

33rd IEEE International Conference on Electronics, Circuits and Systems

FINAL PROGRAM & ABSTRACT BOOK

November 8th-11th, 2026
Thessaloniki, Greece

33rd edition IEEE ICECS



General Chairs Welcome Message

On behalf of the Organizing Committee of the 33rd IEEE International Conference on Electronics, Circuits and Systems (IEEE ICECS 2026), it is our great pleasure to welcome you to Greece and to Thessaloniki.

IEEE ICECS 2026 is the flagship conference of the IEEE Circuits and Systems (CAS) Society in Region 8, covering Europe, Africa and the Middle East. This year, it is being held in Greece, in the historic and vibrant city of Thessaloniki.

Circuits and Systems are at the heart of today's technological evolution and will play an even more critical role in the era of Artificial Intelligence. From the hardware platforms that enable increasingly powerful AI algorithms to energy-efficient, high-performance and intelligent systems at the edge, advances in circuits and systems are essential for transforming new ideas into real-world technologies. The IEEE ICECS 2026 technical program reflects these rapidly evolving trends, bringing together high-quality research papers and innovative approaches addressing emerging challenges and opportunities in electronics, circuits, systems and AI. This year,467 papers were submitted, while 274 papers will be presented, providing a broad overview of the latest developments in the field and an excellent opportunity for researchers, engineers, students and professionals to exchange ideas and explore new directions.

You are in Greece. Is there any special connection between Electronics, Circuits and Systems, the title of the conference, and Greece?

At first glance, perhaps not. Greece is widely regarded as the cradle of democracy, philosophy and rational scientific thought, and its intellectual heritage has profoundly influenced Western civilization. Interestingly, however, this heritage is also reflected in the very language of our field. The word electronics derives from the ancient Greek ēlektron, meaning amber, while systems comes from the Greek systēma, referring to a set of elements brought together into an organized whole. Thus, two of the three words in our conference title—Electronics and Systems—have their roots in the Greek language.

Furthermore, even the word "circuit" evokes a linguistic connection to ancient Greek through the word cycle, which derives from the Greek "kyklos", meaning "circle."

Thessaloniki is the second-largest city in Greece and one of the most historic and vibrant cities in Southeastern Europe. Founded more than 2,300 years ago, its strategic location on the Thermaic Gulf has made it an important commercial, cultural and intellectual centre throughout the centuries.

The city offers a unique combination of history, culture, academic life, gastronomy and a relaxed Mediterranean atmosphere. Its rich heritage is reflected in monuments from the Roman, Byzantine, Ottoman and modern periods. Fifteen Early Christian and Byzantine monuments and sites in Thessaloniki are included in the UNESCO World Heritage List, while the White Tower and the seafront remain among the city's most recognizable landmarks.

Thessaloniki is also renowned for its vibrant gastronomic culture, combining traditional Greek flavors with influences from the many communities that have shaped its history. Its lively university environment and strong academic and research community further contribute to the city's international character, making it an excellent setting for an international scientific conference.

We hope that IEEE ICECS 2026 will give you the opportunity to discover Thessaloniki, experience its history, culture, gastronomy and Mediterranean atmosphere, and enjoy a pleasant and memorable stay.

IEEE ICECS 2026 will take place over four days, from 8 to 11 November, featuring a rich technical and social program. The main conference program comprises 31 technical sessions, 4 special sessions with oral presentations, and 3 poster sessions.

On Sunday, November 8th, six tutorials will run alongside two high-quality satellite events. The program will also feature a joint event by two key CASS groups: Women in Circuits and Systems (WiCAS) and Young Professionals (YP), including an invited talk, two engaging panel discussions, and a dedicated poster session. In addition, the IEEE CASS Region 8 Chapter Leadership Workshop will take place on Monday. The technical program will be further enriched by six demonstrations presented during the poster sessions.

The keynote program of IEEE ICECS 2026 is particularly strong, featuring three distinguished speakers who have kindly accepted our invitations: Professor Themis Prodromakis (University of Edinburgh, UK), Dr. Abu Sebastian (IBM Research Europe / University of Heidelberg), and Professor Kostas Doris (NXP Semiconductors / Eindhoven University of Technology), who will open the keynote program on Monday, November 9th. We sincerely thank our keynote speakers for their participation and for sharing their valuable expertise and insights with the ICECS 2026 community.

I would also like to express my sincere thanks to Dr. Gaia Tomasello, Editor in Chief of Advanced Electronic MaterialsandAdvanced Physics Research (Wiley)who will deliver an invited talk as part of the Women in Circuits and Systems (WiCAS) event.

The conference will conclude on Wednesday, November 11th, with the IEEE ICECS 2026 Awards and Closing Ceremony.

The social program of the conference includes a Welcome Reception on Sunday evening and a Gala Dinner on Tuesday evening, offering participants the opportunity to experience Thessaloniki's gastronomy and enjoy complementary activities that bring them closer to the city's history, culture, and unique character.In addition, options for excursions to significant monuments in the area have been arranged.

IEEE ICECS 2026 would not be possible without the support of many volunteers. We would like to express our sincere thanks to the IEEE Circuits and Systems Society, the IEEE Greece Section, the CAS/SSC joint Chapter in Greece, the Aristotle University of Thessaloniki and Democritus University of Thrace as well as our sponsors for their valuable support and contribution to the organization of this event.

The success of IEEE ICECS 2026 is also the result of the dedication of more than 600 volunteers, including more than 500 reviewers, together with the Organizing Committee, the Track Chairs for Regular and Special Sessions, and the YP/WiCAS organizers.

We would like to extend our special thanks to the hundreds of reviewers who generously devoted their time and expertise to the review process. Their dedication and valuable contribution have been essential in maintaining the high quality of the technical program and making ICECS 2026 possible. Our gratitude also goes to the members of the Organizing Committee and all the volunteers who have contributed, directly or indirectly, to the many tasks and details required to organize a conference of this scale.

We would especially like to thank the TPC Chairs, Vasilis Pavlidis, Alberto Garcia-Ortiz and Ronald Tetzlaff, for their hard work, dedication and initiative in developing the technical program of this edition.

Thank you for participating in IEEE ICECS 2026. We hope that you will enjoy your stay in Thessaloniki and appreciate the technical and social program we have prepared for you.

We also hope that this edition of ICECS will be an opportunity to renew old friendships, establish new ones, exchange ideas and create memories that will last well beyond the conference.

Spiros Nikolaidis & Georgios Sirakoulis
IEEE ICECS 2026 General Chairs



Conference Committees

Conference Committees

General Chairs
Spiros Nikolaidis, Aristotle University of Thessaloniki, Greece
Georgios Sirakoulis, Democritus University of Thrace, Greece

Technical Program Chairs
Vasilis Pavlidis, Aristotle University of Thessaloniki, Greece
Alberto Garcia-Ortiz, University of Bremen, Germany
Ronald Tetzlaff, TU Dresden, Germany

Special Sessions Chairs
Alon Ascoli, Politecnico di Torino, Italy
Joshua Yang, University of Southern California, USA

Tutorials Chairs
Edoardo Bonizzoni, Università di Pavia, Italy
Costas Psychalinos, University of Patras, Greece
Christos Papavasiliou, Imperial College London, UK

Keynote Sessions Chairs
Dimitrios Soudris, NTUA, Greece
Sonal Sherya, Aarhus University, Denmark
Andrei Vladimirescu, University of California, Berkeley, USA

Publicity Chairs
Teresa Serrano Gotarredona, IMSE-CNM, Spain
Hai (Helen) Li, IEEE Fellow, Duke University, USA
Valeri Mladenov, TU Sofia, Bulgaria
Makoto Nagata, Kobe University, Japan
Paul Sotiriadis, IEEE Fellow, NTUA, Greece
Esteban Tlelo-Cuautle, INAOE, Mexico
Yuanqing Cheng, Beihang University, China

Industrial Committee Chairs
Victor Grimblatt, Synopsys
Calliope–Louisa Sotiropoulou, CAST, USA
Emmanuil Zervakis, HETIA, Greece

Publication Chairs
Jose de la Rosa, IMSE-CNM, Spain
Rodrigo Picos-Gaya, Universitat de les Illes Balears, Mallorca, Spain
Christoforos Theodorou, CRNS, France



WiCAS and DEICAS Chairs
Errica Covi, University of Groningen, The Netherlands
Yanqiu Huang, University of Twente, The Netherlands
Maria Papadopoulou, International University of Greece
Ljiljana Trajkovic, Simon Fraser University, Canada

Demo and Student Contest Chairs
Georgios Dimitrakopoulos, Democritus University of Thrace, Greece
Thomas Noulis, Aristotle University of Thessaloniki, Greece
Kostas Siozios, Aristotle University of Thessaloniki, Greece
Ioannis Vourkas, Democritus University of Thrace, Greece

Young Professional Chairs
Vasiliki Gogolou, Aristotle University of Thessaloniki, Greece
Ardalan Najafi, NXP, The Netherlands
Vasilis Ntinas, Aalborg University, Denmark

Award Session Chairs
Sandro Carrara, IEEE Fellow, EPFL, Switzerland
Alkis Hatzopoulos, AUTH, Greece

International Liaison
Jean-Baptiste Bergueret, France
Yehea Ismail, The American University in Cairo, Egypt
Ricardo Reis, UFRGS, Brazil
Thanos Stouraitis, IEEE Fellow, Khalifa University, UAE
Mohamed Tabaa, EMSI Morocco
International Advisory
François RIVET, Université de Bordeaux, France

Finance Chairs
Niki Bai, NBES EVENTS, Greece
Thomas Noulis, Aristotle University of Thessaloniki, Greece

Local Organizing Committee
Ariadne Andreadou, Aristotle University of Thessaloniki, Greece
Vasilios Konstantakos, Aristotle University of Thessaloniki, Greece
Georgios Kousiopoulos, Aristotle University of Thessaloniki, Greece
Vassilis Alimisis, NTUA, Greece



Conference Information

- IEEE ICECS 2026 is a four-day event. It will take place at The Met Hotel, Thessaloniki, Greece.
- Registration will start at 09:00 on Sunday, November8th, and during the conference working hours.
- Six tutorials will take place on Sunday, November 8th, starting at 10:00.
- Two Satellite events will take place on Sunday, 8th November from 13:30 to 17:00.
- The opening ceremony will take place on Monday, November 9th, from 09:00 to 09:30.
- The IEEE CASS Region 8 Chapter Leadership Workshopwill take place on Monday, 9th, November from 14:00 to 17:00.
- The WiCAS panel session is scheduled for Monday, November 9th, starting at 14:00.
- The Young Professionals panel session is scheduled for Tuesday, November 10th, starting at 16:00.
- Coffee breaks will be served at the lobby.
- Lunches will take place at The Met Hotel Restaurant.
- The Welcome reception event will take place at The Met Hotel on Sunday, November 8th, at 18:30-19:30.
- The Gala dinner will take place on Tuesday evening at the Polis convention center. Prior to this, there will be a short tour of the city. Transportation will be provided by bus, departing will start at 18:15-18:30.
- Presentations, tutorials and panels will be held in the following rooms: Maistros ACD, Maistros BEF, Zephyros, Sirocco and Ostria. The specific room for each event will be indicated in the conference program brochure.
- Poster presentations and demonstrations will take place in the hotel lobby area.
- All the attendees are kindly requested to wear their name badge during the conference, which will be provided to them upon registration.
- The official language of the conference is **English.**

Keynote Talks



Prof. Kostas Doris

IFellow NXP Semiconductors
Technical University Eindhoven

Keynote 1: From Automotive Radar to Robotic Perception: Scaling mm-Wave Sensing for responsible robots

Monday November 9th: 09:30 AM -10:30 AM

Room: Maistros ACD
Chair: Calliope-Louiza Sotiropoulou, CAST, USA

Kostas Doris is a Fellow at NXP Semiconductors and part-time Professor at the Technical University of Eindhoven. He pioneered the automotive industry's first mm-wave RFCMOS radar and currently leads NXP's innovations in mm-wave transceivers, silicon processes, and packaging. Kostas has authored numerous publications and patents, and actively collaborates with European universities. He has held leadership roles in ISSCC, RFIC, and EUMW/EURAD.

Abstract:
The rapid evolution of automotive radar sensors operating in the 76-81 GHz band is transforming vehicle perception enabling increasingly advanced autonomous driving functions. Beyond automotive applications, these technologies are emerging as key enablers for autonomous robots, drones, and intelligent machines that require robust environmental awareness under challenging conditions.
Future sensing systems must support high-resolution mapping, localization, and object classification while operating reliably across diverse environments. Meeting these demands requires mm-wave radar sensors capable of generating dense point clouds with high spatial resolution, a wide field of view, and extended range, while remaining compact and energy efficient.
This talk will explore advances in waveform design, antenna technologies, and circuit architectures that continue to push the limits of large-scale MIMO radar in the 76-81 GHz band. It will also discuss opportunities beyond 100 GHz and examine how these developments can complement optical sensing modalities. The convergence of automotive and robotic sensing will be presented as a key driver of the next generation of intelligent perception systems.



2026 IEEE ICECS
November 8th-11th 2026, Thessaloniki, Greece
33rd IEEE International Conference on Electronics, Circuits and Systems

Dr. Abu Sebastian

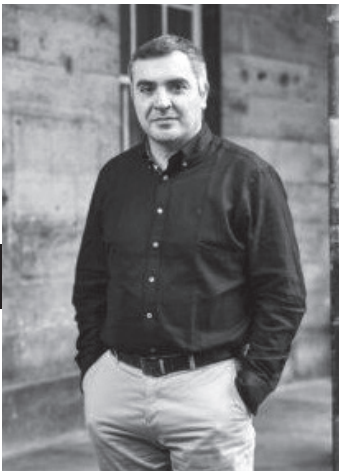
IBM Research Europe, Zurich Research laboratory,
University of Heidelberg

Keynote 2: Neural Principles as a Progressive Path to Efficient Deep Learning Inference Hardware

Tuesday November 10th: 10:30 AM - 11:30 AM
Room: Maistros ACD
Chair: Vasilis Pavlidis, Aristotle University of Thessaloniki, Greece

Abu Sebastian is a Distinguished Scientist at IBM Research – Zurich, where he leads the AI Compute Frontiers group. He is also affiliated to the Kirchhoff-Institut für Physik at Universität Heidelberg. For his contributions to micro- and nanoscale mechatronic systems, he was honored with the IEEE Control Systems Technology Award in 2009 and the IFAC Mechatronic Systems Young Researcher Award in 2013. His research in exploratory memory technologies and emerging computing paradigms has also earned him three European Research Council (ERC) grants: the Consolidator Grant (2015), the Proof-of-Concept Grant (2020), and the Advanced Grant (2025). At IBM, he has been honoured as a Master Inventor, and in 2019 he received the Ovshinsky Lectureship Award for his contributions to phase-change materials for cognitive computing. He is a Fellow of the IEEE.

Abstract:
Deep neural networks, commonly referred to as deep learning, have driven major advances in artificial intelligence over the past decade and are conceptually rooted in principles from computational neuroscience. In contrast, their hardware realizations have largely departed from neuro-inspired design paradigms, giving rise to a pronounced “hardware divide” that manifests as substantial energy inefficiency. Addressing this gap has become a central objective of contemporary classical computing research.
Recent years have seen the gradual incorporation of brain-inspired concepts into AI computing systems. For instance, modern application-specific digital accelerators for deep learning leverage massive parallelism and reduced-precision arithmetic, while in-memory computing (IMC) architectures adopt the principle of co-locating computation and memory to enable in-place processing. In this talk, I will review the current state of the art in IMC-based accelerators and outline my perspective on the additional neuro-inspired principles that must be integrated to further narrow the hardware divide.



Prof. Themis Prodromakis

Regius Chair of Engineering
Centre for Electronics Frontiers,
Institute for Micro Nano Systems,
University of Edinburgh, Edinburgh

Keynote 3: "AI-on-Chip" and "AI-for-Chips": Driving Sustainable Electronics through Bio-inspired Hardware and AI-Enabled Design

Wednesday November 11th: 09:00 AM - 10:00 AM
Room: Maistros ACD
Chair: Costas Psychalinos, University of Patras, Greece

Themis Prodromakis holds the Regius Chair of Engineering at the University of Edinburgh where he established the Centre for Electronics Frontiers (CEF) and is currently the Head of the Institute for Integrated Micro and Nano Systems (IMNS). His work focuses on developing metal-oxide Resistive Random-Access Memory technologies and related applications and leads an interdisciplinary team comprising 60 researchers with expertise across materials process development to electron devices and circuits and systems for embedded applications. He holds an RAEng Chair in Emerging Technologies and is Adjunct Professor at UTS Australia and Honorary Fellow at Imperial College London. He is Fellow of the Royal Society of Chemistry, the British Computer Society, the IET and the Institute of Physics. He served as the Director of the Lloyds Register Foundation International Consortium for Nanotechnology and Co-Director of the UKRI Centre for Doctoral Training in Machine Intelligence for Nano- Electronic Devices and Systems (MINDS). In 2015, he established ArC Instruments Ltd that delivers high-performance testing infrastructure for automating characterisation of novel nanodevices in over 26 countries and in 2025 he founded EVA that is building new power-efficient AI hardware solutions. His contributions in memristive technologies and applications have brought this emerging technology one step closer to the electronics industry for which he was recognised as a 2021 Blavatnik Award UK Honoree in Physical Sciences and Engineering and with the Royal Academy of Engineering Princess Royal Silver Medal in 2025.

Abstract:
The escalating energy demands of modern AI necessitate a fundamental shift from traditional computing paradigms toward hardware that is both energy-efficient and intelligently designed. This plenary talk explores the synergy between AI-on-chip and AI-for-chips as a dual-track solution for a sustainable technological future. We first examine the integration of emerging semiconductor technologies with standard

CMOS to create bio-inspired architectures, enabling high-performance in-memory computing at a fraction of the power cost of von Neumann systems. In parallel, we address the bottlenecks in semiconductor R&D by showcasing AI-driven capabilities to automate and accelerate the design, optimization, and testing of modern electronics. Together, these advancements provide a holistic roadmap for delivering next-generation electronics that are not only inspired by biological efficiency but are also developed through the very intelligence they aim to host.





2026 IEEE ICECS

November 8th-11th 2026, Thessaloniki, Greece

33rd IEEE International Conference on Electronics, Circuits and Systems

Program

SUNDAY NOVEMBER 8TH 2026

<table>
<tr><th>TIME</th><th>ZEPHYROS</th><th>SIROCCO</th><th>OSTRIA AB</th><th>OSTRIA CD</th></tr>
<tr><td>09:00</td><td colspan="4">REGISTRATION</td></tr>
<tr><td>10:30–13:00</td><td colspan="2"></td><td>Tutorial 1</td><td>Tutorial 2</td></tr>
<tr><td>13:00–13:15</td><td colspan="4">Coffee Break</td></tr>
<tr><td>13:15–13:30</td><td colspan="2"></td><td rowspan="3">Tutorial 3</td><td rowspan="3">Tutorial 4</td></tr>
<tr><td>13:30–15:00</td><td>CASS CONNECT Europe</td><td>TL-RNG 2026</td></tr>
<tr><td>15:00–15:15</td><td colspan="2" rowspan="2">Coffee Break</td></tr>
<tr><td>15:15–15:30</td><td colspan="2">Coffee Break</td></tr>
<tr><td>15:30–17:00</td><td>CASS CONNECT Europe</td><td>TL-RNG 2026</td><td rowspan="2">Tutorial 5</td><td rowspan="2">Tutorial 6</td></tr>
<tr><td>17:00–17:30</td><td colspan="2"></td></tr>
<tr><td>18:30–19:30</td><td colspan="4">Welcome Reception (The MET Hotel)</td></tr>
</table>

MONDAY NOVEMBER 9TH 2026

	FOYER				
08:15	REGISTRATION				
	MAISTROS ACD				
09:00–09:30	Welcome Notes				
TIME	**MAISTROS ACD**	**MAISTROS BEF**	**ZEPHYROS**	**SIROCCO**	**OSTRIA**
09:30–10:30	KEYNOTE TALK 1 **Prof. Kostas Doris** **From Automotive Radar to Robotic Perception: Scaling mm-Wave Sensing for responsible robots** Chair: Calliope-Louisa Sotiropoulou				
10:30–11:00	Coffee Break (Foyer)				
11:00–12:30	Chair: Engin Afacan, Thomas Noulis **S01** **Analog/mixed signal/RF circuits 1**	Chair: Nestoras Evmorfopoulos, Cristina Meinhardt **S12** **EDA, Test and Reliability 1**	Chair: Panagiotis Mousouliotis, Athanasios Tziouvaras **S15** **Circuits and Systems for AI Algorithms 1**	Chair: Stefano Gregori, Ioannis Messaris **S22** **Microsystems and Embedded systems**	Chair: Boris Vaisband, Aniket Roy **S28** **Emerging technologies and technology trends 1**
12:30–14:00	Lunch Break				
14:00–15:30	**CASS Chapter Meeting**	**WiCAS**			
15:30–17:00		Coffee Break & Poster Session 1 (Foyer)			
17:00–18:30	Chair: Vasilis Konstantakos, Marius Enachescu **S02** **Analog/mixed signal/RF circuits 2**	Chair: George Lentaris, Apostolos Fournaris **S08** **Digital Circuits and Systems 1**	Chair: Valentino Peluso, Milad Kokhazadeh **S25** **Neural networks, Machine and Deep Learning 1**	Chair: George Athanasiou, Boris Vaisband **S32** **Special Session: Advances in Hardware Security for Electronic Circuits and System**	Chair: Christoforos Theodorou, Hossein Esmailbeygi **S33** **Special Session: Sensor Devices, Circuits, and Systems**

TUESDAY NOVEMBER 10TH 2026

TIME	MAISTROS ACD	MAISTROS BEF	ZEPHYROS	SIROCCO	OSTRIA
08:30	REGISTRATION				
08:30–10:00	Chair: Andreas Demosthenous, Cristina Meinhardt **S06** **Biomedical and Bio-inspired Circuits and Systems 1**	Chair: Valeri Mladenov, Tales Pimenta **S10** **Linear and Non-linear Circuits**	Chair: Milad Kokhazadeh, Athanasios Tziouvaras **S24** **Neural networks, Machine and Deep Learning 2**	Chair: Calliope-Louisa Sotiropoulou, Nikolaos A. Anagnostopoulos **S26** **Signal Processing, Image and Video**	Chair: Vasileios Ntinas, Aniket Roy **S29** **Emerging technologies and technology trends 2**
10:00–10:30	Coffee Break (Foyer)				
10:30–11:30	**KEYNOTE TALK 2** **Room: Maistros ACD** **Dr. Abu Sebastian** **Neural Principles as a Progressive Path to Efficient Deep Learning Inference Hardware** Chair: Vasilis Pavlidis				
11:30–13:00	Chair: Meysam Akbari, Engin Afacan **S03** **Analog/mixed signal/RF circuits 3**	Chair: Valentino Peluso, Manolis Katsaragakis **S09** **Digital Circuits and Systems 2**	Chair: Dimitrios Prousalis, José M. de la Rosa **S18** **AI-driven Circuit Design 1**	Chair: Vasileios Alimisis, Spyridon Vlassis **S20** **Energy Harvesting and Power Management 1**	Chair: Lazaros Papadopoulos, Erica Raviola **S34** **Special Session: Cutting-Edge Research Projects Funded by European and International Programs 1**
13:00–14:00	Lunch Break				
14:00–15:30	Chair: Vasiliki Gogolou, Yasmin Halawani **S04** **Analog/mixed signal/RF circuits 4**	Chair: Lorenzo Benatti, Sahibia Kaur Vohra **S11** **Low-Power/Voltage Design**	Chair: George Floros, Alberto Garcia-Ortiz **S13** **EDA, Test and Reliability 2**	Chair: Panagiotis Mousouliotis, Minas Dasygenis **S16** **Circuits and Systems for AI Algorithms 2**	Chair: Lazaros Papadopoulos, Apostolos Fournaris **S35** **Special Session: Cutting-Edge Research Projects Funded by European and International Programs 2**
15:30–16:00	Coffee Break (Foyer)				
16:00–17:30	**Young Professionals (ACD) & Posters Session 2 (Foyer)**				
18:30–23:00	Mini City Tour & Conference Dinner				

WEDNESDAY NOVEMBER 11TH 2026

TIME	MAISTROS ACD	MAISTROS BEF	ZEPHYROS	SIROCCO	OSTRIA
08:30	REGISTRATION				
09:00–10:00	KEYNOTE TALK 3 Room: Maistros ACD **Prof. Themis Prodromakis** **"AI-on-Chip" and "AI-for-Chips": Driving Sustainable Electronics through Bio-inspired Hardware and AI-Enabled Design** Chair: Costas Psychalinos				
10:00–10:30	Coffee Break (Foyer)				
10:30–12:00	Chair: George Lentaris, Minas Dysagenis **S14** **EDA, Test and Reliability 3**	Chair: George Floros, Manolis Katsaragakis **S19** **AI-driven Circuit Design 2**	Chair: Dimitrios Prousalis, Thomas Noulis **S23** **Sensors and Sensing Systems**	Chair: Nikos Sketopoulos, Nikolaos A. Anagnostopoulos **S27** **VLSI Systems and Applications**	Chair: George Panagopoulos, Mattis Hasler **S31** **Digital Circuits and Systems 3**
12:00–13:15	Lunch Break				
13:15–14:15	Poster Session 3 (Foyer)				
14:15–15:45	Chair: Costas Psychalinos, Vladimir Milovanović **S05** **Analog/mixed signal/RF circuits 5**	Chair: Vasileios Alimisis, Meysam Akbari **S07** **Biomedical and Bio-inspired Circuits and Systems 2**	Chair: Nikos Sketopoulos, Yiorgos Tsiatouhas **S17** **Circuits and Systems for AI Algorithms 3**	Chair: Vasiliki Gogolou, Spyridon Vlassis **S21** **Energy Harvesting and Power Management 2**	Chair: Ioannis Messaris, Iraklis Anagnostopoulos **S30** **Emerging technologies and technology trends 3**
15:45–16:00	Coffee Break (Foyer)				
	MAISTROS ACD				
16:00–16:30	Closing Ceremony				



2026 IEEE ICECS

November 8th-11th 2026, Thessaloniki, Greece

33rd IEEE International Conference on Electronics, Circuits and Systems

Program Booklet



Sunday November 8th 2026

Satellite Events **1:30 PM - 5:00 PM**

CASS-CONNECT Europe

Room: Zephyros

Moderator: Prof. Andreas Demosthenous, UCL, UK

Connecting Academia and Industry in Europe CASS Workshop (CASS CONNECT Europe)

CASS-CONNECT Europe workshop aims to strengthen academia-industry collaboration across the IEEE CASS community. Building on the strong industry engagement demonstrated in recent CASS activities, the workshop will bring together academic and industry researchers to explore forward-looking technologies in circuits and systems. Key themes include healthcare technologies and diagnostics, processor design, circuits and systems for agriculture and advanced sensing/signal processing. The workshop will feature talks from both industry and academia, complemented by interactive discussions and audience participation.

TL-RNG 2026

Room: Sirocco

Moderator: Prof. Hossam O. Ahmed, AUM, Kuwait

TL-RNG 2026 – Trust Level Framework for TRNG & QRNG Systems: Toward the First IEEE Trust Certification Framework for Hardware Random Number Generators

TL-RNG 2026 – Trust Level Framework for TRNG and QRNG Systems is an IEEE CASS Standards Activities Workshop dedicated to advancing the trust, validation, and evaluation of hardware random number generators. The workshop focuses on True Random Number Generators (TRNGs) and Quantum Random Number Generators (QRNGs), which are essential for secure cryptographic systems and emerging security applications. It will bring together international experts from academia and industry to discuss current challenges, evaluation methodologies, and reliability requirements. A key objective is to explore a unified framework for classifying the trust levels of hardware random number generators. Discussions will address entropy assessment, statistical testing, validation procedures, and reproducibility of results. The workshop aims to identify gaps in existing evaluation practices and support the development of a future IEEE standard. Through expert presentations and panel discussions, TL-RNG 2026 seeks to strengthen collaboration between researchers, standardization bodies, and industry practitioners. Ultimately, it aims to establish a foundation for more transparent, reliable, and trustworthy hardware random number generation.

Monday November 9th 2026

S01 | Analog/mixed signal/RF circuits 1 **11:00 AM - 12:30 PM**
Room: Maistros ACD
Chairs: Engin Afacan, Thomas Noulis

1. **Common-mode Tank Impedance Engineering for the Suppression of the 1/f^3 Phase Noise Sideband in Class-C CMOS Oscillators** (47)
   Marco Longo, Marc Tiebout and Andrea Bevilacqua
2. **A Robust 17-21 GHz Frequency Synthesizer in SiGe-BiCMOS for Space** Applications (65)
   Nebojsa Maletic, Frank Herzel, Giovanni Di Pietrantonio and Gunter Fischer
3. **A Low Power Wideband Current-Mode Receiver Frontend for FR3 Band Applications** (101)
   Shiuan-Han Huang, Chien-Nan Kuo and Hwann-Kaeo Chiou
4. **A High-Gain D-Band LNA with Inductively-Tuned Gain-Boosting Technique in 28nm Bulk CMOS** (174)
   Rainier van Dommele, Yiqin Hou, Remco Schalk, Kaijie Ding and Vojkan Vidojkovic
5. **A Highly Linear Autotransformer-Matched S2D LNTA for Sub-6-GHz Receivers in 22-nm FDSOI** (181)
   Meysam Shayestehfar, Harald Pretl and Danilo Manstretta

S12 | EDA, Test and Reliability 1 **11:00 AM - 12:30 PM**
Room: Maistros BEF
Chairs: Nestoras Evmorfopoulos, Cristina Meinhardt

1. **SET: Substrate Extraction Tool for Predicting Substrate Noise in SOI Smart Power ICs** (186)
   Erica Raviola, Gabriele Salzone, Luca Bertolini, Sergio Lecce, Kamel Abouda and Franco Fiori
2. **3MLR: An Investigation of Moment-Matched Mutual Ladder-based Reduction for VLSI Static Timing Analysis** (232)
   Aikaterini Tsilingiri, Dimitrios Tsalapatas, Stavros Simoglou, Stavros Stathoudakis and Christos Sotiriou
3. **A Residual Degree-Guided Heuristic for Delay-Arc Traversal in Combinational Cells Characterization** (268)
   Rita Louro Barbosa and Cristina Meinhardt
4. **Enhancing MBFF Clustering with Continuous Timing Criticality Awareness** (307)
   Lucas Tinoco and José Luis Güntzel
5. **End-to-End GNN-ML Framework for Automated SPICE-to-SystemVerilog Behavioral Modeling** (333)
   Karim Osama, Amr Wael, Kareem Atef, Kareem Abdelraof, Hana Maged, Mariam Ahmed, Mostafa Waseem, Ramy Raafat and Mohamed Abd El Ghany

S15 | Circuits and Systems for AI Algorithms 1 **11:00 AM - 12:30 PM**
Room: Zephyros
Chairs: Panagiotis Mousouliotis, Athanasios Tziouvaras

1. **FPGA and Parallel SNN Inference for RIS-Assisted Transmitter Selection** (134)
   Evangelos Kyparoglou, Minas Dasygenis and Alexandros-Apostolos Boulogeorgos
2. **Efficient Temporal Analog In-Memory Computing via Co-Training of Long- and Short-Term Memory in RRAM Devices** (45)
   Ming-Jay Yang, Dimitrios Spithouris, Johannes Hellwig, Regina Dittmann, Gordon Pipa and John Paul Strachan
3. **Performance Improvement over Recurrent Spiking Neural Networks HW Accelerator for FPGA** (253)
   Diego Avila-Torres, Antonio Perez-Pena, Luis G. Leon-Vega, Jorge Castro-Godínez and Alejandro Linares-Barranco
4. **Co-Optimizing Graph Sparsification and Approximate Computing for Energy-Efficient FPGA-Based GCN Inference** (290)
   Nathaniel Kaye Mellor, Shreejith Shanker and George Floros
5. **SRAM-Based Processing-in-Memory Cell Enabling Binary Neural Networks Through XNOR Multiplication and Charge-Domain Accumulation** (330)
   Cristhopher Rafael Mosquera Montero, Ramiro Taco, José Benjamin Zambrano Chavez, Esteban José Garzón Córdova, Alessandro Catania, Sebastiano Strangio, Giuseppe Iannaccone and Marco Lanuzza

S22 | Microsystems and Embedded systems **11:00 AM - 12:30 PM**
Room: Sirocco
Chairs: Stefano Gregori, Ioannis Messaris

1. **PrefA: Prefetcher as Accelerator for Time Dependent Irregular Workloads on RISC-V** ISA (285)
   Savvina Nastou, Georgios Alexandris, Manolis Katsaragakis, Dimitrios Soudris and Sotirios Xydis
2. **Memristive Slack Time and Event Predictors for Resource-Constrained Embedded Systems** (362)
   Dongchen Li, Zidu Li, Bhaskar Choubey and Roman Obermaisser
3. **A Digital Twin Model and Architecture for Monitoring and Controlling Fleets of Autonomous Unmanned Surface Vehicles** (394)
   Anderson Roberto Pinheiro Domingues, Cássio Jones Dhein da Silva, Francesco da Rocha Lui, Juliano Maia, Rodrigo S. Cenço, César Augusto Mission Marcon and Fernando Gehm Moraes
4. **Exploration Methodology for Resource-aware Multi-model Deployment on Edge Devices** (231)
   Anagnostis Boubardas and Lazaros Papadopoulos
5. **A Hybrid Differentiable Digital Twin for Induction Motors under Deep Magnetic Saturation** (60)
   Xingwei Zhou, Wenshan Hu, Zhongcheng Lei and Jianbin He

S28 | Emerging technologies and technology trends 1 11:00 AM - 12:30 PM
Room: Ostria
Chairs: Boris Vaisband, Aniket Roy

1. **RF GaN technology with enhanced multi-layer Back-End-of-Line (BEOL) process** (119)
   Errikos Lourandakis, Nikolaos Makris, Evangelos Aslanidis, Antonios Stavrinidis, Athanasios Kostopoulos, Eleftherios Sarris, Paolo Fioravanti, George Konstantinidis, Alexandros Gurulian and Giannis Kontogiannopoulos
2. **A 7-Bit Current-Steering DAC With an NRZ Dynamic-Switch Driver in IGZO TFT Process** (131)
   Muhammad Dawood Asghar, Sven Baerten, Mohit Dandekar and Kris Myny
3. **Design of on-chip ESD protection network using a-IGZO TFTs** (163)
   Khang Nguyen, Florian De Roose, Paoline Coulson, Xiaonan Xing, Mauricio Velazquez Lopez, Sebastian Haesler, Marko Simicic, Jan Genoe and Kris Myny
4. **Improvement of IGZO MeSFET Performance by Post-Fabrication Anneal: Electrical Assessment and TCAD Modeling** (323)
   Ivan Lisovyi
5. **Electro-Thermal Analysis and Modeling of Resistor-Transistor Logic in a 200-V E-mode GaN-on-SOI Technology** (365)
   Salvatore Pennisi, Katia Samperi, Andrea Ballo, Urmimala Chatterjee and Stefaan Decoutere

WiCAS Panel Session **2:00 PM - 3:30 PM**
Room: Maistros BEF
Moderator: Prof. Maria Papadopoulou, International Hellenic University, Greece

Invited Talk **2:00 PM - 2:30 PM**

Navigating the GAI Revolution in Scholarly Publishing
Speaker: *Dr. Gaia Tomasello, Editor-in-Chief of Advanced Electronic Materials*

Roundtable Discussion **2:30 PM - 3:30 PM**

Panelists:

- Dr. Chrisasvgi Dre, Deputy Director, Intracom Defense
- Dr. Nathalie Deltimple, Professor, University of Bordeaux
- Dr. Calliope-Louisa Sotiropoulou, Manager, CAST
- Eleni Papageorgiou, PhD Student, Aristotle University of Thessaloniki

This session celebrates the contributions and impact of women in engineering across academia, industry, and entrepreneurship. It will open with an invited talk by Dr. Gaia Tomasello, followed by a roundtable discussion bringing together diverse perspectives on leadership, career development, and the challenges and opportunities facing women in Electronics, Circuits, and Systems.
The discussion will explore personal leadership journeys, strategies for professional success, and the evolving role of women in shaping technological innovation. Audience participation will be encouraged, fostering an open exchange of ideas and experiences.
Open to all conference participants, the session will be particularly valuable for young professionals, women in engineering, aspiring leaders, and anyone interested in innovation, inclusion, and leadership in Electronics, Circuits, and Systems.

IEEE CASS Region 8 Chapter Leadership Workshop **2:00 PM - 5:00 PM**
Room: Maistros ACD

Organizers:

- **Prof. Costas Psychalinos**, IEEE Greece CASS Chapter Chair, University of Patras, Greece.
- **Prof. Maurizio Valle**, R8 Board of Governors (Regional Member at Large), North Italy CASS Chapter Chair, University of Genova, Italy.
- **Prof. Yongfu Li**, IEEE CASS Vice President RAM, Shanghai Jiao Tong University, China.
- **Francois Rivet**, Board of Governors (Member at Large), University of Bordeaux, France.

Time	Duration	Talk	Speaker
14:00 - 14:10	10 min	Opening Remarks & Workshop Objectives	Workshop Organizers
14:10 - 14:30	20 min	Launching CAS, RAS and WiE Student Branch Chapters and Supporting Participation in IEEE CASS Design Competitions	Prof. Erivelton Nepomuceno, Nonlinear Circuits and Systems Technical Committee (NCAS TC), Maynooth University, Ireland
14:30 - 14:50	20 min	Strategic Initiatives, Growth Trajectories, and Operational Outcomes of the IEEE Egypt Section Circuits and Systems Society (CASS) Chapter	Prof. Ahmed H. Madian, IEEE Egypt Section Chair, NILE University, Egypt
14:50 - 15:10	20 min	From Student Engagement to Sustainable CASS	Prof. Mustak Erhan Yalcin, IEEE Türkiye CASS Chapter Chair, Istanbul

		Membership: Lessons from the IEEE CASS Türkiye Chapter	Technical University, Türkiye
15:10 - 15:30	20 min	Connecting Students with Industry through the IEEE CASS Germany Chapter	Prof. Zidu Li, IEEE Germany CASS Chapter Secretary, Universität Siegen, Germany
15:30 – 16.00	30 min	Break	
16:00 - 16:20	20 min	From Circuits to Communities: Empowering the Next Generation through IEEE CASS STEM Outreach	Prof. Abdallah Kassem, IEEE Lebanon CASS Chapter Chair, Notre Dame University-Louisa (NDU), Lebanon
16:20 - 16:40	20 min	Restarting the IEEE Greece CASS Chapter: Good practices and lessons learned	Prof. Costas Psychalinos, IEEE Greece CASS Chapter Chair, University of Patras, Greece.
16:40 - 17:00	20 min	Concluding and closing remarks	Workshop Organizers, Speakers, and IEEE Officers

S02 | Analog/mixed signal/RF circuits 2 **5:00 PM - 6:30 PM**
Room: Maistros ACD
Chairs: Vasilis Konstantakos, Marius Enachescu

1. **A FMIR Radar receiver with 1-ns Step Dynamic Pattern Variable Gain Amplifier** (140)
   Arthur Paviot, Pierre Courouve, Mykhailo Zarudniev, Laurent Ouvry and Gilles Masson
2. **A -57 dBm sensitivity 26 GHz Active Envelope Detector for Wake-Up Receiver Applications** (146)
   Albert Kibiwott, Chadi Jabbour, Reda Mohellebi, Kyriaki Niotaki and Zuo Liu
3. **Broadband Darlington-Cascode Power Amplifier: Analysis, Stability, and Measurement Results** (246)
   Paul François, Thierry Lagutere and Philippe Descamps
4. **GaN-on-SOI Oscillators for Smart Power ICs** (73)
   Paul-Catalin Medinceanu, Florin-Silviu Dumitru and Marius Enachescu
5. **Efficiency-Centric Automated Design Optimization of a Digitally Configurable IR-UWB Transmitter** (238)
   Yağız Kırtaş, Okan Zafer Batur and Burcu Erkmen

S08 | Digital Circuits and Systems 1 **5:00 PM - 6:30 PM**
Room: Maistros BEF
Chairs: George Lentaris, Apostolos Fournaris

1. **A BRAM-Free and DSP-Free On-the-Fly Twiddle Factor Generator for PQC NTT Accelerators** (227)
   Shiva Nejati, Moslem Heidarpur, Lang Lin, Norman Chang and Mitra Mirhassani
2. **MPX: A Unified Systolic Array for Matrix and Polynomial Multiplication** (263)
   George Alexakis, Dimitrios Schoinianakis and Giorgos Dimitrakopoulos
3. **Efficient Dynamic Logic Implementation of Magnitude Comparators** (351)
   Constantinos Efstathiou, John Liaperdos and Yiorgos Tsiatouhas
4. **HPG-SVD High-Performance Generalized Singular Value Decomposition Processor with Adaptive Termination** (401)
   Anjan Banerjee, Souris Sahu, Rohan Purkait, Richard Prajwal, Sasya Merugu Sri, A. Prasanna Laxmi, Sounak Samanta and Acharyya Amit
5. **A Division-Free Fixed-Point Matrix Inversion Architecture Using Normalization-Assisted BI-LUT Reciprocal Approximation** (433)
   Maloth Vishnu Vardhan, E Richard Prajwal Das, Anjan Banerjee, J Sai Maneesh Kumar, Nidhi Verma, P V Phaneendra Kumar and Acharyya Amit

S25 | Neural networks, Machine and Deep Learning 1 **5:00 PM - 6:30 PM**
Room: Zephyros
Chairs: Valentino Peluso, Milad Kokhazadeh

1. **HAMLET: Machine Learning-Driven Placement over Heterogeneous DRAM/NVM Memory Systems** (291)
   Manolis Katsaragakis, Giannis Rokomos and Dimitrios Soudris
2. **Arapaima-LIF: An Oxygen-Regulated Spiking Neuron for Temporal Memory Retention** (308)
   Anupama Santhosh, Sumit Datta, Xinming Shi and Alex James
3. **Continual Learning by Exploiting Adaptive Anti-Class Distributions** (367)
   Christos Tsingiropoulos, Dimitrios Katsikas, Nikolaos Passalis and Anastasios Tefas
4. **EBL: Efficient Broad Learning for Distributed Adaptive Harmonic Analysis** (386)
   Changhong Li, Georgios Floros, Biswajit Basu and Shreejith Shanker
5. **Exploration of a Multilayer Perceptron for Electrical Characterization of Logic Gates** (400)
   Eduarda Moraes dos Santos and Cristina Meinhardt

S32 | Special Session: Advances in Hardware Security for Electronic Circuits System
5:00 PM - 6:30 PM
Room: Sirocco
Chairs: George Athanasiou, Boris Vaisband

1. **Reconfigurability Leading to Reliability and Security: A Charge Trap Transistor Based PUF** (304)
   Fahim Shahriar, Nader Sehatbakhsh and Boris Vaisband
2. **Assessing Runtime Electromagnetic Detection of CPU Hardware Trojans Targeting Kernel Memory** (218)
   Athanasios Moschos, Baki Berkay Yilmaz, Kevin Valakuzhy and Angelos Keromytis
3. **Machine-Learning-Resilient Deca-Mbit Entropy Physical Unclonable Functions in 65 nm CMOS** (314)
   Sabrin Alla, Edoardo Cignoni, Alessandro Bedoya, Marco Lanuzza, Sebastiano Strangio, Alessandro Catania and Giuseppe Iannaccone
4. **Side-Channel Evaluation of a Stealthy Cross-Module Hardware Trojan in a Combined ML-KEM and AES FPGA Design** (187)
   Athanasios Papadimitriou, Nikolaos Athanasios Anagnostopoulos, Mihalis Psarakis and Paris Kitsos
5. **Towards low-power cryptographic circuits in flexible IGZO electronics** (210)
   Jelle Biesmans, Wim Dehaene, Kris Myny and Nele Mentens
6. **Low-Cost Enrollment and Low-Overhead Authentication of SRAM PUFs on STM32 Microcontrollers** (244)
   Stefano Gregori, Kamal Y. Kamal and Esther Bamiduro

S33 | Special Session: Sensor Devices, Circuits, and Systems
5:00 PM - 6:30 PM
Room: Ostria
Chairs: Christoforos Theodorou, Hossein Esmailbeygi

1. **A Time-Domain ADC with Constant-Slope Voltage-to-Time Converter and Conditional Correlated Multiple Sampling for CMOS Image Sensors** (96)
   Jonathan Vrijsen, Wannes Nevens, Thomas Lapauw, Maarten Kuijk, Hans Ingelberts and Thomas Van den Dries
2. **Environment Perception and Multi-Target Detection with a mmW 2D Beamsteering ISAC System** (348)
   Nebojsa Maletic and Markus Petri
3. **Accuracy of Characterization Procedures for Quartz Tuning Fork in QEPAS Sensors as a Function of the Front-End Architecture** (58)
   Vincenzina Zecchino, Luigi Lombardi, Cristoforo Marzocca, Gianvito Matarrese, Pietro Patimisco, Angelo Sampaolo and Vincenzo Spagnolo
4. **A 24 GHz FMCW Radar Sensor System with Embedded Real-Time Processing for Cyclist Dooring Prevention** (204)
   Jonas Gillner, Tristan Thierfelder, Roy Taylor, Leila Sharara, Mohammed Ismail and Klaus Thelen
5. **Optimization of a Cymbal Piezoelectric Transducer Using Finite-Element-Guided ANN and DRL (**205)
   Mina Nazarian Samani, Javad Pourmohammad, Sima A. Alidokht, Heloise Therien-Aubin and Lihong Zhang

Tuesday November 10th 2026

S06 | Biomedical and Bio-inspired Circuits and Systems 1
08:30 AM - 10:00 AM
Room: Maistros ACD
Chairs: Andreas Demosthenous, Cristina Meinhardt

1. **Analog Haar Wavelet Filter Design for Phonocardiogram Signal Analysis** (8)
   Julia Nako, Costas Psychalinos and Ahmed Elwakil
2. **A Low Power, Low Noise Inverter-based Amplifier with High Linearity for Ultrasound Imaging** (157)
   Elham Hatamzadeh, Antonio Aprile, Elisabetta Moisello, Farshad Moradi, Edoardo Bonizzoni, Piero Malcovati and Leonardo Miceli Lara
3. **Parkinsonian Activity Detection in Microelectrode Recordings Using a Memristive Oscillator** (203)
   Nefeli Maria Panagiota Kavouni, Ioannis K. Chatzipaschalis and Georgios Ch. Sirakoulis
4. **A Low-Power Analog IC Artificial Neural Network Architecture for Cardiac Survivability Classification** (247)
   Vassilis Alimisis, Christos Dimas, Alak Majumder, George F. Fragulis and Paul P. Sotiriadis
5. **FPGA Implementation of a Medical Image Watermarking Scheme Based on Region of Interest** (104)
   Abdenour Kifouche and Mohamed Salah Azzaz

S10 | Linear and Non-linear Circuits **08:30 AM - 10:00 AM**
Room: Maistros BEF
Chairs: Valeri Mladenov, Tales Pimenta

1. **Circuit Modeling Using Belevitch Canonical Form and Iterative Least-Squares Optimization** (30)
   Gang Lin
2. **A Sinusoidal Signal Generator based on Second-Order Analog Interpolation with Widely Tunable Amplitude, Mean Value, and Frequency** (133)
   Francesco Gagliardi, Margherita Scognamiglio, Simone Contardi, Iacopo Nannipieri, Massimo Piotto and Paolo Bruschi
3. **p-Optimal Noise-Shaping Digital Δ-Σ Modulators** (178)
   Xu Wang and Patrick Hickey
4. **Memristor Models with Logarithmic-Sigmoidal and Piecewise-Linear Self-Limiting State-Flux Relations** (193)
   Valeri Mladenov and Stoyan Kirilov
5. **A Differential gm/C Oscillator for Mobility-Dependent Frequency** Response (413)
   Tobias Schwanke

S24 | Neural networks, Machine and Deep Learning 2 08:30 AM - 10:00 AM
Room: Zephyros
Chairs: Milad Kokhazadeh, Athanasios Tziouvaras

1. **NPU Deployment of DL-based Spectrum Sensing for Edge-AI Cognitive Radio Systems** (235)
   Andres Rojas and José M. de la Rosa
2. **TE-STDF: Gradient-Guided Structural Pruning for Lightweight Video Enhancement** (150)
   Luis Carlos Rodrigues Linares, Gilberto Kreisler, Daniel Palomino, Bruno Zatt, Attilio Fiandrotti and Guilherme Corrêa
3. **Spiking Neural Network on Microcontroller for Low-cost Speech Enhancement** (151)
   Yusuf Serhat Özkan, Tao Sun and Guangzhi Tang
4. **ActiPercept: Confidence-Guided Active Perception for Robotic Screw Inspection** (260)
   Pavlos Tosidis, Nikolaos Passalis and Anastasios Tefas
5. **Keyword Spotting with Out-of-Vocabulary Rejection for Wearable IoT Devices** (429)
   Yuying Zhuo, Michael Grzybek, Petros Spachos and Stefano Gregori

S26 | Signal Processing, Image and Video 08:30 AM - 10:00 AM
Room: Sirocco
Chairs: Calliope-Louisa Sotiropoulou, Nikolaos Athanasios Anagnostopoulos

1. **Range Utilization and SQNR Decomposition for Dithered Quantizers Below 4 Bits** (59)
   Brian Holford
2. **Generalization of the Gamma/Gammatone Filter Transfer Functions** (152)
   Brent Maundy, Ahmed Elwakil, Costas Psychalinos and Anis Allagui
3. **Using a ring buffer instead of a FIFO in a FIR-DAC significantly improves linearity but increases noise** (251)
   Hanspeter Schmid, Alex Huber, Dominik Hiltbrunner and Mathieu Coustans
4. **A Pareto-Optimized Fused Posit Radix-2 FFT Architecture for FPGA and its validation in FMCW Radar Signal Processing** (275)
   G V N Mokshagna, Kandagatla Venkata Rohith Sai and Madhav Rao
5. **SCARF-DoA: Spectral Coupling ICA and Adaptive Refinement Projection Framework for Coherent and Underdetermined DoA Estimation** (399)
   Rohan Purkait, Souris Sahu, Anjan Banerjee, Nidhi Verma, Amit Kumar Gupta, Sounak Samanta and Acharyya Amit

S29 | Emerging technologies and technology trends 2 **08:30 AM - 10:00 AM**
Room: Ostria
Chairs: Vasileios Ntinas, Aniket Roy

1. **Smart-IMPLY Validation via Custom Integrated Circuits and Hardware-Aware System Simulations** (132)
   Lorenzo Benatti, Tommaso Zanotti, Max Uhlmann and Francesco Maria Puglisi
2. **Analog Area-Type Switching in Ultrathin HfO2 Memristors for Neuromorphic Circuits** (154)
   Cuo Wu, Junjie Zhao, Zipeng Mi, Thomas Mikolajick and Benjamin Max
3. **The impact of memristor I-V asymmetry on charge integration in pseudo-memcapacitive neurotransistors** (199)
   Hampus Hoffman, Eter Mgeladze, Thomas Mikolajick, Richard Schroedter, Ronald Tetzlaff and Stefan Slesazeck
4. **SPICE-Compatible Model for Skyrmionic Devices** (229)
   Georgios Chararas, Aristotelis Tsekouras, Nicolas Reyren, Vincent Cros, Christoforos Moutafis and Vasilis Pavlidis
5. **Topology and Device-Dependent Synaptic Behavior in Memristor Bridges** (416)
   Dimitrios Prousalis, Shamoil Ejaz, Ioannis Messaris, Vasileios Ntinas, Ahmet Samil Demirkol and Ronald Tetzlaff

S03 | Analog/mixed signal/RF circuits 3 **11:30 AM - 1:00 PM**
Room: Maistros ACD
Chairs: Meysam Akbari, Engin Afacan

1. **A 10-b 56-MS/s Calibration-Free SAR ADC Array for Touch Sensing Applications** (127)
   Chin-Wei Cheng
2. **A 5 GS/s 44 dB Burst-Mode ADC for In-Situ Monitoring of Wide-Bandgap Power Devices** (128)
   Tom Herzog, Erik Wehr, Kenny Vohl, Ralf Wunderlich and Stefan Heinen
3. **A Preamplifier-Free Sigma-Delta ADC for fA-Level Current Digitization Using Logarithmic Current Splitting** (190)
   Kerem Yurtseven, Uygar Yildiz, Kemal Ozanoglu and Günhan Dündar
4. **Compensating Excess Loop Delay in CT Incremental Delta-Sigma ADCs using Extended Counting** (332)
   Elias Ruopp, John Kauffman and Maurits Ortmanns
5. **A Method for Designing Decimation Filters for Noise-Shaping SAR ADCs** (346)
   Stan Oosterlynck, Marco Bolognesi, Achim Vandierendonck, Caro Meysmans, Hendrik Meersseman, Brecht Vanlaar and Guy Torfs

S09 | Digital Circuits and Systems 2 **11:30 AM - 1:00 PM**
Room: Maistros BEF
Chairs: Valentino Peluso, Manolis Katsaragakis

1. **RISC-V controlled flexible digital front-end for communications** (88)
   Aleksi Korsman, Otto Simola, Andrei Spelman, Verneri Hirvonen, Jussi Ryynänen and Marko Kosunen
2. **Hardware Architecture for MPEG-5 EVC Intra Prediction Supporting Multiple Coding Unit Sizes** (161)
   Marco Jefferson Mamani Arapa and Claudio Machado Diniz
3. **A Resource-Efficient FPGA Architecture for Numerically Stable Streaming Hermitian Covariance Estimation Using Welford's Online Algorithm** (296)
   Anjan Banerjee, Rohan Purkait, Souris Sahu, Nithin Sai Reddy, Aryaan Mandal, A. Prasanna Laxmi, Sounak Samanta and Amit Acharyya
4. **Core-Lock Scheduling for 5G NR LDPC Decoders** (406)
   Nikolaos Papageorgiou and Vassilis Paliouras
5. **Dynamic Frequency Scaling Evaluation for RTOS in RP2040** (342)
   Ioannis Sofianidis, Vasileios Konstantakos and Spyridon Nikolaidis

S18 | AI-driven Circuit Design 1 **11:30 AM - 1:00 PM**
Room: Zephyros
Chairs: Dimitrios Prousalis, José M. de la Rosa

1. **BABEL: A Bi-fidelity Automatic Bayesian Exploration Loop for Variation-Aware Analog Integrated Circuit Sizing** (83)
   Takayoshi Namura and Nobukazu Takai
2. **Target-Driven Deep RL for Reusable Sizing of Constant gm Biasing Circuit Under Corner Variations** (90)
   Eleni Papageorgiou, Valentina Davidoiu, Andi Buzo, Georg Pelz and Thomas Noulis
3. **Sensitivity-Guided Differential Evolution for Simulation-Based Analog Circuit Sizing** (92)
   Eleni Papageorgiou, Savvas Karipidis, Valentina Davidoiu, Andi Buzo, Georg Pelz and Thomas Noulis
4. **Optimization-Driven Active Learning for High-Yield Dataset Generation in ML-Assisted Design of Sigma-Delta Modulators** (223)
   Francisco J. Gómez-Pulido and Jose M. de La Rosa
5. **SeG-MCTS: Sensitivity-Guided Monte Carlo Tree Search for Device Selection in Analog Circuits** (340)
   Yuto Moriguchi and Nobukazu Takai

S20 | Energy Harvesting and Power Management 1 **11:30 AM - 1:00 PM**
Room: Sirocco
Chairs: Vasileios Alimisis, Spyridon Vlassis

1. **The energy-efficient techniques for high voltage generation of NAND Flash storage using a Fully Integrated Voltage Regulator (FIVR)** (78)
   Minsoon Hwang, Mingi Jeong, Hyunwoo Kim, Sanguk Lee, Jaekwan Kwon, Jongwook Kim, Youngjin Park, Sejin Park, Yoosang Na, Hwang Yoon, Seungyoung Ahn and Dongsop Lee
2. **A 38 nA Quiescent Current, 90 nS Fast Transient Response, Capacitorless Low Dropout Regulator Based on a Class-AB Buffer With Current Limiter and Active Capacitor** (284)
   Jiayang Li, Dai Jiang and Andreas Demosthenous
3. **A 14-mV-Input, 92.9%-Peak-Efficiency Single-Inductor Dual-Output DC–DC Converter for Energy Harvesting Applications** (299)
   Elham Hatamzadeh, Hossein Esmailbeygi and Farshad Moradi
4. **A Load-Independent Hybrid Capacitor-Inductor Interface Circuit for Piezoelectric Energy Harvesting** (319)
   Sima Ghandi, Hamed Nasiri, Lihong Zhang and Mohammad Al Janaideh
5. **Phase-Controlled GHz DC-DC Conversion Using Outphasing Class-E Power Amplifiers** (370)
   Ahilnandan Kabilan and Sakshi Arora

S34 | Special Session: Cutting-Edge Research Projects Funded by European and International Programs 1
11:30 AM - 1:00 PM
Room: Ostria
Chairs: Lazaros Papadopoulos, Erica Raviola

1. **Layer- and Block-wise DVFS for Neural Network Inference on Embedded GPUs: An Empirical Study** (440)
   Quanwen Zheng, Jingjin Li, Heng Yu, George Lentaris, Francky Catthoor and Yajun Ha
2. **Testing Methodology for Nanosatellite On-Board Computer Embedded Software System** (194)
   Panagiotis Karathanasis, Angelos Motsios, Alexandros Tsagkaropoulos, Christoforos Vasilakis, Christos Tsigkanos, George Lentaris and Dionysios Reisis
3. **Carbon-Aware Routing for Function Calling in Edge-Cloud LLM Systems** (82)
   Aikaterini Maria Panteleaki, Varatheepan Paramanayakam, Spyros Tragoudas and Iraklis Anagnostopoulos
4. **Digital Twins of Living Materials: Morphology and Memristive Emulation** (435)
   Georgios Kleitsiotis, Pantelis Fraidakis, Ioannis Tompris, Ioannis K. Chatzipaschalis, Theodoros Panagiotis Chatzinikolaou, Iosif-Angelos Fyrigos and Georgios Ch. Sirakoulis

5. **A Comprehensive Physical-to-Architectural Verification Framework of Memristive Crossbars for Neural Networks** (357)
   Anastasios Liontos, Christina Dilopoulou, Georgia Tsirimokou, Yiorgos Tsiatouhas, Denis Voltmann, Jaroslav Dzuba and Vasileios Tenentes
6. **Compiler-Assisted Radiation Mitigation for Next-Generation Space Computing Systems** (99)
   Federico Reghenzani, Davide Baroffio, Tomas A. Lopez, Emilio Corigliano, Giovanni Alberto Sartorato, Davide Mor, William Fornaciari and Giovanni Agosta

S04 | Analog/mixed signal/RF circuits 4 **2:00 PM - 3:30 PM**
Room: Maistros ACD
Chairs: Vasiliki Gogolou, Yasmin Halawani

1. **On the Equalization Capability of Resonator-Based CTLEs: A Time-Domain Perspective** (52)
   Mayu Zhang, Danilo Manstretta and Andrea Mazzanti
2. **Reconfigurable Integrated RC Networks: Experimental Characterization for Extended Filtering Functions** (69)
   Luis Arturo Martinez Alvarado, Federico Sandoval Ibarra, Miguel Angel García Andrade and Arturo Gutierrez D'Iaz
3. **New Biquadratic Operational Amplifier Based Filters** (153)
   Brent Maundy, Ahmed Elwakil and Abdulrahman Nutfaji
4. **A Calibration-Free GM-C Low-Pass Filter Constructed with Inverter Based OTAs** (292)
   Emre Sayıt, Berkay Çiftci, Adnan Gündel and Hasan Uluşan
5. **Non-Uniform Analog FIR Filtering for Wideband Receivers With Relaxed DAC-Rate Requirements** (312)
   Abhishek Kumar

S11 | Low-Power/Voltage Design **2:00 PM - 3:30 PM**
Room: Maistros BEF
Chairs: Lorenzo Benatti, Sahibia Kaur Vohra

1. **Buffer PSRR improvement through feedforward supply-noise cancellation** (41)
   Spyridon Vlassis, Konstantinos P. Pagkalos, Aikaterini I. Chronopoulou, George Souliotis, Vasileios Kalenteridis and Fotis Plessas
2. **A Synchronous, Latch-Based Digital OTA Operating at 1.25 nW-Power and 0.3V Power Supply** (177)
   Andrea De Gregorio, Danilo Karim Kaddouri and Paolo Stefano Crovetti
3. **An Energy Efficient Precisely Adjustable Hybrid Delay Element in 22-nm FDSOI CMOS** (209)
   Valentin Stadtlander, Florian Protze and Frank Ellinger
4. **Ultra Low Power, CMOS Programmable VRs for Under Voltage Lockout Circuits** (211)
   Giorgio Bersani, Michele Caselli and Andrea Boni

5. **An 83 ppm/ºC Voltage Reference in Flexible TFT Technology** (338)
   Noor Ul Amin, Andrea Ballo, Alfio Dario Grasso and Orazio Aiello

S13 | EDA, Test and Reliability 2 **2:00 PM - 3:30 PM**
Room: Zephyros
Chairs: George Floros, Alberto Garcia-Ortiz

1. **GAVENS: A GPU-Accelerated Value-Exact NPU Simulator for Pre-Silicon Accuracy Validation** (93)
   Chin-Cheng Wei and Chung-Ho Chen
2. **Benchmarking the Open Source Yosys-Nextpnr FPGA Toolflow against Vendor Software** (124)
   Sven Krause, Denis Lebold, Sebastian Wittlich, Tobias Krawutschke, Jens Rettkowski and Hendrik Wöhrle
3. **A Multi-Level Thermal Modeling and FPGA Emulation Framework for Nanophotonic 3D Systems** (197)
   Helen Clara George, Sebastian Fischer and Alberto García-Ortiz
4. **LLM-Assisted Coverage-Driven Verification of RISC-V Processor** (428)
   Zoya Ahmed and Otmane Ait Mohamed
5. **A Physics-Guided Transformer Framework for Electromigration Analysis in Multi-Segment Interconnects** (424)
   Pavlos Stoikos, Anuj Pathania and George Floros

S16 | Circuits and Systems for AI Algorithms 2 **2:00 PM - 3:30 PM**
Room: Sirocco
Chairs: Panagiotis Mousouliotis, Minas Dasygenis

1. **A Rank-8 Rolled Kernel-Grid Accelerator IP Core for CNN Feature Extraction** (143)
   Vinicius Correa, Gabriel Souza and Tales Pimenta
2. **Beyond Multiplication Count: System-Level PPA Trade-offs of Fast-Convolution RTL Accelerators** (182)
   Fernando Moraes and Tarsio Onofrio
3. **Event-Driven Binarized Convolutional Layer for Ultra-Low-Power On-Chip Keyword Spotting** (224)
   Marc Chevalier, Benoit Larras and Antoine Frappé
4. **FPGA-Based Deployment of Quantized DRL Policies for Edge** Robotics (403)
   Irene-Maria Tabakis, Konstantinos Tatas and Minas Dasygenis
5. **ZTA-Q: an Open-source RISC-V Platform for Accurate Quantized CNN Inference** (417)
   Yike Li, Ajay Kumar M, Vishnu P S, Deepu John, Dimitrios S. Nikolopoulos, Bo Ji and Hans Vandierendonck

S35 | Special Session: Cutting-Edge Research Projects Funded by European and International Programs 2
2:00 PM - 3:30 PM
Room: Ostria
Chairs: Lazaros Papadopoulos, Apostolos Fournaris

1. **Aging-Aware Resiliency Analysis for Digital Compute-In-Memory DNN Accelerators** (243)
   Alexis Maras, Panagiotis Chaidos, Georgios Alexandris, Theofilos Spyrou, Anteneh Gebregiorgis, Manil Dev Gomony, Henk Corporaal, Said Hamdioui, Dimitrios Soudris and Sotirios Xydis
2. **6G VERtical trials for SUStainability: towards a Greek pilot utilizing Energy-Efficiency-as-a-Service** (409)
   Aimilios Leftheriotis, Charis Diamantidis, Dimitrios Fragkos, Konstantinos Katsaros, Lazaros Karagiannidis, Spyridon Athanasiadis, Xenofon Vasilakos, Thanasis Douklias, Angelos Amditis, Georgios Theodoridis, George Lentaris and Dimitrios Soudris
3. **FPGA-based acceleration for Homomorphically Encrypted Processing in CoEvolution** (442)
   Stylianos Kotsobolos, Konstantina Karagianni, Vassilis Paliouras and Apostolos Fournaris
4. **Experimental Evaluation of a Low-Power Ultra-Wideband Receiver for Spectrum Sensing** (219)
   Panagiotis Vlachos, Ioannis Bartsiokas, Cedric Dehos, Francois Rivet and George Karachalios
5. **ATHENA Cybersecurity Validation and Testing Environment** (341)
   Lazaros Papadopoulos, Nikos Peppes, Dimitrios Soudris, Ilias Kalouptsoglou, Dimitrios Tsoukalas, Damian Puchalski and Michał Chorąś
6. **SPQR: Smart Precision for Quantum** Resilience (324)
   Simone Perriello, Gerardo Pelosi, Giacomo Lancellotti, Stefano Cherubin, Mohammad B. Rou, Giovanni Agosta and Gunnar Tufte

Young Professionals (YP) Panel Session **4:00 PM - 5:30 PM**
Room: Maistros ACD
Moderator: Dr. Vasileios Ntinas, Assistant Professor, Aalborg University

Roundtable Discussion **4:00 PM - 5:00 PM**

Panelists:

- Dr. Fotis Giannopoulos, R&D Manager, Prisma Electronics
- Dr. Vasiliki Gogolou, Associate Professor, Aristotle University of Thessaloniki
- Kristi Nikolla, Thess IC
- Konstantinos Kapetanos, IEEE AUTH Student Branch
- Thomas Gavris, Master Student, Aristotle University of Thessaloniki

Young Professionals Poster Session **5:00 PM - 5:30 PM**
Room: Foyer

The session will feature a roundtable discussion led by young professionals from academia and industry, focusing on the challenges and opportunities facing young professionals in Electronics, Circuits, and Systems. Topics will include career development, professional growth, and the transformative role of Artificial Intelligence in the evolving technological landscape.
The discussion will provide an opportunity to exchange experiences, explore emerging career paths, and share perspectives on the skills and competencies needed to succeed in a rapidly changing environment.
The session will continue with a poster presentation featuring selected contributions from Young Professionals, showcasing innovative ideas, technical creativity, and practical relevance. ABest YP Poster Award will be announced during the conference closing ceremony.
Following the panel discussion, a dedicated mentorship segment will provide opportunities for informal one-to-one and small-group exchanges between panelists and participants. This activity will encourage knowledge sharing, career guidance, and professional networking, helping young professionals build valuable connections with experienced members of the Electronics, Circuits, and Systems community.

Wednesday November 11th 2026

S14 | EDA, Test and Reliability 3 **10:30 AM - 12:00 PM**
Room: Maistros ACD
Chairs: George Lentaris, Minas Dasygenis

1. **Development of a Radiation-Hardened Standard Cell Library in 7 nm FinFET Technology** (175)
   Aobo Cui, Sarah Azimi and Luca Sterpone
2. **A Fault-Tolerant Serial Peripheral** Interface (265)
   Alex Calinescu and Marius Enachescu
3. **FERRino: A Radiation-Hardened Embedded FPGA in 28 nm CMOS** (322)
   Andrea Mazzanti, Luca Frontini, Alessandra Camplani, Alessandro Andreani, Alberto Stabile and Jafar Shojaii
4. **From Software to Hardware: Cross-Domain Vulnerability Detection in Verilog via Domain Adaptation** (356)
   Anna Maria Voulkopoulou, Ilias Kalouptsoglou, Dimitrios Tsoukalas, Dionysios Kehagias and Dimitrios Tzovaras
5. **Reliability Aware Learning Based Scheduling for Tiled DNN Accelerators** (266)
   Muhammad Zia Hameed and Iraklis Anagnostopoulos

S19 | AI-driven Circuit Design 2 **10:30 AM - 12:00 PM**
Room: Maistros BEF
Chairs: George Floros, Manolis Katsaragakis

1. **Assessing Data Fragmentation and Federated Learning for ML-based Dynamic IR-Drop Analysis** (300)
   Ludovica Deriu, Erich Malan, Valentino Peluso, Andrea Calimera, Stefano Quer and Enrico Macii
2. **Design Space Exploration of Reinforcement Learning Configurations for Post-Legalization Placement** (344)
   Foteini Oikonomou, Maria Pantazi-Kypraiou, George Goudroumanis, Pavlos Stoikos, Athanasios Tziouvaras, Christos Anagnostopoulos, Georgios Karakonstantis and George Floros
3. **GPlaceRL: An Open-Source Graph Reinforcement Learning Framework for Detailed Placement** (423)
   Pavlos Stoikos, Foteini Oikonomou, Christos Poulos, Maria Pantazi-Kypraiou, Athanasios Tziouvaras, Christos Anagnostopoulos, George Floros and Georgios Karakonstantis
4. **Hierarchy-aware Heterogeneous Graph Neural Network based Parasitic RC-risk Screening in Pre-layout Stage** (434)
   Hyemin Ko

S23 | Sensors and Sensing Systems **10:30 AM - 12:00 PM**
Room: Zephyros
Chairs: Dimitrios Prousalis, Thomas Noulis

1. **System-Level Optimization of Signal-to-Noise Ratio in Fiber-Optic Plasmonic Sensing Platforms** (36)
   Seung-Ki Lee

2. **Process Sensors for Neural Network Assisted ADCs** (141)
   Can Tunca, Robert van Veldhoven, Pieter Harpe and Marco Fattori
3. **A Dual-Channel Delay Line Oscillator Readout Circuit for Surface Acoustic Wave-Based Prostate-Specific Antigen Sensing in 22nm FD-SOI Technology** (315)
   Emmanuel John Bonifacio, Nathaniel Renz Domingo, Arcel Leynes, Maria Sophia Ralota, Marc Rosales and Maria Theresa De Leon
4. **A Differential Capacitive Sensor Array with Real-Time ASIC Readout for Dynamic Shear Wave Elastography** (414)
   Maryam Habibollahi, Arthur Jaccottet, Antonio Gomez, Nader Saffari and Andreas Demosthenous
5. **A Parasitic-Tolerant Relaxation Oscillator for GΩ-Range Sensor Interfaces** (419)
   Allen Boston, Matthieu Couriol, Paul Capgras and Pierre-Emmanuel Gaillardon

S27 | VLSI Systems and Applications **10:30 AM - 12:00 PM**
Room: Sirocco
Chairs: Nikos Sketopoulos, Nikolaos Athanasios Anagnostopoulos

1. **AdaptViT: Runtime-Adaptive Vision Transformer Deployment on Custom RISC-V** (397)
   Vishnu P S, Ajay Kumar M, Yike Li, Robert Bogdan Staszewski and Deepu John
2. **Immediate XOR In-Memory Computing in 6T-SRAM** (213)
   Viktor Weinelt and Klaus Hofmann
3. **Implementing Hybrid Homomorphic Encryption on a Silicon Root of Trust** (250)
   Tobias Stelzer, Vishesh Kiran Kumar Bafna and Andreas Zankl
4. **A Gates-Free Approximate 7-2 Adder Compressor for Energy-Efficient Diabetes Classification** (282)
   Leonardo Antonietti, Rodrigo Lopes, Morgana Macedo Azevedo Da Rosa, Eduardo Da Costa and Rafael Soares
5. **Versat-AI: From ONNX to Area-Efficient RISC-V+CGRA Edge AI Subsystems** (398)
   Rúben Teixeira, João Barreiros Rodrigues, Jaime Aguiar and José T. de Sousa

S31 | Digital Circuits and Systems 3 **10:30 AM - 12:00 PM**
Room: Ostria
Chairs: George Panagopoulos, Mattis Hasler

1. **Robust h-Parameter Extraction for RF Oscillators using Multi-Start Optimization and Upper-Bound Constrains** (225)
   Athina Papagianni and Georgios Panagopoulos
2. **A Resonant Low-Pass Ladder Filter for Sigma-Delta Bitstreams on FPGA** (294)
   Sven Finke, Rolf Paulus and Michael Karagounis

3. **A 3-ps-Resolution Wide-Tuning-Range Coarse-Fine Digitally Controlled Delay Line in 180 nm CMOS** (360)
   Lazar Ilić, Đorđe Gačić, Dejan Petković and Vladimir Milovanović
4. **End-to-End Hardware Validation of OSSDM Semantic Waveform for 6G Communications** (372)
   Lahcen Attijani, Panagiotis Vlachos, Mohamed-Amine Hamoura, Emilio Calvanese Strinati and François Rivet
5. **A Low-Power and Area-Efficient Standard-Cell-Based DIGOTA for Flexible Electronics** (388)
   Noor Ul Amin, Andrea Ballo, Alfio Dario Grasso and Orazio Aiello

S05 | Analog/mixed signal/RF circuits 5 **2:15 PM - 3:45 PM**
Room: Maistros ACD
Chairs: Costas Psychalinos, Vladimir Milovanović

1. **Multi-Modulus Frequency Dividers for System-On-Chip Clock Generation in 28-nm CMOS** (97)
   Tze Hin Cheung, Kari Stadius, Aarno Pärssinen and Jussi Ryynänen
2. **Effect of scaling on Shot noise in MOSFETs and Its overall impact on Phase Noise in Ring Oscillators** (138)
   Shubham Sharma, Anshul Verma and Sudeb Dasgupta
3. **A High-Linearity and Size-Optimized Synthesizable Cyclic Time-to-Digital Converter in 180 nm CMOS** (286)
   Dejan Petković and Vladimir Milovanović
4. **An LC VCO with Bias Circuit Noise Rejection using a Pseudo-Resistor-Based Low Pass Filter** (347)
   Mouli Venkata Prakash Pittala and Chithra
5. **Performance Optimization of a Cyclic Fine Interpolator in Nutt TDCs via Delay-Cell Trimming (**408)
   Javad Ahmadi-Farsani, Ion Vornicu and Laurentiu Acasandrei

S07 | Biomedical and Bio-inspired Circuits and Systems 2 **2:15 PM - 3:45 PM**
Room: Maistros BEF
Chairs: Vasileios Alimisis, Meysam Akbari

1. **A Neuromorphic Tactile Sensing System for the Hannes Anthropomorphic Prosthetic Hand** (230)
   Leila Hamadi
2. **Distilling Privileged Information into Efficient Neural Networks for Sepsis Detection** (259)
   Leonidas Paratsokas, Dimitrios Spanos, Nikolaos Passalis and Anastasios Tefas
3. **Improving Calibration in Lightweight Biomedical Deep Learning via Soft Subclass Adaptation** (262)
   Dimitrios Spanos, Nikolaos Passalis and Anastasios Tefas
4. **AcoustEEEcare: A Portable Audio-Based MEMS Patch for On-Device Heart-Rate Monitoring** (363)

Jose Marie Acedillo, Karylle Shainne Cadalzo, Mark Louie Pasturan, Paul Jason Co, Marc Rosales, Nathaniel Orillaza, Ron Louis Nierva and John Richard Hizon

5. **NeuroHub: A Software-Defined Neural Interface for Dynamic Electrode Routing, Biosignal Acquisition, and Four-Wire Impedance Spectroscopy** (412)
   Ebubekir Enes Arslan, Jhan Luke Okkabaz, Gregory Worrell and Nuri Ince

S17 | Circuits and Systems for AI Algorithms 3 **2:15 PM - 3:45 PM**
Room: Zephyros
Chairs: Nikos Sketopoulos, Yiorgos Tsiatouhas

1. **VQ4SNN: Vector Quantization for Memory-Efficient FPGA Spiking Neural** Networks (214)
   Dimitrios Sekertzis and Giorgos Dimitrakopoulos
2. **A Time-Based Readout for Vector-Matrix Multiplication in Fully Analog Memristive SNNs** (148)
   Elia Mateu Barriendos, Álvaro Gómez-Pau, Josep Rius, Daniel Arumí, Rosa Rodríguez-Montañés and Salvador Manich
3. **Train Where You Infer: A System-Level Analysis of On-Device Forward-Forward for IMC Achitectures** (149)
   Neha Ashar, Aymen Romdhane, Kamel-Eddine Sadadou, Johannes Reibold, Guillaume Prenat, Mehdi Tahoori, Jonathan Miquel and Lorena Anghel
4. **Pareto-Optimal Accuracy, Area, and Power in Time-Domain Analog Neural Networks via Hardware-Aware Training** (237)
   Giulio Filippeschi, Giuseppe Iannaccone, Mirko Brazzini, Cristhopher Mosquera, Marco Lanuzza, Alessandro Catania and Sebastiano Strangio
5. **MELIF: Design of Multi-Encoded Leaky-Integrate-and-Fire Neuron in 28 nm CMOS Technology** (361)
   Sahibia Kaur Vohra

S21 | Energy Harvesting and Power Management 2 **2:15 PM - 3:45 PM**
Room: Sirocco
Chairs: Vasiliki Gogolou, Spyridon Vlassis

1. **Waste Heat Recovery from a 3.6 GHz GaN Doherty Power Amplifier using Thermoelectric Generators** (89)
   Lucas Ott, Robert Krämer, Xin Xu, Jens Wagner and Frank Ellinger
2. **Brownout-Aware Energy Management for Wearable Health Patches Using SMM--MDP Optimization** (195)
   Hajer Abdelaziz, Abdullah Yousuf, Shamma Alblooshi, Sara Alessa and Yasmin Halawani
3. **ADRC-Enhanced MPPT With Sensorless Open-Circuit Fault Diagnosis and Reconfiguration for Three-Level PV Boost Converters** (269)
   Abdeldjabar Benrabah, Mouaad Belguedri, Mohamed Amine Yahiaoui, Fayçal Benyamina, Lakhdar Benhamimid and Mohamed Benbouzi

4. **Dual-Mode Ultrasonic System Enabling Simultaneous Energy Harvesting and Data Communication in Water and Ex-Vivo Tissue Channels** (350)
   Ziyi Zhao, Qihan Wang, Chenxue Hou, Zhaoxi Li, Dongdong Chen, Di Li and Yintang Yang

S30 | Emerging technologies and technology trends 3 **2:15 PM - 3:45 PM**
Room: Ostria
Chairs: Ioannis Messaris, Iraklis Anagnostopoulos

1. **Stochastic Resonance for Variability-Tolerant Memristor IMPLY** Logic (327)
   Rana Walied Ahmad, Rosana Rodriguez and Vasileios Ntinas
2. **An arbitrary scale programmable Flash ADC based on analog memristor ladder** (377)
   Diego Guglielmi, Matej Horský, Marián Precner, Sergei Krylov, Tomáš Ščepka, Tibor Izsák and Boris Hudec
3. **Assessment of an 8T Dynamic-Cascode SRAM-CIM Bitcell in IMEC N2 GAA Technology** (381)
   Pio Esposito, Andrea Fasolino, Rosalba Liguori, Luigi Di Benedetto, Alfredo Rubino and Gian Domenico Licciardo
4. **Hybrid CMOS Spin Wave Implementation of Convolutional Neural Networks** (272)
   Pantazis Anagnostou, Arne Van Zegbroeck, Said Hamdioui, Christoph Adelmann, Florin Ciubotaru and Sorin Cotofana
5. **MemrisTec Memristor Model Platform: An Open Educational Framework for Studying Memristive Devices** (443)
   Christian Bruchatz, Dimitrios Prousalis, Ahmed Magdy Ibrahim Abdelsamad, Mohammad Fahim Sultan Anoy, Ahmet Samil Demirkol, Adnan Haidar, Umer Amer Khan, Carsten Knoll, Ioannis Messaris, Sarthak Dharmendra Nandanwar, Pawan Patel, Jonathan Pitzschel, Anagha Raghunath, Nithya Kalyani Ramachandran, Steffen Seitz, Richard Schroedter, Hongliu Yang and Ronald Tetzlaff

Poster Sessions

Poster Session 1 - Monday November 9th 2026, 3:30 PM - 5:00 PM

POSTER #2	Digital
Immediate Boolean IMC – Enabling Area-Efficient SRAM In-Memory Computing	
Viktor Weinelt, Hannes Simon Korn and Klaus Hofmann	

POSTER #9	Digital
A Fully Synthesizable, Technology-Independent Digitally-Controlled Delay Line	
Hannes Simon Korn, Viktor Weinelt and Klaus Hofmann	

POSTER #64	Digital
Frequency Domain Equalization Design for Gigabit Automotive Ethernet System	
Chih-Feng Wu, Muh-Tian Shiue, Yu-Jun Chen, Min-Yuan Tsai and Tong-Da Lee	

POSTER #76	AMS Systems
A 28 Gbps PAM-4 Adaptive Equalizer with Data and Edge ISI Compensation	
Meng-Xuan Lin, Kuo-Hsing Cheng, Wen-Hsuan Hsieh and Hong-Yi Huang	

POSTER #81	Digital
Pipelined Systolic-Ring-Based Digital Oscillator Ising Machine for Highly Dense Graphs	
Anubhav Bhura, Sathish N and Debanjan Bhowmik	

POSTER #87	Digital
Power-Accuracy Design Exploration of Dual-kernel Configurable Gaussian Filter	
Marcio Monteiro, Ismael Seidel, Mateus Grellert, José Luís Güntzel, Leonardo Soares and Cristina Meinhardt	

POSTER #103	Digital
Performance Analysis of Sparse Kernels on a RISC-V Open Source GPU	
Jules Dubois, Eric Guthmuller, Adrian Evans and Frédéric Rousseau	

POSTER #117	AI & applications
Hardware-Aware Acceleration of DNN Inference on Embedded GPUs for Chest X-Ray Classification	
Theodora Sanida, Argyrios Sideris and Minas Dasygenis	

POSTER #118	AI & applications
Hardware–Software Co-Mapping NAS for Heterogeneous Embedded Platforms	
Theodora Sanida, Argyrios Sideris and Minas Dasygenis	

POSTER #123	Digital
Hardware Design of a Scan-Enabled Four-Stage Pipelined SHA-3 Core on FPGA	
Panagiotis Antoniou, Constantinos S. Dionysiou, Panayiotis Pericleous and Harris E. Michail	

POSTER #129	Digital
Predictive and Host I/O–Aware Garbage Collection Optimization in Solid State Drives	
Ojasweeta Singh and Ramanathan Muthiah	

POSTER #137	Technology
A Novel Radiation-Hardened-by-Design Architecture for SRAMs	
Hung-Ming Kuo, Jafar Shojaii, Andrea Mazzanti and Ting-Chang Chang	

POSTER #139	Digital
Challenges of Taping Out a Multi-Tenant MPSoC with Heterogeneous IP Contributions	
Mattis Hasler	

POSTER #145	Technology
A Lightweight Full-Custom FiGaRO TRNG for Resource-Constrained Security in 12 nm FinFET	
Francesca Bagnoli, Ettore Noccetti, Daniele Rossi and Sergio Saponara	

POSTER #170	Technology
Layer-Wise Hardware-Aware Non-Uniform Quantization for Unary-Coded RRAM AI Accelerators	
Ming Cheng, Cao Wang, Xiao-Chun Li, Yanxia Lin and David Keezer	

POSTER #185	Digital
Full-Chip Modeling for Voltage Analysis on the Backside Silicon Substrate	
Rikuu Hasegawa, Kazuki Monta, Takuya Wadatsumi, Alkis Chatzopoulos and Makoto Nagata	

POSTER #191	Technology
Information-Theoretic Causal Modelling of Semiconductor Process Dynamics	
Daniel Sørensen, Giorgio Melchiorre, Sudip Bandyopadhyay, Sandip Halder, Roel Wuyts and Bappaditya Dey	

POSTER #201	Digital
An SVA-to-RTL Transpiler for Synthesizable Hardware Assertion Checkers	
Antonio Giuseppe Maria Strollo, Gennaro Di Meo, Luca Tegazzini, Camillo Perna, Matteo Sommella and Davide De Caro	

POSTER #220	Digital
Evaluating NASA's Core Flight System Fault Detection and Recovery on RISC-V Systems	
Luís França, Fabio Benevenuti, Leonardo R. Gobatto, José Rodrigo F. Azambuja and Fernanda Kastensmidt	

POSTER #226	Digital
Reinforcement Learning-Based Optimization of Workload-Aware Power Delivery Networks	
Oran Hayes, Maria Pantazi, Athanasios Tziouvaras, George Stamoulis, Anuj Pathania, Shreejith Shanker and George Floros	

POSTER #239	Digital
FPGA Congestion Estimation Using Topology and Geometry Aware Graph Neural Networks	
André Cavalcante, Sumanta Chaudhuri and Lirida Naviner	

POSTER #257	Digital
Area-Efficient CKKS Homomorphic Multiplication Design Using Output-Driven Scheduling and Hierarchical Basis Conversion	
Tsu-Hsiung Huang, Yun-Ting Hsu and Ming-Der Shieh	

POSTER #273	AMS Systems
Zero Current Detector for Positive Outputs of a Single-Inductor Triple-Output DC-DC Converter	
Giorgio Rolandi, Alireza Mesri, Elisabetta Moisello, Piero Malcovati and Edoardo Bonizzoni	

POSTER #283	Digital
Design-Space Exploration of Multi-Kernel Multiplierless Gaussian Filters Using 3-2 AxAC	
Rodrigo Lopes, Léo Ribeiro, Leonardo Antonietti, Morgana Macedo Azevedo Da Rosa, Eduardo Da Costa and Rafael Soares	

POSTER #295	Digital
Portable Stimulus Vertical Reuse with VerTrace-Generated Interconnections: A RISC-V Execution Stage Case Study	
Petr Bardonek and Marcela Zachariášová	

POSTER #305	Technology
Self Powered Systems as Security Buttons for Improving Trust in Wildfire Sensor Nodes	
Zane Gunn, Aniket Roy, Samit Hasan, Davis Loose, Megan Gunn, Daniel Truesdell, James Lambert and Benton Calhoun	

POSTER #316	Technology
A 4H-SiC CMOS Op-Amp for High Temperature Applications up to 500°C	
Andrei-Tudor Gheorghe and Marius Enachescu	

POSTER #325	Technology
PCM-enabled energy consumption estimation for language models deployed on edge devices	
Lazaros Papadopoulos, Giorgos Balis, Socratis Nazlidis, Miltiadis Siavvas, Georgios Karagiannidis and Ilias Kalouptsoglou	

POSTER #382	Digital
HOMAxA-2: An Energy-Efficient Hybrid OR–MUX Approximate Adder for Edge AI	
Sagar Bhivsane, Alak Majumder, Vassilis Alimisis and Paul P. Sotiriadis	

POSTER #384	Technology
Micromagnetic Investigation of a Skyrmion-Based Wired AND Logic Device Using Hierarchical VCMA Barriers	
Debasmita Chatterjee	

POSTER #385	Technology
Non-Invasive Hydration Status Detection via Human Body Communication	
Amr N. Abdelrahman, Abdelhay Ali and Ahmed M. Eltawil	

POSTER #393	Digital
A very low cost, fully digital, Clock and Data Recovery circuit for SERDES	
Massimo Ruo Roch and Maurizio Martina	

POSTER #404	Digital
RISC-V for Quantum Control: Evaluation of Soft-Core Architectures on Cryogenic Temperatures	
Gustavo Kanaiama, Rafael Rodrigues, Édney Freitas, Luis Bravo, Gustavo Marcati, Francisco Brito and Bruno Sanches	

POSTER #407	AI & applications
Hamm3R: Optimized 3D Image Matching for Memory-Constraint Environments	
Mert Tas, Halil Durmuş, Bilgi Görkem Yazgaç and Ece Olcay Güneş	

POSTER #425	Digital
FSNIC: A Low-Latency Flow-Based Intrusion Detection Architecture for FPGA SmartNICs	
Nise O'Cuill, Changhong Li, Georgios Floros and Shreejith Shanker	

POSTER #439	Technology
Flexible FPGA Undervolting Test Bench For Research And Teaching	
Lucas Bublitz, Ulf Kulau and Leonard Schauer	

POSTER #405	SP1
KHNUM: A Kernel Hash Network for Unified Multi-function Cryptography	
Mohamed El-Hadedy	

Poster Session 2- Tuesday November 10th 2026, 4:00 PM - 5:30 PM

POSTER #43	AI & applications
Silicon Aware Neural Networks	
Sebastian Fieldhouse and Kea-Tiong Tang	

POSTER #68	AI & applications
A CKKS-Aware Distilled CNN System for Efficient Encrypted Inference	
Babak Moradi, Hassan Rekabi Bana, Moslem Heidarpur and Mitra Mirhassani	

POSTER #71	AI & applications
A Physics-Informed Surrogate Modeling Framework for Real-Time Analog Circuit Performance Prediction: A 55-nm OTA Case Study	
Konstantinos Liakos, Christos Poulos and Fotis Plessas	

POSTER #114	Digital
Efficient VLSI Design for CKKS Key Switching via Adaptive Special Modulus Selection and Approximate Basis Conversion	
Qi-Xian Wu, Tsu-Hsiung Huang, Yun-Ting Hsu and Ming-Der Shieh	

POSTER #155	AI & applications
A Resource-Efficient MobileNetV2 Hardware Engine with Multispectral Fusion on FPGA	
Tsung-Han Tsai, Chun-Tse Lin and Yu Shang-Yin	

POSTER #165	AI & applications
Specific Component Fault Detection in Analog Circuits using the RMS of the Supply-Current Fluctuation	
Vassilis Vassios	

POSTER #176	AI & applications
An End-to-End FPGA-Based SoC Implementation of Quantized Neural Networks for ECG Biometric Recognition	
Claudia Georgiana Barbura, Paul Farago and Sorin Hintea	

POSTER #261	AI & applications
A Latent Ensemble Method for Safe Multimodal Reinforcement Learning in Robot Navigation	
Konstantinos Spyridopoulos, Pavlos Tosidis, Nikolaos Passalis and Anastasios Tefas	

POSTER #270	AI & applications
CuriousRL: A Curiosity-Driven Goal-Conditioned RL Engine for Multi-Objective Analog Circuit Sizing	
Yash Sengupta, Niranjan Gopal and Sakshi Arora	

POSTER #274	AI & applications
Error-Driven Sampling of (,) for Improved Generalization in Multi-Layer Perceptron-Based Compact Modeling	
Yusra Rachidi, Thomas Bedecarrats, Roberto Miorelli and Emeric De Foucauld	

POSTER #289	AI & applications
Acceleration of BiLSTM-based Music Separation Algorithm by Customizing HW/SW on SoC-FPGA	
Sotirios Aias Karioris, Ioannis Voyiatzis and George Lentaris	

POSTER #293	AI & applications
A Pipelined FPGA Architecture for Banded Sparse Matrix–Dense Matrix Multiplication in Longformer	
Phillip Pramberger, Athanasios Tziouvaras, Shreejith Shanker and George Floros	

POSTER #301	Digital
A Framework for Accelerating Transformer Inference on RISC-V for Edge AI	
Ajay Kumar M, Vishnu P S, Yike Li, Shreejith Shanker, Dimitrios S. Nikolopoulos, Bo Ji, Hans Vandierendonck and Deepu John	

POSTER #313	AI & applications
A Technology-Generic ANN Framework for Modeling and Design of Operational Amplifiers	
Engin Afacan, Günhan Dündar and Sefa Çetinkol	

POSTER #358	AI & applications
MIVE: A Minimalist Integer Vector Engine for Softmax LayerNorm and RMSNorm Acceleration	
Kosmas Alexandridis and Giorgos Dimitrakopoulos	

POSTER #375	SP1
Real-time Tunable Probability Distributions From Perimeter-Gated SPAD TRNGs	
Yanbing Icey Chen, Md Sakibur Sajal and Marc Dandin	

POSTER #380	Digital
TCAD-Based Analysis of Single Event Transients in 130 nm Bulk CMOS Technology	
Eleftheria Litou, Pelopidas Tsoumanis, Georgios-Ioannis Paliaroutis, Marko Andjelkovic, Nestor Evmorfopoulos and George Stamoulis	

POSTER #422	Digital
A Layout-Level Single Event Transient Susceptibility Analysis of 4-2 Adder Compressor	
Ana Flávia Reis, Morgana Rosa, Rafael Soares, Eduardo Costa, Cristina Meinhardt and Rafael Schivittz	

POSTER #427	AI & applications
Lightweight Binary Neural Networks for Noise Independent Voice Activity Detection	
Michael Grzybek and Stefano Gregori	

POSTER #421	AMS Systems
Development of Verilog-A Modules for the Acceleration of Time Domain Noise Simulations	
Panagiota Fasili, Christoforos Theodorou and Thomas Noulis	

Young Professionals Poster Session

POSTER #445	YP session
MOF-enabled Capacitive Humidity Sensor for Precision Agriculture	
Seerat Sekhon, Maryam Habibollahi, Vikramjeet Singh and Andreas Demosthenous	

POSTER #446	YP session
A Ka-band 8-channel Beamformer IC for mmWave 5G and Radar Applications	
Ioannis Kousparis, Nikolaos Naskas, Nikolaos Alexiou, Ilias Tsakiridis, Konstantinos Kontaxis, Fanourios Fakoukakis, Georgios Ntounas and Stefanos Horianopoulos	

POSTER #452	YP session
A 0.8-mW Neural Recording Front End for Vagus Nerve Stimulation Applications with 1.8-μVrms Noise and ~1 ms Stimulation-Artifact Recovery	
Trishya Angela Babs	

POSTER #453	YP session
Liquid-Decay Regulation of Switching Sensitivity in a HETUF-Based Memristive Recurrent System	

POSTER #454	YP session
A Complex I/Q Direct-RF Receiver Using Periodic Non-Uniform Sampling.	
Abhishek Kumar	

POSTER #455	YP session
A Comparison of Low-Complexity Iterative Algorithms for Adaptive Active Gate Drivers	
Erica Raviola and Franco Fiori	

POSTER #458	YP session
Energy-Aware NAS for Object Recognition using a Piezoelectric Tactile Sensing System	
Hiba Al Youssef, Razan Khalifeh, Leila Hamadi, Mohamad Raad, Maurizio Valle and Ali Ibrahim	

POSTER #461	YP session
A reconfigurable D-Flip-Flop based on reconfigurable NAND/NOR Logic Gates	
Eric Fehrmann, Juan Martinez, Giulio Galderisi, Thomas Mikolajick and Jens Trommer	

POSTER #463	YP session
Design and Evaluation of a Resource-Efficient Direct Digital Synthesizer	
Dimitra Skinezou and Spyridon Nikolaidis	

POSTER #464	YP session
A Modular Hardware-in-the-Loop Framework for Validating AI-Generated Embedded Firmware	
Bharat Kumar, Harshit Krishna R and Madhav Rao	

POSTER #465	YP session
A Semi-Automated g m /I D -Based Technology Migration Flow for Regenerative Mixed-Signal Circuits	
Ashwin George and Klaus Hofmann	

POSTER #466	YP session
AIHDL: An Agentic AI Based Autonomous Pipeline for Verilog Code Generation, Simulation, Verification and Analysis	
Ashutosh Singh, Anuj Grover and Kunal Iiitd	

Poster Session 3- Wednesday November 11th 2026, 1:15 PM - 2:15 PM

POSTER #17	AMS Systems
A 64×64, 1162 TOPS/W Analog Capacitive In-Memory Computing Array for Vector-Matrix Multiplication	
Fady Sabry Negm, Max Uhlmann, Corrado Carta, Gerhard Kahmen and Philip Ostrovskyy	

POSTER #22	Digital
Explicit Timestamp Propagation for Temporal Processing	
Xavier Lesage, Cristiano Merio, Azzadine Thajte, Luca Sauer de Araujo, Fernando Welzel, Sylvain Engels and Laurent Fesquet	

POSTER #35	AI & applications
Dual Diff Query Based Guiding for LLM Spec Edits for RTL Driven Specification Synchronization	
Rahimullah Shaik	

POSTER #116	AMS Systems
A 4.2-Gb/s Ultra-Low-Power LPDDR4X DQ Receiver Architecture in GF 22-nm FDSOI	
Ashwin George	

POSTER #136	AMS Systems
A High-Performance Single-Stage OTA with 90-dB DC Gain for Nanofarad Capacitive Loads	
Rozhin Salehi Saheb, Erika Covi, Kea-Tiong Tang and Meysam Akbari	

POSTER #142	AMS Systems
A Low-Noise, High-Linearity Three-Stage Amplifier forAutomotive MEMS Gyroscope Capacitance-to-Voltage Converters	
Raphael Maier	

POSTER #147	AMS Systems
A Memristive Synapse for Online STDP Learning and Inference in SNNs	
Elia Mateu Barriendos, Álvaro Gómez-Pau, Daniel Arumí, Rosa Rodríguez-Montañés and Salvador Manich	

POSTER #158	AMS Systems
On the Design of Current Mirrors Using Only Depletion-Mode Transistors in Monolithic GaN Technology	
Fabio Meloni, Sandro Rossi, Elisabetta Moisello, Edoardo Bonizzoni and Piero Malcovati	

POSTER #172	AMS Systems
A 5.26ppm/°C 155-nW CMOS Sub-Bandgap Voltage Reference	
Evi Keramida, George Souliotis, Spyridon Vlassis and Fotis Plessas	

POSTER #183	AMS Systems
A Fractional-N Phase-Locked Loop with a Pseudo-Differential DTC and a Phase-Error Sign Extraction Circuit Based on a Phase-Detector Replica	
Johan Holmstedt, Rainier van Dommele and Vojkan Vidojkovic	

POSTER #202	AMS Systems
Bandwidth Interpolation Capability of NARXNN-Based PA Behavioral Model	
Thuy Pham Trong, Germain Pham, Reda Mohellebi, Pierre Almairac, Carolina Pedrosa and Patricia Desgreys	

POSTER #221	AMS Systems
Patch-Based Spatial Encoding for Memristor Crossbar Array Reservoir Computing	
Karol Bednarz	

POSTER #234	AMS Systems
Fundamental limits of the Pulse Frequency Modulation converter for very high Full-Well Capacity Digital Pixel Sensor	
Martin Lapeyre, Philippe Martin-Gonthier, Olivier Saint-Pe, Thierry Kohl and Vincent Goiffon	

POSTER #240	AMS Systems
A 0.8mW, 22ns Peaking Time, Analog Front-End with Current-Domain Pulse Shaping Filter for High-Density Energy-Resolved Photon-Counting Detector Arrays	
Matteo Petruzzi, Leonardo Gasparini, Stefano D'Amico, Jacopo Quercia and Massimo Gandola	

POSTER #249	AMS Systems
Machine-learning Based Error Correction in Multi-bit Delta-Sigma Digital-to-analog Converters	
Matthew Kinsinger and Arindam Sanyal	

POSTER #281	Digital
Hardware-Friendly FIR Architectures for ECG Denoising and QRS Morphology Preservation	
Morgana Macedo Azevedo Da Rosa, Rodrigo Lopes, Eduardo Da Costa and Rafael Soares	

POSTER #287	Digital
On the Impact of Degradation-Balanced Scheduling in Manycore Systems	
Maria Pantazi-Kypraiou, George Rafael Goudroumanis, George Floros, Anuj Pathania and George Stamoulis	

POSTER #297	AMS Systems
Supercells: Addressing Mismatch Variability in Analog Compute-in-Memory Arrays	
Michael Tomlinson, Richard Niu, Adil Malik, Philippe Pouliquen and Andreas Andreou	

POSTER #298	AMS Systems
Experimental Verification of CMOS-Compatible Capacitive Isolation AC-Switch Driver	
Ming Zhang and Nicolas Llaser	

POSTER #302	AMS Systems
Extended Compact MOSFET Model for Wide Range of Inversion Level	
Daniel do Carmo Silva Ribeiro, Fabian Andrade, Ana Cunha and Edson Pinto Santana	

POSTER #328	AMS Systems
A 273 GOPS and 54.6 TOPS/W FeFET-Based CIM Architecture for Edge AI Inference Applications	
Anisha Sharma, Cheena Singhal, Abhishek Goel and Sudeb Dasgupta	

POSTER #353	AMS Systems
A 1-Gb/s Pulse-Width Modulation Transceiver for 15-m Twisted-Pair Links in 28-nm CMOS	
Johannes Kuhn, Tim Lauber, Tristan Strahler, Ralf Wunderlich and Stefan Heinen	

POSTER #369	AMS Systems
An Efficient Method of Optimal Trimming Pattern Identification for ENOB Maximisation in Analog Front Ends	
Arpit Bal and Laxmeesha Somappa	

POSTER #383	AI & applications
FINNAS: FINN-Guided Hardware-Aware NAS and Pruning for FPGA Jet Substructure Classification	
Eva Chauffour, Changhong Li, Georgios Floros and Shreejith Shanker	

POSTER #387	Digital
SET Mitigation Strategies for a Minimal Standard Cell Library in 7 nm FinFET Technology	
Rafael Neves de Mello Oliveira, Rafael Schivittz and Cristina Meinhardt	

POSTER #402	Digital
Enabling Low-Cost Scalable Cryogenic Quantum Control with FPGA-Based SDR: Experimental Validation in a Cryostat	
Gustavo Kanaiama, Francisco Rouxinol and Bruno Sanches	

POSTER #410	AMS Systems
An Energy Efficient Divider-Free Body-Resonance Transmitter with Direct Oscillator Keying for Human Body Communication	
Amr N. Abdelrahman, Abdelhay Ali and Ahmed M. Eltawil	

POSTER #418	AMS Systems
A Fully Differential Charge Pump Amplifier with High Voltage Gain	
Masoud Askariraad and Stefano Gregori	



2026 IEEE ICECS

November 8th-11th 2026, Thessaloniki, Greece

33rd IEEE International Conference on Electronics, Circuits and Systems

Book of Abstracts

Monday November 9th 2026

S01 | Analog/mixed signal/RF circuits 1 **11:00 AM - 12:30 PM**
Room: Maistros ACD
Chairs: Engin Afacan, Thomas Noulis

Common-mode Tank Impedance Engineering for the Suppression of the 1/f^3 Phase Noise Sideband in Class-C CMOS Oscillators (47)
Marco Longo, Marc Tiebout and Andrea Bevilacqua

A theoretical study of the influence of the common-mode tank impedance on the 1/f noise-to-phase noise upconversion in CMOS harmonic oscillators is presented. We show that the 1/f^3 phase noise sideband can be suppressed by carefully engineering the common-mode tank reactance. In fact, this largely reduces the sensitivity of the Groszkowski shift of the oscillation frequency to fluctuations of the conduction angle of the transistors in a class-C CMOS oscillator, and hence to 1/f noise. The analytical developments are validated by simulations carried out in a 16 nm CMOS technology.

A Robust 17-21 GHz Frequency Synthesizer in SiGe-BiCMOS for Space Applications (65)
Nebojsa Maletic, Frank Herzel, Giovanni Di Pietrantonio and Gunter Fischer

This paper presents a phase-locked loop (PLL)-based frequency synthesizer for space applications. It is designed as part of an intermodulation built-in self-test (BIST) circuit of an analog predistorter block that is driving a solid-state power amplifier (SSPA) or a traveling wave tube amplifier (TWTA). Using two voltage-controlled oscillators (VCOs) with coarse and fine tuning loops, the PLL generates an output signal in the 17-21 GHz frequency range from a 10 MHz reference. The PLL tuning range spans over 32 subbands (16 per VCO), achieving the phase noise of -100 dBc/Hz at 1 MHz offset in the middle of the tuning range (18.82 GHz). The PLL is designed and fabricated in 130 nm SiGe-BiCMOS technology, occupying an area of approximately 6 mm^2 with a power consumption of 750 mW from a nominal 3.3 V supply.

A Low Power Wideband Current-Mode Receiver Frontend for FR3 Band Applications (101)
Shiuan-Han Huang, Chien-Nan Kuo and Hwann-Kaeo Chiou

This paper presents a low-power wideband current-mode receiver (RX) front-end targeting FR3-band applications (7–24 GHz). The proposed low-noise amplifier (LNA) employs a current-reuse topology to minimize power consumption, while a dual-resonance input matching network and multiple inductive peaking techniques are adopted to achieve wideband input matching, high gain, and low noise figure (NF). In addition, the baseband transimpedance amplifier (TIA) adopts a two-stage architecture with feedforward compensation to extend channel bandwidth. By merging the feedforward path with the output stage of the second amplifier stage, the conventional trade-off between bandwidth and power consumption is alleviated. Fabricated in a 90-nm CMOS process, the proposed RX front-end occupies 0.776 mm². Measurement results show a peak conversion gain of 26.0 dB

and a minimum NF of 2.78 dB while consuming only 35.5 mW. The RF bandwidth covers 8.5–18.5 GHz, and a 3-dB channel bandwidth of 1.4 GHz is achieved. The measured in-band IP1dB and blocker P1dB are −31 dBm and −22 dBm, respectively, while the in-band and out-of-band IIP3 are −19 dBm and −13.5 dBm.

A High-Gain D-Band LNA with Inductively-Tuned Gain-Boosting Technique in 28nm Bulk CMOS (174)

Rainier van Dommele, Yiqin Hou, Remco Schalk, Kaijie Ding and Vojkan Vidojkovic

This paper presents the design and implementation of a three-stage gain-boosted low-noise amplifier (LNA) operating in the D-band, which incorporates a novel inductively-tuned negative-resistance gain-boosting technique. To evaluate the effectiveness of the proposed technique, a conventional transformer-coupled three-stage LNA is implemented as well. Both LNAs are implemented in 28nm bulk CMOS and characterized using on-wafer probing. The proposed technique provides a 7.3dB extra gain compared to the three-stage conventional LNA. The power consumption is reduced by 30% with respect to a four-stage conventional LNA that would achieve nearly the same gain. The conventional LNA/gain-boosted LNA provide 11.3dB/18.6dB power gain, -11.7dBm/-13.0dBm input 1-dB compression point at their operating frequencies of 142.5GHz/136GHz respectively. Both LNAs have a noise figure of 7.2dB, occupying a core area of only 0.03mm2. The gain-boosted LNA shows excellent performance compared to the state-of-the-art solutions, which proves the feasibility and effectiveness of the proposed technique.

A Highly Linear Autotransformer-Matched S2D LNTA for Sub-6-GHz Receivers in 22-nm FDSOI (181)

Meysam Shayestehfar, Harald Pretl and Danilo Manstretta

This paper presents an autotransformer - matched single- ended- to- differential (S2D) low-noise transconductance amplifier (LNTA) in 22-nm FDSOI for sub-6-GHz current-mode receivers. Conventional inductor-less CG–CS active baluns are limited by input parasitic poles and the rigid coupling between 50-Ω power matching and noise optimization. The proposed LNTA employs a compact 1:1 on-chip autotransformer to provide broadband pole compensation and relaxation of the noise/power matching trade-off, extending the input match beyond 6 GHz. A cascode-tapped cross-coupled auxiliary-gm core simultaneously enforces differential current balance without load asymmetry and improves linearity through IM3 cancellation. Post-layout simulations of the embedded LNTA predict an NF of 4.9– 5.8 dB, +6.7 to +7.8dBm OOB-IIP3, and less than 1 dB output current amplitude imbalance across 2.7–6 GHz. Full receiver chain measurements demonstrate |S11| < −10 dB from below 2 GHz up to 6.7 GHz, NF of 6.5–7.2 dB, 44 dB of voltage gain, and +5.3 to +6.5dBm OOB-IIP3. The LNTA consumes 7.5mW and has a 0.025-mm2 core- plus- autotransformer footprint.

S12 | EDA, Test and Reliability 1
11:00 AM - 12:30 PM
Room: Maistros BEF
Chairs: Nestoras Evmorfopoulos, Cristina Meinhardt

SET: Substrate Extraction Tool for Predicting Substrate Noise in SOI Smart Power ICs (186)
Erica Raviola, Gabriele Salzone, Luca Bertolini, Sergio Lecce, Kamel Abouda and Franco Fiori

Substrate noise propagation is a critical issue in smart power SOI design, as the switching noise generated by power circuits can couple with sensitive analog circuitry causing malfunctions. This work presents SET, a lightweight tool for extracting SPICE-like netlists to provide a rapid estimation of aggressor-victim coupling through the substrate. By implementing a local mesh refinement, SET reduces the simulation time while preserving accuracy. Validation on elementary test structures shows good agreement with commercial tools in terms of spreading resistance and substrate coupling, making SET a practical solution for early-stage architecture exploration and design definition.

3MLR: An Investigation of Moment-Matched Mutual Ladder-based Reduction for VLSI Static Timing Analysis (232)
Aikaterini Tsilingiri, Dimitrios Tsalapatas, Stavros Simoglou, Stavros Stathoudakis and Christos Sotiriou

Parasitic RC networks extracted from physical layouts grow large at advanced technology nodes, imposing significant memory and runtime overhead on static timing analysis. This paper presents 3MLR, a moment-matched mutual ladder-based reduction algorithm that maps multi-sink RC trees to compact linear ladders via adaptive sink clustering and first-moment matching. A structural complexity score Φ is introduced to identify nets unsuitable for ladder reduction and pass them through unreduced. Evaluated on six OpenCores designs on the ASAP7 7 nm PDK at three parasitic extraction effort levels, 3MLR reduces SPEF node count by up to 23.5% and file size by up to 13.2%, with, on average, 98.3% out of ten thousand paths per design showing <3% arrival error and zero pins in the ≥50% error bin.

A Residual Degree-Guided Heuristic for Delay-Arc Traversal in Combinational Cells Characterization (268)
Rita Louro Barbosa and Cristina Meinhardt

This work formalizes the characterization of combinational cells problem as the Minimum Artificial Transition Cover problem and shows its reduction to the Directed Chinese Postman Problem over a symmetric digraph. We propose a single-pass greedy heuristic that avoids the pre-processing and degree-equalization phases required by exact methods. Over 92 benchmark-circuit instances, the heuristic reaches the optimal solution in about 62% of cases, and on large circuits its additional cost over the optimum averages less than 0.5% of the total arcs to be traversed.

Enhancing MBFF Clustering with Continuous Timing Criticality Awareness (307)

Lucas Tinoco and José Luis Güntzel

Clock networks represent a major source of power consumption in modern VLSI designs. Multi-Bit Flip-Flops (MBFFs) can reduce clock network power by merging groups of Single-Bit Flip-Flops (SBFFs) into shared structures with common clock drivers. However, determining how to cluster SBFFs while preserving timing quality remains challenging. Among the state-of-the-art MBFF clustering techniques, the scalable divide-and-conquer FTray (SFTray) available in OpenROAD platform performs Min-Cost Flow (MCF)-based assignment driven exclusively by Manhattan displacement, without explicitly considering timing criticality until a later Integer Linear Programming (ILP) stage. This paper presents a continuous timing-aware enhancement to the MCF assignment cost formulation within SFTray, by combining a normalized slack-derived criticality score with normalized displacement into a unified edge cost. The experimental results after detailed routing of nine circuits using the ASAP7 enablement showed that our formulation is able to efficiently explore the optimization space in circuits with higher SBFFs ratios, being more effective than SFTray in improving power and preserving timing, while still producing better or similar area and clock buffer reductions.

End-to-End GNN-ML Framework for Automated SPICE-to-SystemVerilog Behavioral Modeling (333)

Karim Osama, Amr Wael, Kareem Atef, Kareem Abdelraof, Hana Maged, Mariam Ahmed, Mostafa Waseem, Ramy Raafat and Mohamed Abd El Ghany

Behavioral modeling for mixed-signal SoC verification remains a manual bottleneck. This paper presents an automated framework translating SPICE netlists into simulation ready SystemVerilog. The system employs a two-stage GraphSAGE GNN with Name Block Voting (NBV) to classify analog blocks into six classes, alongside a three-checkpoint human-in-the-loop (HITL) system to ensure labeling accuracy. Following classification, ML regression and deterministic engines extract parameters to populate SV templates, while a static CMOS extractor and Weisfeiler–Lehman (WL) hash matcher handle the digital domain. The generated SV models are verified through an interactive visual validation workflow. Experimental results demonstrate a 94.32% node-level classification accuracy across 1,306 netlists and 94.92% OTA/BIAS discrimination across 114 netlists. Furthermore, system-level validation of an OTA-RC filter cascade confirms AC response consistency between the generated top-level SV model and the Cadence SPICE reference.

FPGA and Parallel SNN Inference for RIS-Assisted Transmitter Selection (134)
Evangelos Kyparoglou, Minas Dasygenis and Alexandros-Apostolos Boulogeorgos

Reconfigurable Intelligent Surface (RIS)-assisted wireless systems require real-time, low-latency link adaptation, but deploying Spiking Neural Networks (SNNs) for transmitter selection introduces significant computational overhead. This paper presents a hardware-software heterogeneous acceleration framework to optimize SNN-based transmitter selection. The framework evaluates multiple parallel execution strategies, migrating the inference pipeline into optimized sequential C++, OpenMP, MPI, hybrid OpenMP+MPI, and hardware-accelerated FPGA implementations. Experimental results using 200,000 samples demonstrate that while OpenMP delivers the best CPU-side scaling, the FPGA implementation on the AMD Xilinx Kria KR260 provides the highest overall efficiency. The FPGA achieves a 12.24x speedup over the sequential baseline and a 21.39x reduction in energy consumption compared to OpenMP, proving the effectiveness of heterogeneous acceleration for edge inference in next-generation wireless networks.

Efficient Temporal Analog In-Memory Computing via Co-Training of Long- and Short-Term Memory in RRAM Devices (45)
Ming-Jay Yang, Dimitrios Spithouris, Johannes Hellwig, Regina Dittmann, Gordon Pipa and John Paul Strachan

Volatile memristors exhibit intrinsically transient conductance states arising from spontaneous current relaxation, endowing them with rich device-level temporal dynamics. These dynamics naturally encode short-term memory (STM) and time-dependent responses, making volatile memristors promising for temporal and event-driven signal processing. However, a systematic architecture that directly maps the analog temporal kernels of volatile memristors to structured spatio-temporal computation remains lacking. In this work, we leverage the intrinsic temporal kernels of volatile memristors with nonvolatile memristive neural networks. By co-training the network weights and time constants, we show 2-3 orders of magnitude improvement on model parameter efficiency and more than 2 orders of magnitude improvement on power efficiency compared with standard convolution-based and recurrent models running inference on GPU. The results also demonstrate efficient scalability of inference accuracy with increasing numbers of STM taps. When compared to advanced bio-inspired models on memristive neuromorphic chips we also show improved parameter and power efficiency. These improvements suggest further research in algorithm-technology co-design leveraging the analog tunability of memristive devices for analog in-memory computing (AIMC) applications.

Performance Improvement over Recurrent Spiking Neural Networks HW Accelerator for FPGA (253)
Diego Avila-Torres, Antonio Perez-Pena, Luis G. Leon-Vega, Jorge Castro-Godínez and Alejandro Linares-Barranco

Spiking Neural Networks (SNNs) have emerged as a promising paradigm for energy-efficient and biologically inspired computation, particularly in edge and embedded systems. However, their hardware implementation poses significant challenges due to strict timing constraints and the need for efficient data movement. This work presents a performance-oriented enhancement of the ReckOn recurrent spiking neural network (RSNN) accelerator deployed on a FPGA-based MPSoCs. We compare two execution paradigms: PYNQ (Python-based), and CYNQ (C++-based runtime); on an AMD Kria K26 SoM. By migrating the training loop to CYNQ and extending its capabilities with additional IP integration and Python bindings, we achieve substantial improvements in execution efficiency while preserving training accuracy, demonstrating up to 232x speedup in DMA transfers, and an overall reduction in training time of 2.57x, with equivalent accuracy across both frameworks. These findings highlight the effectiveness of CYNQ in reducing software overhead and accelerating neuromorphic workloads.

Co-Optimizing Graph Sparsification and Approximate Computing for Energy-Efficient FPGA-Based GCN Inference (290)
Nathaniel Kaye Mellor, Shreejith Shanker and George Floros

Graph Convolutional Networks (GCNs) have emerged as a powerful framework for learning from graph-structured data, yet their deployment on resource-constrained edge platforms remains challenging due to the computational and memory demands of sparse graph aggregation. This work presents an FPGA-based GCN accelerator that combines DSpar graph sparsification, 8-bit quantization, and approximate multipliers on the AMD Kria KV260. Evaluated on Cora, LastFM Asia, and Amazon Photo, the design explores the interaction between sparsification and approximation across graphs with widely varying densities.
Results show that the effectiveness of approximate arithmetic is governed by accumulation depth within GCN computations. Approximate multipliers are most effective when applied to sparse aggregation operations, while graph sparsification further improves their viability by reducing aggregation depth. The combined approach achieves up to 9.88$\times$ speedup while maintaining 86.6% classification accuracy on Amazon Photo, and 1.52$\times$ speedup with 77.0% accuracy on Cora, with total power consumption below 1 W. These results demonstrate that graph sparsification and approximate computing are complementary techniques whose co-optimization enables efficient low-power GCN inference on edge FPGA platforms.

SRAM-Based Processing-in-Memory Cell Enabling Binary Neural Networks Through XNOR Multiplication and Charge-Domain Accumulation (330)
Cristhopher Rafael Mosquera Montero, Ramiro Taco, José Benjamin Zambrano Chavez, Esteban José Garzón Córdova, Alessandro Catania, Sebastiano Strangio, Giuseppe Iannaccone and Marco Lanuzza

This paper presents a high-efficiency mixed-signal processing-in-memory accelerator for binary neural networks. The architecture uses a novel SRAM bitcell that unifies in-memory XNOR computation with direct analog charge-domain accumulation, thus avoiding power-hungry intermediate digital circuitry. To mitigate hardware non-idealities, the design features a self-biased sensing scheme and a programmable offset-compensation mechanism tailored for batch-normalization folding. Post-layout simulations of a 288 x 388 array in 65-nm CMOS yield up to 23.7 TOPS throughput, 306.8 TOPS/W energy efficiency, and 62.4 TOPS/mm^2 area efficiency. System-level validations on MNIST and CIFAR-10 datasets confirm excellent classification robustness against device mismatch.

S22 | Microsystems and Embedded systems **11:00 AM - 12:30 PM**
Room: Sirocco
Chairs: Stefano Gregori, Ioannis Messaris

PrefA: Prefetcher as Accelerator for Time Dependent Irregular Workloads on RISC-V ISA (285)
Savvina Nastou, Georgios Alexandris, Manolis Katsaragakis, Dimitrios Soudris and Sotirios Xydis

In recent years, memory scaling has become a major bottleneck in the embedded community, overshadowing the computation novelties. Modern ML and HPC workloads can also suffer from this scaling problem, especially when using dynamic workloads with time-dependent irregular memory accesses, which limit their performance. Existing ideas approach this problem by introducing prefetching mechanisms that load blocks into the cache before they are requested by the executable. Prior novel implementations are limited by extra instruction overhead and insufficient acceleration gains when implemented in software, suffer from cache pollution risks, and are constrained by pointer-chasing optimizations when mapped in hardware. Hybrid prefetching came as a middle solution to this design space, providing hardware/software co-design principles in prefetching. Still, this technique cannot guarantee cache pollution avoidance. Additionally, most of the existing systems are only simulated, not assessed for the utilization/performance errors introduced from emulating/synthesizing the Register Transfer Level (RTL). To address these limitations, we introduce PrefA, a non-invasive prefetching accelerator for RISC-V ISAs, using dynamic data-type metadata on external caches and allowing the user to implement strategies that describe how that metadata would be handled and prefetched accordingly in the L1 cache. PrefA uses custom instructions to guide the prefetching and populate the cache with directed parts of dynamic data types, accelerating time-dependent accesses. Evaluated on both microbenchmarks and real-world implementation, PrefA is gaining up to 79% more performance, when in the utilization we showcase 4% lower FF and LUT utilization with 1% URAM overhead when compared to previous SoA, emulated in a ZCU106 FPGA.

Memristive Slack Time and Event Predictors for Resource-Constrained Embedded Systems (362)
Dongchen Li, Zidu Li, Bhaskar Choubey and Roman Obermaisser

Accurate and reliable operation of the Slack Time and Event Prediction Models (SEPM) is critical for deep-learning-based adaptive time-triggered systems. While attention mechanisms and multi-cycle inputs improve prediction accuracy, the resulting increase in inference latency limits their deployment in resource-constrained applications such as advanced driver-assistance systems. Memristive crossbar arrays could provide an efficient solution, but deploying the entire SEPM model is impractical due to its complexity. In this paper, we present a simulation framework that injects memristor-induced noise into different components of SEPM. Results demonstrate that the attention module dominates execution time, reaching up to 73.58\% at a hidden dimension of 100. Moreover, injecting noise exclusively into the attention module causes only minor performance degradation, whereas injecting noise into other components leads to significant performance degradation and task imbalance. These results indicate that the attention module achieves the best trade-off between efficiency and robustness, making it the most suitable candidate for crossbar deployment.

A Digital Twin Model and Architecture for Monitoring and Controlling Fleets of Autonomous Unmanned Surface Vehicles (394)
Anderson Roberto Pinheiro Domingues, Cássio Jones Dhein da Silva, Francesco da Rocha Lui, Juliano Maia, Rodrigo S. Cenço, César Augusto Mission Marcon and Fernando Gehm Moraes

This paper presents Fleet-DT, an open-source, fleet-level Digital Twin (DT) architecture for Autonomous Unmanned Surface Vehicles (AUSVs). The architecture, called Fleet-DT, relies on the Message Queuing Telemetry Transport (MQTT) protocol and integrates the Ardupilot firmware, the WeBots simulator, and several simulation applications into a single system. Fleet-DT dynamically compensates for the wireless communication bottleneck inherent in fleet telemetry by implementing rate regulators in local MQTT brokers. Quantitative latency and link-budget modeling are validated using the Jundiá Project's fleet of AUSVs for environmental DNA sampling. This paper demonstrates that Fleet-DT guarantees Quality of Service (QoS) and maximizes the use of shared wireless links.

Exploration Methodology for Resource-aware Multi-model Deployment on Edge Devices (231)
Anagnostis Boubardas and Lazaros Papadopoulos

Deploying specialized Language Models (LMs) at the edge offers benefits for privacy, security, and localized intelligence. However, executing multi-model workflows on resource-constrained hardware introduces bottlenecks across memory footprint, storage bandwidth, and energy consumption. This work proposes a systematic exploration methodology, supported by an open-source framework, for multi-model deployment on resource-constrained edge devices. Focusing on Language Models (LMs), the methodology identifies trade-offs between quality of model responses, execution latency, memory requirements and energy efficiency for optimizing multi-LM deployments. The framework is validated on two heterogeneous edge architectures under multiple workloads and use cases.

The results demonstrate that through multi-objective optimization, Pareto-optimal configurations are identified that reduce memory usage by up to 59%, while trading off 2.5% to 15% in quality of model responses for up to 21% increase in efficiency.

A Hybrid Differentiable Digital Twin for Induction Motors under Deep Magnetic Saturation (60)
Xingwei Zhou, Wenshan Hu, Zhongcheng Lei and Jianbin He

Differentiable physics operators have recently emerged as a principled foundation for digital twins of electric drives, enabling exact gradient-based parameter identification and trajectory optimization. However, when such purely mechanistic operators are confronted with heavy-duty transients that drive the iron core into deep magnetic saturation, a fundamental failure mode arises: the gradient-descent procedure compensates for unmodeled non-linearities by distorting the structural parameters themselves, yielding physically meaningless estimates that subsequently break field-oriented control. This paper proposes a two-stage hybrid differentiable digital twin that resolves this dilemma. In the first stage, physically interpretable parameters are identified from small-signal data in the linear magnetic regime and then frozen. In the second stage, a C^2-continuous, zero-initialized residual network is embedded in the state-derivative space and trained, via the discrete adjoint method, to absorb the residual dynamics induced by saturation while leaving the physical anchors untouched. A scale-aware normalization is introduced to handle the multi-decade dynamic range of the motor states. Experiments on a 37kW induction motor under transient currents exceeding 600A show that the hybrid operator reduces RMSE on stator currents by more than 95% relative to a saturated mechanistic baseline, and retains close tracking under zero-shot evaluation on an unseen load profile.

S28 | Emerging technologies and technology trends 1 11:00 AM - 12:30 PM

Room: Ostria
Chairs: Boris Vaisband, Aniket Roy

RF GaN technology with enhanced multi-layer Back-End-of-Line (BEOL) process (119)
Errikos Lourandakis, Nikolaos Makris, Evangelos Aslanidis, Antonios Stavrinidis, Athanasios Kostopoulos, Eleftherios Sarris, Paolo Fioravanti, George Konstantinidis, Alexandros Gurulian and Giannis Kontogiannopoulos

This paper presents FORTH's RF GaN process offering depletion mode (D-mode) HEMTs and an improved multi-layer back-end-of-line (BEOL) process with three gold metal layers for designing advanced and high performance passives. Parameter extraction of industry standard HEMT compact model was performed utilizing static and dynamic characterization, setting the framework for circuit design in FORTH's GaN pilot line. The impact of having multiple BEOL metals is discussed by investigating inductive transformer devices used for impedance matching between circuit blocks. As a feasibility study, both a single metal air-bridge transformer and a 3-metal multi-layer transformer have been fabricated and measured by on-wafer probing.

A 7-Bit Current-Steering DAC With an NRZ Dynamic-Switch Driver in IGZO TFT Process (131)
Muhammad Dawood Asghar, Sven Baerten, Mohit Dandekar and Kris Myny

This work presents a 7-bit current-steering (CS) digital-to-analog converter (DAC) fabricated in a 600-nm unipolar (n-type) IGZO thin-film transistor (TFT) technology for flexible and large-area electronic systems. To mitigate linearity degradation and reduce the large device area, an on-chip dynamic-switch driver (DSD) with an inherent level shifter (LS) is introduced. The proposed architecture suppresses the switching skew, while enabling highly compact unit current cells. The prototype achieves a maximum spurious-free dynamic range (SFDR) of 46 dB at 1 MS/s, a settling time of 0.3 µs, and a full-scale output-voltage swing of 1.2 V_Peak-Peak, while consuming 264.4 µW from a 2 V supply. The mean measured INL and DNL are within ±0.5 LSB and ±1 LSB, respectively. The DAC occupies an active area of 1.19 mm2, resulting in improvement >100× in area-efficiency per unit current cell and 2350× enhancement in FoM-I. In addition, FoM-II improves by 139× compared to state-of-the-art (SotA) IGZO-based CS DAC.

Design of on-chip ESD protection network using a-IGZO TFTs (163)
Khang Nguyen, Florian De Roose, Paoline Coulson, Xiaonan Xing, Mauricio Velazquez Lopez, Sebastian Haesler, Marko Simicic, Jan Genoe and Kris Myny

This work demonstrates the design of a whole-chip electrostatic discharge (ESD) protection network for amorphous Indium-Gallium-Zinc-Oxide (a-IGZO) thin-film transistor (TFT) technologies, addressing the limitations of unipolar device operation, without dedicated ESD devices. Diode-connected TFTs are proposed as core protection elements. Optimized local clamp and power clamp topologies are designed and characterized using Human Body Model (HBM) and (Very-Fast) Transmission Line Pulse (VF-TLP) measurements. The ESD robustness of the local clamp against very-fast pulses is improved by the addition of secondary on-chip protection elements. The capacitor-resistor (CR)-triggered transient power clamp provides superior clamping efficiency, higher ESD robustness, and lower static power consumption compared to the static power clamp. The resulting protection network is able to handle long and short ESD events of any polarity and achieves HBM levels up to 900V with static current consumption below 0.1 mA.

Improvement of IGZO MeSFET Performance by Post-Fabrication Anneal: Electrical Assessment and TCAD Modeling (323)
Ivan Lisovyi

This paper presents the development and comprehensive electrical characterization of an In-Ga-Zn-O (IGZO) metal-semiconductor field-effect transistor (MeSFET), where the gate stack consists of a Schottky-forming metal in direct contact with semiconductor without the need for gate dielectric. Functional IGZO MeSFET devices were successfully fabricated and characterized, allowing to extract key electrical parameters to benchmark transistor performance. The device performance was improved through post-fabrication thermal annealing, with a focus on analyzing the effects of oxygen and nitrogen ambients across various temperatures and durations. Furthermore, TCAD modeling of the devices was

conducted, demonstrating good correlation between simulated device physics and measured data under different anneal conditions. Based on these findings, we propose a standardized set of figures of merit for MeSFET device, providing a robust framework for their consistent evaluation and future development.

Electro-Thermal Analysis and Modeling of Resistor-Transistor Logic in a 200-V E-mode GaN-on-SOI Technology (365)
Salvatore Pennisi, Katia Samperi, Andrea Ballo, Urmimala Chatterjee and Stefaan Decoutere

This paper presents an electro-thermal analysis and modeling of Resistor-Transistor Logic (RTL) inverters implemented with low-voltage transistors of a 200-V enhancement-mode (E-mode) GaN-on-SOI monolithic platform. A closed-form analytical model for the rise/fall time, propagation delay and power-delay product is derived, directly linking dynamic performance to the thermal degradation of charge carrier mobility. Experimental characterization across a temperature range from –40 °C to 150 °C reveals a rise time increase from 722 ns to 1.97 ¦s under a capacitive load of around 28.5 pF. Through nonlinear fitting, a mobility degradation exponent ■ ≃ 1.82 is extracted. These results provide an essential physical baseline for designing high-temperature digital circuits on smart GaN-on-SOI platforms.

WiCAS Panel Session **2:00 PM - 3:30 PM**
Room: Maistros BEF
Moderator: Prof. Maria Papadopoulou, International Hellenic University, Greece

Prof. Erivelton Nepomuceno, Launching CAS, RAS and WiE Student Branch Chapters and Supporting Participation in IEEE CASS Design Competitions

This talk shares practical experience from establishing three new IEEE student branch chapters at Maynooth University: Circuits and Systems Society (CAS), Robotics and Automation Society (RAS), and Women in Engineering (WiE). It describes the steps involved, from recruiting student officers and securing faculty advisors to meeting IEEE administrative requirements and planning early activities that build momentum. It shows how running the three chapters together creates shared membership, joint events and mutual support, rather than each competing for student attention. Building on this foundation, the talk presents an ongoing initiative to create a structured process within the Department of Electronic Engineering that helps students take part in IEEE CASS design competitions, with mentoring, access to resources and alignment with academic activities. It closes with lessons learned, common pitfalls, and transferable recommendations for chapter leaders across Region 8 who want to strengthen student engagement and turn CASS strategies into local action.

Prof. Ahmed H. Madian, Strategic Initiatives, Growth Trajectories, and Operational Outcomes of the IEEE Egypt Section Circuits and Systems Society (CASS) Chapter

This presentation highlights the strategic initiatives, growth strategies, and impact outcomes of the IEEE Egypt Section Circuits and Systems Society (CASS) Chapter in advancing microelectronics, integrated circuits (IC), and system-level research across the local and regional engineering ecosystem.

To address the growing demand for specialized hardware design, deep-tech skills, and industry-ready talent, IEEE Egypt CASS implemented a targeted framework focused on three primary pillars: capacity building, academic-industrial collaboration, and student engagement.

Key initiatives include technical workshops on VLSI and analog/digital IC design, expert-led seminar series, and mentorship programs tailored for graduation projects and joint chapter activities.

Quantifiable outcomes demonstrate marked increases in active professional and student memberships, heightened participation in technical design competitions, and a stronger regional pipeline connecting academic research with the semiconductor industry.

These efforts establish a scalable model for technical society chapters aiming to accelerate local technological adoption and professional development within emerging engineering economies.

Prof. Mustak Erhan Yalcin, From Student Engagement to Sustainable CASS Membership: Lessons from the IEEE CASS Türkiye Chapter

Sustainable growth of IEEE and IEEE Circuits and Systems Society (CASS) membership requires more than recruiting established professionals who are already embedded in existing professional and social networks. Our experience in the IEEE CASS Türkiye Chapter suggests that a more sustainable strategy is to build strong relationships with young engineers while they are still students and forming their professional identities and communities. In Türkiye, close cooperation with university student organizations supported the establishment of the IEEE ITU CASS Student Branch Chapter and enabled a growing portfolio of technical trainings, seminars, industry talks, company visits, and student activities. These initiatives helped students experience CASS not only as a technical society, but also as an international professional community.

For students, one of IEEE's strongest differentiators among many university clubs is its global character. This advantage becomes meaningful only when students can actively experience it. We therefore argue that CASS should place greater emphasis on cross-border student engagement through joint events, collaborative projects, competitions, and mobility activities connecting IEEE and CASS Student Branch Chapters in different countries. Dedicated financial support mechanisms are essential to make such participation accessible.

Looking further ahead, IEEE and CASS could also act as a framework for international master's pathways, joint doctoral initiatives, and other structured academic collaborations among universities. Strong relationships formed during the student years can then continue into professional life through technical development, industry engagement, and international peer networks. These long-term professional connections can become an important reason for graduates to retain their IEEE and CASS membership.

Based on lessons from the Türkiye Chapter, this presentation proposes a student-centered membership development model that links recruitment, international belonging, academic collaboration, and professional continuity, with the aim of building a stronger and more sustainable CASS community across Region 8.

Prof. Zidu Li, Connecting Students with Industry through the IEEE CASS Germany Chapter

Building connections between students and local industry can help ease the transition from university to professional life. The IEEE CASS Germany Chapter has organized experience-sharing events that provide students an opportunity directly exchanging with professionals working in different career levels in German companies. It is more helpful for international students, who may also face additional challenges to understand the German industrial environment, workplace culture, career paths, and expectations. Through this direct interaction with industry people, students can gain practical insights, ask questions, and develop a better understanding of preparation for future careers in Germany. This presentation will share our experience in organizing this activity, discuss the value of direct student–industry interaction, and highlight lessons learned. We hope that this approach can provide a practical and transferable model for other CASS Chapters seeking to strengthen connections between students and local industry.

Prof. Abdallah Kassem, From Circuits to Communities: Empowering the Next Generation through IEEE CASS STEM Outreach

This talk presents the experience of the IEEE Lebanon Joint Chapter, IE13/PE31/CAS04/PEL35 (CH08625) in using STEM outreach as a key component of Chapter engagement and community impact. Beyond traditional technical and professional activities, the Chapter has developed initiatives that introduce school students and young learners to circuits, electronics, programming, robotics, renewable energy, and emerging technologies through practical and engaging experiences.
The presentation highlights several initiatives, including Arduino and Raspberry Pi workshops for high-school students, robotics activities for orphaned children using programmable boards, pre-university STEM outreach focused on renewable energy and circuits and systems, and the “ECO SMART CITY” Workshop & Competition, which engages students from public and private high schools in Arduino coding and circuit assembly. The Chapter has also organized a four-day Technological Summer Camp for students aged 12–17, providing hands-on exposure to Arduino programming and electronic systems.
Drawing on these experiences, the talk will share practical approaches for IEEE CASS Chapters to expand their impact through hands-on STEM education, partnerships with schools and other IEEE entities, inclusive community outreach, and activities that connect circuits and systems with real-world challenges. The objective is to demonstrate how Chapter activities can inspire younger generations, strengthen the engineering pipeline, and create lasting connections between IEEE, universities, schools, and local communities.

Prof. Costas Psychalinos, Restarting the IEEE Greece CASS Chapter: Good practices and lessons learned

This talk presents the experience of the IEEE Greece CASS Chapter Chair during the process of the local Chapter restarting. The followed steps are analytically described, including the selection criteria of the officers, the establishment of the Chapter’s site, the chosen activities (such as the organization of Seasonal Schools, Conferences, and lectures within the IEEE CASS Distinguished Lecture Program), as well as the future plans for further development.

In addition, the difficulties which have been faced in these years are discussed, including also the followed strategy for overcoming these obstacles.
The outcome of this talk is intended to be a helpful guide for officers which want to establish a sustainable CASS Chapter. Towards this goal, suggested good practices and, also, the lessons learned, will conclude this talk.

S02 | Analog/mixed signal/RF circuits 2 **5:00 PM - 6:30 PM**
Room: Maistros ACD
Chairs: Vasilis Konstantakos, Marius Enachescu

A FMIR Radar receiver with 1-ns Step Dynamic Pattern Variable Gain Amplifier (140)
Arthur Paviot, Pierre Courouve, Mykhailo Zarudniev, Laurent Ouvry and Gilles Masson

This paper explores a complete FMIR receiver with a dynamic pattern Variable Gain Amplifiers (VGA) featuring 16 reconfigurable gain steps swept at up to 1GHz, enabling adaptive gain patterns optimized for Radar applications with 1GHz bandwidth. The reception chain reaches a time-varying dynamic range of 40dB with 8 bits low power interleaved ADCs. The complete Radar circuit designed in 22-FDSOI achieves a power consumption under nominal operating conditions of no more than 152mW at 0.8V and a compact footprint of less than 6.24mm² while the VGA core area is 0.026mm² and uses 11mW of the total power.

A -57 dBm sensitivity 26 GHz Active Envelope Detector for Wake-Up Receiver Applications (146)
Albert Kibiwott, Chadi Jabbour, Reda Mohellebi, Kyriaki Niotaki and Zuo Liu

This paper presents a pseudo-differential 26 GHz active envelope detector employing a complementary nMOS-pMOS current-reuse stack to maximize the transconductance per unit current, enhancing second-order nonlinearity in the weak inversion region. A co-designed resonant impedance matching network provides passive voltage boosting at the carrier frequency, thereby improving the detector's sensitivity. Implemented in 22 nm GlobalFoundries FD-SOI CMOS technology, the proposed design consumes 2.2 µW from a 0.8 V supply and achieves a sensitivity of -57 dBm and a data rate of 1 Mbps.

Broadband Darlington-Cascode Power Amplifier: Analysis, Stability, and Measurement Results (246)
Paul François, Thierry Lagutere and Philippe Descamps

A broadband GaN/AlGaN Darlington-Cascode power amplifier covering 50 MHz to 800 MHz is presented. The Darlington input stage, combined with dual resistive feedback, keeps a nearly resistive, matchable input impedance across the band, but introduces a negative input resistance, located by a five-node admittance model reduced to a two-port by Kron reduction and suppressed by a gate-drain RC network. The amplifier is a non-grounded multi-transistor stage on which the usual two-port stability tests do not apply. Stability is established by four complementary criteria: Kurokawa, Winslow SMH0/SMY0, Ohtomo, and pole-zero analysis.

Post-layout simulations predict P-1dB = 33-35 dBm and Psat = 38.5-39 dBm, comparable to the state of the art. On-wafer measurements confirm S21 = 20-22 dB up to 800 MHz with no spurious oscillation.

GaN-on-SOI Oscillators for Smart Power ICs (73)
Paul-Catalin Medinceanu, Florin-Silviu Dumitru and Marius Enachescu

While Gallium Nitride (GaN) is a proven enabler for high-efficiency, high-power-density conversion systems, the monolithic integration of robust auxiliary circuits remains a bottleneck for GaN smart power ICs. To address this, this paper proposes four oscillator architectures implemented in a 100~V GaN-on-SOI technology. The first two are ring oscillators designed to provide on-chip process characterization for HEMTs and MIM capacitors. To meet the digital logic and switching converters demands, the third and fourth circuits introduce supply-independent clocking architectures: a three-phase differential ring oscillator and a novel GaN adaptation of a CMOS relaxation oscillator. By eliminating the need for depletion-mode devices, the proposed relaxation oscillator achieves a highly stable frequency response with merely 0.45% variation under a 4V to 6V supply sweep, while consuming a minimal silicon area of 0.034 mm^2.

Efficiency-Centric Automated Design Optimization of a Digitally Configurable IR-UWB Transmitter (238)
Yağız Kırtaş, Okan Zafer Batur and Burcu Erkmen

This paper introduces an efficiency centered design optimization for Impulse Radio Ultra-Wideband (IR-UWB) transmitters. The circuit in the loop CAD design is performed autonomously, utilizing scripted code (Cadence's Ocean Script platform). The proposed Ocean Scripted Autonomous Engine (OSAE) is capable of identifying the ring oscillator circuit parameters Winning}$, W_Pinv, W_TG, and A_Cap that result in peak efficiency in a digitally configurable design. The transmitter is configurable from 0.35 GHz to 1.95 GHz . The design is all digital; thus, the main power consumption is due to the configurable ring oscillator and output stages that drive a 50 ohm antenna. This work focuses on the selection of the ring oscillator design parameters for optimal efficiency in a 4-bit configuration of the ring oscillator. The circuit design is conducted on a 65 nm TSMC MSRF CMOS process with standard transistors and 1.2 Volt supply. The achieved efficiency range for the optimum design using the OSAE method on a 4 bit configuration varies from 33.84 % to 38.28 %. The energy per pulse (EPP) for the optimum design varies from 8.95 pJ/pulse to 47.9 pJ/pulse, and the pulse energy E_P from 3 pJ to 18.28 pJ, depending on the configuration.

S08 | Digital Circuits and Systems 1 **5:00 PM - 6:30 PM**
Room: Maistros BEF
Chairs: George Lentaris, Apostolos Fournaris

A BRAM-Free and DSP-Free On-the-Fly Twiddle Factor Generator for PQC NTT Accelerators (227)
Shiva Nejati, Moslem Heidarpur, Lang Lin, Norman Chang and Mitra Mirhassani

The Number Theoretic Transform (NTT) is the dominant computational kernel in lattice-based post-quantum cryptography. Conventional accelerators store twiddle factors in on-chip BRAM or rely on DSP-based multipliers for on-the-fly generation, incurring overhead that grows with transform size. This paper proposes a BRAM-free and DSP-free on-the-fly twiddle factor generator (TFG) exploiting the carry propagation structure of the bit-reversed counter: each twiddle update reduces to a single multiply by one of W precomputed constants, eliminating both BRAM and DSP. A three-stage registered pipeline achieves 220–230 MHz on Zynq-7000, Kintex UltraScale, and Virtex UltraScale+ with zero BRAM and zero DSP for transform sizes N = 215 and N = 216. The proposed design achieves up to 43× lower equivalent normalized slice (ENS) cost than prior works and is, to the best of the authors' knowledge, the first reported TFG that achieves both zero BRAM and zero DSP simultaneously.

MPX: A Unified Systolic Array for Matrix and Polynomial Multiplication (263)
George Alexakis, Dimitrios Schoinianakis and Giorgos Dimitrakopoulos

Polynomial multiplication is a fundamental kernel in Fully Homomorphic Encryption (FHE) and post-quantum cryptography (PQC) and is commonly accelerated through Number Theoretic Transforms (NTTs). To avoid the cost of designing dedicated cryptographic accelerators, recent efforts have mapped NTT computations onto existing systolic matrix engines, enabling the reuse of AI hardware for cryptographic workloads. In this work, we take the opposite approach. We observe that the wavefront dataflow of systolic arrays naturally aligns with the accumulation pattern of polynomial multiplication and leverage this correspondence to design MPX, a dual-mode systolic array that supports both matrix multiplication and direct polynomial multiplication within the same hardware fabric. Experimental results show that extending a conventional systolic array with this dual-mode capability requires only 20% additional area and introduces negligible power overhead during matrix-multiplication execution. In polynomial-multiplication mode, MPX achieves more than 1.2x lower latency compared to NTT-based polynomial multiplication on systolic matrix engines.

Efficient Dynamic Logic Implementation of Magnitude Comparators (351)
Constantinos Efstathiou, John Liaperdos and Yiorgos Tsiatouhas

High speed comparators are critical blocks in modern computer, communication and digital signal processing systems. Dynamic logic design style provides circuits with the required high-speed performance, although power consumption remains an issue. In this work, an efficient dynamic logic based magnitude comparator architecture is presented that maintains speed performance at reduced (up to 28%) dynamic power consumption and lower silicon area requirements (up to 13.5%) with respect to advanced comparator designs in the literature.

HPG-SVD High-Performance Generalized Singular Value Decomposition Processor with Adaptive Termination (401)

Anjan Banerjee, Souris Sahu, Rohan Purkait, Richard Prajwal, Sasya Merugu Sri, A. Prasanna Laxmi, Sounak Samanta and Acharyya Amit

This paper presents the High-Performance Generalized SVD (HPG-SVD) processor, a multiplier-free FPGA accelerator for real-time MUSIC-based Direction-of-Arrival (DoA) estimation. The proposed architecture combines a Radix-4 CORDIC datapath with Forward Angle Recoding (FAR) to halve microrotation cycles and skip redundant zero-valued rotations. A custom diagonal/off-diagonal BRAM interleaving scheme eliminates two-sided Jacobi read/write collisions, while adaptive early termination recovers idle pipeline cycles without accuracy penalty. Implemented on a Xilinx ZCU104 Zynq UltraScale+ FPGA, the processor achieves 0.046 ms SVD latency for a complex 8 × 8 covariance matrix, consuming only 35 DSP slices and 0.432 W of Programmable Logic dynamic power, outperforming state-of-the-art designs by up to 2.6x.

A Division-Free Fixed-Point Matrix Inversion Architecture Using Normalization-Assisted BI-LUT Reciprocal Approximation (433)

Maloth Vishnu Vardhan, E Richard Prajwal Das, Anjan Banerjee, J Sai Maneesh Kumar, Nidhi Verma, P V Phaneendra Kumar and Acharyya Amit

Matrix inversion is a fundamental operation in many signal processing and communication systems, where hardware efficiency is often limited by reciprocal and division operations. This paper presents a division-free fixed-point matrix inversion architecture based on LDLH decomposition and a normalization-assisted Bipartite Lookup Table (BI-LUT) reciprocal unit. The proposed reciprocal engine replaces iterative Newton–Raphson refinement with direct table-based approximation, enabling fixed-latency operation without DSP-intensive divider hardware. Integrated within a three-stage streaming architecture comprising LDLH decomposition, lower-triangular inversion, and matrix reconstruction, the design supports scalable matrix inversion with continuous throughput. FPGA implementation on a Xilinx ZCU104 achieves a mean relative reciprocal error of 3.70×10^{-4} while reducing LUT utilization by 99.0% and eliminating DSP usage compared with a four-iteration pipelined Newton–Raphson reciprocal architecture. The complete design operates at 333 MHz and is suitable for high-throughput fixed-point matrix inversion applications.

HAMLET: Machine Learning-Driven Placement over Heterogeneous DRAM/NVM Memory Systems (291)
Manolis Katsaragakis, Giannis Rokomos and Dimitrios Soudris

The growing demand for larger memory capacity, combined with the need for cost efficiency, performance, and sustainability, has accelerated the adoption of heterogeneous memory hierarchies that integrate DRAM and Non-Volatile Memory (NVM) technologies. While such architectures provide a promising balance between speed, persistence, and scalability, they also introduce new challenges in workload management and data placement, as applications exhibit diverse and highly dynamic memory access behaviors. In this work, we propose a novel Machine Learning (ML)-driven methodology for workload classification and intelligent task placement in heterogeneous memory systems. Based on this methodology, we introduce HAMLET, a dynamic, class-aware placement policy that optimizes task allocation across DRAM/NVM hierarchies. Our evaluation on real DRAM/NVM hardware demonstrates that HAMLET consistently outperforms state-of-the-art techniques, delivering 4.1x average improvement in execution time while also reducing NVM write operations by 78.2%. In addition, These reductions in write intensity not only improve performance but also reduce energy consumption by 58% on average and extend the effective lifetime of NVM devices, highlighting the practical and sustainable benefits of our approach.

Arapaima-LIF: An Oxygen-Regulated Spiking Neuron for Temporal Memory Retention (308)
Anupama Santhosh, Sumit Datta, Xinming Shi and Alex James

Spiking Neural Networks (SNNs) are attractive for neuromorphic computing due to their event-driven operation and low power properties. The standard Leaky Integrate-and-Fire (LIF) neurons, on the other hand, have limited temporal memory, and the Adaptive LIF (ALIF) neurons mainly depend on the spike-history-dependent threshold adaptation, which is only suitable for short-range temporal retention. Inspired by the oxygen regulation mechanisms in biology, we propose a bio-inspired spiking neuron model, Arapaima-LIF, which features an internal state that is regulated by oxygen and modulates neuronal excitability. The proposed neuron has oxygen-modulated membrane integration and burden-based adaptive threshold dynamics for resource-aware temporal processing and long-term memory retention. To assess long-term temporal learning, Arapaima-LIF is evaluated on a delay-cue benchmark with delays ranging from 20 to 100 timesteps and compared to LIF, ALIF and Oxygen-LIF neurons. Experimental results show that competing models degrade significantly beyond a delay of 60 timesteps. Arapaima-LIF achieves accuracies of 95.17%, 93.75% and 76.42% with delays of 60, 80 and 100 timesteps respectively . In addition, oxygen-dependent excitability acts to suppress excess spike activity, favoring sparse event-driven computation. Arapaima-LIF has simple recursive dynamics and low state overhead, making it suitable for next-generation digital, analog and mixed-signal neuromorphic implementations.

Continual Learning by Exploiting Adaptive Anti-Class Distributions (367)
Christos Tsingiropoulos, Dimitrios Katsikas, Nikolaos Passalis and Anastasios Tefas

Class-incremental learning requires models to learn new classes without forgetting old ones, yet sequential updates often distort class geometry and weaken decision boundaries. This paper proposes an adaptive anti-class supervision objective that regularizes inter-class separation using the model's own confusion patterns, applying stronger repulsion to class pairs that are harder to distinguish. Integrated into a dynamically expandable rehearsal-based framework, the method improves discriminative representation learning without architectural changes or additional training overhead. Extensive experiments on CIFAR-100, ImageNet-100, and ImageNet-1000 show strong overall gains over previous state-of-the-art methods, demonstrating that confusion-guided anti-class regularization is an effective mechanism for mitigating catastrophic forgetting in class-incremental continual learning.

EBL: Efficient Broad Learning for Distributed Adaptive Harmonic Analysis (386)
Changhong Li, Georgios Floros, Biswajit Basu and Shreejith Shanker

Renewable energy systems and electrified transport have found widespread adoption in recent years. The integration of these non-linear loads, dominated by electric vehicle (EV) charging, however, has introduced severe harmonic distortion into the power grid, impacting efficiency and lifetime of substation equipments, and switch gear in the distribution network.
Rapid and high-precision harmonic analysis has hence become a prerequisite for effective harmonic control at the source of injection.
This paper proposes an Efficient Broad Learning (EBL) framework for distributed adaptive harmonic estimation.
As the first edge-deployed deep learning harmonic estimation acceleration framework, it offers high accuracy estimation with half-cycle input, reconfigurable flexibility enabled by the FPGA-based implementation, and ultra-low latency, achieving 17.4 $\times$ faster predictions than the current state-of-the-art methods.
In the harmonic prediction across multi-scenario charging and discharging nodes, the online transfer learning based on a closed-form solution rather than back propagation in EBL demonstrates rapid adaptability.
By exploiting bespoke quantisation and sparsity, the approach also reduces resource usage, consuming 5.9\% LUTs on the Zynq Ultrascale+ ZU7EV FPGA, and $\approx$ 82\% lower resources than the state-of-the-art FPGA-accelerated estimator in the literature.

Exploration of a Multilayer Perceptron for Electrical Characterization of Logic Gates (400)
Eduarda Moraes dos Santos and Cristina Meinhardt

Traditional standard cell characterization is hardly dependent of exhaustive electrical simulations across numerous Process, Voltage, and Temperature (PVT) corners. As semi-conductor manufacturing advances deep into the nanoscale, the exponential growth of these corners, driven by process variability, environmental fluctuations, it creates an unsustainable computational bottleneck. To address this challenge and mitigate foundries' reluctance to share confidential transistor model cards, this work explores a lightweight, machine learning-based black-box flow for the electrical characterization of logic gates. Using a 32 nm bulk CMOS inverter cell as a case study, we show the results for a trained Multilayer Perceptron (MLP) architecture to predict critical performance features, specifically propagation delays (tpHL and tpLH) and energy consumption (E), under varying operational parameters and transistor sizing. The proposed MLP model is evaluated against traditional Decision Tree and Random Forest regression models. The experimental results demonstrate that the proposed MLP framework achieves high predictive accuracy, matching the state-of-the-art Random Forest model with a coefficient of determination (R2) of 0.9983 for energy consumption prediction. These findings underscore the potential of neural-network-based black-box models to preserve foundry confidentiality while dramatically accelerating early- stage technology and cell-level electrical evaluations.

S32 | Special Session: Advances in Hardware Security for Electronic Circuits and System

5:00 PM - 6:30 PM
Room: Sirocco
Chairs: George Athanasiou, Boris Vaisband

Reconfigurability Leading to Reliability and Security: A Charge Trap Transistor Based PUF (304)
Fahim Shahriar, Nader Sehatbakhsh and Boris Vaisband

Physically unclonable functions (PUFs) provide hardware intrinsic fingerprints for secure device authentication. Conventional PUFs however, rely solely on fabrication mismatch, resulting in static entropy sources that cannot be modified once deployed. Consequently, if a cryptographic key is compromised, the hardware identity becomes permanently vulnerable. Furthermore, PUFs that are solely based on process variations are not robust under voltage and temperature variations. In this work, a reliable and reconfigurable PUF architecture based on twin charge trap transistors (CTTs) is proposed, enabling post fabrication reliability enhancement and secure identity reconfiguration. CTTs are standard CMOS process MOSFETs that utilize intentional charge trapping within the gate dielectric to enable modulation of the threshold voltage by applying gate pulses. Designed in GF 22 nm FDSOI technology with a supply voltage of 0.8 V, the proposed approach achieves an intra hamming distance of 0.21% and an inter hamming distance of 49%, similar to the state-of-the-art (SOTA). The PUF cell exhibits a compact area of 1.99 μm^2/bit (1.4-93.2X smaller than SOTA) and low energy consumption of 5.47 fJ/bit (2.4-553.5X lower than SOTA), demonstrating that CTTs are useful for low-power, compact, CMOS-compatible PUFs with enhanced reliability and reconfigurability.

Assessing Runtime Electromagnetic Detection of CPU Hardware Trojans Targeting Kernel Memory (218)

Athanasios Moschos, Baki Berkay Yilmaz, Kevin Valakuzhy and Angelos Keromytis

Side-channels are a promising approach for hardware trojan detection, as they can passively reveal malicious hardware activity without requiring additional circuitry or destructive analysis of the device under test.
This work investigates the use of electromagnetic (EM) emanations for runtime detection of hardware trojan attacks.
Using an open-source hardware trojan that implements an arbitrary memory-write primitive, we tamper with the kernel memory of a Linux system running on a RISC-V microarchitecture and attempt a real-time baseline detection using the processor's EM emissions.
Our results indicate that, under certain conditions, hardware trojans can be detected indirectly through the anomalous software behavior they enable.

Machine-Learning-Resilient Deca-Mbit Entropy Physical Unclonable Functions in 65 nm CMOS (314)

Sabrin Alla, Edoardo Cignoni, Alessandro Bedoya, Marco Lanuzza, Sebastiano Strangio, Alessandro Catania and Giuseppe Iannaccone

We present a two-stage physical unclonable function (PUF) fabricated with a 65 nm CMOS process exhibiting an entropy of 47 Mbit and resilience to machine-learning (ML) attacks. The PUF is based on an analog entropy source exploiting threshold-voltage mismatch in large arrays of minimum-size transistors biased in weak inversion. The entropy source serves as the modular element of a multi-stage architecture, in which the stabilized response of a first PUF stage is used as a hidden input for a second stage, thus obscuring the direct relation between challenge and response. Electrical characterization in a broad range of operating conditions shows that BCH correction of five errors enables reliable operation of the two-stage PUF from -35°C to +80°C and in the 1-1.4 V supply voltage range with a false negative rate smaller than 2×10^{-3}.

Side-Channel Evaluation of a Stealthy Cross-Module Hardware Trojan in a Combined ML-KEM and AES FPGA Design (187)

Athanasios Papadimitriou, Nikolaos Athanasios Anagnostopoulos, Mihalis Psarakis and Paris Kitsos

We present a cross-module hardware Trojan architecture targeting a combined ML-KEM-512 key encapsulation and Boolean masking AES-128 encryption system implemented on a Xilinx Artix-7 FPGA. The Trojan is distributed in both the ML-KEM encapsulation logic and the AES accelerator, creating a design which is more stealthy against single-module functional testing and side-channel characterization. Initially, the attacker forces a specific value into the encapsulation key input of the encaps implementation, and a random subset of the key (8, 16, 32, 64 or 128 bits) is used to arm the trigger. Then the following normal encapsulation triggers the payload through the existing key handoff logic that connects the ML-KEM and the AES-128 implementations. Once triggered, the payload effectively cancels the masking to allow retrieval of the secret key by means of a first-order side-channel

attack. Additionally, we perform a sidechannel detection approach using the Welch's t-test against the comparator responsible for arming the trigger of the Trojan, and we compare it to the combinatorial testing approach. Results show that the t-test approach manages to detect the Trojan when arming is performed with 16 bits of the encapsulation key or more, whereas for the 8-bit arming size, combinatorial testing is more effective.

Towards low-power cryptographic circuits in flexible IGZO electronics (210)
Jelle Biesmans, Wim Dehaene, Kris Myny and Nele Mentens

Flexible IGZO electronics are attractive for low-cost, battery-powered applications, but their lack of p-type devices forces the use of unipolar logic styles with always-on pull-up networks. As a result, commonly used standard cells suffer from large static power consumption, making the integration of on-chip cryptographic accelerators challenging.
This paper presents a differential standard-cell topology for IGZO circuits that eliminates the always-on pull-up network. The topology exploits the low threshold voltage of IGZO transistors to limit the voltage loss across each stage, which is then recovered through capacitive boosting. Measurements on ring oscillators show that the proposed logic style shifts the dominant power from static leakage to dynamic switching power. For the same voltage levels, the energy-delay product is improved by up to one order of magnitude compared to two-stage resistive-load logic.
To demonstrate its potential for cryptographic applications, the proposed differential logic is used to implement a circuit based on the Ascon S-box. At low frequency and with fixed inputs, the power consumption is reduced from 221 µW to 2.16 µW at the same operating voltage. These results demonstrate that low-power cryptographic circuits can be implemented in flexible IGZO electronics through the use of unipolar differential logic that is dominated by dynamic rather than static power consumption.

Low-Cost Enrollment and Low-Overhead Authentication of SRAM PUFs on STM32 Microcontrollers (244)
Stefano Gregori, Kamal Y. Kamal and Esther Bamiduro

This paper presents a low-cost framework for enrolling and characterizing SRAM physical unclonable functions (PUFs) on 25 STM32 microcontroller units (MCUs). The setup captures the SRAM2 power-up state before firmware initialization using an STM32 Nucleo-L412KB development board, an Arduino Uno for automated cold-boot cycling, and a Python-based host application for data logging and analysis. During enrollment, the initial SRAM2 power-up data from 25 STM32 MCUs were acquired across 100 cold-boot readouts. The results show an average of 50,830 stable bits per MCU, representing 77.56% of SRAM2 bits. Authentication testing across 1,000 additional SRAM2 readouts further shows that the SRAM2 region can provide both a large stable-bit pool during enrollment and high reliability during authentication.

S33 | Special Session: Sensor Devices, Circuits, and Systems
5:00 PM - 6:30 PM
Room: Ostria
Chairs: Christoforos Theodorou, Hossein Esmailbeygi

A Time-Domain ADC with Constant-Slope Voltage-to-Time Converter and Conditional Correlated Multiple Sampling for CMOS Image Sensors (96)
Jonathan Vrijsen, Wannes Nevens, Thomas Lapauw, Maarten Kuijk, Hans Ingelberts and Thomas Van den Dries

This paper presents a 10-bit column-parallel Time-Domain Analog-to-Digital Converter (TD-ADC) for CMOS Image Sensors (CISs). It implements a Constant-Slope Voltage-to-Time Converter (CS-VTC) and a 4x Conditional Correlated Multiple Sampling (CCMS) scheme to improve the low-light performance of the imager without introducing additional row-time or power-overhead compared to conventional approaches.
The prototype ADC is implemented in a 1P4M 0.35 µm process for a 320x240 pixel array with 15 µm column pitch. Measurements show that the proposed CCMS scheme does not introduce additional Differential Non-Linearity (DNL). The input-referred read noise is decreased by a factor of 1.46 through the CCMS scheme and is measured to be 204 µVrms.

Environment Perception and Multi-Target Detection with a mmW 2D Beamsteering ISAC System (348)
Nebojsa Maletic and Markus Petri

Integrated Sensing and Communications (ISAC) is considered a key feature of next-generation (6G) mobile networks. This technology is expected not only to enhance network efficiency but also to ensure reliable communication, particularly in the millimeter-wave (mmW) frequency bands. This paper presents a mmW ISAC platform featuring two-dimensional (2D) beam steering. The system combines a Xilinx RFSoC ZCU111 board with commercial dual 60 GHz beamforming modules to create a compact ISAC solution. A proprietary software-defined radio (SDR) solution enables the evaluation of many ISAC waveforms. Utilizing an OFDM waveform, a series of experiments is conducted to demonstrate the platform's sensing capabilities regarding environmental perception and object detection in both azimuth and elevation.

Accuracy of Characterization Procedures for Quartz Tuning Fork in QEPAS Sensors as a Function of the Front-End Architecture (58)
Vincenzina Zecchino, Luigi Lombardi, Cristoforo Marzocca, Gianvito Matarrese, Pietro Patimisco, Angelo Sampaolo and Vincenzo Spagnolo

Quartz-Enhanced Photoacoustic Spectroscopy (QEPAS) enables ppb-level trace gas detection by leveraging the high-quality factor resonance of a quartz tuning fork (QTF) as an acoustic sensor. Achieving peak sensitivity requires precise excitation of the QTF at its resonance frequency, thus an accurate characterization procedure must be used to determine this parameter. In this work, we provide an experimental demonstration of the frequency mismatch that exists between conventional electrical characterization and the actual resonance frequency observed during normal sensor operation, which can be obtained using optical

characterization. Furthermore, we show that this frequency shift strongly depends on the front-end readout electronics. To investigate this effect, three dedicated preamplifier architectures, voltage, transimpedance, and charge amplifiers, were physically implemented and tested. Results reveal that the sensor-to-electronics interface shifts the resonance peak by introducing different circuit loading effects, with the charge amplifier providing the minimum mismatch. These results confirm the critical role of preamplifier selection and sensor-circuit co-design in ensuring accurate frequency alignment and optimizing the overall performance of QEPAS sensors.

A 24 GHz FMCW Radar Sensor System with Embedded Real-Time Processing for Cyclist Dooring Prevention (204)

Jonas Gillner, Tristan Thierfelder, Roy Taylor, Leila Sharara, Mohammed Ismail and Klaus Thelen

Cyclist dooring accidents require timely detection of approaching cyclists near parked vehicles. This paper presents the development of a 24 GHz FMCW radar sensor for cyclist dooring prevention. The system integrates a directional 4-by-4 microstrip patch antenna array, RF PCB implementation, and embedded real-time signal processing. The antenna array is implemented on a hybrid Rogers/FR4 PCB stack-up, while the RF front-end performs analog I/Q filtering and high-gain amplification for FMCW beat-signal extraction. Target range and radial velocity are estimated using embedded FFT-based processing with stationary-clutter suppression via reference spectrum subtraction. Cyclist approach experiments evaluate detection performance over distances between 2.5 and 15 m. The system achieves update rates supporting decision latencies below 10 ms. The presented prototype demonstrates the potential of compact FMCW radar sensors for automotive door safety assistance systems.

Optimization of a Cymbal Piezoelectric Transducer Using Finite-Element-Guided ANN and DRL (205)

Mina Nazarian Samani, Javad Pourmohammad, Sima A. Alidokht, Heloise Therien-Aubin and Lihong Zhang

Piezoelectric energy harvesting can convert traffic-induced mechanical energy into electrical energy. This study presents a cymbal-type piezoelectric transducer for asphalt pavement energy harvesting under vehicular loading. The proposed design uses a PZT-5H element bonded between steel end caps to improve strain transfer, protect the piezoelectric layer, and enhance electromechanical conversion under repeated compressive loading. A finite-element model is developed to evaluate the coupled mechanical and electrical response of the embedded transducer. The generated simulation data are then used to train artificial neural network surrogate models for rapid prediction of stored electrical energy and pavement-displacement response. To improve the design process, a proximal policy optimization-based reinforcement learning framework is applied to search for a geometry that increases energy harvesting performance while maintaining pavement compatibility. The optimized design is further fabricated and experimentally tested under vertical compressive loading. The results confirm that the proposed ANN-assisted PPO framework can effectively guide the design of cymbal piezoelectric harvesters and support the development of pavement-integrated energy harvesting systems.

Tuesday November 10th 2026

S06 | Biomedical and Bio-inspired Circuits and Systems 1
08:30 AM - 10:00 AM
Room: Maistros ACD
Chairs: Andreas Demosthenous, Cristina Meinhardt

Analog Haar Wavelet Filter Design for Phonocardiogram Signal Analysis (8)
Julia Nako, Costas Psychalinos and Ahmed Elwakil

In this work a multi-resolution analysis of a Phono-cardiogram (PCG) signal is implemented by employing an analog Haar wavelet. The Laplace transform of the Haar wavelet is approximated using the curve-fitting method with a resulting rational transfer function of fourth-order. The decomposition of the PCG signal is realized for different values of the characteristic time value T of the Haar wavelet, and as the resulting transfer functions from MATLAB have the same form, a programmable design is employed. The results are confirmed by simulation through Cadence IC Design suite using real PhysioNet records, where Operational Transconductance Amplifiers (OTAs) are employed at a multi-feedback structure, offering electronic tuning capability.

A Low Power, Low Noise Inverter-based Amplifier with High Linearity for Ultrasound Imaging (157)
Elham Hatamzadeh, Antonio Aprile, Elisabetta Moisello, Farshad Moradi, Edoardo Bonizzoni, Piero Malcovati and Leonardo Miceli Lara

This paper presents a constraint-driven low-noise amplifier (LNA) for PMUT-based ultrasound receivers. The LNA employs a fully differential inverter-based capacitive-feedback amplifier (CFA) that achieves ultra-low input-referred noise, low power consumption, and high linearity without compromising any of these performance aspects. Implemented in 130-nm CMOS, it achieves an input-referred noise density of 1.84 nV/√Hz at 3 MHz while consuming only 7.76 mW from a 3.3-V supply, showing superior noise–power efficiency. The second and third harmonic distortions (HD2 and HD3) remain below -50 dB at 3 MHz with an output swing of 2.82 Vpp. In addition, the design provides programmable gain settings of 9, 12, and 15 dB, and achieves a power supply rejection ratio (PSRR) of -78.5 dB.

Parkinsonian Activity Detection in Microelectrode Recordings Using a Memristive Oscillator (203)
Nefeli Maria Panagiota Kavouni, Ioannis K. Chatzipaschalis and Georgios Ch. Sirakoulis

Parkinson's disease (PD) is associated with pathological neural synchronization and burst-like activity, whose real-time detection remains challenging for low-power embedded biomedical systems. This work presents a memristor-based neuromorphic interface for circuit-level detection of Parkinsonian burst activity in microelectrode recording (MER) signals. The proposed architecture combines analog signal conditioning with a memristive relaxation oscillator based on the

nonlinear, state-dependent dynamics of a unipolar CBRAM device. Instead of relying on digital preprocessing or handcrafted feature extraction, temporally dense neural activity is mapped into clustered memristive switching events, which are subsequently integrated and thresholded to generate a binary burst-detection output. Transient simulations under synthetic excitation, neural-like stimulation, and clinical MER recordings demonstrate selective sensitivity to burst-related activity. With a peak power consumption of 123.8 μW and a time-averaged power of 0.14 μW, the proposed system provides a compact and energy-efficient circuit-level approach for future closed-loop neurostimulation interfaces.

A Low-Power Analog IC Artificial Neural Network Architecture for Cardiac Survivability Classification (247)
Vassilis Alimisis, Christos Dimas, Alak Majumder, George F. Fragulis and Paul P. Sotiriadis

This work presents a power-efficient, fully integrated analog classifier for real-time post-infarction patient survivability prediction. The architecture employs current-mode Squarer circuits to compute squared feature deviations from trained reference currents, followed by electronically tunable Sigmoid Function Circuits for smooth bounded activation, and a current correlator stage performing harmonic-mean-type aggregation within each Neuron Cell. Cascode current mirrors sum the per-cell outputs into a class score, and a high-precision current comparator renders the final binary decision against a programmable threshold. Implemented in TSMC 65 nm CMOS and verified via Cadence IC Suite with layout, the design achieves 91.65% mean classification accuracy (vs. 92.85% in software) at 843 nW and 350K classifications per second, making it well suited for ultra-low-power, portable biomedical diagnostics.

FPGA Implementation of a Medical Image Watermarking Scheme Based on Region of Interest (104)
Abdenour Kifouche and Mohamed Salah Azzaz

Data protection, authenticity, and confidentiality of medical images are major challenges in online healthcare and telemedicine environments, where data is widely exchanged over networks. In order to guarantee the authenticity of images while preserving their diagnostic quality, this paper presents an FPGA implementation of a blind digital watermarking system. The proposed method embeds Patient Information PI (a binary watermark) into MRI (Magnetic Resonance Imaging) images of the brain. It is based on a predefined mask, a discrete cosine transform, and a spread spectrum. The embedding and extraction components are implemented using Vivado tool for Artix-7 XC7A100T FPGA, resulting in low processing latencies and low resource usage. Hardware experiments, including encryption, embedding, extraction, and decryption, validate the proper functioning of our system. In-depth evaluations demonstrate high imperceptibility with a peak signal-to-noise ratio (PSNR) of 31 dB and structural similarity (SSIM) close to 0.92, indicating higher image similarity. The bit error rate (BER) value is close to 0, indicating less distortion in the watermark. In addition, high robustness is achieved against common signal processing attacks (compression, scaling, filtering, etc.), and an acceptable payload for MRI images with an insertion capacity of 0.004 bpp. These results confirm that the proposed architecture is a well-suited solution for enabling real-time image watermarking on resource-constrained boards.

Circuit Modeling Using Belevitch Canonical Form and Iterative Least-Squares Optimization (30)
Gang Lin

Accurate circuit modeling from measured impedance data is essential for modern high-frequency and semiconductor design. Conventional approaches typically require preselecting a circuit topology and solving a highly nonlinear parameter optimization problem, which can lead to convergence issues and suboptimal solutions. In this paper, we propose a circuit modeling method based on the Belevitch canonical form of the scattering matrix. The proposed approach reduces the reliance on explicit circuit topology specification by parameterizing the network through the Belevitch canonical form. To address the resulting high-order nonlinear optimization problem, we develop an iterative least-squares algorithm inspired by vector fitting, in which the denominator polynomial is updated using a modified weighting and offset scheme. The resulting formulation improves numerical robustness and enables efficient rational model estimation. Validation using synthetic and measured data demonstrates accurate impedance matching and robust convergence.

A Sinusoidal Signal Generator based on Second-Order Analog Interpolation with Widely Tunable Amplitude, Mean Value, and Frequency (133)
Francesco Gagliardi, Margherita Scognamiglio, Simone Contardi, Iacopo Nannipieri, Massimo Piotto and Paolo Bruschi

Recent applications in the biomedical, biochemical, and battery management fields require miniaturized devices able to perform electrical impedance spectroscopy (EIS). In this work, we present a single-chip sinusoidal signal generator (SSG) for EIS, combining direct digital synthesis and the first waveform reconstruction method based on analog second-order interpolation (SOI). Measurement results of a 0.18-µm CMOS prototype highlight wide tunability of the frequency (1-100 kHz), mean value (0.55-2.75 V), and amplitude (0.5-3 Vpp) under modest oversampling ratio, equal to 8, and spurious-free dynamic range of 33.6-36.4 dB. These findings appoint the proposed SOI-SSG as a solution of potential interest for miniaturized EIS devices.

p-Optimal Noise-Shaping Digital Δ-Σ Modulators (178)
Xu Wang and Patrick Hickey

Digital Δ-Σ modulators (DDSM's) are widely used in phase locked loops to control the instantaneous divide ratio to achieve fractional-N frequency synthesis. Compared to conventional Multi-stAge noise-SHaping (MASH) DDSMs, recently advanced Irwin-Hall DDSMs optimally shape the statistical distribution of the quantization error (QE) to achieve maximum order of spur immunity for a given QE range in the presence of nonlinear distortion within the loop—this is called

“p-optimality”. However, the spectral shaping of the QE of the reported IrwinHall DDSMs is typically suboptimal and limited to maximum second order, leading to excess in-band noise contribution. This work introduces a novel p-optimal noise-shaping (POSH) family of Irwin-Hall DDSMs that features p-optimal spur immunity whilst simultaneously applying a programmable (up to 5th) order of spectral shaping for the QE; together, these result in minimized jitter contribution by the DDSM-based fractional-N modulation.

Memristor Models with Logarithmic-Sigmoidal and Piecewise-Linear Self-Limiting State-Flux Relations (193)
Valeri Mladenov and Stoyan Kirilov

Owing to their very low power consumption, nano-sizes and applicability in neural networks and memory high-volume arrays, memristors are very important and perspective electronic components. A significant step in engineering memristor-based circuits is their modeling. In this paper, simple and fast-operating memristor models are proposed. The first one is with a logarithmic-sigmoidal state-flux relation, while the next one is with a piecewise-linear self-limiting characteristics. The proposed models are analyzed in MATLAB and LTSPICE simulators. A comparison of the proposed memristor models with some of the commonly used standard and existing modified memristor models is conducted, paying attention on their simplicity, fast operation, low simulation time and applicability for analysis of complex memristor-based electronic circuits and devices. The present work is a step in memristor’ and memristor device modeling.

A Differential gm/C Oscillator for Mobility-Dependent Frequency Response (413)
Tobias Schwanke

This paper presents a fully integrated 180nm CMOS gm/C differential ring oscillator featuring inherent large-signal amplitude self-stabilization without auxiliary control loops. Validated via silicon measurements, the dual-mode architecture leverages an on-chip biasing network to address two distinct applications. In frequency-reference mode, a proportional-to-absolute-temperature (PTAT) bias counteracts thermal degradation, yielding a residual temperature coefficient of 218 ppm/K. Alternatively, a constant-Vov biasing scheme forces the output frequency to precisely replicate the intrinsic carrier mobility $\mu(T)$. By operating in this secondary mode, the architecture successfully tracks the actual, process-dependent trajectory of the relative mobility change over temperature. This enables the direct on-chip characterization of how process variations modulate the thermal transport dynamics of the carriers, making the device highly relevant for advanced technology diagnostics and variability analysis.

NPU Deployment of DL-based Spectrum Sensing for Edge-AI Cognitive Radio Systems (235)
Andres Rojas and José M. de la Rosa

Spiking Neural Networks (SNNs) have emerged as a promising paradigm for energy-efficient and biologically inspired computation, particularly in edge and embedded systems. However, their hardware implementation poses significant challenges due to strict timing constraints and the need for efficient data movement. This work presents a performance-oriented enhancement of the ReckOn recurrent spiking neural network (RSNN) accelerator deployed on a FPGA-based MPSoCs. We compare two execution paradigms: PYNQ (Python-based), and CYNQ (C++-based runtime); on an AMD Kria K26 SoM. By migrating the training loop to CYNQ and extending its capabilities with additional IP integration and Python bindings, we achieve substantial improvements in execution efficiency while preserving training accuracy, demonstrating up to 232x speedup in DMA transfers, and an overall reduction in training time of 2.57x, with equivalent accuracy across both frameworks. These findings highlight the effectiveness of CYNQ in reducing software overhead and accelerating neuromorphic workloads.

TE-STDF: Gradient-Guided Structural Pruning for Lightweight Video Enhancement (150)
Luis Carlos Rodrigues Linares, Gilberto Kreisler, Daniel Palomino, Bruno Zatt, Attilio Fiandrotti and Guilherme Corrêa

Spatio-Temporal Deformable Fusion (STDF) is a Video Quality Enhancement (VQE) architecture that effectively reduces video compression artifacts but is too computationally demanding for resource-constrained consumer devices. Previous optimization efforts often rely on semi-structured techniques that fragment feature maps, break memory contiguity, and fail to deliver hardware acceleration. To address these limitations, we propose TE-STDF, an accelerated framework compressed through Taylor-series-based importance estimation. The method removes redundant convolutional channels according to their impact on the global loss function, producing an architecture that preserves topological constraints and remains natively compatible with hardware accelerators. Experiments on the MFQEv2 dataset show that TE-STDF reduces the parameter count by 68.27% while decreasing overall ΔPSNR by only 0.06 dB. Measurements on an NVIDIA RTX A6000 GPU demonstrate a 64.1% inference speedup and a 39.0% reduction in per-frame energy consumption.

Spiking Neural Network on Microcontroller for Low-cost Speech Enhancement (151)
Yusuf Serhat Özkan, Tao Sun and Guangzhi Tang

Speech enhancement on wearable and voice-facing devices increasingly has to run on small battery-powered microcontrollers rather than in the cloud. Spiking neural networks (SNNs) are an attractive candidate: they process temporal audio with

sparse spike events and a compact recurrent state, but their reported efficiency is almost always measured on dedicated neuromorphic silicon, and whether a time-domain spiking speech model can run on a low-cost MCU is largely untested. In this work, we deploy a spiking convolutional network (SCNN) end-to-end onto an MCU (STM32U5 with ARM Cortex-M33). We lift every piece of cross-frame state out of the stateless inference graph as explicit recurrent tensors and reformulate the model as a single-frame streaming pass that carries this state from frame to frame. The deployed network reaches 16.70 dB SI-SNR on VoiceBank-DEMAND and runs at 2.715 ms per frame. The result shows that an SNN can serve as a practical temporal processing block for low-cost embedded AI.

ActiPercept: Confidence-Guided Active Perception for Robotic Screw Inspection (260)
Pavlos Tosidis, Nikolaos Passalis and Anastasios Tefas

Object detectors used in robotic inspection are typically applied from fixed or externally defined viewpoints, which can be restrictive for small objects whose confidence depends on viewpoint, scale, occlusion, and local image resolution. This paper proposes \emph{ActiPercept}, a lightweight active perception framework for confidence-guided screw/noscrew verification. The method closes the loop between perception and sensing by transforming detector outputs into deterministic, bounded camera-control actions: detections are assigned to selected inspection regions, a stable target is selected using holder and screw/noscrew detections, the target is centered in the image, and the camera approaches it until the selected-region confidence exceeds a threshold. The rule-based formulation is interpretable, safety-aware, and does not require policy training, reward engineering, or action annotations. In Webots simulation, ActiPercept achieves 100.0% verified success, i.e., correct selected-region screw/noscrew predictions above the confidence threshold before timeout, in nominal conditions, and 96.7% verified success across six illumination levels, with 0.0% false verification. These results substantially improve over static inference, random local search, multi-view voting, and a no-cache ablation.

Keyword Spotting with Out-of-Vocabulary Rejection for Wearable IoT Devices (429)
Yuying Zhuo, Michael Grzybek, Petros Spachos and Stefano Gregori

This paper presents a voice command interface designed for resource-constrained wearable Internet of Things (IoT) devices. The proposed system performs all processing locally on the ESP32-DevKitC V4 platform using dynamic time warping (DTW) for speech recognition. Building upon previous work in our laboratory, this study systematically investigates how DTW rejection thresholds influence recognition accuracy and F1 score of the system. The DTW thresholds are employed to reject out-of-vocabulary words and ambient noise, enhancing its robustness in noisy environments. The proposed system achieves a total end-to-end Flash footprint of approximately 430 KB and a peak RAM consumption of 34 KB, integrating the full pipeline from raw acoustic acquisition to keyword classification. Its reliable performance and low processing latency makes it highly suitable for battery-powered wearable applications.

Range Utilization and SQNR Decomposition for Dithered Quantizers Below 4 Bits (59)
Brian Holford

Below 4 bits, dither headroom consumes more than 25% of the hardware input range, yet no closed-form budget exists for the resulting signal-to-quantization-noise ratio (SQNR) penalty. We decompose the effective SQNR into three terms, hardware SQNR, utilization loss η, and noise floor increase, revealing that at N = 2 the combined dither cost is 5.5 dB (RPDF) or 10.8 dB (TPDF), and that the two standard step-size conventions diverge by 3.5–6.0 dB because they clip at different boundaries. Monte Carlo simulations validate all predictions for N = 1–8.

Generalization of the Gamma/Gammatone Filter Transfer Functions (152)
Brent Maundy, Ahmed Elwakil, Costas Psychalinos and Anis Allagui

A study of the behavior of filters derived from the Gamma kernel function is performed in this work. A generalized filter function derived by employing the Mittag-Leffler function is initially derived and some special cases of filters are presented including low-pass, band-pass, and Gammatone filters. The approximation transfer functions are obtained through an efficient approximation of the fractional-order Laplacian operator. The findings of this work are supported by simulation results.

Using a ring buffer instead of a FIFO in a FIR-DAC significantly improves linearity but increases noise (251)
Hanspeter Schmid, Alex Huber, Dominik Hiltbrunner and Mathieu Coustans

Finite-impulse-response digital-to-analog converters (FIR-DACs) can be built with a number of unit current sources (in our example: 128) that are switched such that their sum averages a digital sigma-delta modulated (SDM) bitstream.
In the literature, the digital control of the current sources is always done using a first-in first-out (FIFO) register: every time a new SDM bit arrives, all bits are shifted. It can happen that all 128 current sources switch.
We propose using a ring-buffer instead: every time a new SDM bit arrives, it overwrites the control bit of one current source. Therefore, at most one current switches.
In this paper we show by simulation and measurements of two chips fabricated in 0.18um CMOS that using a ring buffer gives much better linearity at the expense of more noise, leading to a 3dB SNR increase in one application.
We also discuss the sources of non-linearity and noise: non-linearity comes from non-ideal settling (hence less switching activity gives less non-linearity), whereas additional noise in the ring-buffer case comes from the mismatch of the current sources mixing some of the out-of-band SDM noise into the signal band.

A Pareto-Optimized Fused Posit Radix-2 FFT Architecture for FPGA and its validation in FMCW Radar Signal Processing (275)
G V N Mokshagna, Kandagatla Venkata Rohith Sai and Madhav Rao

This paper presents a fully parameterized Fused Posit Radix-2 FFT architecture for FPGA, targeting the elim-nation of redundant intermediate encode–decode operations that afflict conventional Posit butterfly datapaths. By operating entirely in the raw unpacked domain and applying Round-to-Nearest-Even (RNE) rounding only at the final output boundary, the proposed design reduces the Codec count from 20 decoders and 10 encoders to 6 decoders and 4 encoders per butterfly. Pareto optimization over Normalized RMSE and Normalized Power for FFT architecture, identifies Posit(10,1) as the optimal configuration. Implemented on a Xilinx Zynq UltraScale+ MPSoC FPGA, the fused architecture achieves an 8.8% LUT savings over IEEE FP16 and a 37.5% reduction over the non- fused Posit baseline, with a 41.4% improvement in maximum operating frequency, while maintaining less than 2% difference in on-chip power. At equivalent bit-width, Posit(16,1) outperforms FP16 by +6.65 dB SNR on real-world FMCW radar signals, confirming superior numerical fidelity. The hardware design files are made freely available for further usage to the designers and researchers' community.

SCARF-DoA: Spectral Coupling ICA and Adaptive Refinement Projection Framework for Coherent and Underdetermined DoA Estimation (399)
Rohan Purkait, Souris Sahu, Anjan Banerjee, Nidhi Verma, Amit Kumar Gupta, Sounak Samanta and Acharyya Amit

Accurate Direction-of-Arrival (DoA) estimation is critical for radar, sonar, and acoustic sensing applications such as source localisation and target tracking. Conventional subspace methods like MUSIC and ESPRIT often fail under coherent-source conditions and in underdetermined scenarios, where the number of sources exceeds the available sensors. This work presents a six-stage co-array DoA estimation framework combining Spectral Coupling Independent Component Analysis (SC-ICA) and Adaptive Fractional Projection Refinement (AFPR) to address these challenges. The proposed method resolves 8 sources using only 6 sensors, surpassing the classical M-1 limit by 60%, even with coherent source pairs. SC-ICA enhances source separation via eigenspace decorrelation and AFPR adjusts the noise subspace with adaptive Stiefel-manifold projection. Monte-Carlo simulations with 1000 trials show that the resolution probability is improved and the RMSE is reduced for different SNR levels as compared to MUSIC and ESPRIT. Hardware validation on the Xilinx Zynq UltraScale+ ZCU104 reveals complete alignment with software simulation findings, demonstrating the framework's practical suitability for embedded DoA systems.

Smart-IMPLY Validation via Custom Integrated Circuits and Hardware-Aware System Simulations (132)
Lorenzo Benatti, Tommaso Zanotti, Max Uhlmann and Francesco Maria Puglisi

Among the Logic in Memory (LiM) approaches enabled by memristors, the material implication (IMPLY) logic scheme is one of the most explored. Although the traditional implementation suffers from several design challenges and reliability issues, a recent Smart-IMPLY (SIMPLY) solution has been proposed to overcome these limitations while improving energy efficiency. While the SIMPLY architecture has been demonstrated in simulation or prototypical works, here a first integrated solution is presented. Measurement results for memristors switching and operational circuitry are presented and then adopted for hardware-aware simulations, demonstrating the validity and scalability of the scheme.

Analog Area-Type Switching in Ultrathin HfO2 Memristors for Neuromorphic Circuits (154)
Cuo Wu, Junjie Zhao, Zipeng Mi, Thomas Mikolajick and Benjamin Max

Analog neuromorphic circuits require synaptic devices that combine continuously tunable conductance with low-disturb readout and predictable pulse-based weight update. However, many area-type memristive devices suffer from parasitic charging, zero-crossing shift, and history-dependent drift, which directly degrades analog weight accuracy. In this work, we investigate ultrathin TiN/HfO2/Pt memristors as compact analog synaptic elements for circuit-level neuromorphic applications. The 5.08 nm HfO2 devices exhibit gradual area-type switching with negligible zero-crossing shift, no pronounced rectification, and highly repeatable conductance modulation. Low-bias measurements reveal a stable read window below ±0.5 V, while pulse-width-dependent programming shows controllable potentiation and depression with an onset around 8 µs. Temperature-dependent and electrical analyses are consistent with TAT-dominated steady-state transport, while gradual switching may originate from field-driven changes in trap occupancy. These results suggest that ultrathin HfO2 memristors provide a circuit-compatible analog weight element for neuromorphic and in-memory computing systems.

The impact of memristor I-V asymmetry on charge integration in pseudo-memcapacitive neurotransistors (199)
Hampus Hoffman, Eter Mgeladze, Thomas Mikolajick, Richard Schroedter, Ronald Tetzlaff and Stefan Slesazeck

In this work, a pseudo-memcapacitive neurotransistor is experimentally realized using a nonvolatile and analog Ti/Al2O3/NbOx/Ti resistive random-access memory (RRAM) device connected to the gate of a transistor. The neurotransistor acts as a leaky integrate-and-fire artificial neuron input stage with the synaptic element already part of the device, and we show how the analog nature of the RRAM device translates to an analog pseudo-memcapacitive state change that can be used to

gradually tune the charge integration behavior. Additionally, through a calibrated compact model of the neurotransistor circuit implemented in LTspice, we investigate how the asymmetry in the I-V characteristics of the memristor affects the charge integration. We reveal that the conductance in the low-resistive and high-resistive states of the RRAM independently impact the time constants for charging and discharging. These results highlight the potential of the neurotransistor as a flexible implementation of an artifical neuron where co-design with the RRAM technology allows for a wide range of neurotransistor characteristics.

SPICE-Compatible Model for Skyrmionic Devices (229)
Georgios Chararas, Aristotelis Tsekouras, Nicolas Reyren, Vincent Cros, Christoforos Moutafis and Vasilis Pavlidis

Recent strides in the fabrication of skyrmionic devices have enabled realistic implementations of deterministic control over magnetic skyrmions. These magnetic quasiparticles can now be nucleated, transported, and deleted in a controlled manner, facilitating the deployment of skyrmionic devices in a broad spectrum of applications, including memory, logic, and neuromorphic and edge computing. To date, micromagnetic simulation has been extensively employed to predict skyrmion behavior. Although highly accurate, this approach is computationally expensive and hinders adequate large-scale device and circuit level exploration. Since these devices should be integrated with CMOS technologies, there is a need to co-simulate skyrmion-CMOS circuits with standard commercial circuit simulators. A compact model is presented for skyrmion generation, transport, and detection. The model is calibrated against reported skyrmionic device measurement data. Emphasis is placed on the stochastic and thermally assisted skyrmion nucleation process. Efficient simulation of skyrmion-based devices is thereby enabled while preserving the key stochastic and experimentally observed features of skyrmionic devices. Results demonstrate several orders of magnitude reduction in simulation time compared to micromagnetic simulation, while the error remains below 8%.

Topology and Device-Dependent Synaptic Behavior in Memristor Bridges (416)
Dimitrios Prousalis, Shamoil Ejaz, Ioannis Messaris, Vasileios Ntinas, Ahmet Samil Demirkol and Ronald Tetzlaff

Memristor bridge synapses provide a compact cir- cuit element for realizing bipolar synaptic weights in neural networks. While early studies relied on idealized memristor mod- els, practical devices follow physics-based dynamics with thresh- old behavior, asymmetric SET/RESET kinetics, and topology- dependent state evolution. In this work, we present a circuit- level study of the two-memristor–two-resistor (2M–2R) bridge, focusing on how its synaptic weight can be written, read, and tuned when the memristive devices follow such physics-based dynamics. Using SPICE simulations of the JART Valence- Change Mechanism (VCM) and Stanford RRAM models across the four 2M–2R configurations, we characterize the transfer characteristic, identify writing and reading regimes, and quantify the role of design parameters such as series resistance and pulse width. The results show that topology and device physics jointly shape the synaptic behavior.

A 10-b 56-MS/s Calibration-Free SAR ADC Array for Touch Sensing Applications (127)
Chin-Wei Cheng

This paper presents a 10-bit 56-MS/s dual channel SAR ADC array in 90-nm CMOS for capacitive touch sensing. To meet the demands for multi-touch accuracy and rapid response times, a calibration-free architecture is proposed. The design integrates a subranging SAR structure to reduce switching energy and mitigate the impact of DAC capacitor mismatch. To enhance linearity and achieve offset cancellation, the system employs Residue Oversampling (ROSR) combined with a comparator input-swapping technique. Additionally, a reference swapping mechanism intrinsically eliminates inter-channel gain mismatch. At a 56-MS/s sampling rate, with 2-MHz input frequency, the measured peak signal-to-noise-and-distortion ratio (SNDR) and spurious-free dynamic range (SFDR) are 56.83 dB and 68.88 dB, respectively. Measurement results show that the channel-to-channel static offset match with a difference of only 0.26 LSB. Furthermore, with the reference swapping mechanism enabled, the average channel-to-channel gain error mismatch is successfully reduced from 3.51 LSB to 1.73 LSB. The total power consumption of the dual channels is 4.34 mW under a 1.1 V supply.

A 5 GS/s 44 dB Burst-Mode ADC for In-Situ Monitoring of Wide-Bandgap Power Devices (128)
Tom Herzog, Erik Wehr, Kenny Vohl, Ralf Wunderlich and Stefan Heinen

This work presents a 5 GS/s burst-mode analog-to-digital converter (ADC), co-integrated with an adaptive gate driver, that enables high-speed in-situ monitoring of wide- bandgap (WBG) power transistors. The system is implemented in a 180 nm high-voltage silicon-on-insulator (HV-SOI) technology and occupies a total area of 1.11 mm2 on a 10 mm2 gate driver IC. It uses a burst-mode architecture to achieve a maximum sampling rate of 5 GS/s despite the modest CMOS node size. Measurement results demonstrate a best-in-class 5 GS/s sampling rate for this architecture and a peak signal-to-noise-and-distortion ratio (SNDR) of 44.41 dB, the highest reported to date.

A Preamplifier-Free Sigma-Delta ADC for fA-Level Current Digitization Using Logarithmic Current Splitting (190)
Kerem Yurtseven, Uygar Yildiz, Kemal Ozanoglu and Günhan Dündar

Precision ionization chamber radiometry demands fA-to-pA current digitization from a high-impedance near-DC source, a regime where conventional resistive TIAs and hard-reset capacitive integrators fall short in leakage suppression, noise, and integration density. This work presents a first-order current-mode preamplifier-free Sigma-Delta ADC in 180 nm CMOS utilizing a leakage-compensated integrator with split-switch resetting and a current DAC with logarithmic current splitting. Transistor-level simulations demonstrate 11.09 ENOB over a ±1 pA input range with

sub-fA INL while consuming 2.26 mW average power. Results demonstrate robust operation across process, voltage, and temperature corners, validating the proposed architecture as a compact solution for precision fA-level sensing.

Compensating Excess Loop Delay in CT Incremental Delta-Sigma ADCs using Extended Counting (332)

Elias Ruopp, John Kauffman and Maurits Ortmanns

Continuous-time (CT) I-DS modulators (I-DSMs) offer the unique potential to realize high-resolution analog-to-digital converters (ADCs) operating at Nyquist-rate with a CT input. Among the reported techniques to enhance the resolution of I-DSMs is a two-step conversion approach referred to as extended counting (EC) or extended range (ER). However, state of the art (SoA) publications have not discussed the EC performance degradation due to excess loop-delay (ELD) and its necessary compensation, which is a known non-ideality in CT I-DSMs. This paper presents an analysis of the hitherto unreported systematic performance loss of EC CT I-DSMs in the presence of ELD, and it proposes an approach to partially compensate for it.

A Method for Designing Decimation Filters for Noise-Shaping SAR ADCs (346)

Stan Oosterlynck, Marco Bolognesi, Achim Vandierendonck, Caro Meysmans, Hendrik Meersseman, Brecht Vanlaar and Guy Torfs

This paper presents a method for designing decimation filters for noise-shaping (NS) SAR ADCs. A novel equalization filter at Nyquist rate based on an improved Modulated-Sinc-Sum Filter (MSSF) is proposed that has a linear phase response and is fully defined by the preceding decimation filters, leading to simple calculation of the coefficients. Together with a 10th order Cascaded-Integrater-Comb (CIC) and a 47 tap halfband filter (HBF), the example decimation filter achieves more than 80dB stop-band attenuation and less than ±0.05dB in-band ripple for a 19 tap MSSF. It relaxes the requirements on the analog anti-alias filter by greatly attenuating the close out of band blockers that may reside between the Nyquist and sampling rate of oversampling ADCs. Measurement data from a 4th order NS SAR ADC that was designed and fabricated in 180nm is used to validate the filter designs. The ADC achieves 70.5dB SNDR for a bandwidth of 125kHz and power consumption of 1.33mW, leading to a figure of merit of 150.2dB.

RISC-V controlled flexible digital front-end for communications (88)
Aleksi Korsman, Otto Simola, Andrei Spelman, Verneri Hirvonen, Jussi Ryynänen and Marko Kosunen

As wireless communication systems evolve toward 5G and beyond, transceiver front-ends must deliver both high throughput and flexibility to support dynamically changing standards and spectrum environments. This paper presents a RISC-V–controlled digital down/up-converter (DDUC) accelerator that reuses a single sampling rate converter and CORDIC for both transmit and receive operation modes that is configured at run-time with a RISC-V processor through a software-controllable crossbar. This structure allows the same silicon to be allocated dynamically, supporting workload adaptation under changing uplink/downlink demand. The system was implemented in 22 nm FD-SOI technology and verified with QAM-64 test signals. Measurement results demonstrate accurate signal conversion with an error vector magnitude (EVM) well below 1% with stopband attenuation of 80 dB while achieving a sample rate of 560 MS/s and total system power consumption of 133 mW. The proposed design achieves the efficiency of dedicated hardware with software-like flexibility, offering a scalable foundation for next-generation hardware-accelerated SDR systems.

Hardware Architecture for MPEG-5 EVC Intra Prediction Supporting Multiple Coding Unit Sizes (161)
Marco Jefferson Mamani Arapa and Claudio Machado Diniz

With the emergence of the MPEG-5 Essential Video Coding (EVC) standard, a new trade-off between compression efficiecy and licensing transparency has been established. While the EVC Baseline profile is royalty-free, the Main profile introduces advanced tools that significantly enhance coding efficiency at the cost of increased computational complexity. This paper presents a high-performance hardware implementation for the Intra Prediction module of the EVC Main profile. The design was synthesized for GPDK 45nm CMOS technology and the experimental results demonstrate that our implementation is able to encode Full HD videos at 30 frames per second running at 106.3 MHz with a power dissipation of 129 mW. To the best of the author's knowledge, this is the first work in the literature to report a hardware implementation for the MPEG-5 EVC Main profile Intra Prediction.

A Resource-Efficient FPGA Architecture for Numerically Stable Streaming Hermitian Covariance Estimation Using Welford's Online Algorithm (296)
Anjan Banerjee, Rohan Purkait, Souris Sahu, Nithin Sai Reddy, Aryaan Mandal, A. Prasanna Laxmi, Sounak Samanta and Amit Acharyya

High-throughput edge signal processing is frequently bottlenecked by real-time Hermitian covariance matrix estimation. Traditional batch-based FPGA designs rely on exten-sive snapshot buffering, which incurs heavy memory overhead and exposes fixed-point arithmetic to catastrophic cancellation. We eliminate these bottlenecks by implementing a reconfigurable hardware pipeline that executes Welford's online algorithm for single-pass, streaming covariance updates. Fabricated on a Xilinx Zynq UltraScale+ ZCU104 SoC using Q20.12 precision, the architecture integrates a 45-lane Gauss three-multiplier core that cuts DSP48E2 consumption by 25%, alongside a packed upper- triangular storage scheme that drops BRAM utilization by 44% for a 9×9 matrix. Operating at 150 MHz, the accelerator delivers a continuous throughput of 14.42 MVec/s and a deterministic 63.7 µs pipeline latency at just 0.328 W dynamic power. Board- level testing against double-precision software references vali- dates that element-wise errors remain strictly within theoretical quantization limits, rendering this core an ideal, power-efficient front-end for next-generation adaptive array systems.

Core-Lock Scheduling for 5G NR LDPC Decoders (406)
Nikolaos Papageorgiou and Vassilis Paliouras

Accelerating iterative LDPC decoding convergence while preserving error-correction reliability is a critical challenge in latency-constrained 5G New Radio (NR) and non-terrestrial networks. This paper proposes a hardware-efficient scheduling mechanism based on a localized core-lock criterion that monitors sub-iteration convergence. By evaluating the syndrome weight of only the systematic core sub-matrix, the mechanism selectively bypasses extension-layer updates. With an optimal threshold of θ=0.20, the proposed method reduces decoding latency by up to 31.3% when integrated with advanced static schedules, and by up to 45.5% compared to alternative dynamic schemes utilizing the full parity-check matrix, without block error rate (BLER) degradation. Crucially, compared to full-matrix check alternatives, the core-only approach yields a 76% reduction in shifter network area and a 73% reduction in XOR gate count. For decoders with existing early termination, the propose mechanism requires only a compact adder tree and a comparator, demanding an area of only 3,539 µm2 in 28-nm technology.

Dynamic Frequency Scaling Evaluation for RTOS in RP2040 (342)
Ioannis Sofianidis, Vasileios Konstantakos and Spyridon Nikolaidis

Energy efficiency remains a challenge in the proliferation of battery-powered Internet of Things devices, where embedded microcontrollers must sustain operation for months or even years without intervention. Dynamic frequency scaling appears promising for energy reduction, yet its practical benefit is tied to the nature of the executed workload, a trade-off that becomes particularly subtle in real-time operating system (RTOS) environments where tasks interleave and block. This work investigates whether the complexity of active frequency scaling can be bypassed by a single, benchmarked fixed frequency, without sacrificing energy efficiency. Using an RP2040 microcontroller running FreeRTOS, we compare both

strategies across a representative mixed workload comprising fixed-duration audio acquisition, computational signal processing (FIR filtering and FFT), and Wi-Fi communication with a typical blocking delay. Our experimental findings reveal that the optimal policy is not universal: while dynamic scaling proves advantageous for short and long processing bursts, a fixed frequency can outperform it in intermediate regimes with a narrow region where a fixed frequency proves competitive, the extent of which depends on processing duration. Furthermore, we demonstrate that RTOS-driven task switching during communication delays yields energy savings over bare-metal busy-waiting.

S18 | AI-driven Circuit Design 1 **11:30 AM - 1:00 PM**
Room: Zephyros
Chairs: Dimitrios Prousalis, José M. de la Rosa

BABEL: A Bi-fidelity Automatic Bayesian Exploration Loop for Variation-Aware Analog Integrated Circuit Sizing (83)
Takayoshi Namura and Nobukazu Takai

Developing automated sizing methods that simultaneously address manufacturing variations and performance trade-offs is crucial for efficient analog integrated circuit design. However, existing techniques typically treat these challenges separately and lack a unified optimization framework. This paper proposes an efficient automated sizing methodology based on multi-fidelity, multi-objective Bayesian optimization to concurrently handle variations and trade-offs within a single framework. Experimental results on an operational amplifier design demonstrate that the proposed method reduces execution time by approximately 40% compared to standard Bayesian optimization in a TSMC 180 nm process, while achieving superior Pareto fronts across multiple process nodes.

Target-Driven Deep RL for Reusable Sizing of Constant gm Biasing Circuit Under Corner Variations (90)
Eleni Papageorgiou, Valentina Davidoiu, Andi Buzo, Georg Pelz and Thomas Noulis

Automated sizing of biasing circuits remains a bottleneck in EDA due to high-dimensional design spaces and PVT robustness requirements. While current RL-based methods focus on satisfying performance thresholds for amplifiers, they often fail to achieve the high-precision convergence required for biasing. This paper proposes a target-driven Deep RL framework for the optimization of flattened biasing circuits. Our approach introduces a composite reward function that balances accuracy, noise, and stability across process corners. Notably, the agent demonstrates generalization, enabling adaptation to new voltage targets without model retraining. Validated on a circuit with 27 design parameters, the framework significantly reduces the manual "design-and-simulate" loop. Results indicate that our method outperforms traditional meta-heuristics, such as Differential Evolution and PSO, by providing superior convergence precision and reduced computational overhead during the inference phase.

Sensitivity-Guided Differential Evolution for Simulation-Based Analog Circuit Sizing (92)
Eleni Papageorgiou, Savvas Karipidis, Valentina Davidoiu, Andi Buzo, Georg Pelz and Thomas Noulis

This paper introduces Sensitivity-Guided Differential Evolution (SG-DE), an adaptive optimization framework for simulation-based analog circuit sizing. The proposed method enhances the conventional DE algorithm with a per-dimension mutation scaling mechanism driven by Spearman rank-correlation sensitivity estimates, derived from a violation-weighted population archive. To maintain computational efficiency, an event-driven stagnation detector invokes sensitivity recomputations only when optimization progress stalls, eliminating the overhead associated with fixed-interval updates. The algorithm was evaluated on a low-dimensional operational amplifier and a production-level TSMC 65nm bandgap reference circuit featuring 22 design parameters and nine simultaneous performance constraints. Experimental results indicate that while standard DE remains efficient for small-scale designs, SG-DE provides superior convergence stability and faster constraint satisfaction in high-dimensional search spaces. Notably, the sensitivity-aware mutation strategy achieves these improvements without requiring surrogate model training, gradient information, or additional simulations beyond the standard DE budget.

Optimization-Driven Active Learning for High-Yield Dataset Generation in ML-Assisted Design of Sigma-Delta Modulators (223)
Francisco J. Gómez-Pulido and José M. de La Rosa

This paper presents an optimization-driven approach for the generation of high-yield datasets to enhance machine-learning-based automation and system-level optimization of Continuous-Time Sigma-Delta Modulators (CT-SDMs).
The proposed methodology integrates a continuous surrogate-retraining strategy with efficient behavioral simulations to dynamically explore the complex, multi-dimensional design space.
This framework is applied to various Active-RC CT-SDM architectures, systematically mapping both optimal performance basins and physical instability boundaries across different loop-filter orders and topologies. Extensive simulation results demonstrate the efficiency of the proposed method in resolving the data sparsity bottleneck inherent in traditional sampling methods.
Specifically, the framework autonomously generates novel, validated configurations, demonstrating a significant enhancement in the general metrics obtained across different topologies.
Compared with prior approaches, this work enables the high-throughput generation of maximum-performance data, providing a robust foundation for its application in AI-driven synthesis of analog circuits.

SeG-MCTS: Sensitivity-Guided Monte Carlo Tree Search for Device Selection in Analog Circuits (340)
Yuto Moriguchi and Nobukazu Takai

Automated analog circuit design methods have mainly focused on circuit sizing with fixed device types. However, device selection from a process design kit is a categorical optimization problem with discontinuous circuit responses and expensive evaluations. This paper proposes SeG-MCTS, a sensitivity-guided Monte Carlo tree search method for analog circuit device selection. Sobol-based global sensitivity analysis estimates variable sensitivities and higher-order interactions from a surrogate model, which are used to determine the decision order, group interacting variables, and prune device candidates. The method was evaluated on a BGR circuit designed in a TSMC180nm CMOS process with 186,624 device configurations. Over 10 random seeds, the selected configurations ranked within the top 0.02% on average. Compared with categorical Bayesian optimization, SeG-MCTS achieved better search performance while reducing the mean algorithmic runtime from 510 to 0.5 min, corresponding to a 1,020$\times$ speedup.

S20 | Energy Harvesting and Power Management 1 11:30 AM - 1:00 PM

Room: Sirocco
Chairs: Vasileios Alimisis, Spyridon

The energy-efficient techniques for high voltage generation of NAND Flash storage using a Fully Integrated Voltage Regulator (FIVR) (78)
Minsoon Hwang, Mingi Jeong, Hyunwoo Kim, Sanguk Lee, Jaekwan Kwon, Jongwook Kim, Youngjin Park, Sejin Park, Yoosang Na, Hwang Yoon, Seungyoung Ahn and Dongsop Lee

NAND stack heights increase in high-performance SSDs and mobile storage, improving Word Line (WL) rising time, which is critical for performance. To avoid increasing internal pump size, which constrains die area, this paper proposes an On-Package Fully Integrated Voltage Regulator (FIVR) for mobile devices (UFS/eMMC) that lack external VPP pins. Fabricated using a 55nm PMIC process with a magnetic core inductor, the FIVR converts a 2.5V input into a 12V VPP. It delivers over 50 mA with 78.4% power efficiency and 80.4% with an external inductor, significantly higher than the 40–55% efficiency of traditional internal charge pumps. When integrated into a UFS package, the FIVR reduces read time (tR) by 5 us and program time (tPROG) by 30 us. Consequently, energy efficiency improves by 11% for reads and 16% for writes (pJ/bit). The embedded silicon inductor can make the UFS PKG less than 1mm. By employing Peak Power Management (PPM) to minimize inductor current requirements, this approach substantially enhances the PPAC (Performance, Power, Area, and Cost) competitiveness of mobile storage solutions.

A 38 nA Quiescent Current, 90 nS Fast Transient Response, Capacitorless Low Dropout Regulator Based on a Class-AB Buffer With Current Limiter and Active Capacitor (284)
Jiayang Li, Dai Jiang and Andreas Demosthenous

This paper presents a 38 nA, ultra-low quiescent current, capacitorless low dropout regulator (LDO) based on a class-AB buffer, fabricated in a 130 nm CMOS technology. An active capacitor structure is introduced to improve the overall response speed and reduce the voltage ripples during load transitions. Combined with adaptive and dynamic bias techniques, the LDO achieves a measured maximum response time of 90 ns and a maximum ripple of 210 mV for load current transition from 0 to 20 mA. Additionally, a current limiter is introduced to set an upper limit for the quiescent current of the class-AB output stage, enabling low power operation while supporting up to 20 mA load current.

A 14-mV-Input, 92.9%-Peak-Efficiency Single-Inductor Dual-Output DC–DC Converter for Energy Harvesting Applications (299)
Elham Hatamzadeh, Hossein Esmailbeygi and Farshad Moradi

This paper presents a single-inductor dual-output DC–DC converter for MTJ-based RF energy harvesting, operating in discontinuous conduction mode (DCM). To address the limited input power harvested by MTJ devices, the converter employs a selectable MTJ input array, along with loss-reduction circuit techniques, to enhance input power and enable efficient voltage boosting. The converter is implemented in a 180 nm CMOS technology. Post-layout simulation results show that the converter sustains operation at an input voltage as low as 14 mV. In addition, it achieves a peak efficiency of 92.9% and maintains efficiency above 83.1% over a load-power range from 92 μW to 5.75 mW.

A Load-Independent Hybrid Capacitor-Inductor Interface Circuit for Piezoelectric Energy Harvesting (319)
Sima Ghandi, Hamed Nasiri, Lihong Zhang and Mohammad Al Janaideh

This paper presents a hybrid piezoelectric energy harvesting interface circuit that integrates Synchronized Switch Harvesting on Capacitor–Inductor (SSHCI) with Synchronous Electric Charge Extraction (SECE). By combining efficient voltage inversion with load-independent charge extraction, the proposed circuit structure simultaneously improves energy harvesting capability and operating range. A capacitor–inductor reservoir is utilized to enhance voltage flipping and increase the harvested energy amplitude. Experimental results demonstrate a rectifier peak output power (RPOP) improvement of 4.3x compared with a conventional full-bridge rectifier (FBR), while the optimal rectified voltage range (ORVR) extended by approximately 3x compared with a standalone SSHCI interface. The results confirm that the proposed hybrid interface structure provides both enhanced power extraction and improved operational robustness for piezoelectric energy harvesting applications.

Phase-Controlled GHz DC-DC Conversion Using Outphasing Class-E Power Amplifiers (370)
Ahilnandan Kabilan and Sakshi Arora

Conventional switched-inductor DC–DC converters face increasing challenges in meeting the power-density and transient-response requirements of emerging wearable and edge-AI systems due to passive-component size and bandwidth limitations. To address these challenges, this paper presents a GHz-frequency DC–DC converter that controls output voltage through outphasing-based phase modulation. The proposed architecture comprises two Class-E power-amplifier branches, a Chireix combiner, a synchronous rectifier, and a compact output filter. A two-stage control scheme uses a lookup-table (LUT)-based controller to place the converter near the target operating point, while a proportional–integral (PI) feedback loop corrects the residual voltage error. Power delivery is controlled through the relative phase between the two Class-E branches, enabling regulation without conventional PWM duty-cycle modulation.
The architecture is designed and validated through simulation for two applications: a wearable power supply and an edge-AI accelerator core rail. A 2.4-GHz wearable-power implementation converts 2.0 V to 1.2 V and delivers 40 mW with a simulated output-voltage ripple of 1.81 mVpp and an efficiency of 72.06%, including modeled GaN-switch and passive losses. A 1.8-GHz edge-AI core-supply implementation converts 3.3 V to 1.2 V across a 0.5–1.4 W load range, achieving 65.21% simulated efficiency and sub-10 ns settling times. The proposed converter trades peak efficiency against passive component volume and transient bandwidth, realizing passives in the 0.1–20 nH and 0.2–30 pF range while supporting ns-scale dynamic response.

S34 | Special Session: Cutting-Edge Research Projects Funded by European and International Programs 1

11:30 AM - 1:00 PM
Room: Ostria
Chairs: Lazaros Papadopoulos, Erica Raviola

Layer- and Block-wise DVFS for Neural Network Inference on Embedded GPUs: An Empirical Study (440)
Quanwen Zheng, Jingjin Li, Heng Yu, George Lentaris, Francky Catthoor and Yajun Ha

Neural network inference behaviors on embedded GPUs change rapidly between compute-bound and memory-bound execution phases, while conventional dynamic voltage and frequency scaling (DVFS) governors react only to coarse operating system metrics. This mismatch makes it difficult to select an energy-efficient GPU frequency at the granularity where inference behavior changes. This paper presents an empirical layer- and block-wise DVFS framework evaluated on the NVIDIA Jetson Nano. It profiles latency and power for each network layer, classifies layers by measured energy sensitivity, and generates frequency schedules before deployment. Noting that a key practical constraint is the GPU frequency transition stall, adjacent layers with similar energy behavior are aggregated into blocks with sufficient dwell time, reducing unnecessary DVFS transitions while preserving most

of the energy benefit of fine-grained control. Across four CNNs and two MobileViT models, block-wise DVFS reduces energy by up to 52.83\% over the best single static frequency and consistently outperforms layer-wise DVFS at batch size 1. Batch-size experiments further show that the best control granularity is workload-dependent: larger workloads can amortize transition cost and make finer layer-wise control preferable.

Testing Methodology for Nanosatellite On-Board Computer Embedded Software System (194)

Panagiotis Karathanasis, Angelos Motsios, Alexandros Tsagkaropoulos, Christoforos Vasilakis, Christos Tsigkanos, George Lentaris and Dionysios Reisis

The proliferation of the "New Space" paradigm necessitates agile yet rigorous testing methodologies for nanosatellite flight software, particularly given the stringent time-to-launch constraints and the reliance on Commercial-off-the-Shelf (COTS) components in this mission class. This study presents the software testing methodology employed for the ERMIS mission, a constellation of 6U/8U CubeSats designed for 5G-IoT, laser optical downlink and hyperspectral earth observation. Addressing the challenges of architectures where the On-Board Computer (OBC) constitutes a single point of failure, we propose an emulator-first testing strategy that decouples functional verification from physical hardware availability. The methodology entails compiling embedded AVR flight software for an x86-64 host environment, employing link-time dependency wrapping to mock hardware interfaces and isolate unit logic. We analyze the implementation of a dual-configuration testing regime, for testing both "Golden" (survivability) and "Nominal" (operational) firmware images and discuss technical solutions and according lessons learned.

Carbon-Aware Routing for Function Calling in Edge-Cloud LLM Systems (82)

Aikaterini Maria Panteleaki, Varatheepan Paramanayakam, Spyros Tragoudas and Iraklis Anagnostopoulos

Large Language Models (LLMs) with function-calling capabilities are becoming critical for modern agentic AI systems. Nevertheless, current deployments typically route inferences to powerful cloud-based models, incurring significant energy use and carbon emissions. We address this sustainability challenge with a carbon-aware routing framework that distributes function-calling queries across a three-tier edge-cloud architecture, combining edge and cloud LLMs on heterogeneous hardware. At its core, a lightweight k-NN predictor operating in a unified semantic-lexical embedding space estimates query-specific accuracy, delay, and power consumption on each edge tier. These predictions are then combined with real-time marginal grid carbon intensity to route every query to the lowest-emission tier capable of executing it successfully. Evaluated on state-of-the-art function-calling benchmarks and LLM families, our framework matches cloud-level accuracy while reducing operational carbon emissions by up to 4x.

Digital Twins of Living Materials: Morphology and Memristive Emulation (435)

Georgios Kleitsiotis, Pantelis Fraidakis, Ioannis Tompris, Ioannis K. Chatzipaschalis, Theodoros Panagiotis Chatzinikolaou, Iosif-Angelos Fyrigos and Georgios Ch. Sirakoulis

Engineered Living Materials (ELMs) offer sustainable, adaptive material systems, yet their biological variability and environmental sensitivity challenge reproducibility and control. We present the hardware-accelerated digital twin framework developed under the EU EIC Pathfinder project FUNGATERIA (Grant No.\ 101071145), targeting mycelium-based ELMs grown from a synthetic fungal--bacterial co-cultivation system. A layered architecture progresses from a Digital Twin Prototype to a Digital Twin Aggregate, both deployed on FPGAs. Morphology is captured by a hardware-compatible reaction--diffusion Cellular Automata model, with growth parameters fitted to experimental images by a genetic algorithm; the electrical dimension uses a reconfigurable memristive spiking grid on SiO$_x$ CBRAM devices to emulate mycelial spiking; a separate reservoir computing project built on these spiking-grid dynamics reaches 97.09% on MNIST and is validated against Parkinson's disease EEG. A custom 2-bit-per-cycle SAR ADC closes the biohybrid loop. Though demonstrated on fungal mycelium, the framework applies to any living material coupling morphological dynamics with bioelectric signalling, including biofilms, organoids, and plant tissues.

A Comprehensive Physical-to-Architectural Verification Framework of Memristive Crossbars for Neural Networks (357)

Anastasios Liontos, Christina Dilopoulou, Georgia Tsirimokou, Yiorgos Tsiatouhas, Denis Voltmann, Jaroslav Dzuba and Vasileios Tenentes

Memristive crossbar arrays offer immense potential for accelerating Artificial Intelligence (AI) through highly efficient, parallel in-memory computing. However, bridging the gap between physical devices and the deployment of high-level neural networks remains a significant engineering challenge. This paper presents a comprehensive, end-to-end verification framework to transition from physical memristors to functional AI hardware accelerators. It supports training and mapping of neural networks to crossbar arrays together with a tuning technique to bridge model to silicon discrepancies. To validate the framework through SPICE, we design a 4x3 crossbar together with essential peripheral drivers, including a 7-bit R-2R DAC and a high-speed 6-bit Flash ADC driven by a TIA-based I-V converter. On a trained neural network for the IRIS dataset, the framework provides identical results with SPICE simulation in terms of both response values and classification accuracy. Furthermore, using an in-house characterization board, we collect characteristics from an actual Bizzcom memristive crossbar, and we use the verification framework to evaluate it on executing AI workload.

Compiler-Assisted Radiation Mitigation for Next-Generation Space Computing Systems (99)
Federico Reghenzani, Davide Baroffio, Tomas A. Lopez, Emilio Corigliano, Giovanni Alberto Sartorato, Davide Mor, William Fornaciari and Giovanni Agosta

The increasing computational demand of modern satellite telecommunication payloads is driving the adoption of high-performance Commercial-Off-The-Shelf (COTS) devices in space systems. However, their susceptibility to radiation-induced faults still limits their use in critical missions. This paper presents the ASTRAEUS project, which investigates software- driven radiation mitigation techniques to enable reliable and cost-effective COTS-based computing platforms for space environments. The proposed approach combines Software-Implemented Hardware Fault Tolerance (SIHFT), targeting SDR-based on-board processors for multicarrier regenerative payloads. A key contribution is the automatic injection of fault-tolerance mechanisms through ASPIS, an LLVM-based compiler framework implementing data protection and control-flow checking at the intermediate representation level. The framework is evaluated through fault injection campaigns on benchmark applications and a FreeRTOS-based workload, comparing ASPIS-only protection with CRC checkpointing and TMR-based recovery. Results show that compiler-assisted SIHFT can improve system resilience while preserving programmability and limiting implementation overhead, supporting its use as a complementary strategy for future reliable space and high-performance computing systems.

S04 | Analog/mixed signal/RF circuits 4
2:00 PM - 3:30 PM
Room: Maistros ACD
Chairs: Vasiliki Gogolou, Yasmin Halawani

On the Equalization Capability of Resonator-Based CTLEs: A Time-Domain Perspective (52)
Mayu Zhang, Danilo Manstretta and Andrea Mazzanti

Resonator-based continuous-time linear equalizers (CTLEs) have emerged as a key topology for high-speed SerDes receivers operating at 112~Gb/s and beyond. While these architectures achieve high peaking gain through high Q-factors, the relationship between peaking gain and actual equalization capability has not been rigorously established. This paper presents a time-domain analytical framework that decomposes the CTLE impulse response into a derivative term and a base response term. The analysis reveals that both the zero frequency ω_z and Q-factor can increase the depth of the first post-cursor $|c_1|$, but through fundamentally different mechanisms: ω_z provides equalization without side effects, whereas Q simultaneously introduces self-induced ISI through ringing at subsequent post-cursors. Peaking gain, as a composite metric dominated by Q in the resonator regime, cannot distinguish between these two contributions and therefore fails as a reliable proxy for equalization capability. Behavioral-level simulations quantify the cost-benefit trade-off of Q and demonstrate that, for a given peaking gain, an optimal balance between ω_z and Q exists that maximizes net equalization performance --- a design choice entirely invisible to the peaking gain metric alone.

Reconfigurable Integrated RC Networks: Experimental Characterization for Extended Filtering Functions (69)
Luis Arturo Martinez Alvarado, Federico Sandoval Ibarra, Miguel Angel García Andrade and Arturo Gutierrez D'Iaz

This paper presents the experimental characterization of a reconfigurable integrated passive RC structure fabricated using an academic thin-film process. A single on-chip layout is configured as either a low-pass or high-pass network by modifying only the external interconnection, enabling direct comparison between complementary filtering responses under identical physical conditions. Frequency response measurements based on lock-in techniques are employed to evaluate the impact of layout-dependent parasitic effects on the electrical behavior of the structure. The experimental results reveal that the effective capacitance is strongly influenced by parasitic contributions, leading to significant deviations from nominal geometric expectations. Nevertheless, the measured responses exhibit predictable first-order behavior and are accurately described by an equivalent RC model derived from experimentally extracted parameters. The obtained results demonstrate that reliable prediction of integrated passive RC structures can be achieved through experimental characterization, even under parasitic-dominated conditions and in the absence of detailed process models. Moreover, the proposed reconfigurable structure provides a practical basis for extended integrated passive filtering approaches without requiring modifications to the physical layout. Index Terms—Integrated passive filters, RC filters, layout parasitics, experimental characterization, frequency response, reconfigurable filters.

New Biquadratic Operational Amplifier Based Filters (153)
Brent Maundy, Ahmed Elwakil and Abdulrahman Nutfaji

In this paper we propose new biquadratic operational amplifier filters that have comparable or better responses than that of the ubiquitous Kerwin, Huelsman, and Newcomb (KHN) filter. By using a T like grounded RC structure connected to an inverter, three circuits are proposed. The corresponding gains, quality factor and center frequency gains are compared and contrasted. Simulation and experimental results are also presented.

A Calibration-Free GM-C Low-Pass Filter Constructed with Inverter Based OTAs (292)
Emre Sayıt, Berkay Çiftci, Adnan Gündel and Hasan Uluşan

This paper presents a calibration-free 5th order Gm–C low-pass filter that maintains its cutoff frequency and transfer shape against process, voltage, and temperature (PVT) variations without requiring a reference clock or auxiliary tuning loop. The filter is based on inverter-based OTAs and an active-gyrator realization of a fifth-order ladder prototype. Robustness is achieved by combining strong degeneration and feedback capacitance, which reduces sensitivity to lossy components and transconductance variations. Measurements from fabricated prototypes demonstrate a nominal 5 MHz cutoff frequency with <4.8% sample-to-sample variation while preserving the intended filter response. The intrinsic open-loop structure of GM-C topology and good high frequency behavior of inverters make this architecture highly suitable for wideband communication transceivers.

Non-Uniform Analog FIR Filtering for Wideband Receivers With Relaxed DAC-Rate Requirements (312)
Abhishek Kumar

Analog finite-impulse-response (AFIR) filtering provides sharp, programmable blocker rejection, but a conventional-tap implementation requires coefficient updates at DAC =. Lowering this rate compresses the response and brings aliases and stopband flyback toward the wanted band. Time-approximation filtering (TAF) encodes a multilevel waveform as gated pulses, but finite timing resolution creates a nonmonotonic stopband. Non-uniform TAF (NU-TAF) spreads coherent aliases, although timing-induced flyback remains. This paper proposes a non-uniform AFIR (NU-AFIR) that combines multilevel analog FIR weighting with periodically non-uniform sampling. Fourier-series analysis shows that the coherent alias at is redistributed among spectral branches, while a first-order timing-error bound explains the flyback advantage of amplitude weighting. At a fixed 2.8 GHz DAC rate, a 13-tap, 10-bit behavioral model increases the contiguous bandwidth satisfying 20 dB alias rejection from 413 MHz to 815 MHz. Frequency scaling then gives the same 413 MHz bandwidth at 1.42 GHz, a 49.3% DAC-rate reduction. Across 80 MHz to 600 MHz, the mean folded-alias power is reduced by 6.47 dB, and blocker-limited 64-QAM EVM improves from 9.54% to 3.74%.

S11 | Low-Power/Voltage Design **2:00 PM - 3:30 PM**
Room: Maistros BEF
Chairs: Lorenzo Benatti, Sahibia Kaur Vohra

Buffer PSRR improvement through feedforward supply-noise cancellation (41)
Spyridon Vlassis, Konstantinos P. Pagkalos, Aikaterini I. Chronopoulou, George Souliotis, Vasileios Kalenteridis and Fotis Plessas

This paper presents a simple and effective technique for improving the PSRR of a low-voltage buffer. The operating principle is theoretically analyzed and verified through simulation results. The circuit is implemented in 65 nm TSMC technology, operates from a 1.2 V supply, and drives a 10 mA load current with a 200 pF load capacitor. The proposed topology achieves an approximately 10 dB improvement in low-frequency PSRR and extends the flat-PSRR bandwidth to about 330 kHz. These enhancements are achieved without additional circuitry, extra power consumption, or degradation of the error amplifier stability.

A Synchronous, Latch-Based Digital OTA Operating at 1.25 nW-Power and 0.3V Power Supply (177)

Andrea De Gregorio, Danilo Karim Kaddouri and Paolo Stefano Crovetti

This work presents a novel synchronous Latch Based (LB) Digital Operational Transconductance Amplifier (LB-DIGOTA) for ultra-low voltage (ULV) and ultra-low power (ULP) biosensors. The circuit has been designed in 130nm CMOS technology and demonstrated with Post-Layout Simulations (PLS) at 300mV supply voltage while driving a 150pF capacitive load. Its operation is validated varying the clock frequency fclk from 5kHz to 500kHz. The proposed circuit achieves a minimum power consumption of 1.25 nW at fclk=5kHz, while increasing the DC gain up to 31.17 dB and reaching a minimum input-referred noise of 11.02 µVrms at fclk=500kHz. Simulations reveal an excellent small-signal (large signal) Figure of Merit FOMGBW (FOMSR) of 246.24 kHz pF/nW (0.096 V/µs pF/nW) at 5 kHz clock frequency, making it well suited for ULV and ULP constrained biosensing applications.

An Energy Efficient Precisely Adjustable Hybrid Delay Element in 22-nm FDSOI CMOS (209)

Valentin Stadtlander, Florian Protze and Frank Ellinger

In this paper, a hybrid delay element (DE) with both analog and digital control is investigated and presented. The DE is based on two cascaded CMOS inverters. The coarse digital tuning is implemented using a discrete capacitive digital-to-analog converter (CDAC), while fine analog biasing is achieved by exploiting the strong body-bias capabilities of the utilized 22FDX® process. The back-gate (BG) voltages of the CMOS inverter transistors can be adjusted independently, enabling delay tuning with step sizes below 100 fs. This allows using the advantages of conventional analog controlled DEs (ACDEs), particularly fine-grained delay time tuning, while avoiding their typical drawbacks, such as increased power consumption, reduced area efficiency, and input-data-dependent delay. To prove the concept, the proposed DE was implemented inside a delay line (DL) of a time-to-digital converter (TDC) and fabricated in the 22FDX® process. The DE has a tuning range of 11.8–22.6 ps, consumes 292.9 µW from an 800 mV supply at a 20 Gb/s input data rate, and achieves an energy per delay of only 14.64 fJ.

Ultra Low Power, CMOS Programmable VRs for Under Voltage Lockout Circuits (211)

Giorgio Bersani, Michele Caselli and Andrea Boni

This paper presents the design of an UVLO system that is composed by 3 main blocks, a PVREF (Programmable Voltage Reference), a flipped-designed VR and a comparator. The main core of this work relies in the design phase where, using our OptiGap CAD Toolbox, the Voltage References are designed, while the comparator is hand-sized. The entire system works with a quiescient current of 3.5 nA in post-layout simulations at room temperature, while keeping a programmability, thanks to the PVREF voltage taps, of the shut-off voltage for the circuit, and the TC of the shut-off voltage swept in the temperature interval [-40,80]°C can be as low as 38ppm/°C.

An 83 ppm/ºC Voltage Reference in Flexible TFT Technology (338)
Noor Ul Amin, Andrea Ballo, Alfio Dario Grasso and Orazio Aiello

This paper proposes a low power voltage reference realized in Indium-Gallium-Zinc Oxide (IGZO)- based thin film transistor (TFT) technology using the 600-nm Pragmatic Process Design Kit. Unlike conventional silicon technologies, IGZO TFT technology lacks mature device options such as PN junctions and complimentary transistor structures, posing challenges for voltage reference realization. An NMOS-only voltage reference topology with a trimming network is proposed to enable compact implementation and calibration capability. The proposed circuit is designed to operate over a 2- 5 V supply voltage range and generates a 751-mV reference voltage, achieving a minimum temperature coefficient of 83 ppm/°C at 2 V supply over the 0–60 °C temperature range. Post-layout simulations demonstrate stable performance under both temperature and supply variations, highlighting the feasibility of compact, low power voltage references toward more mature TFT-based circuit implementations.

S13 | EDA, Test and Reliability 2
2:00 PM - 3:30 PM
Room: Zephyros
Chairs: George Floros, Alberto Garcia-Ortiz

GAVENS: A GPU-Accelerated Value-Exact NPU Simulator for Pre-Silicon Accuracy Validation (93)
Chin-Cheng Wei and Chung-Ho Chen

Deploying quantized neural networks on neural processing units (NPUs) requires a numerical reference that is value-exact with respect to the target NPU arithmetic datapath. We build VENS, a CPU value-exact simulator, abstracted from the instruction-set simulator of the target NPU, and use it to measure how far fake-quant baselines, including FQ-ViT and Microsoft Microxcaling, deviate.
Their final-logit relative L2 distance to VENS is 13--56% across the vision models, exceeding their distance to the floating-point reference and showing that fake-quant can miss datapath-level numerical behavior. Yet VENS is too slow to evaluate the full dataset, so its results cannot drive hardware-software co-optimization before tapeout. We present GAVENS, a GPU-accelerated version of VENS in which each stage is verified for value equivalence against the CPU reference. To bridge the quantization-scheme mismatch between NPU silicon datapaths and GPUs, it maps the exact-integer datapath onto a matrix core, e.g., NVIDIA's Tensor Core, via Zero-Point Compensation, ZPC, and Block-Aware Padding, BAP, complemented by compiler-assisted kernel fusion. Across two NPU variants and four architectures, GAVENS's best-observed production configuration reaches up to 8,733× the throughput of the unchanged CPU VENS oracle; on DeiT-Tiny, it matches fake-quant speed and reduces a full ImageNet validation run to less than two minutes. This minutes-scale feedback lets designers revise the quantization scheme or datapath and re-validate the resulting accuracy before tapeout, making hardware-software co-optimization practical.

Benchmarking the Open Source Yosys-Nextpnr FPGA Toolflow against Vendor Software (124)

Sven Krause, Denis Lebold, Sebastian Wittlich, Tobias Krawutschke, Jens Rettkowski and Hendrik Wöhrle

In the wake of global supply chain instabilities and the push for \textbf{technological sovereignty}, the European research landscape is increasingly prioritizing open-source EDA tools.
Transitioning from academic niches to robust ecosystems, the combination of \textbf{Yosys} for synthesis and \textbf{Nextpnr} for place-and-route (PnR) offers an emerging alternative to proprietary suites like Vivado or Diamond.
Despite this momentum, there are few empirical performance benchmarks for mid-range FPGA architectures.
This paper addresses this gap by systematically evaluating the Yosys-Nextpnr toolchain against established industry software.
Using a set of benchmarks, such as \textbf{RISCV-core (VEXRISCV)} and DSP-intensive designs like FIR-filters, we measure resource utilization, maximum clock frequency (f_{max}), and compilation time across a hardware spectrum including AMD, Lattice, Gowin, and Cologne Chip.
Our findings reveal that the open-source toolchain consistently achieves faster compilation times and competitive, often higher f_{max} compared to vendor tools, useful for Design Space Exploration and rapid prototyping.
However, this performance usually comes at the cost of higher resource utilization.
Furthermore tool maturity correlates with architectural documentation; less well-documented architectures occasionally exhibit inconsistent performance.
The results provide an empirical foundation for assessing the maturity of FOSS-EDA, demonstrating that the toolchain can be a viable alternative for modern hardware development.

A Multi-Level Thermal Modeling and FPGA Emulation Framework for Nanophotonic 3D Systems (197)

Helen Clara George, Sebastian Fischer and Alberto García-Ortiz

Thermal sensitivity is a key challenge for the incorporation of nanophotonic computation and communication technologies in heterogeneous 3D architectures. In particular, analyzing the impact of high-level parameters—such as task scheduling—on the thermal profile of the system and, consequently, on MRR is especially challenging. This paper develops a comprehensive multi-stage modeling framework and an FPGA -based emulation methodology for temperature-aware nanophotonic heterogeneous 3D architectures, reproducing the full physics-based thermal-response workflow with an average RMSE of 1.355% while running at hardware speed. The framework proceeds from physics-level thermal simulations of the microring resonator (MRR) to Lorentzian fitting of the transmission spectra, yielding reliable temperature–wavelength relationships. These fitted models expose how temperature variations distort modulated data and characterize the nonlinear response of the MRR, motivating an analytical model that captures the temperature-induced transformation of optical modulation outputs. The resulting model is implemented at the RTL level using fixed-point arithmetic, enabling efficient FPGA-based emulation of thermal effects. To demonstrate system-level implications, the model is integrated into the RTL implementation of a photonic transformer accelerator with dynamic photonic tensor cores, enabling direct evaluation of how temperature-induced transmittance variations propagate through photonic AI computation.

LLM-Assisted Coverage-Driven Verification of RISC-V Processor (428)
Zoya Ahmed and Otmane Ait Mohamed

The verification of RISC-V processor cores remains challenging due to the increasing complexity and stringent quality requirements of modern hardware systems. This paper presents VeriLLM, an industry-aligned LLM-assisted verification framework for RISC-V processors that generates fully functional assembly programs as test stimuli, simplifying post-simulation debugging and result processing. VeriLLM introduces three key contributions: (1) it distinguishes hint-before-solving prompting for self checking, deterministic test generation from iterative, zero-shot prompting for coverage-driven, constrained-random test generation; (2) it generates RISC-V ISA-compliant, self-checking deterministic tests and coverage-driven constrained-random tests for architectural and microarchitectural verification; and (3) it implements a coverage-guided iterative prompt refinement loop for constrained-random tests, accelerating coverage closure. Evaluation on the production-quality 32-bit Ibex RISC-V core demonstrates that VeriLLM accelerates coverage closure and improves verification efficiency, achieving a functional coverage gain of 22.68% over conventional constrained-random testing and 48.97% over an existing LLM-based framework, while maintaining a modular design easily adaptable to other open-source processor cores.

A Physics-Guided Transformer Framework for Electromigration Analysis in Multi-Segment Interconnects (424)
Pavlos Stoikos, Anuj Pathania and George Floros

As technology scales to smaller nodes, increasing current densities make electromigration (EM) one of the dominant reliability challenges in on-chip interconnects. Accurate transient stress analysis is needed to identify wires susceptible to EM degradation, but applying physics-based solvers across many interconnects remains computationally expensive. This paper proposes a physics-guided transformer framework for fast EM stress prediction in multi-segment interconnect lines. The framework converts each line into geometry- and DC-aware segment tokens and uses transformer attention to capture line-level context. A lightweight query decoder then predicts stress at selected locations and time instants. The model is trained with an objective that combines normalized supervised regression, linewise relative-L2 loss, and physics-guided continuity and terminal-flux terms. Experiments on IBM power grid benchmarks show that the proposed model achieves relative-L2 error below 8% and reaches up to 2459.68x speedup compared with the matrix exponential solver.

S16 | Circuits and Systems for AI Algorithms 2
2:00 PM - 3:30 PM
Room: Sirocco
Chairs: Panagiotis Mousouliotis, Minas Dasygenis

A Rank-8 Rolled Kernel-Grid Accelerator IP Core for CNN Feature Extraction (143)
Vinicius Correa, Gabriel Souza and Tales Pimenta

Hardware acceleration of deep learning inference can be approached in two ways: using automated tools to convert high-level models into RTL descriptions or designing circuits from scratch using hardware description languages such as Verilog. Although the first approach accelerates development, it reduces the designer's control over the generated RTL and is often restricted to heavily quantized networks. The second enables fine-grained optimizations through hardware-oriented techniques such as pipelining. Following Amdahl's Law, which states that optimization should focus on the most computationally intensive stage, we adopt the latter approach and propose a kernel-grid feature extraction unit: a pipelined MAC architecture that performs 64 cross-correlations in parallel on an input image, extracting 64 CNN feature maps simultaneously, directly mapping the behavior of the first convolutional layer of a Super-Resolution Convolutional Neural Network (SRCNN). Our IP module demonstrated significant performance gains through parallel feature map extraction and pipelined execution, achieving higher efficiency than a nonpipelined implementation, at the cost of increased resource utilization when synthesized using a 65 nm CMOS standard-cell technology.

Beyond Multiplication Count: System-Level PPA Trade-offs of Fast-Convolution RTL Accelerators (182)
Fernando Moraes and Tarsio Onofrio

Fast convolution reduces the number of multiplications required by small-kernel CNN layers, but these arithmetic savings do not automatically translate into system-level gains once control, storage, and memory bandwidth are included. This paper evaluates this effect through a parameterizable RTL fast-convolution accelerator synthesized in a 28~nm technology node. The system integrates the convolution core, control unit, and memory interface under a common PPA methodology, and varies both the number of instantiated multipliers and the effective bus width. Results show that a one-sample-per-cycle bus limits cycle reduction to 40-51% relative to a direct convolution baseline, indicating a memory-bound regime. When the bus width is matched to the input tile size, cycle reduction increases to 70--2%. However, the fastest configurations are not necessarily the most energy-efficient: TC_2x2 achieves the highest energy reduction, 62%, while IF_3x3 with 18 multipliers matches the best observed cycle reduction, 82%, with 57% energy reduction and lower overhead than the comparable TC_4x4 throughput point. These results show that multiplication count alone is not a reliable proxy for system-level efficiency. Fast-convolution accelerators should be selected by jointly considering algorithmic structure, multiplier count, memory bandwidth, and datapath utilization.

Event-Driven Binarized Convolutional Layer for Ultra-Low-Power On-Chip Keyword Spotting (224)

Marc Chevalier, Benoit Larras and Antoine Frappé

Ultra-low-power processing for in-sensor edge-AI classification typically combines event-driven feature extraction with always-on neural network classifiers operating in the digital, clocked domain. This transition to clocked processing compromises the energy efficiency and speed of event-driven approaches by neglecting the sparse-event nature of input data. To overcome these limitations, we introduce a fully clockless, event-driven convolutional neural network layer that leverages in-memory computing to enable continuous-time processing with near-instantaneous latency. At the same time, the use of binarized weights further simplifies the design and significantly reduces power consumption. Focusing on keyword spotting, we compare our continuous-time model with a discrete-time approach to demonstrate its value.

FPGA-Based Deployment of Quantized DRL Policies for Edge Robotics (403)

Irene-Maria Tabakis, Konstantinos Tatas and Minas Dasygenis

Deep reinforcement learning (DRL) is increasingly used for autonomous decision-making in edge robotic systems, but deploying DRL policies under strict latency and energy constraints remains challenging. This paper evaluates the deployment of lightweight multi-layer perceptron (MLP) policies for coverage path planning on an FPGA-based edge platform using the AMD/Xilinx Deep Processing Unit (DPU). Rather than proposing a custom accelerator, the work studies how a vendor-supported DPU overlay executes compact non-convolutional DRL workloads. The policy is exported through ONNX, quantized to INT8 precision, compiled with the Vitis AI flow, and deployed on the AMD/Xilinx Kria KR260 platform. Under batch-one inference, the FPGA-DPU implementation achieves 3.2 ms latency and 0.9 mJ per inference. Profiling shows that compact DRL policies underutilize the DPU, with compute utilization between 42% and 58%, and that execution is dominated by memory-related activity. These results show that vendor DPU overlays can support low-power edge robotic inference, but also motivate hardware-aware model selection for lightweight DRL workloads.

ZTA-Q: an Open-source RISC-V Platform for Accurate Quantized CNN Inference (417)

Yike Li, Ajay Kumar M, Vishnu P S, Deepu John, Dimitrios S. Nikolopoulos, Bo Ji and Hans Vandierendonck

Low-precision inference is widely adopted in edge AI to reduce computational cost and memory footprint. However, existing open-source accelerator platforms provide limited end-to-end support for CNNs following the standard TensorFlow Lite integer inference scheme. This paper presents ZTA-Q, an open-source RISC-V-based platform that enables accurate deployment of TensorFlow Lite INT8 models. In addition to extending operator support, ZTA-Q provides a configurable post-processing datapath for studying how circuit-level approximations, including reduced multiplier precision, shared shift scaling, and simplified rounding, affect model accuracy. The proposed system is implemented on a Digilent Arty A7-100T

FPGA and operates at 83.3 MHz. Evaluations on representative CNN models show that with LUT, register, and DSP overheads of 26.3%, 12.6%, and 150%, respectively, ZTA-Q limits the degradation in both top-1 and top-5 accuracy to within 0.1% while introducing no additional inference latency.

S35 | Special Session: Cutting-Edge Research Projects Funded by European and International Programs 2

2:00 PM - 3:30 PM
Room: Ostria
Chairs: Lazaros Papadopoulos, Apostolos Fournaris

Aging-Aware Resiliency Analysis for Digital Compute-In-Memory DNN Accelerators (243)

Alexis Maras, Panagiotis Chaidos, Georgios Alexandris, Theofilos Spyrou, Anteneh Gebregiorgis, Manil Dev Gomony, Henk Corporaal, Said Hamdioui, Dimitrios Soudris and Sotirios Xydis

As Deep Neural Networks (DNNs) shift toward edge applications, SRAM-based Digital Compute-In-Memory (DCIM) architectures have emerged as a vital solution to alleviate memory transfer overheads and meet tight power budgets. However, as these accelerators scale to advanced technology nodes, long-term hardware reliability becomes a major challenge due to device wear-out from mechanisms like Bias Temperature Instability. To accurately evaluate these reliability concerns, this paper presents a cross-layer study, an aging-aware fault-injection framework developed under the EU project CONVOLVE. We demonstrate that accurately assessing the impact of aging on DCIM architectures relies heavily on the collaborative efforts of both system-level application mapping and low-level circuit/material degradation analysis. Our results reveal an inherent structural resilience in DNNs: the critical Most Significant Bits naturally center around a stable, symmetrical signal probability near 0.5, protecting them from severe degradation. Consequently, the evaluated networks maintain a resilient relative inference accuracy, proving that SRAM-based DCIM topologies provide an exceptionally rugged foundation for long-term Edge AI deployments.

6G VERtical trials for SUStainability: towards a Greek pilot utilizing Energy-Efficiency-as-a-Service (409)

Aimilios Leftheriotis, Charis Diamantidis, Dimitrios Fragkos, Konstantinos Katsaros, Lazaros Karagiannidis, Spyridon Athanasiadis, Xenofon Vasilakos, Thanasis Douklias, Angelos Amditis, Georgios Theodoridis, George Lentaris and Dimitrios Soudris

As the telecommunications landscape transitions towards the sixth generation (6G) of mobile networks, sustainability and energy efficiency have emerged as critical design principles. The 6G-VERSUS project aims to address these challenges by deploying a novel triadic application architecture across various European testbeds. This paper focuses on the Greek pilot, which explores collaborative robotics for Public Protection and Disaster Response scenarios, specifically wildfire detection and search-and-rescue. Given the strict energy constraints of battery-powered edge devices operating in the field, we detail the initial integration of an Energy Efficiency as a Service (EEaaS) framework. By accurately modeling the hardware power

consumption of AI inference workloads at the edge, EEaaS lays the foundational monitoring infrastructure necessary to develop operational energy-saving profiles, ultimately enabling human-in-the-loop orchestrators to make critical trade-offs between mission performance and battery life.

FPGA-based acceleration for Homomorphically Encrypted Processing in CoEvolution (442)

Stylianos Kotsobolos, Konstantina Karagianni, Vassilis Paliouras and Apostolos Fournaris

The CoEvolution project aims to provide tools against evolving AI cybersecurity threats. In this context Homomorphic Encryption is an emerging quantum-resistant technique which allows processing of encrypted data without decrypting them first, however, only at the cost of excessive computational requirements. We demonstrate the integration of a hardware accelerator into an application and quantify the performance achieved. By actual measurements on two setups, results reported here show that the end-to-end execution time can be substantially reduced, even when data transfer and library integration are considered. The developed system integrates hardware acceleration into OpenFHE and supports ring dimensions and ciphertext moduli of sufficient.

Experimental Evaluation of a Low-Power Ultra-Wideband Receiver for Spectrum Sensing (219)

Panagiotis Vlachos, Ioannis Bartsiokas, Cedric Dehos, Francois Rivet and George Karachalios

Walsh-sequence-based receiver architectures offer an alternative approach for the reception and reconstruction of multiple simultaneous RF signals over wide bandwidths. While previous works have focused on the architecture and theoretical operation of Walsh-domain processing, limited experimental results have been reported so far on hardware operation under realistic conditions. This paper presents the experimental evaluation of a Walsh-sequence-based receiver prototype. A measurement campaign was conducted under various signal placement, interference, and receiver configuration scenarios. Reconstructed spectra, Walsh-lane power spectral densities (PSDs), and Error Vector Magnitude (EVM) measurements were used to investigate Walsh-domain signal representation, interference effects, and receiver performance. Finally, dense- spectrum experiments involving up to ten simultaneously active signals demonstrate the potential of the architecture for low power wideband spectrum sensing applications.

ATHENA Cybersecurity Validation and Testing Environment (341)

Lazaros Papadopoulos, Nikos Peppes, Dimitrios Soudris, Ilias Kalouptsoglou, Dimitrios Tsoukalas, Damian Puchalski and Michał Choraś

The rapidly shifting regulatory landscape in the European Union is marked by the introduction of frameworks such as the Cyber Resilience Act (CRA), NIS2 Directive, and Data Act. These frameworks place an operational and compliance burden on software and hardware manufacturers. This challenge is especially relevant for resource-constrained SMEs, which have limited access to specialized cyber-engineering personnel and testing infrastructure. This paper presents an overview of ATHENA, a 36-month Horizon Europe Innovation Action project, which is designed to lower these barriers through a novel, flexible, and automated

cybersecurity platform. The ATHENA platform is an end-to-end security engineering and validation environment and its architecture is organized around five pillars: security engineering-by-design, AI-driven security services, automated vulnerability detection and repair toolsets, regulatory compliance mapping, and a multi-tier continuous validation environment using virtualized cyberranges. The evaluation of the market-readiness of the ATHENA platform will be based on six real-world industrial use cases spanning critical infrastructure, public safety, autonomous robotics, and supply chain logistics. ATHENA provides European manufacturers with an automated solution towards achieving continuous security-by-design and formal EU compliance certification.

SPQR: Smart Precision for Quantum Resilience (324)

Simone Perriello, Gerardo Pelosi, Giacomo Lancellotti, Stefano Cherubin, Mohammad B. Rou, Giovanni Agosta and Gunnar Tufte

The rapidly shifting regulatory landscape in the European Union is marked by the introduction of frameworks such as the Cyber Resilience Act (CRA), NIS2 Directive, and Data Act. These frameworks place an operational and compliance burden on software and hardware manufacturers. This challenge is especially relevant for resource-constrained SMEs, which have limited access to specialized cyber-engineering personnel and testing infrastructure. This paper presents an overview of ATHENA, a 36-month Horizon Europe Innovation Action project, which is designed to lower these barriers through a novel, flexible, and automated cybersecurity platform. The ATHENA platform is an end-to-end security engineering and validation environment and its architecture is organized around five pillars: security engineering-by-design, AI-driven security services, automated vulnerability detection and repair toolsets, regulatory compliance mapping, and a multi-tier continuous validation environment using virtualized cyberranges. The evaluation of the market-readiness of the ATHENA platform will be based on six real-world industrial use cases spanning critical infrastructure, public safety, autonomous robotics, and supply chain logistics. ATHENA provides European manufacturers with an automated solution towards achieving continuous security-by-design and formal EU compliance certification.

Young Professionals (YP) Panel Session **4:00 PM - 5:30 PM**
Room: Maistros ACD
Moderator: Dr. Vasileios Ntinas, Assistant Professor, Aalborg University

The session will feature a roundtable discussion led by young professionals from academia and industry, focusing on the challenges and opportunities facing young professionals in Electronics, Circuits, and Systems. Topics will include career development, professional growth, and the transformative role of Artificial Intelligence in the evolving technological landscape.
The discussion will provide an opportunity to exchange experiences, explore emerging career paths, and share perspectives on the skills and competencies needed to succeed in a rapidly changing environment.
The session will continue with a poster presentation featuring selected contributions from Young Professionals, showcasing innovative ideas, technical creativity, and practical relevance. ABest YP Poster Award will be announced during the conference closing ceremony.
Following the panel discussion, a dedicated mentorship segment will provide opportunities for informal one-to-one and small-group exchanges between panelists and participants. This activity will encourage knowledge sharing, career guidance, and professional networking, helping young professionals build valuable connections with experienced members of the Electronics, Circuits, and Systems community.

Wednesday November 11th 2026

S14 | EDA, Test and Reliability 3 **10:30 AM - 12:00 PM**
Room: Maistros ACD
Chairs: George Lentaris, Minas Dasygenis

Development of a Radiation-Hardened Standard Cell Library in 7 nm FinFET Technology (175)
Aobo Cui, Sarah Azimi and Luca Sterpone

As technology scales to advanced FinFET nodes, digital circuits become increasingly susceptible to radiation-induced soft errors, creating a major reliability concern for safety-critical and space applications. Conventional hardening often requires extensive circuit-level vulnerability analysis, leading to substantial design time and effort. This paper presents a radiation-hardened (RH) standard cell library for 7 nm FinFET technology based on the ASAP7 PDK, aiming to improve radiation resilience while limiting both overhead and design time. The library is developed through transistor-level fault injection and 3D layout-aware particle-strike simulation to derive customized hardening strategies for individual cells. Experimental results show that the RH cells achieve significantly higher threshold LET and shorter error duration than standard counterparts. Case studies on a D flip-flop and a C432 benchmark further demonstrate significantly reduced SEE cross-section with modest area, power, and delay overheads, confirming a practical and scalable approach to SEE-resilient design in advanced FinFET circuits.

A Fault-Tolerant Serial Peripheral Interface (265)
Alex Calinescu and Marius Enachescu

Serial Peripheral Interface (SPI) Controller IP that can avoid critical failures during operation is presented. The IP features Triple Modular Redundancy (TMR), Signal Qualifiers (SQ), digital filters and watchdog timers for fault detection, along with a strict state machine transition policy dedicated to fault detection. The IP was implemented in a commercial 130nm BCD CMOS process, yielding an area of 0.029 mm2, a maximmum frequency above 100 MHz and draining up to 0.688 mW when running at 50 MHz with a supply voltage of 1.5 V.

FERRino: A Radiation-Hardened Embedded FPGA in 28 nm CMOS (322)
Andrea Mazzanti, Luca Frontini, Alessandra Camplani, Alessandro Andreani, Alberto Stabile and Jafar Shojaii

Radiation-tolerant embedded FPGAs (eFPGAs) offer a promising solution for introducing post-fabrication reconfigurability into ASICs deployed in space and high-energy physics environments. This work presents FERRino, a radiation-hardened eFPGA developed within the framework of the FERRad project (Radiation-Hardened Embedded FPGA for harsh-environment Applications) and fabricated in a commercial 28 nm CMOS technology using an open-source design flow based on the FABulous framework. The architecture incorporates Triple

Modular Redundancy (TMR) in the user logic, configuration circuitry, and peripheral control modules, while the configuration memory employs self-correcting TMR storage elements to mitigate Single Event Upsets (SEUs). The fabricated prototype occupies an area of 1 mm x 1 mm and contains a 4 x 4 LUT-based fabric with 16 programmable GPIOs. Heavy-ion irradiation tests were performed using nine ion species with LET values in the range from 1.15 to 37 MeV cm2/mg. No SEUs were observed in the TMR-protected test structures, and no functional interruptions were detected while executing firmware within the eFPGA. These results demonstrate the effectiveness of the proposed hardening techniques and validate the use of open-source eFPGA frameworks for radiation-tolerant reconfigurable systems.

From Software to Hardware: Cross-Domain Vulnerability Detection in Verilog via Domain Adaptation (356)

Anna Maria Voulkopoulou, Ilias Kalouptsoglou, Dimitrios Tsoukalas, Dionysios Kehagias and Dimitrios Tzovaras

Vulnerability detection has become increasingly effective in software systems, yet hardware description languages remain quite underexplored due to the scarcity of labeled vulnerability datasets. In this paper, we propose a domain adaptation technique that transfers vulnerability detection knowledge from C/C++ to Verilog code. We fine-tune CodeBERT, a pre-trained Transformer model, using the BigVul dataset as the source domain, while simultaneously aligning feature representations with unlabeled Verilog samples from a hardware vulnerability dataset, using Maximum Mean Discrepancy (MMD). The total training objective combines supervised cross-entropy loss in the source domain with a weighted MMD term in order to optimize the classification performance, while reducing the gap between source and target domains. We evaluate the proposed method on the target Verilog dataset and further assess its capabilities on the BugWhisperer dataset, which contains a variety of known hardware vulnerabilities. Results demonstrate that domain adaptation with MMD enables effective transfer of vulnerability detection capabilities from C/C++ to Hardware Description Languages.

Reliability Aware Learning Based Scheduling for Tiled DNN Accelerators (266)

Muhammad Zia Hameed and Iraklis Anagnostopoulos

Deep neural network (DNN) accelerators continue to grow in size and complexity. As accelerator scale increases, inter layer scheduling becomes important for utilization, energy efficiency, and resource balancing. Existing scheduling frameworks mainly focus on energy delay product (EDP) efficiency. However, these methods overlook tile aging imbalance and stress accumulation. Analysis of activity factor distributions reveals non-uniform tile utilization and aging imbalance. To address this problem, this work proposes a reinforcement learning based reliability aware inter layer scheduling framework for tiled DNN accelerators. The framework incorporates tile fatigue tracking, queue level fatigue propagation, and tile rotation to redistribute hardware stress. This work develops two reliability aware proximal policy optimization agents for EDP and delay optimization. The agents learn adaptive scheduling policies from fatigue aware observations and workload characteristics.

Experimental results show improved aging balance while maintaining competitive energy and latency.

S19 | AI-driven Circuit Design 2 **10:30 AM - 12:00 PM**
Room: Maistros BEF
Chairs: George Floros, Manolis Katsaragakis

Assessing Data Fragmentation and Federated Learning for ML-based Dynamic IR-Drop Analysis (300)
Ludovica Deriu, Erich Malan, Valentino Peluso, Andrea Calimera, Stefano Quer and Enrico Macii

Machine learning-based approaches to IR-drop prediction have emerged as a fast, scalable alternative to computationally intensive sign-off analysis. However, most existing methods assume access to centralized datasets collected across multiple designs (an assumption that is often unrealistic in practice due to intellectual property restrictions and organizational data silos). In this work, we examine how data fragmentation affects the accuracy of ML-based dynamic IR-drop prediction and explore Federated Learning as a privacy-preserving framework for collaborative model development. We analyze realistic deployment scenarios where design type, operating frequency, or design–frequency combinations dictate data distribution across clients.
Experiments on the CircuitNet-N28 benchmark reveal that isolated training can substantially degrade predictive performance, resulting in accuracy losses of up to 25.96% relative to an ideal centralized setting. By contrast, Federated Learning consistently recovers most of this lost performance, achieving prediction quality close to that of centralized training while enabling effective collaboration without sharing proprietary data.

Design Space Exploration of Reinforcement Learning Configurations for Post-Legalization Placement (344)
Foteini Oikonomou, Maria Pantazi-Kypraiou, George Goudroumanis, Pavlos Stoikos, Athanasios Tziouvaras, Christos Anagnostopoulos, Georgios Karakonstantis and George Floros

Graph Neural Networks (GNNs) and Reinforcement Learning (RL) have recently been explored for VLSI placement optimization, but their effectiveness for post-legalization placement refinement remains unclear. This paper presents a systematic design space exploration of a GNN-Proximal Policy Optimization (GNN-PPO) framework across three ISPD2015 benchmarks. We evaluate 18 configurations covering reward functions, GNN architectures, action spaces, candidate-set sizes, and training strategies. We analyse the results we obtain to identify pitfalls and to assess the prospects and opportunities of the RL methods for post-legalization placement. On top of that we propose a design classification approach to guide the design decisions of RL models.

GPlaceRL: An Open-Source Graph Reinforcement Learning Framework for Detailed Placement (423)

Pavlos Stoikos, Foteini Oikonomou, Christos Poulos, Maria Pantazi-Kypraiou, Athanasios Tziouvaras, Christos Anagnostopoulos, George Floros and Georgios Karakonstantis

Reinforcement learning (RL) has emerged as a promising approach for placement optimization, particularly when combined with graph neural networks (GNNs) that capture circuit connectivity. However, most learning-based placement approaches focus on floorplanning, macro placement, or global placement, while detailed placement refinement remains relatively unexplored. In this paper, we present GPlaceRL, an open-source graph reinforcement learning framework for detailed placement refinement. GPlaceRL represents legalized placements as graphs and provides a modular environment for studying graph encoders, policy architectures, reward formulations, and local placement actions. To demonstrate the capabilities of GPlaceRL, we conduct a systematic evaluation of proximal policy optimization (PPO) policies with graph attention network (GAT) encoders in a per-design optimization setting. Across five placement benchmarks, the best greedy evaluation results achieve HPWL improvements ranging from 3.27% to 32.87%. The results highlight the importance of compact GAT architectures and flexible local action spaces for placement optimization. Overall, GPlaceRL provides a reproducible and extensible framework for systematic research on RL-based detailed placement refinement.

Hierarchy-aware Heterogeneous Graph Neural Network based Parasitic RC-risk Screening in Pre-layout Stage (434)

Hyemin Ko

Early awareness of interconnect parasitics before placement and routing is crucial for reducing design iterations and achieving efficient timing and power closure. In this paper, we present a hierarchy-aware heterogeneous graph neural network (HGNN) that predicts parasitic RC risk at the pre-layout stage using only netlists, without physical geometry or routing information. The circuit is represented as a hierarchical heterogeneous graph composed of net, pin, instance, and module nodes. Relation-wise graph attention captures electrical connectivity and hierarchical dependencies, while attention-based global context fusion incorporates design-level information into net representations. Compared with an HGNN without module-level hierarchy, the proposed model improves RC estimation and risk ranking, achieving a top-10\% recall of 66.4\%, and a variant-consistent recall of 80.4\%. These results indicate that our method is effective for identifying potential interconnect-risk candidates in the early design stage.

System-Level Optimization of Signal-to-Noise Ratio in Fiber-Optic Plasmonic Sensing Platforms (36)
Seung-Ki Lee

Fiber-optic plasmonic sensors have emerged as promising candidates for compact, label-free, and remote sensing applications. While prior studies have mainly focused on improving intrinsic sensitivity through nanostructure design, practical sensing performance is often limited by system-level factors, including optical coupling efficiency, detector noise characteristics, and interface stability.
In this work, we present a comprehensive system-level investigation of signal-to-noise ratio (SNR) in fiber-optic plasmonic sensing platforms. The proposed study integrates experimental evaluation and system-level analysis to examine the effects of optical routing architecture, source–detector selection, air-gap conditions, index-matching interfaces, and fiber-core geometry.
Experimental results show that LED–photodetector configurations reduce noise by up to a factor of four compared to spectrometer-based systems, while multi-wavelength averaging partially compensates spectrometer noise. Furthermore, air-gap distances below 30 μm maintain near-contact-level SNR, whereas larger gaps introduce both transmission loss and fluctuation. The use of index-matching gel reduces noise by approximately 40%, and large-core fibers (400 μm) provide up to 40× higher SNR than small-core fibers (105 μm).
These results demonstrate that SNR in fiber-optic plasmonic sensors is governed by the interplay of optical, electronic, and mechanical subsystems. The proposed framework establishes practical design guidelines for high-performance, compact sensing systems and highlights the importance of system-level optimization beyond intrinsic sensor sensitivity.

Process Sensors for Neural Network Assisted ADCs (141)
Can Tunca, Robert van Veldhoven, Pieter Harpe and Marco Fattori

A programmable CMOS differential ring oscillator designed for use as a process sensor is described. Sensor output frequency shows strong dependence to process variations of NMOS, PMOS, and polysilicon resistances but has low sensitivity to temperature and supply voltage. Process shifts are detected using threshold voltage and carrier mobility variations on MOS devices and sheet resistance variations of polysilicon resistances. The circuit is designed using a 0.16μm CMOS process and operates with 1.8 V ±10% power supply (1.62 - 1.98 V) in temperature range of −40◦C to 120◦C.

A Dual-Channel Delay Line Oscillator Readout Circuit for Surface Acoustic Wave-Based Prostate-Specific Antigen Sensing in 22nm FD-SOI Technology (315)
Emmanuel John Bonifacio, Nathaniel Renz Domingo, Arcel Leynes, Maria Sophia Ralota, Marc Rosales and Maria Theresa De Leon

This work presents a dual-channel delay line oscillator readout circuit for surface acoustic wave (SAW)-based Prostate-Specific Antigen (PSA) biosensing, implemented in 22nm FD-SOI technology. The architecture utilizes parallel sensing and reference oscillation loops to introduce frequency shifts with a range of 500 Hz-8 kHz from a resonant frequency of 173.4 MHz. With the dual-channel delay-line approach, common-mode noise and temperature variation are suppressed. The readout circuit determines these frequency shifts through the use of a two-stage noise cancellation low-noise amplifier (LNA), a high-swing folded Gilbert-cell mixer, and a 4th-order all-backgate input low-pass Gm-C filter. Using the periodic steady-state analysis, the system achieves a phase noise of -81.89~dbc/Hz at 10kHz, a worst-case jitter of $\pm$27.32~Hz for 1-second averaging time, and a maximum of 2\% error within the range of expected frequency shifts.

A Differential Capacitive Sensor Array with Real-Time ASIC Readout for Dynamic Shear Wave Elastography (414)

Maryam Habibollahi, Arthur Jaccottet, Antonio Gomez, Nader Saffari and Andreas Demosthenous

Capacitive sensing arrays for dynamic shear wave elastography (SWE) demand high resolution, wide dynamic range, and compact, low-power implementations suitable for transluminal applications. Existing discrete readout systems present limitations in size, power consumption, and real-time configurability, restricting their suitability for in-vivo deployment. This paper presents the verification and system-level implementation of a mixed-signal capacitance-to-time (C2T) ASIC interfaced with a differential capacitive sensor array for dynamic shear sensing. An HDL-based verification framework and MATLAB post-processing pipeline are described, alongside a real-time serial data transfer protocol enabling multi-sensor array operation. Simulation results and preliminary system-level verification demonstrate sub-fF resolution with adaptive tuning across a wide input range, supporting dynamic shear detection over 0–1000 Hz. The proposed system advances toward a miniaturised, low-power readout platform for transluminal SWE-based tissue characterisation.

A Parasitic-Tolerant Relaxation Oscillator for GΩ-Range Sensor Interfaces (419)
Allen Boston, Matthieu Couriol, Paul Capgras and Pierre-Emmanuel Gaillardon

Resistive sensors in the GΩ range are often constrained by parasitic capacitances (Cpara) that severely limit the throughput of conventional relaxation-oscillator interfaces. While state-of-the-art designs often utilize DC biasing to mitigate these parasitic effects, such techniques are unsuitable for emerging sensor technologies prone to baseline drift. This paper presents a parasitic-tolerant relaxation-oscillator interface for wide-range resistive sensing (100 kΩ to 50 GΩ) that maintains non-DC operation while suppressing the influence of Cpara. The proposed architecture introduces a mitigation technique that decouples the oscillation period from the sensor parasitic capacitance, enabling higher output rates and reduced silicon area compared to prior non-DC-biased topologies. Validated through post-parasitic extraction (post-PEX) Monte Carlo and noisy transient simulations in 180 nm CMOS, the interface demonstrates output rates up to 100 Hz at high resistance values with a maximum error below 0.62% and a total power consumption of 2.65 mW per channel.

S27 | VLSI Systems and Applications **10:30 AM - 12:00 PM**
Room: Sirocco
Chairs: Nikos Sketopoulos, Nikolaos Athanasios Anagnostopoulos

AdaptViT: Runtime-Adaptive Vision Transformer Deployment on Custom RISC-V (397)
Vishnu P S, Ajay Kumar M, Yike Li, Robert Bogdan Staszewski and Deepu John

Deploying Vision Transformers (ViTs) on low-power edge devices is challenging due to high computational demands. Conventional pruning frameworks require a separate compiled binary for each sparsity level, increasing storage overhead and limiting runtime adaptability. This paper presents an end-to-end deployment pipeline that transforms pretrained ViTs into a single runtime-configurable binary, enabling dynamic compute-budget switching on embedded CPUs. This is achieved by restructuring generated C kernels with modified loop bounds and binary-mask control logic, allowing execution to switch across discrete sparsity levels via compact external configuration files. Compared to multi-binary deployment, the proposed runtime-adaptive approach reduces on-device storage by up to 4.86×, requiring only 163 MB for ViT-Base instead of nearly 800 MB. To maximize pruning efficiency, we introduce a hardware-aligned block pruning strategy for Multi-Layer Perceptron (MLP) layers. In addition, a custom ISA extension is proposed to exploit input-reuse patterns in linear projection kernels. On a Synopsys TRV32P3FX RISC-V processor, the full system achieves up to 2.8× speedup at 65% MLP and 50% attention-head pruning for ViT-Base. The ISA extension alone provides a 1.56× speedup and 33% lower inference energy, with a 24% area overhead in a TSMC 28 nm implementation.

Immediate XOR In-Memory Computing in 6T-SRAM (213)
Viktor Weinelt and Klaus Hofmann

This paper introduces an innovative method for performing XOR operations directly within a 6T-SRAM array. The approach combines an external operand with a word stored in SRAM, where the stored operand is overwritten during execution, effectively implementing an immediate XOR (IXOR) instruction. Unlike conventional bit-line computing, this method eliminates undesirable charge flow between SRAM cells during computation, thereby enhancing stability and reducing the risk of unintended state changes.
The design is validated in a test setup based on a 16×16 SRAM array implemented using unmodified SRAM cells from the GlobalFoundries 22FDX+ PDK, demonstrating a simple setup and performance at up to 1 GHz. The proposed IXOR SRAM enables efficient in-memory computing for applications such as cryptographic operations and algebraic In-Memory computing.

Implementing Hybrid Homomorphic Encryption on a Silicon Root of Trust (250)
Tobias Stelzer, Vishesh Kiran Kumar Bafna and Andreas Zankl

Hybrid Homomorphic Encryption (HHE) makes Fully Homomorphic Encryption (FHE) practical for resource-constrained clients by replacing expensive FHE encryption with a lightweight symmetric cipher, eliminating ciphertext expansion and heavy computation on the client side. However, hardware acceleration for HHE remains largely unexplored, and no prior work has realized HHE through a HW/SW codesign approach on a real-world security platform. We close this gap and present the first HW/SW codesign of the HHE cipher PASTA on the OpenTitan Root of Trust (RoT) by extending its cryptographic coprocessor, the OpenTitan Big Number Accelerator, with 64- bit SIMD instructions supporting moduli beyond 32 bits. A design-space exploration yields five microarchitectural variants trading off area against latency. Our best design in terms of ATP achieves a 13.5x speedup over the software baseline at 3.8% area overhead, demonstrating that a silicon RoT can serve as a practical platform for privacy-preserving data transmission via HHE.

A Gates-Free Approximate 7-2 Adder Compressor for Energy-Efficient Diabetes Classification (282)
Leonardo Antonietti, Rodrigo Lopes, Morgana Macedo Azevedo Da Rosa, Eduardo Da Costa and Rafael Soares

This paper proposes a new gates-free approximate 7-2 adder compressor (Ax7-2 AC) targeting energy-efficient arithmetic circuits. The proposed architecture approximates the K least significant bits using direct input-to-output mappings, while the remaining most significant bits are computed with exact 7-2 compressors. An exhaustive design space exploration evaluates all 94 = 6,561 possible mappings between the nine compressor inputs and the four output signals for each approximation level. The selected configurations are implemented and synthesized in a 65 nm CMOS technology at 100 MHz. In addition to the arithmetic-level analysis, the proposed Ax7-2 AC is evaluated in an AI-based case study for diabetes classification using the PIMA Indians Diabetes dataset. In this case study, seven

selected clinical features are processed through a multiplierless perceptron datapath based on shift operations and a 7-2 compressor reduction stage. The synthesis results show that the proposed Ax7-2 AC reaches energy savings of up to 93.62% and area savings of up to 78.61% at the compressor level. When applied to the diabetes classifier, the proposed architecture achieves 90.76% classification accuracy and an F2-score of 0.8871 with 88.81% energy savings for K = 5.

Versat-AI: From ONNX to Area-Efficient RISC-V+CGRA Edge AI Subsystems (398)
Rúben Teixeira, João Barreiros Rodrigues, Jaime Aguiar and José T. de Sousa

Deploying neural network inference on resource-constrained embedded systems requires efficient hardware and software. This paper presents Versat-AI, an open-source frame-work that converts ONNX models into a RISC-V + CGRA subsystem for running inference using single-precision floating-point as a first implementation, with fixed-point quantization as planned future work. Versat-AI co-generates the RISC-V application software and a companion hardware accelerator. Versat-AI is built on top of Versat (CGRA compiler) and Py2HWSW (Python framework for codesigning hardware and software). Py2HWSW generates a compact and energy-efficient SoC; the Versat accelerator alone uses fewer than 10k LUTs and FFs, 4 DSPs, and 200 BRAMs on a Xilinx Ultrascale FPGA. Inference results on the MLPerf Tiny benchmark indicate speedups of 1.7–18× over ONNX2C.

S31 | Digital Circuits and Systems 3 **10:30 AM - 12:00 PM**
Room: Ostria
Chairs: George Panagopoulos, Mattis Hasler

Robust h-Parameter Extraction for RF Oscillators using Multi-Start Optimization and Upper-Bound Constrains (225)
Athina Papagianni and Georgios Panagopoulos

Translating oscillator datasheet specifications into unified power-law noise coefficients is a persistent challenge in circuit modeling. Traditional symmetric optimization treats maximum datasheet limits as precise goals and often generates models that violate the manufacturer's specifications in order to minimize the aggregate error. This paper presents an automated extraction methodology using a Multi-Start optimization algorithm with an asymmetric Upper-Bound cost function. By mathematically penalizing outputs that exceed limits, the proposed algorithm ensures strict adherence to worst-case performance bounds. The tool incorporates a dual analytical and crossing point seeding strategy for reliable convergence across diverse oscillator profiles, from precision quartz-based oscillators to atomic clocks. Validation against commercial components is necessary to demonstrate the tool's ability to generate high fidelity models and to perform reliable parameter extraction for electronic design automation.

A Resonant Low-Pass Ladder Filter for Sigma-Delta Bitstreams on FPGA (294)
Sven Finke, Rolf Paulus and Michael Karagounis

We present a sigma-delta bitstream realization of a resonant low-pass ladder filter based on a continuous-time reference model. Processing in the bitstream domain enables a multiplierless FPGA implementation. The high oversampling ratio keeps both the per-stage discretization error and the phase error from the global feedback delay small, allowing the continuous-time topology to carry over directly. The architecture cascades four first-order low-pass stages with global feedback to provide simultaneous 6, 12, 18, and 24 dB/oct outputs with runtime-tunable cutoff, resonance, and input gain. The Artix-7 implementation uses 566 LUTs and 318 flip-flops, with no DSP blocks or BRAM. Hardware measurements at the 24 dB/oct output match the reference cutoff attenuation within 0.3 dB across seven octaves and demonstrate operation up to stable self-oscillation.

A 3-ps-Resolution Wide-Tuning-Range Coarse-Fine Digitally Controlled Delay Line in 180 nm CMOS (360)
Lazar Ilić, Đorđe Gačić, Dejan Petković and Vladimir Milovanović

This paper presents a standard-cell-based coarse-fine digitally controlled delay line (DCDL) architecture, which aims to establish a path toward an agile, automated design flow. The proposed design features a two-stage topology utilizing a NAND-based coarse tuning stage (CTS) and a NOR-based digitally controlled varactor (DCV) fine tuning stage (FTS) to decouple power, performance, and area (PPA) trade-offs. Efficient mathematical delay approximation models are explored to enable rapid pre-layout estimations with accuracy comparable to SPICE simulations. The prototype is fabricated in a baseline 180 nm CMOS technology, occupies 13125 µm2, and consumes 0.6 to 3.16 mW. Measurements indicate a CTS resolution of 143.12 ps, an FTS resolution of 3.49 ps, and a maximum tunable delay range of 5.52 ns. The maximum measured DNL/INL values are 0.1/0.29 LSB for the CTS and 0.2/1.09 LSB for the FTS. These results confirm that all-digital architectures can deliver fine temporal resolution and wide tuning range without major linearity, power, and area penalties, proving the design methodology feasible.

End-to-End Hardware Validation of OSSDM Semantic Waveform for 6G Communications (372)
Lahcen Attijani, Panagiotis Vlachos, Mohamed-Amine Hamoura, Emilio Calvanese Strinati and François Rivet

Future networks are evolving toward semantic communications, where preserving task-relevant information becomes more important than ensuring bit-level reliability. Recent work introduced Orthogonal Semantic Sequency Division Multiplexing (OSSDM), a Walsh-based semantic waveform that outperformed semantic OFDM (S-OFDM) in simulation. However, its end-to-end performance over a real wireless link remained unverified. This paper presents, to the best of the authors' knowledge, the first hardware demonstration of an OSSDM-based semantic communication system using a Universal Software Radio Peripheral (USRP) X310 software-defined radio (SDR) at 2.45 GHz. CIFAR-100 images are transmitted through OSSDM and an S-OFDM baseline, and evaluated using

classification accuracy as the task-level metric. While both systems achieve a similar Tx-side accuracy ceiling of about 76%, OSSDM consistently preserves higher accuracy after wireless transmission. These results provide the experimental validation that the robustness gains of OSSDM observed in simulation translate into improved task-level performance over a real wireless channel.

A Low-Power and Area-Efficient Standard-Cell-Based DIGOTA for Flexible Electronics (388)
Noor Ul Amin, Andrea Ballo, Alfio Dario Grasso and Orazio Aiello

This paper presents a digital operational transconductance amplifier (DIGOTA) implemented using Indium-Gallium-Zinc Oxide (IGZO)-based flexible substrate technology. The DIGOTA is realized using standard cells provided by the Pragmatic PDK. Post-layout simulations result in a DC gain of 51.7 dB, a phase margin of 57.3º, and a gain-bandwidth product (GBW) of 114.4 kHz. For a capacitive load of 200 pF, the circuit exhibits an average slew rate of 8.8 V/ms. Compared with state-of-the-art flexible DIGOTAs and analog OTAs, the DIGOTA consumes a low power of 18.7 µW and occupies a compact area of 0.0074 mm2. Furthermore, the small-signal figure of merit is improved by 3.8x-301x, while maintaining comparable large-signal figure of merit performance.

S05 | Analog/mixed signal/RF circuits 5 **2:15 PM - 3:45 PM**
Room: Maistros ACD
Chairs: Costas Psychalinos, Vladimir Milovanović

Multi-Modulus Frequency Dividers for System-On-Chip Clock Generation in 28-nm CMOS (97)
Tze Hin Cheung, Kari Stadius, Aarno Pärssinen and Jussi Ryynänen

Multi-modulus frequency dividers (MMD) are utilized in frequency synthesizers for LO-signal generation and as a part of the clock tree within digital circuits. In this paper, we compare three different variations of frequency dividers with alternating architectures and logic families. Two of the designs are based on current-mode logic, and the third one is a true-single-phase-clock (TSPC) logic divider. The paper presents the circuit structures and design procedures in detail. The performance of the MMDs is extensively characterized with post-layout PVT simulations and measurement results are provided for comparison and decision making. A prototype circuit that includes these three MMD variants within a single overall chip is implemented on 28-nm bulk CMOS process. The proposed designs were measured to cover an operating frequency range of 0.3–7.1 GHz while consuming 5.4–17.9 mW at a 0.9V supply. The measurements determined the TSPC option to be the most efficient when considering overall frequency range and power consumption together, while current-mode logic provides the highest operation frequency.

Effect of scaling on Shot noise in MOSFETs and Its overall impact on Phase Noise in Ring Oscillators (138)
Shubham Sharma, Anshul Verma and Sudeb Dasgupta

Shot noise present in MOSFET channels is often excluded for the Phase Noise (PN) analysis, which may lead to inaccurate PN predictions for advanced-node CMOS-based Ring Oscillators (ROs). To resolve this concern, an analytical approach for estimation of phase noise specifically considering shot noise (PNSN) in advanced-node CMOS-based ROs is presented. It has been showed that with scaling of technology nodes, shot noise and its overall impact on white noise increase, which makes its analysis more crucial as we shift towards more advanced technology nodes. The temperature dependency of various parameters which can affect PNSN has been reviewed. The proposed work is verified and validated in 22 nm FDSOI process technology for the temperature range down to 233 K, with an accuracy of about 96%.

A High-Linearity and Size-Optimized Synthesizable Cyclic Time-to-Digital Converter in 180 nm CMOS (286)
Dejan Petković and Vladimir Milovanović

A fully synthesizable cyclic time-to-digital converter (TDC) in 180 nm CMOS is presented. The described architecture employs an inverter-based delay ring with dual cycle counters for metastability-robust operation. An optimized delay-ring length minimizes area and hardware cost while enabling low-complexity digital post-processing. Linearity is improved by selectively inserting dummy inverter capacitive loads in the fastest delay paths using back-annotated timing information. A straightforward calibration scheme is used to estimate the average delay-cell propagation. The prototype occupies 4500 µm2 and typically dissipates 2.7 mW at 1.8 V. It achieves 52 ps resolution over a 6.6 ns full-scale range, with measured DNL/INL of 0.38/0.7 LSB and single-shot precision of up to 16 ps/0.31 LSB. Owing to its small size, good linearity, and standard-cell-only implementation simplicity, the proposed TDC is well suited for compact digital-intensive systems requiring many integrated timing channels.

An LC VCO with Bias Circuit Noise Rejection using a Pseudo-Resistor-Based Low Pass Filter (347)
Mouli Venkata Prakash Pittala and Chithra

In this work, we propose a technique to filter noise from the bias circuit of the tail current source in an LC VCO, which is a dominant contributor to the output phase noise. This work demonstrates the suppression of the bias circuit noise using an on-chip 0.5Hz low-pass filter using a pseudo-resistor. Post-layout simulations show an FoM improvement by 19.5dB at 1kHz offset in the proposed technique compared to an LC VCO without filtering, and an improvement of 1.1dB compared to an LC VCO which is aggressively optimized for bias circuit noise. For similar phase noise performance, the proposed LC VCO consumes 15.3% less power than the VCO with the bias circuit optimized for noise.

Performance Optimization of a Cyclic Fine Interpolator in Nutt TDCs via Delay-Cell Trimming (408)
Javad Ahmadi-Farsani, Ion Vornicu and Laurentiu Acasandrei

In high-resolution Nutt Time-to-Digital Converters (TDCs), the precision of the fine interpolation stage is crucial. This paper investigates the optimization of a cyclic fine interpolator by calibrating the trimming voltage of delay cells within its integrated time amplifier (TA). Building upon previous work that characterized a standalone TA with a gain of two, this study embeds the TA directly into a Nutt TDC architecture to evaluate system-level performance enhancements. Fabricated in a 180nm CMOS process and operating with a 100 MHz reference clock, the TDC achieves its highest precision at an optimal trimming voltage of 900 mV, reducing the standard deviation (σ) of the 10-bit fine counter output to 4.781 LSB. In contrast, shifting this calibration voltage to ground (0V) or the supply voltage (1.8V) significantly degrades the standard deviation to 7.36 LSB and 6.784 LSB, respectively. These experimental results validate that delay-cell trimming effectively mitigates conversion errors, aligning closely with standalone component characterizations.

S07 | Biomedical and Bio-inspired Circuits and Systems 2 2:15 PM - 3:45 PM
Room: Maistros BEF
Chairs: Vasileios Alimisis, Meysam Akbari

A Neuromorphic Tactile Sensing System for the Hannes Anthropomorphic Prosthetic Hand (230)
Leila Hamadi

Efficient artificial touch in prosthetics demands two key factors: low-power computing and robustness to uncontrolled interactions, where contact force and sliding velocity vary unpredictably. This work presents a neuromorphic tactile sensing system integrated into the Hannes prosthetic hand, targeting texture classification under such realistic conditions. The system employs eight flexible piezoelectric sensors mounted on the prosthetic index finger, capturing spatio-temporal tactile signals while the prosthetic hand is manually guided over texture surfaces. The acquired signals are converted into spike trains using a sigma-delta modulator and processed by a Recurrent Spiking Neural Network (RSNN). We systematically analyze the effect of encoding threshold and spike polarity on the tradeoff between classification accuracy and computational cost. The proposed system achieves a peak accuracy of 92.05% using combined ON/OFF encoding, while an OFF-only configuration reaches 91.39% accuracy with only 85.77K synaptic operations, offering a potential candidate for low-power neuromorphic de ployments. These results pave the way toward the deployment on neuromorphic hardware, enabling real-time and low-power tactile texture processing for prosthesis.

Distilling Privileged Information into Efficient Neural Networks for Sepsis Detection (259)

Leonidas Paratsokas, Dimitrios Spanos, Nikolaos Passalis and Anastasios Tefas

Sepsis is a time-critical condition whose accurate severity stratification often depends on clinical attributes that are slow or costly to acquire, yet point-of-care deployment in embedded systems demands fast predictions from only routinely available measurements. Existing deep models assume that the full attribute set used in training remains available at inference, which fails in time-constrained settings where such privileged attributes are absent. We propose a privileged knowledge distillation framework that transfers knowledge from a teacher trained on privileged information to a lightweight student restricted to readily available attributes at inference, concentrating the privileged supervision on the cases where it is decisive and transferring it at both the output and intermediate layers. On a real clinical dataset, the proposed method improves binary and multiclass sepsis classification over all baselines while operating at a 5x reduction in parameters and FLOPs relative to the privileged model. These results confirm privileged knowledge distillation as an effective strategy for accurate, lightweight deployment in real-time, resource-constrained biomedical systems.

Improving Calibration in Lightweight Biomedical Deep Learning via Soft Subclass Adaptation (262)

Dimitrios Spanos, Nikolaos Passalis and Anastasios Tefas

Deep neural networks are increasingly deployed in embedded biomedical decision-support systems, where dependable operation rests not only on accuracy but on well-calibrated confidence, since miscalibrated outputs can trigger unsafe clinical actions. Training with rigid one-hot labels, however, ignores the latent structure within each diagnostic class and assumes every sample is equally certain, which degrades calibration. We propose Soft Subclass Adaptation (SSA), a training scheme that lets a network refine its own targets from latent subclass structure and expert clinical priors instead of fixed labels. Per-class prototypes yield soft subclass targets that are adapted online by an auxiliary similarity network, the predicted subclass is aligned across network depth, and clinically implausible labels are rewritten using domain priors. On a real sepsis dataset, SSA improves accuracy, F-score, and three calibration metrics over strong baselines. Crucially, SSA introduces zero overhead during inference, as all added machinery is confined to training. The deployed model is an unchanged lightweight backbone, meaning inference latency and memory match those of a baseline network — an essential property for time-critical, resource-constrained clinical hardware.

AcoustEEEcare: A Portable Audio-Based MEMS Patch for On-Device Heart-Rate Monitoring (363)

Jose Marie Acedillo, Karylle Shainne Cadalzo, Mark Louie Pasturan, Paul Jason Co, Marc Rosales, Nathaniel Orillaza, Ron Louis Nierva and John Richard Hizon

Wearable acoustic sensing offers a single-sensor route to cardiac vital signs, but most systems offload computation to a phone or server, breaking the promise of a self-contained patch. This paper presents AcoustEEEcare, a skin-conformable acoustic sensing chain that estimates heart rate (HR) entirely on a microcontroller. The front end couples an analog MEMS microphone to the skin through a flexible printed circuit board and a biocompatible film. Audio is digitized on a Seeed Studio XIAO nRF52840 Sense and processed on-device through band-pass filtering, Mel-Frequency Cepstral Coefficient (MFCC) extraction, and a quantized int8 two-dimensional convolutional neural network (2D-CNN) deployed using TensorFlow Lite Micro. The deployed model occupies 28.5~kB of flash with a 104~kB tensor arena. Preliminary results show on 650 recordings from five participants across three auscultation sites, chest placement yields an HR mean absolute error (MAE) of 2.44–2.88 bpm against pulse-oximeter and smartwatch references, within the ANSI/AAMI EC13 tolerance, showing clinically acceptable cardiac monitoring is achievable entirely at the edge.

NeuroHub: A Software-Defined Neural Interface for Dynamic Electrode Routing, Biosignal Acquisition, and Four-Wire Impedance Spectroscopy (412)

Ebubekir Enes Arslan, Jhan Luke Okkabaz, Gregory Worrell and Nuri Ince

Modern electrophysiology and neuromodulation studies increasingly rely on high-density electrode arrays comprising tens to hundreds of contacts. However, existing acquisition and stimulation systems often provide limited flexibility for dynamically accessing large electrode populations while simultaneously supporting electrophysiological recording and tissue impedance monitoring. We present NeuroHub, a programmable neural interface platform that integrates a solid-state analog routing matrix, high-resolution biosignal acquisition, and two-point and four-point impedance spectroscopy within a compact modular architecture. The current prototype dynamically routes 32 electrode inputs to four configurable analog pathways using software-controlled switching, enabling recording, stimulation, and impedance measurements through a shared electrode access network. The system combines an ADS1299-based biosignal acquisition front end, an AD5940-based impedance spectroscopy module, and an ESP32-C6 controller for real-time routing and system management. Bench validation demonstrated relay-switched ECG acquisition with a 0.43 µV RMS error relative to a direct connection, sequential routing across 32 channels without observable waveform degradation, accurate biosignal acquisition, and impedance spectroscopy measurements consistent with analytical RC-network models. These results establish the feasibility of dynamically configurable electrode access for integrated electrophysiology and impedance monitoring. The proposed architecture provides a scalable foundation for future 256- to 1024-channel research platforms supporting bedside neuromodulation, impedance-guided neural monitoring, and the development of next-generation implantable neurotechnologies.

VQ4SNN: Vector Quantization for Memory-Efficient FPGA Spiking Neural Networks (214)
Dimitrios Sekertzis and Giorgos Dimitrakopoulos

Spiking Neural Networks (SNNs) offer an energy-efficient paradigm for edge AI, making them attractive for hardware acceleration. However, deploying dense SNNs on FPGAs is constrained by limited on-chip memory for synaptic weight storage. To address this bottleneck, we propose VQ4SNN, a hardware-aware architecture that reduces memory requirements through Vector Quantization (VQ)-based weight sharing. To the best of our knowledge, this is the first application of VQ to pipelined spatial-dataflow SNN accelerators on FPGAs. VQ4SNN replaces conventional weight storage with a two-level memory organization consisting of compact pointers and a shared codebook of quantized weight vectors. The proposed design integrates FPGA-aware memory mapping with analytical VQ parameter selection, enabling efficient deployment on such accelerators while preserving inference accuracy. The experimental results show a reduction of 52-61% in the total number of BRAMs compared to the state-of-the-art uncompressed FPGA SNNs without increasing overall logic utilization.

A Time-Based Readout for Vector-Matrix Multiplication in Fully Analog Memristive SNNs (148)
Elia Mateu Barriendos, Álvaro Gómez-Pau, Josep Rius, Daniel Arumí, Rosa Rodríguez-Montañés and Salvador Manich

Artificial neural networks rely on vector–matrix multiplications (VMMs), whose implementation in von Neumann architectures is dominated by costly data movement between memory and processing units. Spiking neural networks (SNNs) mitigate this bottleneck by performing in-memory, analog VMMs using memristive crossbar arrays. However, conventional current-mode readout circuits incur significant area and power overhead.
This work proposes a fully analog readout architecture based on voltage-to-time conversion of the VMM output. By sensing the column voltage, the proposed approach avoids current-mode summing and scaling circuitry, improving area and energy efficiency. Post-layout results in 130nm CMOS validate the proposed readout scheme for fully analog memristive SNNs.

Train Where You Infer: A System-Level Analysis of On-Device Forward-Forward for IMC Achitectures (149)
Neha Ashar, Aymen Romdhane, Kamel-Eddine Sadadou, Johannes Reibold, Guillaume Prenat, Mehdi Tahoori, Jonathan Miquel and Lorena Anghel

In-Memory Computing (IMC) promises to address computing and memory bottleneck of AI workloads. However, IMC adoption is hindered by low memory density which necessitates heavy model quantization, and analog non-idealities leading to computational inaccuracies, a problem usually tackled with hardware-in-the-loop training. As traditional back propagation is too resource-intensive for edge training, designers are forced to rely on static inference. To address these hardware challenges, we propose using the forward-forward (FF) algorithm to enable direct on-device model training at the edge. This approach allows models to adapt dynamically to analog non-idealities without requiring prior hardware knowledge. We propose a complete system-level architecture that maps the FF algorithm onto an edge device paired with analog IMC crossbars. Our implementation requires 390ms and 314µJ per training step (128 samples). We show that the dynamic FF model tuning significantly outperforms static inference on MNIST, CIFAR-10 and Google Speech Commands datasets.

Pareto-Optimal Accuracy, Area, and Power in Time-Domain Analog Neural Networks via Hardware-Aware Training (237)
Giulio Filippeschi, Giuseppe Iannaccone, Mirko Brazzini, Cristhopher Mosquera, Marco Lanuzza, Alessandro Catania and Sebastiano Strangio

Silicon-based analog neural networks suffer from dynamic non-idealities that degrade inference accuracy. We present a hardware-aware (HWA) training framework for timedomain (TD) vector-matrix multipliers (VMM) using singletransistor floating-gate memory cells. Unlike conventional models for TD approaches, our method discretizes operations into adaptive time slots to capture instantaneous distortions such as bitline voltage drop and layout-dependent crosstalk. By integrating this model directly into the training loop, it becomes possible to mitigate the loss in classification accuracy caused by nonidealities. For example, the classification error of a VGG-style CNN on CIFAR-10 can be reduced from 29.68% to 6.37%. While conservative analog designs mitigate crosstalk through larger layouts and preserve bit-line voltage accuracy using high-gain and power-hungry amplifiers, our Pareto-optimal strategy instead tolerates both effects and compensates for them through HWA retraining, enabling aggressive asymmetric layouts, achieving a 39% area reduction and eliminating the power overhead typically required for design-level compensation.

MELIF: Design of Multi-Encoded Leaky-Integrate-and-Fire Neuron in 28 nm CMOS Technology (361)
Sahibia Kaur Vohra

Designing a neural encoder that can translate signals into spikes is essential for applying input to spiking neuromorphic systems. The most suitable encoding scheme depends on the application, architecture, and requirements of the system. In this work, we present a compact CMOS circuit of a multi-encoded leaky-integrate-and-fire (MELIF) neuron capable of translating the input sensory data into rate-encoded and temporal time-to-first-spike (TTFS)-encoded signal. The MELIF neuron exhibiting TTFS encoding is further extended to demonstrate inter-spike-interval (ISI) encoding. In addition to its multi-encoding property, the neuron circuit offers various adjustable features such as threshold, pulse width and refractory period control which facilitates the use of the proposed neuron for various applications at different frequencies. Simulations of the designed neuron circuit in 28\,nm technology verify its multi-encoded property and adjustable features.

S21 | Energy Harvesting and Power Management 2 **2:15 PM - 3:45 PM**
Room: Sirocco
Chairs: Vasiliki Gogolou, Spyridon Vassis

Waste Heat Recovery from a 3.6 GHz GaN Doherty Power Amplifier using Thermoelectric Generators (89)
Lucas Ott, Robert Krämer, Xin Xu, Jens Wagner and Frank Ellinger

This paper investigates waste-heat recovery from a radio-frequency (RF) power amplifier (PA) using thermoelectric generators (TEGs). An electrothermal equivalent model is used to analyze the interaction between the PA, the TEG, and the heat sink, revealing the existence of an optimal number of thermoelectric leg-pairs for maximum recoverable power. To validate the approach, a 3.6 GHz gallium nitride (GaN) Doherty PA is used as a heat source. Measurements demonstrate a recovered power of 89 mW, corresponding to a system-level efficiency improvement of maximum 0.74%. Although the recovered power represents only a small fraction of the dissipated energy, it is sufficient to supply low-power peripheral transceiver modules. The presented analysis further provides design guidelines for thermally optimized TEG-based harvesting systems in RF applications.

Brownout-Aware Energy Management for Wearable Health Patches Using SMM--MDP Optimization (195)
Hajer Abdelaziz, Abdullah Yousuf, Shamma Alblooshi, Sara Alessa and Yasmin Halawani

Wearable health patches are usually evaluated using average-power battery lifetime, although practical failures can occur earlier when optical sensing or wireless-transmission pulses collapse the battery voltage below the usable range. This paper presents a brownout-aware energy-management framework that explicitly separates physiological workload dynamics, wireless transmission events, average energy depletion, and voltage-headroom feasibility. A three-state Semi-Markov Model (SMM) is derived from MMASH sleep and activity annotations using empirical SLEEP, LOW, and HIGH dwell-time distributions, while transmission is modeled as a device-level overlay process rather than as a physiological state. The SMM is coupled with a conservative rate-derated battery model and a brownout guard that checks whether each battery can sustain action-dependent sensing and TX peak currents. A Markov decision process (MDP) then selects among practical firmware modes, including nominal sensing, feature processing, batching, event-alert transmission, and adaptive duty-sensing with LED-current reduction. Monte Carlo validation confirmed internal model consistency, with event-level simulated lifetimes matching analytical estimates within 1.81\% across the evaluated battery cases. Relative to nominal sense-send operation, adaptive duty-sensing extended useful monitoring lifetime by approximately (54–60)\% in the reported battery cases while maintaining or reducing voltage-warning rates. In the constrained CR1632 primary coin-cell case, brownout-aware MDP optimization further increased lifetime by 67\% under the evaluated QoS setting by avoiding overly aggressive high-load actions.

ADRC-Enhanced MPPT with Sensorless Open-Circuit Fault Diagnosis and Reconfiguration for Three-Level PV Boost Converters (269)
Abdeldjabar Benrabah, Mouaad Belguedri, Mohamed Amine Yahiaoui, Fayçal Benyamina, Lakhdar Benhamimid and Mohamed Benbouzi

This paper presents a robust fault-tolerant control framework for a three-level boost converter in photovoltaic (PV) systems. An Active Disturbance Rejection Control (ADRC)-based maximum power point tracking (MPPT) strategy is proposed to improve tracking performance under irradiance variations and disturbances. A sensorless open-circuit fault diagnosis method is developed using existing control variables, eliminating the need for additional sensors. Upon fault detection, the converter is automatically reconfigured from three-level to two-level operation, ensuring continuous power transfer without extra hardware. Simulation results validate the effectiveness of the proposed approach, demonstrating improved MPPT robustness, reliable operation, and uninterrupted operation under varying operating conditions.

Dual-Mode Ultrasonic System Enabling Simultaneous Energy Harvesting and Data Communication in Water and Ex-Vivo Tissue Channels (350)
Ziyi Zhao, Qihan Wang, Chenxue Hou, Zhaoxi Li, Dongdong Chen, Di Li and Yintang Yang

Ultrasonic links are attractive for communication through liquid and tissue-like media, but a receiver that also extracts power must balance waveform fidelity, harvested energy, and circuit complexity. This paper presents a dual-mode ultrasonic link in which the received signal is split into a communication branch and an AC-DC/DC-DC energy-management branch. In the high-rate mode, orthogonal frequency-division multiplexing (OFDM) is used with the energy branch inactive. In the energy-harvesting mode, binary frequency-shift keying (2FSK) provides robust low-rate data transfer while the rectifier and buck converter generate a regulated 3.3 V supply for low-power acquisition/control electronics. The operating frequency is selected from impedance/phase measurements and a water-channel pulse-echo sweep, yielding a 1.02 MHz operating point. Experiments over a 14 cm path in water and ex-vivo beef tissue demonstrate FSK communication up to 80 kbps, an OFDM physical-layer rate of 241.67 kbps, a regulated 3.3 V output, and maximum rectified powers of 12.6 mW in water and 114.302 mW in beef tissue.

S30 | Emerging technologies and technology trends 3 **2:15 PM - 3:45 PM**
Room: Ostria
Chairs: Ioannis Messaris, Iraklis Anagnostopoulos

Stochastic Resonance for Variability-Tolerant Memristor IMPLY Logic (327)
Rana Walied Ahmad, Rosana Rodriguez and Vasileios Ntinas

Switching variability is the main barrier to reliable in-memory computing with memristive material-implication (IMPLY) gates, because one fixed programming pulse overdrives some cells and fails to switch others. We show that deliberately adding small, zero-mean voltage noise rescues the gate, a circuit instance of stochastic resonance in which the nonlinear switching kinetics turn added noise into a reliability gain. Using the JART-TUD compact model of HfO_x valence-change memristors under realistic variability, we study this assist and find that the point at which the noise enters the gate strongly shapes its effect. Noise on the drive rail rescues the slow cells but eventually disturbs the cells that must hold, which caps the gain. Noise injected at the node resistor instead appears as a common-mode signal that the stored memristance cancels at the gate node in the hold case, so it widens the usable bias window threefold and exceeds the best noise-free correctness at one fifth of its energy when combined with a shorter pulse, while drive-rail injection reaches moderate correctness at the lowest energy per operation. Across the whole correctness-versus-energy trade-off, the noise assist improves on every operating point of the noise-free design.

An arbitrary scale programmable Flash ADC based on analog memristor ladder (377)

Diego Guglielmi, Matej Horský, Marián Precner, Sergei Krylov, Tomáš Ščepka, Tibor Izsák and Boris Hudec

Analog memristors are emerging semiconductor devices which offer precise programming of their resistance and can be fabricated using back-end-of-line CMOS compatible process. In this study we have fabricated and adopted HfO2/TiO2-based analog memristors (instead of resistors) for an experimental realization of an arbitrary scale programmable 3-bit Flash ADC. By differently programming this memristor ladder we demonstrate ADC functionality with six different non-linear transfer functions, namely exponential linear unit (ELU), rectified linear unit (ReLU), hyperbolic tangent (tanh), exponential (exp), hyperbolic sine (sinh), and logarithmic (log).

Assessment of an 8T Dynamic-Cascode SRAM-CIM Bitcell in IMEC N2 GAA Technology (381)

Pio Esposito, Andrea Fasolino, Rosalba Liguori, Luigi Di Benedetto, Alfredo Rubino and Gian Domenico Licciardo

SRAM-based computing-in-memory (CIM) reduces data movement in edge-AI accelerators, but its accuracy is limited by PVT variations. This work evaluates the migration of a dynamic-cascode 8T SRAM-CIM bitcell from 65 nm Bulk CMOS to IMEC N2 GAA technology. The N2 cell is sized and compared with the original implementation in terms of read stability, write-ability, and CIM operation. Read stability is assessed using Butterfly Curves and Read N-curves metrics, while write-ability is evaluated through Write N-curves. At VDD = 0.7 V, the N2 implementation shows comparable voltage-domain read stability and higher current-domain read/write margins. Bitcell's current robustness is then analyzed through 2000-sample Monte Carlo simulations under voltage-programmed and current-programmed Input Activations. Current programming strongly reduces current dispersion under Global process variation. However, Global-plus-Local variability in the N2 implementation significantly degrades this benefit, identifying local variability as a key limiter for current-programmed analog SRAM-CIM in deeply scaled GAA technologies.

Hybrid CMOS Spin Wave Implementation of Convolutional Neural Networks (272)

Pantazis Anagnostou, Arne Van Zegbroeck, Said Hamdioui, Christoph Adelmann, Florin Ciubotaru and Sorin Cotofana

The rapid growth of Deep Neural Networks (DNNs) has raised serious concerns over their energy cost, motivating the search for ultra-low power computing paradigms. Among these, Spin Wave (SW)-based computing is particularly promising, as SW propagation is considered energy-free. Although an ultra-low power SW-based $2{\times}2$ convolution device has been demonstrated, its extension to larger kernels, which is essential for non-toy Convolutional Neural Networks (CNNs) implementations, requires prohibitively time-consuming micromagnetic simulations. In light of this, in this paper, we first introduce a Multiply-Accumulate (MAC) SW-device that provides the means for weighted-sum calculations, which are fundamental for convolution evaluation. Building on this, we propose a framework

that scales software-trained weights and activations within the SW device dynamic ranges to enable highly power-efficient convolutional layer inference calculation within the SW domain. We verify the approach by emulating convolutional layers under standard image-processing kernels and achieved a Root Mean Square Error (RMSE) of \num{0.0026} for a Gaussian Blur kernel, and \num{0.1385} across 1000 randomly generated kernels. Finally, we briefly discussed CNNs' potential hybrid CMOS-SW implementation by means of the proposed approach.

MemrisTec Memristor Model Platform: An Open Educational Framework for Studying Memristive Devices (443)
Christian Bruchatz, Dimitrios Prousalis, Ahmed Magdy Ibrahim Abdelsamad, Mohammad Fahim Sultan Anoy, Ahmet Samil Demirkol, Adnan Haidar, Umer Amer Khan, Carsten Knoll, Ioannis Messaris, Sarthak Dharmendra Nandanwar, Pawan Patel, Jonathan Pitzschel, Anagha Raghunath, Nithya Kalyani Ramachandran, Steffen Seitz, Richard Schroedter, Hongliu Yang and Ronald Tetzlaff

Memristive devices are promising candidates for next-generation non-volatile memory, in-memory computing, and neuromorphic systems, and the number of compact models published every year keeps growing. Keeping track of all of these models, and understanding how they actually behave, is becoming increasingly hard. This paper presents the MemrisTec Memristor Model Platform, an open educational framework developed at TUD Dresden University of Technology that gathers recent compact memristor models in a unified Python environment and lets website visitors directly interact and play with them. Rather than offering yet another written review, the platform invites students and interested newcomers to obtain hands-on experience with the published models, deepen their intuition for the underlying device physics, and become aware of the diversity of approaches that exist in the literature. The platform is built on three pillars: a curated Memristor Model Table, an interactive Model Simulator, and a Model Repository hosted on a publicly reachable GitLab. As an illustrative example, we demonstrate the simulator using the simplified JART model for VCM devices. A future extension, the Circuits Simulator, will allow users to study the behaviour of memristive circuits utilising different models.

POSTER SESSION 1

Monday November 9th, 2026, 3:30 PM - 5:00 PM

Immediate Boolean IMC - Enabling Area-Efficient SRAM In-Memory Computing. (2)

Viktor Weinelt, Hannes Simon Korn and Klaus Hofmann

Abstract. This paper presents a novel approach to integrating the Boolean instructions AND and OR in 6T SRAM, enabling in-memory computing. Unlike in conventional state-of-the-art implementations, computing is performed during write access and combines an external operand with an operand stored in SRAM. The result of the instruction is available in the same SRAM cells immediately after the instruction. By executing the instruction during write access, rather than during read access, typical issues arising from conflicting design targets can be avoided, such as static noise margin (SNM), write ability margin (WAM), read current, and layout area. Furthermore, a writeback of the result into memory is no longer required, and with a capable word-line driver, the same immediate instruction can be applied to multiple words simultaneously. The concept is proven, using two SRAM test setups based on modified versions of available SRAM standard cells in GlobalFoundries 22FDX+ technology. These operate at a maximum frequency of 1 GHz. Additional required layout area needed, solely originates from the need for more complex word-line drivers, compared to SRAM without IMC capability. As the cells themselves remain the same size as before integrating the IMC capability, this approach accomplishes high area efficiency.

A Fully Synthesizable, Technology-Independent Digitally-Controlled Delay Line. (9)

Hannes Simon Korn, Viktor Weinelt and Klaus Hofmann

Abstract. Digitally-Controlled Delay Lines (DCDL) are key components in integrated circuits. They allow for precise timing of signals, such as centering a clock within a received data eye. Conventional DCDL implementations predominantly rely on full- or semi-custom designs that incorporate analog elements, such as variable capacitors or tri-state drivers. These approaches require significant effort for manual schematic design and layout, resulting in considerable development costs and limited portability across technology nodes. This work addresses these limitations by introducing a fully synthesizable, all-digital DCDL architecture. By utilizing automated digital design flows, the proposed methodology eliminates labor-intensive manual steps, significantly reducing design time and enhancing scalability. Furthermore, the presented workflow facilitates seamless migration to advanced technology nodes, ensuring rapid adaptation to evolving process technologies.

Frequency Domain Equalization Design for Gigabit Automotive Ethernet System. (64)

Chih-Feng Wu, Muh-Tian Shiue, Yu-Jun Chen, Min-Yuan Tsai and Tong-Da Lee

Abstract. In this paper, a frequency domain equalization scheme is presented for IEEE 802.3bp 1000BASE-T1 automotive Ethernet transceiver. The frequency domain equalization scheme is composed of a feed-forward equalizer (FFE) in the frequency domain and a feed-backward equalizer (FBE) in the time domain. The simulation results show that the frequency domain equalization can effectively suppress the inter-symbol interference (ISI) with the improvement of signal-to-interference ratio (SINR) around 0.2 dB compared with the conventional time domain equalization (TDE) regardless of channel lengths 15- and 40-m. Considering the BER less than 10^(-6), the given signal-to-noise ratio (SNR) for channel lengths 15- and 40-m only requires larger than 23 and 28 dB, respectively.

A 28 Gbps PAM-4 Adaptive Equalizer with Data and Edge ISI Compensation. (76)

Meng-Xuan Lin, Kuo-Hsing Cheng, Wen-Hsuan Hsieh and Hong-Yi Huang

Abstract. This paper presents an innovative transmission gate decision feedback equalizer (TG-DFE) designed for 28 Gbps four-level pulse amplitude modulation (PAM-4) data. The proposed TG-DFE eliminates both data and edge inter-symbol interference (ISI) while alleviating the jitter induced by excessive loop delay in conventional decision feedback equalizers (DFEs). Furthermore, an adaptive system is used to adjust the weights of the continuous time linear equalizer (CTLE) and the TG-DFE under different channel loss conditions. The proposed adaptive equalizer is fabricated by TSMC 40 nm (TN40G) 1P10M CMOS process with 0.9 V supply voltage. When the channel loss is 12 dB, the peak-to-peak jitter of the equalized data is 50.4 ps and the root mean square (RMS) jitter is 14.7 ps. When the channel loss is 16 dB, the peak-to-peak jitter of the equalized data is 51.6 ps and the RMS jitter is 15.2 ps. The total power consumption is 95.9 mW, the chip area is 1.15 mm2 and the core area is 0.09 mm2.

Pipelined Systolic-Ring-Based Digital Oscillator Ising Machine for Highly Dense Graphs. (81)

Anubhav Bhura, Sathish N and Debanjan Bhowmik

Abstract. Solving combinatorial optimization problems related to dense graphs of size N has various applications in edge AI and involves exponential time complexity on accurate classical solvers and quadratic or higher-degree-polynomial (O(N^2) or above) time complexity on approximate classical solvers, run on CPU. Analog or digital oscillator Ising machine (OIM) implementations with all-to-all connectivity between the oscillators have O(log(N)) time complexity. But spatial resources grow as O(N^2), which can't be avoided in extremely dense graphs through sparsification. Here, we propose, design, and implement (on a FPGA board) a digital OIM in which coupled non-linear differential equations corresponding to N oscillators, corresponding to N nodes of the graph, are solved numerically in parallel through a systolic ring architecture. By solving MaxCut on fully connected weighted graphs up to N = 400, we report experimentally that number of cycles per iteration grows as O(N) in the designed OIM and overall time complexity is O(Nlog(N)), still a significant improvement over classical solvers on CPU. Also, in our designed OIM, the growth of spatial resources is restricted to O(N), avoiding quadratic scaling unlike all-to-all connected OIM implementations. Furthermore, taking advantage of pipelining while implementing the non-linear function calculation, we ensure that number of cycles per iteration only grows as c1(N)+c2 with c1 as low as 1. This makes the actual computational time of the designed OIM on the FPGA board extremely low.

Power-Accuracy Design Exploration of Dual-kernel Configurable Gaussian Filter. (87)

Marcio Monteiro, Ismael Seidel, Mateus Grellert, José Luís Güntzel, Leonardo Soares and Cristina Meinhardt

Abstract. Gaussian filtering is a widely used image preprocessing technique in power-restricted applications, such as edge computing and smart cities, demanding a constant balance between power consumption and image quality. In this work, we present a design exploration of a dual-kernel size-configurable Gaussian filter architecture. This architecture allows the runtime configuration of the kernel's operation mode, switching between a 3×3 and a 5×5 kernel according to application requirements. The power saving is achieved by replacing exact full adders with four types of approximate adders. By combining the design-time configuration of approximate adders with runtime kernel selection, the proposed architecture achieves power reductions of up to 48% for 3×3 mode, with quality varying from 21 dB to 57 dB. For 5×5 mode, our approach reduced total power by up to 74%, while quality ranged from 20 dB to 58 dB.

Performance Analysis of Sparse Kernels on a RISC-V Open Source GPU. (103)

Jules Dubois, Eric Guthmuller, Adrian Evans and Frédéric Rousseau

Abstract. Graphics Processing Units (GPUs) are the workhorse in many computational domains including AI and HPC. When processing regular data streams, the effective performance is typically close to the roofline. For workloads based on sparse matrices, where the vast majority of entries are zero, the irregular distribution of the non-zero elements stresses the memory system resulting in performance being bottlenecked by the memory bandwidth. On these parallel architectures, workload imbalances also limit performance. GPUs running sparse matrix kernels using state-of-the-art Basic Linear Algebra Subsystem (BLAS) libraries are central in modern HPC systems. Although RISC-V application processors are gaining in performance, RISC-V based open-source GPUs are in an early stage of development. We benchmark sparse kernels, both on modern HPC-grade GPUs and on Vortex, an open-source RISC-V GPU that is gaining adoption. We analyse their performance under memory-bound workloads and report the gaps in software and hardware required to enable efficient sparse BLAS processing on RISC-V GPUs.

Hardware-Aware Acceleration of DNN Inference on Embedded GPUs for Chest X-Ray Classification. (117)

Theodora Sanida, Argyrios Sideris and Minas Dasygenis

Abstract. Deep learning-based inference for chest X-ray classification on embedded GPU platforms is a promising direction for real-time edge intelligence in medical imaging. However, deploying deep neural networks (DNNs) on resource-constrained devices, such as NVIDIA Jetson platforms, remains challenging due to computational and energy constraints. This paper presents a hardware-aware low-complexity DNN optimization methodology for chest X-ray inference on NVIDIA Jetson embedded GPU platforms, specifically AGX Xavier and AGX Orin. The proposed approach combines structured pruning with post-training quantization (PTQ) and quantization-aware training (QAT) to reduce model complexity while preserving classification performance. Experimental results show that moderate sparsity improves generalization, with the 15\% pruned model achieving 93.55\% accuracy, surpassing the FP32 baseline. In addition, INT8 quantization reduces model size by up to 69.10\%, while QAT-based FP16 and INT8 models maintain accuracy above 93\%. These findings demonstrate that pruning and quantization, when integrated within a hardware-aware design methodology aligned with the architectural characteristics of embedded GPU platforms, provide an effective solution for accurate, compact, and efficient real-time chest X-ray classification.

Hardware–Software Co-Mapping NAS for Heterogeneous Embedded Platforms. (118)

Theodora Sanida, Argyrios Sideris and Minas Dasygenis

Abstract. Neural Architecture Search (NAS) automates the design of convolutional neural networks (CNNs), but conventional NAS methods often ignore hardware constraints. This limitation is critical in heterogeneous embedded systems, where performance depends on both network topology and runtime resource allocation.

This paper proposes a hardware--software co-design framework for embedded GPU/DLA platforms. The proposed method introduces a DLA-native search space that jointly explores CNN architecture, numerical precision, and hardware mapping. A parametric multi-objective scheduler dynamically assigns MobileNet-style CNN blocks to GPU, DLA$_0$, or DLA$_1$ based on computational load, transfer overhead, workload balance, energy cost, and runtime constraints.

The framework is implemented using Optuna-based optimization and evaluated on NVIDIA AGX Xavier and NVIDIA AGX Orin platforms.

The selected architecture achieved 98.05\% test accuracy.

Experimental results show that AGX Orin achieved up to 17.50\% latency reduction, 21.14\% throughput improvement, and 18.62\% energy reduction compared with the static scheduling configuration.

The obtained results demonstrate that jointly optimizing CNN architecture and runtime scheduling significantly improves execution efficiency and resource utilization in heterogeneous embedded AI systems.

Hardware Design of a Scan-Enabled Four-Stage Pipelined SHA-3 Core on FPGA. (123)

Panagiotis Antoniou, Constantinos S. Dionysiou, Panayiotis Pericleous and Harris E. Michail

Abstract. Cryptographic hash functions are key components in modern security systems, since they are widely used in applications that require authentication, integrity verification, and secure data processing. In many of these applications, hardware implementations are preferred because they can provide higher processing speed and better efficiency than software-only solutions. At the same time, modern hardware designs are becoming increasingly complex, while their reliability must be ensured both after implementation and during controlled test operation in demanding environments. These requirements motivate the use of Design for Testability (DFT) techniques, where scan paths provide an effective method for improving the controllability and observability of the internal sequential elements of a circuit. In this paper, a four-stage pipelined SHA-3 FPGA architecture is investigated with and without embedded scan-path support. The proposed pipelined design divides the 24 rounds of the Keccak-f[1600] permutation into four stages of six rounds each, while the scan-enabled version replaces the four 1600-bit state registers and the registers of the four local 3-bit control counters with scan-capable registers. A complete description of the conventional pipelined core and the scan-enabled architecture is presented. Both implementations are synthesized and evaluated under the same conditions on three FPGA devices, and their characteristics are compared in terms of area, maximum operating frequency and throughput. The obtained results show that scan-path insertion enables test-mode access to the internal pipeline state and control registers, with an area increase of about 9.02% to 11.23% and a small maximum-frequency reduction of less than 0.7%. Therefore, this work provides a quantitative evaluation of the cost of scan-path integration in a four-stage pipelined SHA-3 FPGA core and highlights the trade-off between improved testability and implementation efficiency.

Predictive and Host I/O–Aware Garbage Collection Optimization in Solid State Drives. (129)

Ojasweeta Singh and Ramanathan Muthiah

Abstract. Abstract—Garbage collection (GC) is a fundamental background operation in solid state drives (SSDs), required to reclaim space from invalid data and sustain write performance. Existing GC mechanisms are largely reactive, relying on static thresholds such as free block count or valid fragment count, without considering ongoing host access behavior or future invalidation trajectories. This often results in unnecessary data relocation, increased write amplification, degraded quality of service (QoS), and reduced flash endurance. This paper presents a unified garbage collection optimization framework that integrates host I/O–aware source block selection with predictive Auto Free–aware GC deferral. The proposed approach leverages real time analysis of host access locality, temporal write characteristics, and block level invalidation dynamics to guide GC decisions. Blocks associated with active host access are de-prioritized for relocation, while blocks exhibiting a high likelihood of imminent self cleaning (automatic block reclamation) are temporarily deferred from GC. An adaptive feedback mechanism continuously refines deferral thresholds based on observed results, enabling robust operation across diverse workloads without host intervention or hardware modification. Experimental evaluation using firmware level modeling and workload replay demonstrates reductions of up to 70–75% in unnecessary GC operations, 10–20% reduction in write amplification, and significant improvements in tail latency under overwrite dominated and mixed workloads.

A Novel Radiation-Hardened-by-Design Architecture for SRAMs. (137)

Hung-Ming Kuo, Jafar Shojaii, Andrea Mazzanti and Ting-Chang Chang

Abstract. This work addresses the Multiple-Node Upset (MNU) vulnerability of radiation-hardened Static Random-Access Memory (SRAM) cells, particularly the Dual-Interlocked Storage Cell (DICE). We propose a three-stage Radiation-Hardening-By-Design (RHBD) methodology and two novel SRAM cells derived from the conventional six-transistor (6T) and DICE architectures: a Pass-Transistor-Based Six-Transistor Cell (6TPT) and a Pass-Transistor-Based Dual-Interlocked Storage Cell (DICEPT). The flow combines pass-transistor-based feedback gating, sizing-assisted node restoration, and layout-oriented SEU optimization to suppress MNU. Schematic-level SEU simulations show that the proposed 6TPT SRAM cell achieves Single-Node Upset (SNU) immunity; furthermore, the proposed DICEPT SRAM cell preserves the stored bit under Double-Node Upset (DNU) conditions, extending upset tolerance beyond the SNU-level capability of conventional DICE. Layout-level SEU simulations further confirm that strategic spatial adjustment of sensitive nodes significantly reduces sensitive areas in both proposed 6TPT and DICEPT compared to their conventional counterparts, improving robustness in practical radiation environments where multiple nodes can be affected simultaneously.

Challenges of Taping Out a Multi-Tenant MPSoC with Heterogeneous IP Contributions. (139)

Mattis Hasler

Abstract. Multi-tenant MPSoCs offer a cost-effective path for small research teams to participate in silicon fabrication by sharing tapeout resources. However, combining intellectual property (IP) from diverse contributors into a single chip introduces substantial integration challenges spanning environment compatibility, platform adaptation, physical design, and project management. In this paper we report our experience taping out a ~5 MGate, 12 mm² MPSoC built on the M3 platform within one year. The chip incorporates four external IP contributions — a RISC-V core cryptography extensions, a post-quantum cryptography accelerator, a ring-oscillator entropy source, and a micro-AI speech-recognition accelerator—each supplied under a different collaboration model. We describe the integration workflow, the concrete technical obstacles encountered at each level, and the platform tools we developed or extended to ease future collaborative tapeouts. The resulting GDSII passed all sign-off checks and has been submitted for production.

A Lightweight Full-Custom FiGaRO TRNG for Resource-Constrained Security in 12 nm FinFET. (145)

Francesca Bagnoli, Ettore Noccetti, Daniele Rossi and Sergio Saponara

Abstract. This work presents ultra-lightweight True Random Number Generator (TRNG) engines based on Fibonacci (FiRO), Galois (GaRO), and hybrid FiGaRO architectures. Implemented at the schematic level in 12 nm LP+ FinFET technology, these engines target resource-constrained hardware-security applications by ensuring high area and power efficiency. Mutually coprime delay elements are introduced along the feedback paths to sustain chaotic dynamics and mitigate periodic-collapse phenomena. Gaussian power-noise injection is employed to model realistic supply fluctuations, which shows improved entropy extraction and reduced output bias; a possible on-chip entropy-enhancement mechanism based on controlled switching activity on a virtual power rail is also discussed. Standalone FiRO and GaRO cores remain partially dominated by deterministic dynamics, while the FiGaRO architecture significantly improves spectral randomization and suppresses limit-cycle behavior. A dual FiGaRO configuration combined through an XOR stage achieves high-entropy raw-bit generation at 500 Mbps, passing the selected NIST statistical tests under preliminary schematic-level evaluation, making it a promising candidate for integrated hardware-security primitives in constrained systems.

Layer-Wise Hardware-Aware Non-Uniform Quantization for Unary-Coded RRAM AI Accelerators. (170)

Ming Cheng, Cao Wang, Xiao-Chun Li, Yanxia Lin and David Keezer

Abstract. RRAM-based AI accelerators exploit crossbar arrays for efficient in-memory multiply-accumulate operations, but device variation degrades inference accuracy. Unary coding improves variation tolerance by distributing weight significance over multiple cells, yet it introduces substantial storage and operational-unit redundancy. This paper proposes a layer-wise hardware-aware non-uniform quantization (LWNQ) framework for unary-coded RRAM accelerators. LWNQ combines Hessian-based layer sensitivity, K-Means initialized quantization levels, an integer-programming bit-width allocation, and learnable lookup-table quantization-aware training. The framework assigns each layer only the unary precision it needs while adapting non-uniform levels to the layer weight distribution. On CIFAR-10, LWNQ achieves 92.33% accuracy on VGG16 and 92.68% on ResNet20, outperforming the evaluated fixed quantizers. Under a 70% unary hardware-resource reduction proxy, it maintains 91.86% and 92.46% accuracy, respectively. These results show that layer-dependent unary redundancy can be substantially reduced with limited accuracy degradation.

Full-Chip Modeling for Voltage Analysis on the Backside Silicon Substrate. (185)

Rikuu Hasegawa, Kazuki Monta, Takuya Wadatsumi, Alkis Chatzopoulos and Makoto Nagata

Abstract. Voltage fluctuations in integrated circuits are important for chip-level noise analysis and future physical leakage evaluation. In flip-chip packages, the backside silicon substrate is exposed and can be used as an observation surface for voltage fluctuations caused by internal circuit operation. In this work, we extend a full-chip chip power model (CPM)-based modeling method to analyze voltage fluctuations on the backside silicon substrate of a flip-chip IC. The target chip was fabricated in a TSMC 180 nm CMOS process and contains six identical AES cores. Among them, two AES cores are selected as operating targets, and the backside voltage distributions are evaluated by both simulation and measurement. The simulation model divides the P-substrate into multiple layers to represent voltage propagation through the silicon substrate. The simulated backside voltage heatmaps are compared with measured heatmaps obtained by direct backside probing. The results show that the simulation and measurement exhibit similar spatial trends when different AES cores are operated. This suggests that the extended CPM-based modeling approach can capture the main tendency of backside silicon substrate voltage fluctuations and can be used as a basis for future physical leakage evaluation.

Information-Theoretic Causal Modelling of Semiconductor Process Dynamics. (191)

Daniel Sørensen, Giorgio Melchiorre, Sudip Bandyopadhyay, Sandip Halder, Roel Wuyts and Bappaditya Dey

Abstract. With the progress of the semiconductor industry toward increasingly complex compute devices and tighter process tolerances, advanced process control has become crucial. This work explores a novel framework to infer the underlying dynamics of semiconductor processes, directly from raw equipment log-file time-series data. By modelling the tool dynamics as a stochastic dynamical system comprising (a) a deterministic component and (b) a stochastic component, we estimate entropy transfer rates between variables through the Liang-Kleeman and Pires formalism. Preliminary results indicated that 7.5% of the inferred dependencies were known, 36.0% were plausible, 17.5% represented previously uncharacterized relationships, and 39.0% were inconsistent with established process knowledge. These findings demonstrate the framework's capability to uncover novel causal insights, while motivating further improvements to reduce inconsistent findings.

An SVA-to-RTL Transpiler for Synthesizable Hardware Assertion Checkers. (201)

Antonio Giuseppe Maria Strollo, Gennaro Di Meo, Luca Tegazzini, Camillo Perna, Matteo Sommella and Davide De Caro

Abstract. This paper presents an open-source transpiler that generates synthesizable RTL hardware checkers from a subset of SystemVerilog Assertions. The proposed flow uses the Slang frontend to extract assertion information from the elaborated SystemVerilog abstract syntax tree; assertions are then translated into hardware by combining a forward pipeline for antecedent evaluation with a checkpointed lookback engine for consequent synthesis. The lookback engine decomposes consequents into prefixes and evaluates endpoint matches at fixed checkpoints. This simplifies hardware implementation while accepting bounded failure latency. The current prototype supports sampled-value functions, bounded delays and repetitions, goto-based antecedents, and throughout conditions. Generated checkers were validated through formal equivalence-style analysis. A synchronous FIFO case study demonstrates that non-trivial assertions can be synthesized on FPGA with moderate resource overhead.

Evaluating NASA's Core Flight System Fault Detection and Recovery on RISC-V Systems. (220)

Luís França, Fabio Benevenuti, Leonardo R. Gobatto, José Rodrigo F. Azambuja and Fernanda Kastensmidt

Abstract. This paper presents a fault injection campaign to evaluate the effectiveness of the NASA Core Flight System (cFS) Health Services (HS) application as an active fault detection and recovery mechanism in satellite On-Board Computers (OBCs). The proposed approach integrates HS with the Critical Data Store (CDS) to enable autonomous system monitoring and state preservation, leveraging a hardware watchdog timer to detect failures and trigger recovery from fatal system crashes. The methodology is validated on a RISC-V System-on-Chip running FreeRTOS, where more than 100,000 data memory bit-flips are injected to emulate radiation-induced faults. Experimental results demonstrate that the proposed architecture achieves 1,292 successful recoveries, significantly reducing I/O timeout failuresfrom previously reported levels of 79% to only 0.8%.

Reinforcement Learning-Based Optimization of Workload-Aware Power Delivery Networks. (226)

Oran Hayes, Maria Pantazi, Athanasios Tziouvaras, George Stamoulis, Anuj Pathania, Shreejith Shanker and George Floros

Abstract. Power Delivery Networks (PDNs) are critical components of modern VLSI chips, providing stable voltage levels while satisfying electromigration (EM) and IR-drop constraints. Conventional PDN design methodologies typically rely on worst-case assumptions, often resulting in over-provisioned networks and inefficient use of resources. This paper presents a reinforcement learning-based framework for the optimization of workload-aware PDNs. The proposed methodology first generates workload-aware PDNs using architectural power traces obtained from system-level simulations. These power traces are mapped to spatial power density distributions, enabling adaptive allocation of PDN resources according to local current demand. A reinforcement learning agent then performs wire-width optimization to minimize PDN area while maintaining EM and voltage integrity constraints. Electrical and reliability metrics are obtained using SPICE-based circuit analysis and EM lifetime estimation. Experimental evaluation is performed on a dataset of workload-aware PDNs generated from 4-, 8-, and 16-core multiprocessor floorplans using PARSEC and SPLASH-2 benchmark workloads. Furthermore, the proposed DQN-based optimizer reduces the average normalized PDN area by 47% while satisfying all EM and IR-drop constraints. Compared to simulated annealing, the proposed approach achieves comparable optimization quality while providing approximately 26$\times$ faster optimization.

FPGA Congestion Estimation Using Topology and Geometry Aware Graph Neural Networks. (239)

André Cavalcante, Sumanta Chaudhuri and Lirida Naviner

Abstract. Wire demand estimation is used for optimizing placement of FPGA circuits. Here we propose a learning methodology for topology/geometry aware graph neural networks used for the problem of wire demand estimation at placement time. Our system considers a graph in which nodes represent FPGA fabric elements, including communication channels and switches. Furthermore, we propose dedicated pre-processing pipelines for each graph feature. Finally, we propose a loss function composed of three metrics in order to reduce the deterioration of information about least and most used wire channels. The proposed method improves Pearson correlation from 0.81 to 0.90 and Spearman correlation from 0.82 to 0.86 compared to the standard method.

Area-Efficient CKKS Homomorphic Multiplication Design Using Output-Driven Scheduling and Hierarchical Basis Conversion. (257)

Tsu-Hsiung Huang, Yun-Ting Hsu and Ming-Der Shieh

Abstract. To speedup CKKS fully homomorphic encryption (FHE), existing hardware accelerators often rely on extensive parallelism and large on-chip memory capabilities, resulting in high hardware overhead. This work presents an area-efficient CKKS homomorphic multiplication (HMult) design with hybrid key-switching (HKS) process. By analyzing the limb-level fan-in cones of the primary outputs, we propose an Output-Driven Scheduling (ODS) scheme that avoids redundant computations and exploits shared intermediate data for reuse across multiple outputs, thereby reducing off-chip data movement. Moreover, lightweight moduli and fast hierarchical basis conversion (FHBC) are employed to reduce computational complexity and area overhead. Compared with a related 7-nm ASIC design, the proposed design, implemented in a 40-nm process, achieves up to 1.59× higher effective throughput per unit area after technology scaling. These results demonstrate the effectiveness of our design for CKKS-based FHE.

Zero Current Detector for Positive Outputs of a Single-Inductor Triple-Output DC-DC Converter. (273)

Giorgio Rolandi, Alireza Mesri, Elisabetta Moisello, Piero Malcovati and Edoardo Bonizzoni

Abstract. This paper presents a zero current detector (ZCD) circuit for positive outputs of a single-inductor triple-output bipolar DC-DC converter. This converter is used in spindle and voice coil motor controllers intended for Hard Disk Drive (HDD) applications. From input voltage of VM, it provides a negative voltage of -10V (VN), and two positive outputs of 30V (VP), and VM+5V (VGH). The proposed circuit is implemented in 110nm STMicroelectronics BCD technology and tracks the inductor current by monitoring the drain source voltage (VDS) of highside switch (HS) during inductor discharge phase. The ZCD uses the source terminal of the power transistor as its first input and a signal analogous to the drain voltage, shifted by a small offset, as its second input. This offset follows the corners of the HS in order to reduce the ZCD threshold. The circuit is able to operate in the high-voltage domain and is protected against the negative voltage reached by the high-side switching node during VN regulation. Simulations confirm the expected performance and the robustness of the circuit across process corners and in the presence of device mismatch, as verified by Monte Carlo analysis.

Design-Space Exploration of Multi-Kernel Multiplierless Gaussian Filters Using 3-2 AxAC. (283)

Rodrigo Lopes, Léo Ribeiro, Leonardo Antonietti, Morgana Macedo Azevedo Da Rosa, Eduardo Da Costa and Rafael Soares

Abstract. This work proposes approximate multiplierless Gaussian-filter architectures using approximate 3-2 adder compressors (3-2 AxAC) in the accumulation trees. We explore 3 × 3, 5 × 5, and 7 × 7 kernels using three AxAC from the literature, namely GFAC, LOPES, and KONERU, along with the gate-free TRUNC 3-2 adder compressor (AC) proposal, applied to different numbers of least significant columns. The proposed design-space exploration comprises 136 approximate configurations for 8- bit grayscale images. All architectures were synthesized using Cadence Genus with a 65nm CMOS technology, a 1.25 supply voltage, and target frequencies of 63 and 249 MHz. The image-quality analysis employs the average normalized cross-correlation (NCC) over all evaluated images and the peak signal-to-noise ratio. The results show area- and energy-savings of up to 97.91% and 98.32%, respectively, compared with the corresponding exact architectures. The GFAC configurations maintain an average NCC above 0.99 and achieve energy savings of 137.37× compared with related Gaussian-filter architectures.

Portable Stimulus Vertical Reuse with VerTrace-Generated Interconnections: A RISC-V Execution Stage Case Study. (295)

Petr Bardonek and Marcela Zachariášová

Abstract. As hardware verification faces increasing pressure from growing design complexity and greater automation in hardware development, reuse in verification becomes increasingly important. The Portable Stimulus Standard (PSS) supports hardware verification by capturing verification intent in reusable Portable Models (PMs), from which tests can be generated. However, vertical reuse of PMs across design hierarchy levels remains difficult because their realization bindings must be adapted when block-level interfaces become internal, indirectly controlled, or constrained after integration. This paper presents an end-to-end case study of static-analysis-based PSS vertical reuse, implemented in the prototype VerTrace toolchain, on the execution (EX) stage of the RISC-V processor RI5CY. Existing PMs created for block-level verification of integrated modules are reused as modeling elements in the EX-stage PM. VerTrace derives connectivity information from the register-transfer-level design description and generates interconnections and supporting artifacts that bind reused PMs to the EX-stage integration context. Tests generated from the EX-stage PM are executed in a simulation environment, with waveforms showing value propagation to interfaces of integrated modules and collected functional coverage reporting the behavior exercised after integration. The case study also exposes integration-induced adaptations, including constant-tied handshake signals and constrained interaction between components. These necessary adaptations clarify how generated realization artifacts support PM reuse while allowing verification engineers to focus on creating verification scenarios for the higher-level integration context.

Self Powered Systems as Security Buttons for Improving Trust in Wildfire Sensor Nodes. (305)

Zane Gunn, Aniket Roy, Samit Hasan, Davis Loose, Megan Gunn, Daniel Truesdell, James Lambert and Benton Calhoun

Abstract. Wildfire sensor nodes are increasingly vulnerable to both physical and network-layer attacks, yet replacing existing hardware to address these vulnerabilities incurs prohibitive cost and risk. This paper presents two modeled self-powered security ”button” systems designed as low-cost, easily deployable, add-ons to existing wireless sensor nodes (WSNs) for use in wildfire monitoring. The capabilities of each button are tailored to location. The system power of both buttons in an attack scenario are modeled, and both systems are shown to be feasible using outdoor solar harvesting consuming only $\sim$10–20 μW average power. A combined configuration is also proposed, demonstrating how buttons can collaborate to provide the full complement of defenses from both systems.

A 4H-SiC CMOS Op-Amp for High Temperature Applications up to 500°C. (316)

Andrei-Tudor Gheorghe and Marius Enachescu

Abstract. Silicon carbide (SiC) CMOS technologies enable integrated electronics for high-temperature environments. This work presents the design of a two-stage operational amplifier, employing a folded-cascode OTA followed by a class-AB output stage, in 2 µm 4H-SiC CMOS. Supplied at 20V and driving a 10 kΩ, 100 pF load, the amplifier is evaluated across a temperature range from 25 ◦C to 500 ◦C. Simulation results confirm operational capability over the entire range, achieving a distinguishing maximum DC gain of 96.6 dB and maintaining a high gain range of 79.6 dB to 96.6 dB, remaining robustly above 90 dB for temperatures between 50 ◦C and 430 ◦C. This design demonstrates considerably higher peak gain compared to recent state-of-the-art designs which report maximum gains of 71 dB and 51 dB respectively. The design further achieves a minimum gain-bandwidth product of 63.6 kHz. The results confirm the technology's potential and design methodology suitability for high-temperature op-amp based applications.

PCM-enabled energy consumption estimation for language models deployed on edge devices. (325)

Lazaros Papadopoulos, Giorgos Balis, Socratis Nazlidis, Miltiadis Siavvas, Georgios Karagiannidis and Ilias Kalouptsoglou

Abstract. Deploying language models on resource-constrained edge devices is challenging due to performance, memory and energy constraints. Accurate power models can effectively drive DVFS techniques and OS governors to optimize energy consumption. In this work, we propose a lightweight, portable, hybrid energy consumption estimation model that combines analytical approaches with data-driven techniques. The model was applied and evaluated on a CPU-based Raspberry 5 edge device, demonstrating an error rate of under 5.47\%. Moreover, we demonstrate that leveraging our model to guide direct CPU frequency scaling on language models achieves energy savings of up to 14.78\% compared to the default RPi-OS governor.

HOMAxA-2: An Energy-Efficient Hybrid OR–MUX Approximate Adder for Edge AI. (382)

Sagar Bhivsane, Alak Majumder, Vassilis Alimisis and Paul P. Sotiriadis

Abstract. Approximate adders are widely used in edge-AI accelerators, where addition dominates multiply--accumulate operations. This paper proposes HOMAxA-2, a hybrid OR--MUX lower-part approximate adder employing a two-tier error-control strategy: bits 0--2 use XOR-based partial sums with a shared bit-2 generate signal, bit 3 uses a 2-to-1 MUX-based carry-select controlled by the bit-3 generate condition, and a single-AND-gate boundary carry ensures accurate upper-part carry propagation, limiting the worst-case error to 7. Exhaustive evaluation over all 65,536 input combinations using the carry-kept 9-bit convention yields an MED of 2.125, normalized MED of 0.00417, error rate of 57.81\%, and MRED of 0.0113, outperforming LOA, OLOCA, M-HERLOA, KMECA, and ETA-I in MED, worst-case error, and MRED. On the Lena benchmark, HOMAxA-2 achieves a PSNR of 40.70~dB and SSIM of 0.9855. ASIC implementation using the OpenROAD SkyWater 130~nm flow produces a DRC-clean layout that, under identical constraints, attains a 17.1\% shorter critical-path delay and 42.3\% lower power than an exact ripple-carry adder.

Micromagnetic Investigation of a Skyrmion-Based Wired AND Logic Device Using Hierarchical VCMA Barriers. (384)

Debasmita Chatterjee

Abstract. Magnetic skyrmions have emerged as promising candidates for emerging spintronic computing technologies due to their nanoscale stability, non-volatility, and ultra-low energy consumption. In this work, a hierarchical skyrmion-based wired AND logic architecture is proposed using a customized branched nanotrack geometry integrated with staged voltage-controlled magnetic anisotropy (VCMA) barriers. Logic states are encoded through the presence or absence of skyrmions, while spin–orbit torque (SOT) drives skyrmion transport through the nanotrack structure.The proposed device exploits collective skyrmion interaction and threshold-assisted depinning dynamics to realize wired AND-like functionality. Micromagnetic simulations performed using MuMax3 demonstrate progressive enhancement in skyrmion transport capability with increasing active inputs, where the 111 configuration successfully propagates through the hierarchical VCMA barriers while lower-input states remain partially pinned. The proposed skyrmion wired AND exhibits an average energy consumption on the order of 10^−17 J, demonstrating significantly lower switching energy compared to conventional CMOS wired logic implementations. Furthermore, the hierarchical branched geometry enables scalable multi-input operation without requiring standardized skyrmion logic layouts, thereby providing flexibility for future compact and energy efficient emerging spintronic computing architectures.

Non-Invasive Hydration Status Detection via Human Body Communication. (385)

Amr N. Abdelrahman, Abdelhay Ali and Ahmed M. Eltawil

Abstract. Dehydration is a prevalent yet under-detected condition impacting cognitive and physical performance. Clinical assessments are invasive and unsuitable for continuous monitoring, motivating non-invasive alternatives. This work proposes human body communication (HBC) as a novel sensorless, noninvasive hydration-sensing modality, where the body channel itself serves as the sensing element without dedicated chemical or optical transducers. Using dry electrodes across the full arm-to-arm path, a pilot study is conducted with a single acquisition instrument probing five frequency channels (1 kHz, 100 kHz, 1 MHz, 5 MHz, 10 MHz) under hydrated and fasting conditions on seven healthy subjects. Time-, frequency-, and time–frequency-domain features are extracted, and 15 classifiers across three families are evaluated under leave-one-subject-out (LOSO)-cross-validation (CV) with within-subject rebalancing. Linear support vector machine (SVM) on 89 features achieves the highest balanced accuracy, with a four-band configuration{1 kHz, 100 kHz, 1 MHz, 10 MHz} matching peak performance while eliminating the detrimental 5 MHz band, offering a direct path to hardware simplification. These findings establish HBC as a viable, physics-grounded hydration-sensing modality and motivate validation on larger cohorts.

A very low cost, fully digital, Clock and Data Recovery circuit for SERDES. (393)

Massimo Ruo Roch and Maurizio Martina

Abstract. The paper presents a novel approach for clock and data recovery from a serial stream of bits. The proposed solution is based on fully digital circuits, without the need of PLLs, DLLs, and VCOs at all. It can recover clock and data informations either from continuous bit streams, or from data packets with a short preamble. Due to the exclusive usage of simple gates or flip-flops, it can be implemented on a wide variety of digital technologies, starting from fully digital FPGAs, up to standard cell libraries, or full custom designs. Moreover, no oversampling of the input data stream is needed, allowing to optimally use frequency characteristics of the target technology. A description of the architecture is shown with experimental results obtained through a proof-of-concept implementation on a low-cost, low-power FPGA.

RISC-V for Quantum Control: Evaluation of Soft-Core Architectures on Cryogenic Temperatures. (404)

Gustavo Kanaiama, Rafael Rodrigues, Édney Freitas, Luis Bravo, Gustavo Marcati, Francisco Brito and Bruno Sanches

Abstract. The scalability of quantum computing architectures is currently constrained by the thermal and cabling bottle necks associated with connecting room-temperature electronics to cryogenic stages. Integrating control electronics directly into the cryostat can mitigate these limitations, provided that power dissipation remains within the system's strict thermal budget. This paper evaluates and compares three soft-core processor architectures: Intel Altera Nios II/e, Nios II/f, and an open source RISC-V core synthesized on a Cyclone V FPGA for localized quantum control applications. The architectures were benchmarked based on area utilization, computational latency, and power consumption across three signal processing kernels: sine generation, Gaussian pulse modeling, and a Radix-4 Fast Fourier Transform (FFT). Experimental results demonstrate that the RISC-V core achieves a 50% reduction in register utilization compared to proprietary Nios II alternatives and exhibits the highest power efficiency, with a full system average dissipation of approximately 1.25 W (including voltage Regulators, Clock generation, PLL, FPGA and 4 Microwave front-end channels up to 6 GHz). Although the Nios II/f provides the lowest raw latency, the RISC-V architecture offers a superior performance-to-power ratio and Instruction Set Architecture (ISA) extensibility. These findings suggest that RISC-V is a highly promising candidate for deployment at the 4 K to 50 K cooling plates of modern dilution based quantum computing cryostats, enabling more localized, efficient and scalable cryogenic control.

Hamm3R: Optimized 3D Image Matching for Memory-Constraint Environments. (407)

Mert Tas, Halil Durmuş, Bilgi Görkem Yazgaç and Ece Olcay Güneş

Abstract. State-of-the-art 3D image matching architectures require a lot of memory and processing power, which makes it very hard to use them on resource constrained devices. This study introduces Hamm3R, an optimized framework that modifies the resource-intensive MASt3R model for efficient execution on a single-core CPU. We present a hardware-efficient, training-free approach that integrates post-training quantization of the feature extractor backbone and Hamming-FastNN, a binary-domain reciprocal matching technique that utilizes efficient bitwise operations. Evaluations on the CO3Dv2 dataset show that Hamm3R can speed up inference by up to 10.86 times and cut the size of the model by 5.07 times, all while maintaining highly competitive pose estimation accuracy.

FSNIC: A Low-Latency Flow-Based Intrusion Detection Architecture for FPGA SmartNICs. (425)

Nise O'Cuill, Changhong Li, Georgios Floros and Shreejith Shanker

Abstract. Modern data centres require high-performance networking alongside effective real-time security.

Traditional Intrusion Detection Systems (IDS) commonly rely on general-purpose processors and often struggle to inspect high-speed traffic at line rate without introducing latency or performance bottlenecks.

Smart Network Interface Cards (NICs) provide an alternative by enabling computation directly within the network data plane.

This work presents a machine learning-based IDS implemented within an FPGA-based SmartNIC pipeline.

The system integrates P4-based packet parsing with a LogicNets IDS model implemented in RTL, enabling deterministic, low-latency inference.

Compared with traditional stateless packet-level classifiers, the proposed stateful flow-based IDS introduces minimal state by aggregating features across packets, capturing behavioural patterns not observable at the packet level.

Experimental results on the UNSW-NB15 dataset show that the flow-based IDS improves detection accuracy from 86.92\% to 97.68\% compared with stateless packet-level classification.

We also evaluate the proposed IDS on CICIDS2017 and compare its real-time hardware performance with prior FPGA-based IDS designs. Through hardware-software co-design, the proposed SmartNIC implementation achieves up to 1,113× lower latency and 55.9× lower LUT usage compared with existing FPGA IDS accelerators.

Flexible FPGA Undervolting Test Bench For Research And Teaching. (439)

Lucas Bublitz, Ulf Kulau and Leonard Schauer

Abstract. Low-power FPGAs are increasingly being used in small embedded systems and sensors as processing units or accelerators, thanks to its flexible hardware design. But in terms of energy efficiency especially under fluctuating load conditions, which are common in sensor applications they are significantly less efficient than ASICs or microcontrollers. One approach to improving this efficiency is undervolting. To facilitate student engagement with this research and integrate it into the lecture, a development environment is required that supports both undervolting and evaluation while being intuitive and user-friendly. A test bench is presented that is based on an FPGA as the control unit and communicates directly via USB. It offers a high degree of flexibility regarding both software and hardware.

KHNUM: A Kernel Hash Network for Unified Multi-function Cryptography. (405)

Mohamed El-Hadedy

Abstract. Constrained edge and space platforms often need several cryptographic services at once, yet dedicating a separate hardware block to each one raises area, power, and integration cost. This paper presents KHNUM, a Kernel Hash Network built on an ASCON-CXOF-chain kernel reused as a shared substrate for five embedded-security workloads: hardware-rooted random generation, key derivation, message authentication, Merkle-node hashing, and WOTS+ chain computation. The kernel is signed off as a SKY130 tile, where a framed byte-wide boundary replaces a wide host-datapath multiplexer and closes timing across all reported PVT corners at 50 MHz in about 0.30 mm² of standard-cell area. We then replicate the same kernel structure into a 20-lane PYNQ-Z2 fabric driven through the framed interface; across all five workloads the datapath is unchanged, and only the customization string, request format, and chain depth vary. The main systems result is that the kernel scales while end-to-end throughput is bounded by framed control and 32-byte result readback: WOTS+ chains reach a two-worker ceiling near 1.88×, while per-operation services remain near framed-request latency. KHNUM therefore pairs a circuit-level CXOF-chain signoff with a multi-function fabric and identifies the control and readback path as the next target for silicon.

POSTER SESSION 2

Tuesday November 10th, 2026, 4:00 PM - 5:30 PM

Silicon Aware Neural Networks. (43)

Sebastian Fieldhouse and Kea-Tiong Tang

Abstract. Recent work in the machine learning literature has demonstrated that deep learning can train neural networks made of discrete logic gate functions to perform simple image classification tasks at very high speeds on CPU, GPU and FPGA platforms. By virtue of being formed by discrete logic gates, these Differentiable Logic Gate Networks (DLGNs) lend themselves naturally to implementation in custom silicon - in this work we present a method to map DLGNs in a one-to-one fashion to a digital CMOS standard cell library by converting the trained model to a gate-level netlist. We also propose a novel loss function whereby the DLGN can optimize the area, and indirectly power consumption, of the resulting circuit by minimizing the expected area per neuron based on the area of the standard cells in the target standard cell library. Finally, we also show for the first time an implementation of a DLGN as a silicon circuit in simulation, performing layout of a DLGN in the SkyWater 130nm process as a custom hard macro using a Cadence standard cell library and performing post-layout power analysis. We find that our custom macro can perform classification on MNIST with 97% accuracy 41.8 million times a second at a power consumption of 83.88 mW.

A CKKS-Aware Distilled CNN System for Efficient Encrypted Inference. (68)

Babak Moradi, Hassan Rekabi Bana, Moslem Heidarpur and Mitra Mirhassani

Abstract. CKKS-based encrypted CNN inference is con strained by rotations, ciphertext products, rescaling, key support, and multiplicative depth. This paper presents a CKKS-aware distilled CNN system that uses offline knowledge distillation to reduce the online encrypted execution graph while preserving valid convolution, square activation, average pooling, and a no bootstrapping path. On CIFAR-10 with Microsoft SEAL-Python, the student reaches 70.31% accuracy versus 72.59% for the teacher, while reducing parameters by 31.47% and FLOPs by 31.15%. Under CKKS inference, runtime decreases from 355.81 to 240.33 s/image, a 32.46% reduction without added depth. Memory does not scale with CNN size: the 0.57 GB Galois key footprint is only about 7% of the 8.22 GB peak RAM, indicating that live ciphertext state, temporaries, and library Overhead dominates peak memory. The result is a measurement driven CAS design study separating width-induced arithmetic savings from memory behavior in depth-constrained encrypted inference.

A Physics-Informed Surrogate Modeling Framework for Real-Time Analog Circuit Performance Prediction: A 55-nm OTA Case Study. (71)

Konstantinos Liakos, Christos Poulos and Fotis Plessas

Abstract. Fast and reliable analog circuit evaluation is a central bottleneck in design-space exploration, because each candidate sizing normally requires repeated Spectre analyses across corners and operating conditions. This paper presents a physics-informed machine-learning surrogate for real-time forward prediction of a two-stage Miller operational transconductance amplifier (OTA) in 55-nm CMOS. The surrogate is trained on a 10,000-sample Cadence/Spectre dataset generated by the closed-loop-aware methodology described in the companion work [1]. The proposed pipeline transforms 14 user-controlled design, process-corner, and operating-condition inputs into a leakage-safe 45-feature representation, including geometry ratios, current-density proxies, compensation-related quantities, and one-hot process-corner encodings. Target-specific Light Gradient Boosting Machine (LightGBM) regressors and classifiers are trained under feasibility-aware NaN masking, where non-physical or oscillatory simulation regions are excluded from regression targets. The surrogate predicts 46 outputs, consisting of 41 regression targets and five feasibility-classification labels. On a held-out 2,000-sample test set, the regressors achieve a mean coefficient of determination R2 of 0.942, with 30 out of 41 targets exceeding $R2 \geq 0.95$, while all feasibility classifiers reach receiver-operating-characteristic area-under-curve (ROC-AUC) values of at least 0.975. A single forward evaluation over all outputs completes in under 5 ms in the deployed online interface, compared to several minutes for a full Spectre run. The trained models are implemented as a deployable prototype for rapid pre-Spectre design-space inspection.

Efficient VLSI Design for CKKS Key Switching via Adaptive Special Modulus Selection and Approximate Basis Conversion. (114)

Qi-Xian Wu, Tsu-Hsiung Huang, Yun-Ting Hsu and Ming-Der Shieh

Abstract. Fully homomorphic encryption (FHE) enables privacy-preserving computation on encrypted data. CKKS, a promising FHE scheme for AI applications, suffers from key-switching as a performance bottleneck. Conventional special modulus configurations are not optimized across multiplicative levels, leading to computational inefficiency. To address this, we proposed an adaptive special modulus selection (ASMS) technique that dynamically determines the optimal special modulus based on the required multiplicative level, further incorporating the impact of rescaling operation. In addition, we present a systematic key-switching architecture that integrates the approximate basis conversion (ABC) and single-port merged-bank-based number theoretic transform (NTT) approaches for area efficiency. Experimental results demonstrate that the proposed accelerator with ASMS reduces the average key-switching computation time by 35.7% over the baseline, thereby improving the practicality of CKKS accelerators in real-world applications.

A Resource-Efficient MobileNetV2 Hardware Engine with Multispectral Fusion on FPGA. (155)

Tsung-Han Tsai, Chun-Tse Lin and Yu Shang-Yin

Abstract. This paper presents an energy-efficient, full-integer MobileNetV2 hardware accelerator implemented on the Xilinx Zynq MPSoC platform for multispectral crop classification. A Discrete Wavelet Transform (DWT)-based fusion scheme is adopted to combine RGB and NIR imagery to maximize accuracy, while Quantization-Aware Training (QAT) compresses the model to 8-bit integer (INT8) precision to alleviate memory bottlenecks. Experimental results on the RGB/NIR Aerial Crop Dataset show that the fused imagery achieves a high accuracy of 96.8% under INT8 precision, with a negligible 0.2% degradation compared to the FP32 baseline. Implemented on the ZCU106 FPGA at 140 MHz, the customized architecture delivers a throughput of 76.8 GOPs at 3.7 W. Utilizing only 55.9k LUTs and 75 DSPs, the design achieves an exceptional power efficiency of 20.6 GOPs/W, outperforming existing implementation. FPGA, Hardware Description language, CNN Accelerator, SoC System, Multispectral System

Specific Component Fault Detection in Analog Circuits using the RMS of the Supply-Current Fluctuation (165)

Vassilis Vassios

In this paper an enhanced approach for accurate identification, localization and classification of catastrophic faults impacting analog integrated circuits is proposed. This approach is based on previous work and relies on observing the fluctuations of the power supply current. The measurements are used to calculate new metrics for the good or faulty circuit under test. The mean and fundamental rms values of the supply current are calculated. From these calculations, 12 points of interest are obtained. The under/overshooting, the damping and the relaxation point. The general behavior of the algorithm in hard fault detection was studied in the authors previous works and this paper is an extension on the faulty component detection capabilities of the algorithm. The considered faults affect only resistor and capacitor values. The algorithm has been tested in simulations and actual circuit measurements and the results are given. The results obtained show that all tested CUTs have high average accuracy of classification. The proposed method has achieved a high level of classification accuracy which is achieved with the use of classifiers and with much less complexity and processing power as compared to previous research studies.

An End-to-End FPGA-Based SoC Implementation of Quantized Neural Networks for ECG Biometric Recognition. (176)

Claudia Georgiana Barbura, Paul Farago and Sorin Hintea

Abstract. This paper presents an end-to-end FPGA-based system for ECG biometric recognition using quantized neural networks (QNNs). The proposed solution is implemented on a Zynq-7000 SoC (Zybo Z7-20) using the FINN framework and supports real-time inference under embedded hardware constraints. Three neural network architectures, namely MLP, 1D CNN, and 2D CNN, are designed and quantized to 4-bit weights and activations (4W4A) using quantization-aware training. The models are evaluated in terms of classification accuracy, inference latency, power consumption, and FPGA resource utilization. The deployed models achieve classification accuracies above 90%, demonstrating the effectiveness of low-precision inference on FPGA hardware. Among the evaluated architectures, the quantized 1D CNN provided most efficient trade-off between classification performance and hardware efficiency, achieving 93% accuracy with a power consumption of 1.765 W and an inference latency of 2.22 ms per sample. The results demonstrate the feasibility of deploying QNN-based ECG biometric recognition systems on resource-constrained FPGA platforms using automated dataflow synthesis tools.

A Latent Ensemble Method for Safe Multimodal Reinforcement Learning in Robot Navigation. (261)

Konstantinos Spyridopoulos, Pavlos Tosidis, Nikolaos Passalis and Anastasios Tefas

Abstract. Deep reinforcement learning (DRL) enables mobile robots to learn complex navigation policies, but learned controllers may become unsafe when deployed outside the distribution observed during training. This paper presents a lightweight safety layer that transforms latent-space ensemble disagreement into a real-time anomaly score. A frozen Dreamer-V3 world model fuses RGB and LiDAR observations into a compact latent representation, while an ensemble of small transition heads predicts the next latent state. The variance among these predictions is aggregated over a sliding window and classified as nominal or anomalous. The resulting score supports real-time intervention through a MAPE-K feedback loop in Webots and is also used for adaptive encoder switching in visually corrupted Atari Pong. Experiments in empty and furnished Webots environments show that the proposed latent disagreement signal detects RGB and LiDAR faults effectively, reaching ROC-AUC values up to 0.94 in cluttered scenes. The Atari proof-of-concept further indicates that the same uncertainty signal can improve robustness under severe visual corruption.

CuriousRL: A Curiosity-Driven Goal-Conditioned RL Engine for Multi-Objective Analog Circuit Sizing. (270)

Yash Sengupta, Niranjan Gopal and Sakshi Arora

Abstract. Analog transistor sizing is a major bottleneck in analog design, requiring the designer to balance many competing specifications over a large parameter space. Reinforcement learning (RL) is attractive because a single goal-conditioned policy can be reused across target specifications, but standard RL scales poorly as the number of design variables grows, confining it to small problems with few objectives. We present a goal-conditioned RL formulation that augments a Soft Actor-Critic agent with an Intrinsic Curiosity Module (ICM) to improve exploration of large, sparsely feasible design spaces. We apply it to a three-stage capacitor-less low-dropout regulator (LDO) with 26 design variables and 6 competing objectives, training a single policy that sizes the circuit for any requested specification. Evaluated across a wide range of target specifications, the curiosity-driven agent achieves a 100% feasible-design hit rate, compared with 87.2% for an identical agent without curiosity, and a substantially higher hypervolume (1.2375 vs. 0.4471). These results show that a single intrinsic-curiosity term, rather than a larger model, is enough to make complex, many-objective analog sizing tractable.

Error-Driven Sampling of (,) for Improved Generalization in Multi-Layer Perceptron-Based Compact Modeling. (274)

Yusra Rachidi, Thomas Bedecarrats, Roberto Miorelli and Emeric De Foucauld

Abstract. This work studies the impact of gate length (L) and width (W) sampling on training neural network (NN) models for drain (Id) and gate (Ig) currents prediction in FD-SOI transistors. Existing datasets are usually designed for physics-based model parameter extraction and may be sub-optimal for training NN models. We show that these L-W grids lead to poor generalization capabilities and we propose a methodology to build a more suitable grid. The optimized dataset improves generalization while reducing the number of devices required for training. Although demonstrated on FD-SOI technology, the proposed method can be extended to other device technologies.

Acceleration of BiLSTM-based Music Separation Algorithm by Customizing HW/SW on SoC-FPGA. (289)

Sotirios Aias Karioris, Ioannis Voyiatzis and George Lentaris

Abstract. Specialized hardware accelerators for AI/ML applications enable embedded devices to meet real-time constraints with reduced power consumption. This work proposes a HW/SW co-design approach with VHDL acceleration on SoC-FPGA for music source separation. We rely on the established CrossNet Open-Unmix algorithm, which utilizes computationally intensive BiLSTM functions, to extract the audio signals of individual instruments from a single input stream. We methodically partition the operations into multi-core CPU, SIMD extensions, and FPGA circuits, which we then customize with increased parallelization and tightly-coupled co-processing. We develop parametric VHDL for dense vector-matrix multiplications and LSTM gate calculations. We develop C/C++ code on SIMD for general matrix operations, exponentiations, and floating-point conversions. On our Zynq XCZU7EV implementation, four VHDL cores are integrated via AXI interconnect with four ARM cores and their corresponding NEON units. The design achieves 42.2x acceleration with 3.1 GOP/s processing rate to decrease one inference latency at 46.3 ms. Moreover, the energy efficiency improves by approx. 9x with the SoC-FPGA consuming only 167 mJ per inference.

A Pipelined FPGA Architecture for Banded Sparse Matrix–Dense Matrix Multiplication in Longformer. (293)

Phillip Pramberger, Athanasios Tziouvaras, Shreejith Shanker and George Floros

Abstract. Sparse attention mechanisms have become increasingly important for transformer models processing long input sequences due to their lower computational and memory complexity compared to full self-attention. Longformer achieves this through a sliding-window attention mechanism that produces a structured banded sparse attention matrix. However, existing sparse transformer accelerators primarily target attention generation or unstructured sparsity, leaving sparse matrix--dense matrix multiplication (SpMM) for structured sparse attention largely unexplored. This paper presents a pipelined FPGA architecture for accelerating banded SpMM in Longformer. The proposed design exploits the predictable sparsity pattern of Longformer's attention matrix through a custom row-wise storage scheme with implicit indexing, eliminating the overhead of conventional sparse matrix formats while enabling regular memory accesses. The architecture employs parallel processing elements, pipelined adder trees, and a dual-path computation strategy to maximize throughput and hardware utilisation. Implemented in Verilog and evaluated on an RFSoC platform using Vivado 2024.2, the accelerator sustains one complete dot-product result per clock cycle after an initial latency of 11 cycles while maintaining power consumption below 2.9 W. Operating at 100 MHz, the design achieves over 100 million dot-product outputs per second, demonstrating the effectiveness of directly exploiting structured sparsity for sparse transformer acceleration.

A Framework for Accelerating Transformer Inference on RISC-V for Edge AI. (301)

Ajay Kumar M, Vishnu P S, Yike Li, Shreejith Shanker, Dimitrios S. Nikolopoulos, Bo Ji, Hans Vandierendonck and Deepu John

Abstract. This work presents a framework for accelerating transformer-based language models (LMs) on resource-constrained IoT devices. The framework targets compact LMs: BERT-Tiny (B-Ty), MobileBERT (M-Bt), MiniLM (M-Lm), Electra (E-Lt) and DeBERTa (D-Bt) — selected for their architectural diversity and use in edge inference scenarios. The proposed flow derives lightweight instruction set extensions tailored to the non- obvious computational patterns of these models. In addition, a custom instruction is introduced to accelerate the address generation stage of batch matrix multiplication, achieving a performance improvement of 15.39–21.74% with modest ASIC overheads of 6.79% in area and 2.33% in power. To further enhance performance without incurring additional processor core hardware cost, an optional compiler-directed loop unrolling strategy is employed, trading increased code size for overall reduced execution time. Evaluation on the Synopsys trv32p3f RISC-V core demonstrates inference speedups of up to 2.19×, and FPGA implementation on the AMD Zynq UltraScale+ ZCU102 shows a 32.94% area overhead at 75 MHz, whereas the ASIC implementation using the TSMC 28 nm library incurs a 21.08% area overhead while operating at 250 MHz.

A Technology-Generic ANN Framework for Modeling and Design of Operational Amplifiers. (313)

Engin Afacan, Günhan Dündar and Sefa Çetinkol

Abstract. Over the last decade, artificial intelligence (AI) has revolutionized engineering fields, with AI-driven systems performing tasks ranging from simple automation to complex optimization. The integrated circuit (IC) industry has adapted to this shift by utilizing AI across various design steps. However, while progress in digital ICs has been rapid, the case for analog/RF circuit design remains challenging due to the high sensitivity of analog circuits to device nonlinearity. Although recent studies have explored AI-driven approaches with promising results, a crucial bottleneck remains: technology dependency. Current ML models developed for the sizing and inverse-sizing of analog circuits cannot transfer their acquired knowledge to different process nodes. To address this limitation, this paper investigates technology-generic model development for operational amplifiers. Specifically, we present a systematic framework that combines a structurally dataset generation approach with AI-based modeling to achieve cross-technology portability. Experimental results indicate that the proposed model can mimic SPICE simulations with an average accuracy of 98.3% among key performance metrics across 5 different technology nodes.

MIVE: A Minimalist Integer Vector Engine for Softmax LayerNorm and RMSNorm Acceleration. (358)

Kosmas Alexandridis and Giorgos Dimitrakopoulos

Abstract. The rapid growth of Large Language Models (LLMs) has intensified the need for specialized hardware accelerators that can satisfy stringent inference latency and power constraints. Although matrix multiplications dominate the overall computational workload, non-linear vector normalization operations, such as LayerNorm, RMSNorm and Softmax can become critical hardware bottlenecks. Existing accelerators typically implement these functions using dedicated hardware blocks, leading to duplicated resources and inefficient silicon utilization. To address this limitation, we propose a Minimalist Integer Vector Engine (MIVE), a programmable architecture capable of executing all three operations within a unified datapath. By exploiting common computational patterns across LayerNorm, RMSNorm and Softmax the proposed vector engine maximizes hardware sharing while reducing implementation overhead. Physical ASIC implementation results show that MIVE provides comprehensive multi-function support while achieving higher area and hardware efficiency than most state-of-the-art standalone accelerators.

Real-time Tunable Probability Distributions From Perimeter-Gated SPAD TRNGs. (375)

Yanbing Icey Chen, Md Sakibur Sajal and Marc Dandin

Abstract. Binary true random number generators (TRNGs) based on the dark noise of perimeter-gated single-photon avalanche diodes (pg-SPADs) naturally produce outputs that follow a tunable Bernoulli distribution. For applications requiring non-binary TRNGs, multiple such units can be coupled, or the inherent Poisson statistics of the accumulated dark events from a single pg-SPAD can be exploited. In this work, we present the principles for extending pg-SPAD-based TRNGs from the binary to the non-binary regime and investigate the impact of the perimeter gate voltage on the underlying probability distributions. The proposed concepts are experimentally validated using a prototype 64 × 64 pg-SPAD chip, fabricated in a standard 0.35 µm CMOS process. Collectively, we show that common non-binary distributions such as Poisson, exponential, and approximated Gaussian can be achieved and tuned in real-time using pg-SPAD TRNGs.

TCAD-Based Analysis of Single Event Transients in 130 nm Bulk CMOS Technology. (380)

Eleftheria Litou, Pelopidas Tsoumanis, Georgios-Ioannis Paliaroutis, Marko Andjelkovic, Nestor Evmorfopoulos and George Stamoulis

Abstract. The increasing susceptibility of modern CMOS technologies to radiation effects necessitates accurate modeling of Single Event Transients (SETs). This work presents a TCAD-based study of radiation-induced SETs in 130 nm bulk CMOS using Synopsys Sentaurus. Both device-level and mixed-mode simulations are employed to analyze NMOS and PMOS transistors, as well as CMOS logic gates, including inverter, NAND, and NOR. Furthermore, the impact of ionizing particle key parameters such as strike location, angle of incidence, Linear Energy Transfer (LET), and device configuration is investigated. Both the generated transient current and the resulting SET pulse characteristics are examined. Overall, a comprehensive analysis of SET generation across device and circuit levels is provided, offering insight into radiation-induced behavior in CMOS technologies. The results reveal a strong dependence of SET characteristics on LET, transistor type, logic-cell topology, and strike angle. More oblique ion strikes (e.g., 30° and 150°) produce higher transient currents, while NMOS strikes generally generate higher peak currents and PMOS strikes produce wider SET pulses.

A Layout-Level Single Event Transient Susceptibility Analysis of 4-2 Adder Compressor. (422)

Ana Flávia Reis, Morgana Rosa, Rafael Soares, Eduardo Costa, Cristina Meinhardt and Rafael Schivittz

Abstract. This paper presents a layout-level Single Event Transient (SET) susceptibility analysis of a 4–2 adder compressor designed using the IHP 130 nm CMOS technology. Post-layout simulations evaluate the fault-susceptible area, propagation delay, power consumption, and total layout area. The results show that XOR gates concentrate the most susceptible active regions, while logical masking in MUX-based outputs reduces the effective susceptibility for a subset of input vectors. The main contribution of this work is the layout-level implementation of the 4–2 adder compressor and a comprehensive characterization of its construction, performance, and robustness under realistic post-layout conditions.

Lightweight Binary Neural Networks for Noise Independent Voice Activity Detection. (427)

Michael Grzybek and Stefano Gregori

Abstract. Voice activity detection (VAD) is a critical front-end component in always-on speech-processing systems, enabling energy savings by activating computationally intensive functions only when speech is present. Although neural-network-based VAD systems improve robustness in noisy environments, their computational and memory requirements limit their applications for resource-constrained systems. Binary neural networks (BNNs) reduce these costs through binary-valued weights and activations; however, their performance is often limited when the input layer is fully binarized, particularly for VAD across diverse contexts. This work proposes a hardware-efficient input expansion method that applies multiple learnable thresholds to each input feature, preserving magnitude information without requiring a full precision input layer. The proposed method achieves an AUROC of 89.64% while requiring only 2.181 KB of memory. This performance is comparable to a full precision input layer BNN, which requires nearly twice the memory. These results demonstrate that learnable binary input expansion can improve noise-independent BNN-based VAD while retaining operations suitable for low-power hardware implementation.

Development of Verilog-A Modules for the Acceleration of Time Domain Noise Simulations. (421)

Panagiota Fasili, Christoforos Theodorou and Thomas Noulis

Abstract. This work presents an exploration of various Verilog-A based modeling methods that aim at the acceleration of noise-aware transient circuit simulations. Thermal noise, 1/f noise and gate noise are all addressed, using 65nm TSMC technology devices as test vehicles, and various implementation ideas for each. These include white noise sources combined with low-pass or high-pass filters, or probability-based approaches closer to the trap physics of the device. All simulation results using the developed models are benchmarked against the standardized simulator-provided Transient Noise method, revealing an acceleration advantage varying from 2 up to 38%, depending on the type of noise and simulation specifics.

MOF-enabled Capacitive Humidity Sensor for Precision Agriculture. (445)

Seerat Sekhon, Maryam Habibollahi, Vikramjeet Singh and Andreas Demosthenous

Abstract. Distributed humidity monitoring requires low-cost sensors with stable baselines for low-power readout. We present a flexible differential capacitive humidity sensor based on UiO-66-NH_2-coated interdigitated electrodes. The exposed IDE tracks humidity-dependent changes in the dielectric properties of the sensing film, while a sealed reference IDE rejects common-mode drift arising from temperature, ageing, and parasitic effects. At a 1 kHz operating frequency, the sensor exhibits a monotonic increase in capacitance from 30 to 49 pF over the 20–90% relative humidity range, with low hysteresis and a sensitivity of up to 1 pF/%RH above 70% RH. The proposed architecture is well suited for integration with compact capacitance-to-digital readout circuits in battery- or solar-powered wireless sensing nodes.

A Ka-band 8-channel Beamformer IC for mmWave 5G and Radar Applications. (446)

Ioannis Kousparis, Nikolaos Naskas, Nikolaos Alexiou, Ilias Tsakiridis, Konstantinos Kontaxis, Fanourios Fakoukakis, Georgios Ntounas and Stefanos Horianopoulos

Abstract. A Ka-band, 24-31GHz 8-channel dual-polarization bulk CMOS Beamformer IC (BFIC) is presented. The circuit achieves a Tx/Rx gain of 30dB/17dB, while employing both amplitude and phase controllability. Tx output 1dB compression point is greater than 10dBm per element and Rx noise figure (NF) per element is 5.4dB including a Time Division Duplex (TDD) switch. Chip size is 5.2x4.4mm2. It is intended for 5G communication applications and mmWave Radar Active Electronically Scanned Array (AESA) systems. Key simulation results are presented herein, verifying its suitability for future 5G and Radar systems.

A 0.8-mW Neural Recording Front End for Vagus Nerve Stimulation Applications with 1.8-µVrms Noise and ~1 ms Stimulation-Artifact Recovery. (452)

Trishya Angela Babs

Abstract. Reliable recording of tiny neural signals in the presence of large stimulation artifacts is a major challenge in implantable vagus nerve stimulation systems. Since stimulation artifacts can be much larger than the neural signals of interest, they may saturate the recording front end and mask early neural responses. To address this issue, this work presents a neural recording system designed to suppress stimulation artifacts and reduce amplifier recovery time to approximately 1 ms. The chip is designed in 180 nm CMOS technology and simulated using an electrode–nerve model that reproduces stimulation-induced artifacts at the recording input. The recorder achieves low input referred noise of 1.8 µVrms, a programmable gain of 49-60 dB, and a bandwidth of 300 Hz to 10 kHz. The recorder supports a maximum input range of ±5 mV while consuming 0.8 mW.

Liquid-Decay Regulation of Switching Sensitivity in a HETUF-Based Memristive Recurrent System. (453)

Bilgi Görkem Yazgaç and Mert Tas

Abstract. A state-dependent liquid-decay mechanism inspired by liquid time-constant networks (LTCs) is introduced into a hybrid exponential Tanh unit function (HETUF)-driven extension of a two-neuron memristive recurrent system. Relative to unit-decay baseline, it shifts the operational onset of robust conditional trajectory separation from k=5.125 to k=5.489 (+7.1%) and reduces the operating-point standard deviations of x and y by 9.6% and 8.1%, respectively. Regularizing the finite HETUF jump at $|u|=2$ reduces the separation rate from 1.1815 to 0.0022 while preserving approximately the same switching-surface crossing rate and stroboscopic geometry. The conclusions are unchanged under the tanh(x) memristor-state drive and support liquid-decay regulation of discontinuity-induced sensitivity rather than a broadband chaotic response.

A Complex I/Q Direct-RF Receiver Using Periodic Non-Uniform Sampling. (454)

Abhishek Kumar

Abstract. Periodic non-uniform sampling (PNUS) is an established alias-spreading technique for direct-RF receivers. This paper proposes an I/Q-compatible receiver interface for PNUS: the real-sampled IF is reconstructed, digitally mixed with cosine and sine local oscillators, and RRC-matched filtered before symbol recovery. The periodic sampling pattern redistributes a blocker fold before this complex reconstruction. For a 50-MHz 16-QAM channel with a 0-dBc alias-band blocker, the proposed receiver reduces RMS EVM from 76.75% with uniform sampling to 5.56%.

A Comparison of Low-Complexity Iterative Algorithms for Adaptive Active Gate Drivers. (455)

Erica Raviola and Franco Fiori

Abstract. Adaptive Active Gate Drivers (AGDs) improve the trade-off between switching losses and transient overshoots by dynamically shaping the gate current. However, their implementation requires the optimization controller to run on limited computational resources while coping with fast load variations. This work compares three low-complexity algorithms for adaptive AGDs, analyzing their computational cost and convergence behavior. The algorithm performance are evaluated on a real AGD test case, showing that steepest gradient descend provides the best trade-off between speed and complexity.

Energy-Aware NAS for Object Recognition using a Piezoelectric Tactile Sensing System. (458)

Hiba Al Youssef, Razan Khalifeh, Leila Hamadi, Mohamad Raad, Maurizio Valle and Ali Ibrahim

Abstract. This paper proposes a multi-state, energy variance-aware neural architecture search (NAS) framework for 1D-CNN based object recognition using a piezoelectric tactile sensing system. Candidate models are evaluated using an Energy Variance-Aware metric that balances accuracy, latency, computational efficiency, and robustness across energy states. The framework is validated on tactile signals acquired from 32 sensing elements during the manipulation of six daily-life objects. The selected models achieve up to 96.56% test accuracy while requiring at most 12.35 KiB of Flash, 9.25 KiB of RAM, 81,192 MACCs, and 0.555 mJ per inference. These results demonstrate that incorporating multi state feasibility and performance variance into the search process enables the identification of energy- efficient architectures that maintain high performance under resource-constrained operating conditions.

A reconfigurable D-Flip-Flop based on reconfigurable NAND/NOR Logic Gates. (461)

Eric Fehrmann, Juan Martinez, Giulio Galderisi, Thomas Mikolajick and Jens Trommer

Abstract. The co-integration of CMOS and reconfigurable field-effect transistors is a promising approach to realize adaptable digital circuits. This work presents a fully reconfigurable D flip-flop based on three-independent-gate RFETs that combines two NAND/NOR-reconfigurable logic gates with a programmable inverter/follower stage, enabling the dynamic reconfiguration of the memory storage.

Design and Evaluation of a Resource-Efficient Direct Digital Synthesizer. (463)

Dimitra Skinezou and Spyridon Nikolaidis

Abstract. DDS (Direct Digital Synthesis) chips are widely employed in modern electronic systems due to their digital programmability, ability to generate multiple waveforms through ROM-based lookup tables (LUTs), fine frequency resolution, and fast frequency switching. Practical DDS implementations involve design trade-offs between area utilization and spectral performance. This work presents the design and experimental evaluation of a resource-efficient FPGA-based Direct Digital Synthesizer using quarter-wave ROM compression, phase truncation, and linear interpolation. The impact of these techniques on waveform quality is examined through theoretical analysis and experimental evaluation in both the time and frequency domains. The design was implemented on a Xilinx Spartan 7 board.

A Modular Hardware-in-the-Loop Framework for Validating AI-Generated Embedded Firmware. (464)

Bharat Kumar, Harshit Krishna R and Madhav Rao

Abstract. Large Language Models (LLMs) are increasingly used to generate microcontroller firmware, creating a critical need for automated validation beyond static analysis and software simulation. Existing evaluation methods fail to expose hardware-specific defects, such as peripheral misconfigurations, protocol violations, and communication failures that manifest only during execution on physical hardware. This paper presents a modular Hardware-in-the-Loop (HIL) framework for validating AI-generated embedded firmware. By treating generated firmware as a black-box System Under Test (SUT), the framework integrates a reusable Hardware Abstraction Layer (HAL), protocol-aware validation suites, and software-based fault injection to enable automated functional verification. A prototype validates GPIO, SPI, and I2C implementations using an STM32 Nucleo-F401RE, an ESP32 SPI slave, and an MPU6050 sensor. Experimental evaluation demonstrates that the framework successfully validates functionally correct implementations, accurately detects representative implementation defects, and isolates failures to their corresponding subsystems, providing a practical methodology for validating AI-generated embedded firmware on physical hardware.

A Semi-Automated g m /I D -Based Technology Migration Flow for Regenerative Mixed-Signal Circuits. (465)

Ashwin George and Klaus Hofmann

Abstract. Migration to newer CMOS nodes creates a growing need to reuse verified analog and mixed-signal IP while reducing redesign effort and verification time. Unlike digital logic, analog circuits cannot be retargeted by geometric scaling because device electrostatics, current density, intrinsic gain, and parasitic capacitances evolve non-uniformly. This problem is especially severe for regenerative circuits governed by nonlinear, time-varying large-signal dynamics. This paper presents a semi-automated schematic-level migration flow that preserves transistor inversion level and branch current using technology-specific g m /I D –I D /W characteristics. Regeneration speed is additionally assessed through g m /C eff. The flow is demonstrated by migrating a PMOS-input StrongARM comparator and regenerative latchfor an LPDDR4X PHY from GF22FDX to GF12LP+ FinFET. Schematic and post-layout simulations confirm multi-Gb/s operation without iterative manual resizing, demonstrating a practical first-pass approach to technology-aware reuse of regenerative mixed-signal IP.

AIHDL: An Agentic AI Based Autonomous Pipeline for Verilog Code Generation, Simulation, Verification and Analysis. (466)

Ashutosh Singh, Anuj Grover and Kunal Iiitd

Abstract. This paper presents an agentic AI system for automated Verilog hardware description language (HDL) generation, simulation, verification and analysis. The proposed framework implements a complete end-to-end pipeline that autonomously generates HDL code from textual specifications or state diagram images, a feature that enhances its versatility. It compiles the generated code, simulates it, and analyzes simulation outputs through a self-correcting feedback loop. The system works with combinational, sequential, and finite state machine (FSM) designs. It orchestrates multiple large language models (LLMs) with specialized capabilities, using Gemini for visual understanding of hardware diagrams, an interchangeable code model for Verilog generation, and Gemma for simulation output analysis. When presented with a circuit diagram or FSM specification, the system interprets the visual input, generates compliant Verilog implementation code including testbenches, executes simulations using industry-standard tools (Icarus Verilog), and analyzes waveform outputs to verify correctness. When a compilation or an output comparison fails, the controller builds a refined prompt from the compiler log and the output mismatch report, then regenerates the design until a satisfactory implementation is reached or the iteration budget is exhausted. Across different standard design tasks the pipeline reaches a 94% pass rate on our directed testbenches within two feedback iterations, against 68% for one shot generation. This framework reduces development time while maintaining quality standards and opens new avenues for collaborative human-AI hardware design workflows.

POSTER SESSION 3

Wednesday November 11th, 2026, 1:15 PM - 2:15 PM

A 64×64, 1162 TOPS/W Analog Capacitive In-Memory Computing Array for Vector-Matrix Multiplication. (17)

Fady Sabry Negm, Max Uhlmann, Corrado Carta, Gerhard Kahmen and Philip Ostrovskyy

Abstract. This paper explores a novel fully analog in-memory computing unit (IMCU) architecture for performing vector-matrix multiplication (VMM) operations. The proposed design is highly energy-efficient because it uses only the inherent energy of the applied input voltage. In this architecture, input vectors are encoded as voltages, with synaptic weights implemented by programmable capacitor banks controlled by dedicated SRAM cells. We performed a complete evaluation of the unit's performance in a 22 nm Fully Depleted Silicon on Insulator (FD-SOI) CMOS technology, focusing specifically on area and energy efficiency. The designed 64 × 64 IMCU runs at 0.52 TOPS, while demonstrating an area efficiency of 1.82 TOPS/mm2, and an energy efficiency of 1587 TOPS/W (4-bit input and weight precision) when clocked at 250 MHz.

Explicit Timestamp Propagation for Temporal Processing. (22)

Xavier Lesage, Cristiano Merio, Azzadine Thajte, Luca Sauer de Araujo, Fernando Welzel, Sylvain Engels and Laurent Fesquet

Abstract. The issue of time representation in digital circuits is often overshadowed by considerations of periodic implementations. Nevertheless, a non-uniform representation has the potential to eliminate temporal redundancies and reduce the number of computations, thereby enhancing system performance. In this contribution, we present a distributed processing system design that relies on an explicit time representation and transactional units. The additional complexity resulting from the transport and computation of timestamps is mitigated by modular arithmetic, limiting the area overhead. To demonstrate the advantages of this method we implemented an array of 256 LIF neurons using 28nm technology. The results show a considerable reduction in dynamic power consumption by a factor of 220, with an area overhead of 60%. Future work includes optimizing static power consumption, increasing the locality of unit synchronization, and exploiting the explicit timestamp values for new temporal processing capabilities.

Dual Diff Query Based Guiding for LLM Spec Edits for RTL Driven Specification Synchronization. (35)

Rahimullah Shaik

Abstract. Specification drift between micro architecture specification (MAS) documents and register transfer level (RTL) code is a significant source of inefficiency during debug, verification, and broader development phases in modern system-on-chip (SoC) programs. This paper introduces a large language model (LLM) assisted approach that automatically proposes localized MAS updates from RTL diffs while preserving document structure and minimizing unintended modifications. The proposed system extracts dual change signals from each diff: Q+ to capture newly added semantics and Q− to capture removed semantics. These signals are utilized to assign one of four targeted update operations to replace, delete, insert or context only prior to generation. The LLM is subsequently prompted with action specific constraints to produce minimal edits and prevent the alteration of unrelated content. We implement this approach within a production grade workflow compatible with Perforce based version control and document repositories, incorporating token budgeted retrieval for lengthy specifications. We further evaluate the generated edits using an LLM as a judge protocol, employing multiple judge models to characterize scoring stability under varying context capacities. Our approach demonstrates higher edit precision, fewer spurious modifications, and improved reviewer acceptance compared to conventional single query retrieval augmented generation (RAG) baselines. Overall, the results indicate that this dual change cataloging strategy offers a highly deployable method for synchronizing specifications with RTL modifications in practical circuit and system design environments.

A 4.2-Gb/s Ultra-Low-Power LPDDR4X DQ Receiver Architecture in GF 22-nm FDSOI. (116)

Ashwin George

Abstract. This work presents a 4.2-Gb/s ultra-low-power LPDDR4X DQ receiver architecture implemented in GF 22nm FDSOI technology for energy-efficient memory-controller interfaces. LPDDR4X reception is challenging because the reduced voltage swing and low input common-mode limit the effectiveness of conventional receiver front-ends. To address this, the proposed receiver employs a PMOS-input StrongARM sense-amplifier comparator for improved sensitivity under low common-mode operation, followed by a regenerative latch for rail-to-rail data restoration and a compact 1:8 deserializer for digital back-end interfacing. Robust double-edge sampling is achieved using complementary DQS strobes, enabling interleaved data detection with a simple front-end structure. Post-layout simulation results confirm correct operation at data rates up to 4.266Gb/s with an effective clock frequency of 2.133GHz while consuming less than 300µW of average power. These results demonstrate that the proposed mixed-signal receiver is well suited for low-power LPDDR4X PHY implementations in mobile and embedded systems.

A High-Performance Single-Stage OTA with 90-dB DC Gain for Nanofarad Capacitive Loads. (136)

Rozhin Salehi Saheb, Erika Covi, Kea-Tiong Tang and Meysam Akbari

Abstract. A single-stage operational transconductance amplifier is realized using a flipped-voltage-follower (FVF) adaptive biasing scheme combined with a conjugated current-mirror architecture. The amplifier can drive nanofarad capacitive loads without requiring a separate slew rate enhancer circuit. To maximize DC gain, the conjugated current mirrors create high impedance internal nodes, which substantially increase the effective transconductance while maintaining stability through a local negative feedback loop. For large signal operation, the slew rate is dynamically boosted by FVF based variable current sources that strongly drive the p-channel differential pair, together with quasi variable current sources that sink the n-channel input devices. Implemented in a standard 0.18-µm CMOS technology, the proposed amplifier occupies an active area of 0.0125 mm². Operating from a 1.8-V supply, post-layout simulations confirm consistent performance across process, voltage, and temperature variations. The design attains a gain-bandwidth product of 1.6 MHz while driving a differential capacitive load of 2×2.5 nF, along with an open-loop gain of 90.83 dB and an average slew rate of 1.49 V/µs.

A Low-Noise, High-Linearity Three-Stage Amplifier for Automotive MEMS Gyroscope Capacitance-to-Voltage Converters. (142)

Raphael Maier

Abstract. This work presents a fully differential three-stage amplifier for automotive MEMS gyroscope capacitance-to-voltage (C/V) converters in 40nm CMOS. The amplifier combines a current-reuse input stage, a class-AB output stage, and hybrid feed-forward/cascode compensation to achieve high gain, low noise, high linearity, and low temperature drift. Post-layout simulations show a gain of 130dB, an input-referred noise floor of 20nV/$\sqrt{\text{Hz}}$, and a worst-case total harmonic distortion (THD) of −78dB. Across −40 ◦C to 125 ◦C, the standalone OTA gain drift remains below 25m%, while the complete C/V gain drift remains below 190m%. The phase drift is below 192 µrad across the same temperature range. The circuit occupies 0.0256mm2 and consumes 42.5 µW.

A Memristive Synapse for Online STDP Learning and Inference in SNNs. (147)

Elia Mateu Barriendos, Álvaro Gómez-Pau, Daniel Arumí, Rosa Rodríguez-Montañés and Salvador Manich

Abstract. This work presents a fully analog memristive synaptic circuit for online spike-timing-dependent plasticity (STDP) learning in spiking neural networks (SNNs). The proposed synapse integrates a local STDP circuit generating gradual timing-dependent conductance updates directly from pre- and post-synaptic spikes. Learning occurs during normal network operation without requiring external digital control or explicit STDP waveform synthesis.

Post-layout simulations of the memristive synapse implemented in a 130nm CMOS technology show timing-dependent synaptic adaptation during normal SNN operation. A 2x2 SNN simulation further illustrates online neuron specialization through unsupervised learning.

On the Design of Current Mirrors Using Only Depletion-Mode Transistors in Monolithic GaN Technology. (158)

Fabio Meloni, Sandro Rossi, Elisabetta Moisello, Edoardo Bonizzoni and Piero Malcovati

Abstract. This paper proposes three current mirror architectures implemented in a monolithic depletion-mode GaN technology from STMicroelectronics: Diodes + Resistor (DR), Diodes + Current Generator (DCG) and Resistors + Current generator (RCG). Due to the normally-on behavior of depletion GaN HEMTs and the absence of a gate oxide, dedicated biasing strategies based on source degeneration are required.

The designs are evaluated through simulations under identical device dimensions and bias conditions in Cadence Virtuoso at 5~V single-supply operation. Performance is assessed in terms of current transfer error, output resistance, and minimum operating output voltage.

Simulation results show near-linear current transfer for all architectures. The DCG mirror achieves the lowest error and a minimum operating voltage of 1.75~V, while the RCG topology provides the highest output resistance (5.7~MΩ) at the cost of increased headroom (2.2~V). The DR solution shows higher error with reduced complexity.

These results highlight the trade-off between current transfer accuracy, output resistance, and voltage headroom in depletion-mode GaN current mirrors.

A 5.26ppm/°C 155-nW CMOS Sub-Bandgap Voltage Reference. (172)

Evi Keramida, George Souliotis, Spyridon Vlassis and Fotis Plessas

Abstract. A CMOS-only voltage reference with a very low temperature coefficient (TC) is presented in this paper. The proposed design relies exclusively on CMOS devices and an operational amplifier (op-amp) to generate the reference voltage, eliminating the need for bipolar transistors commonly used in conventional bandgap reference circuits. The architecture is simple and employs a symmetrical op-amp configuration. The proposed topology has been implemented in a 65-nm CMOS technology. Simulation results demonstrate a TC of 5.26 ppm/°C under typical conditions and 8.92 ppm/°C in the worst-case scenario. The circuit achieves a power supply rejection ratio (PSRR) of −46 dB at 100 Hz while consuming only 172 nA over a supply voltage range of 0.9 V to 1.3 V. The generated reference voltage is 504.4 mV.

A Fractional-N Phase-Locked Loop with a Pseudo-Differential DTC and a Phase-Error Sign Extraction Circuit Based on a Phase-Detector Replica. (183)

Johan Holmstedt, Rainier van Dommele and Vojkan Vidojkovic

Abstract. Abstract— An analog fractional-N sampling phase-locked loop (PLL) is presented. The PLL achieves 80-fs rms jitter in integer-N mode and 400-fs rms jitter in fractional-N mode, integrated from 10 kHz to 100 MHz, and attains a −253-dB figure of merit (FoM) in integer mode and -238 dB in fractional-N mode using a 100-MHz reference clock. The integer spur is -61 dBc. The proposed PLL employs a pseudo-differential digital-to-time converter (DTC) together with background DTC-gain calibration. A key contribution of this work is a phase-error sign extraction circuit based on a replica sampling phase detector (SPD), which enables simple and robust DTC calibration. Furthermore, the paper introduces a simple method to improve robustness against delay mismatch between the divider path and the synchronization flip-flop path. The PLL consumes 8.7 mW and occupies 0.16 mm2.

Bandwidth Interpolation Capability of NARXNN-Based PA Behavioral Model. (202)

Thuy Pham Trong, Germain Pham, Reda Mohellebi, Pierre Almairac, Carolina Pedrosa and Patricia Desgreys

Abstract. This paper presents a compact behavioral modeling approach for wideband RF power amplifiers (PAs) using a Nonlinear AutoRegressive with eXogenous inputs neural network (NARXNN). The objective is to evaluate whether NARXNN can maintain modeling accuracy when the signal bandwidth varies, which is a key requirement for bandwidth-agile RF transmitters.

The proposed model is validated using measured 5G NR input-output data from a 2.6 GHz LDMOS PA over bandwidths from 20 MHz to 100 MHz. Comprehensive model ex-tractions on wideband measurement results are conducted across three validation schemes: standard single-band, standard multi- band, and interpolation.

In the standard single-band setting, NARXNN achieves an average NMSE of −48.60 dB, improving over the GMP reference by approximately 10 dB with moderate model complexity. In the multi-band setting, NARXNN reaches −42.35 dB with 642 parameters, outperforming LSTM, BiLSTM, GRU, and ARVTD baselines. Under interpolation, NARXNN maintains NMSE below −40 dB across intermediate bandwidths, confirming its robustness for bandwidth-scalable PA modeling.

Patch-Based Spatial Encoding for Memristor Crossbar Array Reservoir Computing. (221)

Karol Bednarz

Abstract. Memristive devices are promising physical substrates for reservoir computing owing to their intrinsic nonlinearity and history-dependent conductance. This paper introduces a patch-based spatial encoding strategy that exploits the parallel structure of a memristor crossbar array. Each input image is decomposed into overlapping patches whose rows simultaneously drive the crossbar inputs, allowing each column to aggregate nonlinear contributions from all rows via Kirchhoff's current law in a single forward pass. This preserves local two-dimensional spatial correlations in the input and produces a high-dimensional feature representation. The approach is validated on MNIST digit classification using an 8×8 Self-Directed Channel memristor crossbar, both in simulation with the Mean Metastable Switch model and experimentally on physical hardware. The patchbased reservoir achieves 91 percent classification accuracy in simulation and 87 percent on physical hardware with a linear ridge-regression readout, a 30.6 percentage point improvement over the row-wise encoding baseline in simulation. These results demonstrate that crossbar-level parallel processing of spatially coherent inputs significantly enhances the discriminative capacity of memristive reservoirs.

Fundamental limits of the Pulse Frequency Modulation converter for very high Full-Well Capacity Digital Pixel Sensor. (234)

Martin Lapeyre, Philippe Martin-Gonthier, Olivier Saint-Pe, Thierry Kohl and Vincent Goiffon

Abstract. In this paper, a Pulse Frequency Modulation (PFM) based Digital Pixel Sensor (DPS) Analog-to-Digital Converter (ADC), 100 frames/s full-frame global shutter, prototype is proposed. Developed in 180 nm/1.8V CMOS process node, this DPS has a maximum Full Well Capacity (FWC) of 12 Ge- per frame and can handle photocurrent beyond 100 nA per pixel, driven by MWIR and LWIR applications. Consequently, it benefits from future technology scaling, which both imposes lower supply voltages and drives increasingly digital implementations. A 96x64 pixel-array with several variants of pixels is tested to gain insight into key electrical parameters. Experimental results are compared with simulation, providing a clear understanding of the modulator operated under very high photocurrent and enables the identification of the major architectural and electrical limitations of the proposal system.

A 0.8mW, 22ns Peaking Time, Analog Front-End with Current-Domain Pulse Shaping Filter for High-Density Energy-Resolved Photon-Counting Detector Arrays. (240)

Matteo Petruzzi, Leonardo Gasparini, Stefano D'Amico, Jacopo Quercia and Massimo Gandola

Abstract. This work presents an analogue front-end (AFE) designed in 180-nm CMOS technology for high-rate and high- density X-ray spectroscopy applications, incorporating a current domain-Pulse shaping filter. The shaping filter comprises a cascade of biquadratic stages operating in the current domain to synthesise a pseudo-Gaussian response, achieving peaking times in the range of 10–27 ns. In the context of current-mode operations, it is possible to bypass the use of operational amplifiers, which are known to have significant bandwidth and total power consumption requirements per channel. As demonstrated by the simulation results, the proposed AFE out-stands the state of the art in term of power efficiency, with a total consumption confined between 0.5 mW and approximately 0.8 mW per channel, across a count rate ranging from 100 kpcs to 1 Mpcs. This feature well meets the requirement for a scalable solution for high-density 3D- stacked detector arrays. Furthermore, an equivalent noise charge (ENC) of approximately 24 e^- (rms) is maintained, even if the peaking time is as low as 27 ns, thus guaranteeing an adequate energy resolution.

Machine-learning Based Error Correction in Multi-bit Delta-Sigma Digital-to-analog Converters. (249)

Matthew Kinsinger and Arindam Sanyal

Abstract. This work presents a machine-learning (ML) technique for correction of both static and dynamic errors in multi-bit delta-sigma digital-to-analog converters (DACs). Static errors in delta-sigma DAC arise due to mismatch between unit DAC elements, while dynamic inter-symbol interference (ISI) errors arise due to mismatch between rise and fall times of input pulses driving the DAC elements. Compared to existing dynamic element matching (DEM) techniques, the proposed ML technique achieves the best signal-to-noise-and-distortion ratio (SNDR) with a simple artificial neural network while simultaneously suppressing both static mismatch and dynamic ISI errors.

Hardware-Friendly FIR Architectures for ECG Denoising and QRS Morphology Preservation. (281)

Morgana Macedo Azevedo Da Rosa, Rodrigo Lopes, Eduardo Da Costa and Rafael Soares

Abstract. This paper presents a comparative analysis of hardware-friendly finite-impulse response filters for electrocardiogram preprocessing. Three low-complexity architectures are evaluated: a Hanning-inspired binomial filter, a hardware-friendly Savitzky-Golay filter, and an exact Moving Average filter. The analysis considers both signal-quality preservation and hardware implementation metrics for low-power wearable and embedded biomedical systems. The evaluation focuses on denoising capability, QRS morphology preservation, frequency response, area, and timing performance. Results show that the Moving Average filter achieved the best denoising performance, improving the QRS-local normalized cross-correlation by about 34% and reducing the R-peak amplitude error by more than 85% compared to the noisy ECG signal. The Hanning-inspired filter presented the lowest hardware cost due to its reduced arithmetic complexity. In contrast, the Savitzky-Golay filter offered a balanced trade-off between smoothing and preservation of diagnostic ECG features.

On the Impact of Degradation-Balanced Scheduling in Manycore Systems. (287)

Maria Pantazi-Kypraiou, George Rafael Goudroumanis, George Floros, Anuj Pathania and George Stamoulis

Abstract. Electromigration degradation in manycore systems depends strongly on how runtime workload is distributed across cores. Task placement affects power, temperature, and current density, which in turn shape the spatial distribution of degradation over time. However, a scheduler that ranks cores only by their current reliability state may not always distinguish among candidate cores, especially when reliability values are close at decision time. In this case, scheduling decisions can be driven by secondary factors rather than by long-term degradation balance. This paper studies degradation-balanced scheduling and proposes a Reliability-Balanced (RB) policy that predicts the full-task effect of each candidate assignment. RB penalizes added imbalance in electromigration exposure, stress, active time, and task count while guarding the predicted Rvalue floor. On randomized continuous mixed workloads, the proposed scheduler preserves completed tasks while improving average Rvalue by 0.37% and minimum Rvalue by 4.35%. It reduces Rvalue, exposure, stress, and active-time standard deviation by 10.84%, 9.39%, 8.35%, and 16.38%, respectively.

Supercells: Addressing Mismatch Variability in Analog Compute-in-Memory Arrays. (297)

Michael Tomlinson, Richard Niu, Adil Malik, Philippe Pouliquen and Andreas Andreou

Abstract. Compute-in-Memory (CiM) refers to an alternative computer architecture that is seeing increasing interest as a low-power, low-latency solution to the memory bottleneck present in standard Von Neumann architectures. CiM designs typically co-locate state holding elements (memory) with analog processing, minimizing data movement. These implementations tend to suffer from cell-to-cell mismatch, arising from process variation and limited computational accuracy. In this paper, we propose grouping CiM bitcells into calibrated Supercells to address mismatch. We first decompose compute variability in an existing CiM platform into spatial and temporal components. We then experimentally implement Supercells on this platform, practically eliminating mismatch with 4 calibration bits per Supercell.

Experimental Verification of CMOS-Compatible Capacitive Isolation AC-Switch Driver. (298)

Ming Zhang and Nicolas Llaser

Abstract. In this paper, the experimental test of a CMOS-compatible capacitively isolated (cap-iso) AC-switch driver is presented. Unlike a DC-switch, an AC-switch can conduct an alternative current bidirectionally. To do so, a floating gate driving voltage source is mandatory. To generate such a floating voltage source, an electrical isolation is necessary. Taking into account of the context of our application (medical therapy under MRI monitoring), capacitive isolation is the only choice among the others. The proposed AC-switch driver with capacitive isolation is not only verified by simulation results in CMOS HV 0.35µm technology but confirmed by experimental results as well. A PCB design based on the similar integrated circuit (IC) operation principle was implemented and tested. The main focus of this paper is on the PCB design, the experimental setup and experimental results.

Extended Compact MOSFET Model for Wide Range of Inversion Level. (302)

Daniel do Carmo Silva Ribeiro, Fabian Andrade, Ana Cunha and Edson Pinto Santana

Abstract. The mobility variation with gate voltage is modeled inside the specific current of the Improved ACM model, but the expected behavior is not well reproduced in weak inversion. A more accurate representation of this phenomenon is here proposed to expand the inversion range applicability of the compact model in analog circuit design. The proposed model, thus named Extended ACM model, is applied to the design of simple current sources and sinks and, compared to the Basic and Improved ACM models, proves to give the best simulation results for most design specifications.

A 273 GOPS and 54.6 TOPS/W FeFET-Based CIM Architecture for Edge AI Inference Applications. (328)

Anisha Sharma, Cheena Singhal, Abhishek Goel and Sudeb Dasgupta

Abstract. The rise in demand for AI and machine learning systems has driven the development of Compute-In-Memory (CIM) architectures that solve the von-Neumann bottleneck. The proposed work demonstrates a reconfigurable CIM system architecture with FeFET devices for carrying out logical as well as MAC operations. In this paper, a novel FeFET based CIM architecture is proposed for NAND, NOR, XOR, and Half Adder operations through the exploitation of the nonvolatility, high ON/OFF ratio, and independent read/write paths of FeFETs. In addition, a novel architecture based on 2C-1C weight categorization is proposed for 8-bit MAC operations using a 128×128 FeFET CIM array along with an analog shift & add technique to reduce the number of ADCs. The proposed FeFET memory array completes 4096 8-bit MAC operations in four clock cycles with a latency of 15ns in 45nm CMOS technology and operating voltage of 1V, the throughput of the proposed architecture is 273 GOPS with 54.6 TOPS/W energy efficiency, 5 mW power dissipation and 20 mV signal margin.

A 1-Gb/s Pulse-Width Modulation Transceiver for 15-m Twisted-Pair Links in 28-nm CMOS. (353)

Johannes Kuhn, Tim Lauber, Tristan Strahler, Ralf Wunderlich and Stefan Heinen

Abstract. This paper presents a fully time-modulated wireline transceiver fabricated in a 28 nm technology node, optimized for low-latency gigabit communication over twisted-pair cables in automotive and industrial applications. As technology scaling reduces voltage margins, traditional amplitude-modulated schemes suffer from signal-to-noise ratio degradation. To address this, the proposed design exploits the enhanced time-domain resolution of advanced nodes by implementing high-order pulse-width modulation (PWM-4 and PWM-16). The core of both the digital-to-time and time-to-digital converters relies on a delay-locked loop architecture that inherently cancels systematic nonlinearities and provides robust operation across environmental variations. Operating at a data rate of 1 Gb/s, the transceiver achieves a clock-embedded link over single twisted-pair channels up to 15 m. Experimental results demonstrate a bit error rate of less than 10^-12 for short cable lengths up to 2 m and 10^-10 up to 15 m, offering a low-complexity, highly portable physical layer solution for advanced system-on-chip integration.

An Efficient Method of Optimal Trimming Pattern Identification for ENOB Maximisation in Analog Front Ends. (369)

Arpit Bal and Laxmeesha Somappa

Abstract. To improve the process, voltage and temperature (PVT) robustness of analog integrated circuits trimming can be used, which achieves this by adjusting the key electrical parameters of a chip. With an increase in the number of trimming bits, it is quite challenging and time-consuming for the designers to identify the optimal bit pattern to achieve the best possible performance. Our work proposes an efficient method to identify the trimming pattern to obtain the optimal effective number of bits (ENOB) for any Analog Front End (AFE). Through simulation, it was shown that this approach can reduce the required transient simulation time by 30X using commercial EDA tools, achieving 99.5% accuracy in identifying the optimal trimming code within an error tolerance of 0.5 bits.

FINNAS: FINN-Guided Hardware-Aware NAS and Pruning for FPGA Jet Substructure Classification. (383)

Eva Chauffour, Changhong Li, Georgios Floros and Shreejith Shanker

Abstract. FPGAs are well suited to deploying quantised neural networks (QNNs)

under strict accuracy, latency, and resource constraints; however, identifying efficient model-accelerator combinations commonly requires extensive manual design-space exploration and repeated hardware synthesis. This paper presents FINNAS, a FINN-guided hardware-aware evolutionary neural architecture search framework. FINNAS jointly searches quantised MLP depth, width, and global precision settings, and ranks candidates using proxy validation accuracy together with FINN-estimated LUT usage and latency under a fully parallel mapping. Selected finalists are fully retrained, subjected to post-search unstructured pruning, and validated using RTL simulation and Vivado out-of-context synthesis. On the CERNBox jet substructure classification task, the searched implementations expose competitive accuracy-resource trade-offs. Compared with a manually optimised dense FINN accelerator, a compact FINNAS design improves accuracy from 73.78\% to 74.36\%, while reducing LUT usage by \(8.5\times\) and RTL-simulation latency by \(1.77\times\). Unstructured pruning further provides consistent LUT and FF reductions across the fully parallel finalists.

SET Mitigation Strategies for a Minimal Standard Cell Library in 7 nm FinFET Technology. (387)

Rafael Neves de Mello Oliveira, Rafael Schivittz and Cristina Meinhardt

Abstract. This work evaluates the effectiveness of transistor sizing and DCELL insertion to mitigate Single-Event Transients (SETs) in a subset of a 7~nm FinFET standard-cell library. Cells with minimum size have the circuit susceptibility increased, more than doubling the number of sensitive node/vector combinations compared with the baseline library. DCELL insertion partially mitigates this degradation, reducing the number of vulnerable conditions by up to 37%. The most expressive improvements are obtained with the 3-fin implementation, while the combined 3-fin + DCELL approach achieves the best overall robustness, reducing the total number of sensitive cases from 108 to 50.

These results demonstrate that the 3-fin implementation is the dominant contributor to SET hardening, whereas selective DCELL insertion provides complementary improvements for advanced FinFET technologies.

Enabling Low-Cost Scalable Cryogenic Quantum Control with FPGA-Based SDR: Experimental Validation in a Cryostat. (402)

Gustavo Kanaiama, Francisco Rouxinol and Bruno Sanches

Abstract. The scalability of quantum computers requires operating control electronics within cryogenic environments to reduce noise, connection complexity and increase system integration. This paper evaluates the operational stability, power dissipation, and signal integrity of a software-defined radio (SDR) based qubit control platform under high-vacuum (10^{-6} Torr) and cryogenic conditions down to $-195.156\,^{\circ}C$. The system architecture comprises of an expandable set of Altera Cyclone V FPGAs and an AD9361 RFICs, which was characterized using a custom firmware routine and implementation of digital signal processing tasks, including Fast Fourier Transforms (FFT) and modulated Gaussian control pulse generation up to 6 GHz. Experimental results demonstrate robust stability in the system clock and Digital-to-Analog Converter (DAC) performance, remaining sufficient compensated and invariant across the thermal gradient. Signal integrity analysis, performed via Hilbert transform envelope fitting, yielded a coefficient of determination (R2) of 0.9943, confirming high fidelity pulse generation at cryogenic temperatures. Although an average power dissipation of 2.82 W was observed during transmission (fully active, with half RF output amplitude), the results pave the way for deploying the proposed system at the 50 K to 4 K level of commercial cryostats, offering a scalable and modular solution for cryogenic quantum control.

An Energy Efficient Divider-Free Body-Resonance Transmitter with Direct Oscillator Keying for Human Body Communication. (410)

Amr N. Abdelrahman, Abdelhay Ali and Ahmed M. Eltawil

Abstract. Human body communication (HBC) exploits the human body as a guided signal medium, offering orders-ofmagnitude lower path loss than conventional radio frequency (RF) links for wearable body-area networks. Within the HBC spectrum, the 40–120 MHz body-resonance band provides a quasi-flat, wideband channel uniquely suited to high-throughput, non-coherent on-off keying (OOK) transmission. Achieving µWclass operation, however, remains challenging, as carrier generation in this band has typically required crystal-referenced phaselocked loops (PLLs) or injection-locked synthesizers whose static local oscillator (LO) overhead often dominates the total transmitter (TX) power budget. This paper presents a divider-free body resonance (BR)-HBC TX in TSMC 65 nm general purpose (GP) CMOS at 1 V, combining a thermometric-compensated currentstarved ring oscillator (CSRO), a digital frequency-locked loop (FLL) with 9 µs process, voltage, and temperature (PVT) lock time, and direct oscillator keying (DOK) to tie average power directly to data activity. Post-layout simulation confirms 7.8– 12.3 µW core power across 40–120 MHz, supporting 8–24 Mb/s at 0.51–0.98 pJ/bit core efficiency. At 40 MHz and 8 Mb/s, the total TX — including the tapered body driver for driving skinelectrode load — draws 161.8 µW at 20.2 pJ/bit at 50 % activity, with a core-only overhead of 7.8 µW, occupying a compact active area of 0.0021 mm2 in 65 nm CMOS.

A Fully Differential Charge Pump Amplifier with High Voltage Gain. (418)

Masoud Askariraad and Stefano Gregori

Abstract. This paper presents a novel differential charge pump amplifier with an exponential gain of 3^N, where N is the number of stages. The proposed circuit achieves the highest reported gain among published charge pump amplifier topologies for a given number of stages. It exploits the symmetry of exponential charge pumps and employs a modified bootstrapping technique to enable rail-to-rail input voltage. The impact of parasitic capacitances on conversion ratio and power dissipation is analyzed and compared with a conventional series–parallel charge pump amplifier. Simulation results in a 28-nm CMOS technology show that the proposed amplifier achieves higher power efficiency and reduced gain sensitivity to parasitic capacitances, making it a promising solution for low-power and compact voltage amplifiers.



Authors Appendix

Authors Appendix

A. Prasanna Laxmi
Aarno Pärssinen
Abdeldjabar Benrabah
Abdelhay Ali
Abdenour Kifouche
Abdullah Yousuf
Abdulrahman Nutfaji
Abhishek Goel
Abhishek Kumar
Acharyya Amit
Achim Vandierendonck
Adil Malik
Adnan Gündel
Adnan Haidar
Adrian Evans
Ahilnandan Kabilan
Ahmed Elwakil
Ahmed M. Eltawil
Ahmed Magdy Ibrahim Abdelsamad
Ahmet Samil Demirkol
Aikaterini I. Chronopoulou
Aikaterini Maria Panteleaki
Aikaterini Tsilingiri
Aimilios Leftheriotis
Ajay Kumar M
Alak Majumder
Albert Kibiwott
Alberto García-Ortiz
Alberto Stabile
Alejandro Linares-Barranco
Aleksi Korsman
Alessandra Camplani
Alessandro Andreani
Alessandro Bedoya
Alessandro Catania
Alex Calinescu
Alex Huber
Alex James
Alexandros Gurulian
Alexandros Tsagkaropoulos
Alexandros-Apostolos Boulogeorgos
Alexis Maras
Alfio Dario Grasso
Alfredo Rubino
Ali Ibrahim
Alireza Mesri
Alkis Chatzopoulos
Allen Boston
Álvaro Gómez-Pau
Amit Acharyya
Amit Kumar Gupta
Amr N. Abdelrahman
Amr Wael
Ana Cunha
Ana Flávia Reis
Anagha Raghunath
Anagnostis Boubardas
Anastasios Liontos
Anastasios Tefas
Anderson Roberto Pinheiro Domingues
Andi Buzo
André Cavalcante
Andrea Ballo
Andrea Bevilacqua
Andrea Boni
Andrea Calimera
Andrea De Gregorio
Andrea Fasolino
Andrea Mazzanti
Andreas Andreou
Andreas Demosthenous
Andreas Zankl
Andrei Spelman
Andrei-Tudor Gheorghe
Andres Rojas
Angelo Sampaolo
Angelos Amditis
Angelos Keromytis
Angelos Motsios
Aniket Roy
Anis Allagui
Anisha Sharma
Anjan Banerjee
Anna Maria Voulkopoulou
Anshul Verma
Anteneh Gebregiorgis
Antoine Frappé
Antonio Aprile
Antonio Giuseppe Maria Strollo
Antonio Gomez
Antonio Perez-Pena
Antonios Stavrinidis
Anubhav Bhura
Anuj Grover
Anuj Pathania
Anupama Santhosh
Aobo Cui
Apostolos Fournaris
Arcel Leynes
Argyrios Sideris
Arindam Sanyal
Aristotelis Tsekouras

Arne Van Zegbroeck
Arpit Bal
Arthur Jaccottet
Arthur Paviot
Arturo Gutierrez D'Iaz
Aryaan Mandal
Ashutosh Singh
Ashwin George
Athanasios Kostopoulos
Athanasios Moschos
Athanasios Papadimitriou
Athanasios Tziouvaras
Athina Papagianni
Attilio Fiandrotti
Avi Wayzer
Aymen Romdhane
Azzadine Thajte
Babak Moradi
Baki Berkay Yilmaz
Bappaditya Dey
Benjamin Max
Benoit Larras
Benton Calhoun
Berkay Çiftci
Bharat Kumar
Bhaskar Choubey
Bilgi Görkem Yazgaç
Biswajit Basu
Bo Ji
Boris Hudec
Boris Vaisband
Brecht Vanlaar
Brent Maundy
Brian Holford
Bruno Sanches
Bruno Zatt
Burcu Erkmen
Camillo Perna
Can Tunca
Cao Wang
Caro Meysmans
Carolina Pedrosa
Carsten Knoll
Cássio Jones Dhein da Silva
Cedric Dehos
César Augusto Mission Marcon
Chadi Jabbour
Changhong Li
Charis Diamantidis
Cheena Singhal
Chenxue Hou
Chien-Nan Kuo
Chih-Feng Wu
Chin-Cheng Wei
Chin-Wei Cheng
Chithra
Christian Bruchatz
Christina Dilopoulou
Christoforos Moutafis
Christoforos Theodorou
Christoforos Vasilakis
Christoph Adelmann
Christos Anagnostopoulos
Christos Dimas
Christos Poulos
Christos Sotiriou
Christos Tsigkanos
Christos Tsingiropoulos
Chung-Ho Chen
Chun-Tse Lin
Claudia Georgiana Barbura
Claudio Machado Diniz
Constantinos Efstathiou
Constantinos S. Dionysiou
Corrado Carta
Costas Psychalinos
Cristhopher Mosquera
Cristhopher Rafael Mosquera Montero
Cristiano Merio
Cristina Meinhardt
Cristoforo Marzocca
Cuo Wu
Damian Puchalski
Daniel Arumí
Daniel do Carmo Silva Ribeiro
Daniel Palomino
Daniel Sørensen
Daniel Truesdell
Daniele Rossi
Danilo Karim Kaddouri
Danilo Manstretta
David Keezer
Davide Baroffio
Davide De Caro
Davide Mor
Davis Loose
Debanjan Bhowmik
Debasmita Chatterjee
Deepu John
Dejan Petković
Denis Lebold
Denis Voltmann
Di Li
Diego Avila-Torres
Diego Guglielmi
Dimitra Skinezou
Dimitrios Fragkos
Dimitrios Katsikas
Dimitrios Prousalis
Dimitrios S. Nikolopoulos

Dimitrios Schoinianakis
Dimitrios Sekertzis
Dimitrios Soudris
Dimitrios Spanos
Dimitrios Spithouris
Dimitrios Tsalapatas
Dimitrios Tsoukalas
Dimitrios Tzovaras
Dionysios Kehagias
Dionysios Reisis
Dominik Hiltbrunner
Dongchen Li
Dongdong Chen
Dongsop Lee
Đorđe Gačić
E Richard Prajwal Das
Ebubekir Enes Arslan
Ece Olcay Güneş
Édney Freitas
Edoardo Bonizzoni
Edoardo Cignoni
Edson Pinto Santana
Eduarda Moraes dos Santos
Eduardo Costa
Eduardo Da Costa
Eleftheria Litou
Eleftherios Sarris
Eleni Papageorgiou
Elham Hatamzadeh
Elia Mateu Barriendos
Elias Ruopp
Elisabetta Moisello
Emeric De Foucauld
Emilio Calvanese Strinati
Emilio Corigliano
Emmanuel John Bonifacio
Emre Sayıt
Engin Afacan
Enrico Macii
Eric Fehrmann
Eric Guthmuller
Erica Raviola
Erich Malan
Erik Wehr
Erika Covi
Errikos Lourandakis
Esteban José Garzón Córdova
Esther Bamiduro
Eter Mgeladze
Ettore Noccetti
Eva Chauffour
Evangelos Aslanidis
Evangelos Kyparoglou
Evi Keramida
Fabian Andrade
Fabio Benevenuti
Fabio Meloni
Fady Sabry Negm
Fahim Shahriar
Fanourios Fakoukakis
Farshad Moradi
Fayçal Benyamina
Federico Reghenzani
Federico Sandoval Ibarra
Fernanda Kastensmidt
Fernando Gehm Moraes
Fernando Moraes
Fernando Welzel
Florian De Roose
Florian Protze
Florin Ciubotaru
Florin-Silviu Dumitru
Foteini Oikonomou
Fotis Plessas
Francesca Bagnoli
Francesco da Rocha Lui
Francesco Gagliardi
Francesco Maria Puglisi
Francisco Brito
Francisco J. Gómez-Pulido
Francisco Rouxinol
Francky Catthoor
Franco Fiori
Francois Rivet
François Rivet
Frank Ellinger
Frank Herzel
Frédéric Rousseau
G V N Mokshagna
Gabriel Souza
Gabriele Salzone
Gang Lin
Gennaro Di Meo
Georg Pelz
George Alexakis
George F. Fragulis
George Floros
George Goudroumanis
George Karachalios
George Konstantinidis
George Lentaris
George Rafael Goudroumanis
George Souliotis
George Stamoulis
Georgia Tsirimokou
Georgios Alexandris
Georgios Ch. Sirakoulis
Georgios Chararas
Georgios Floros
Georgios Karagiannidis

Georgios Karakonstantis
Georgios Kleitsiotis
Georgios Ntounas
Georgios Panagopoulos
Georgios Theodoridis
Georgios-Ioannis Paliaroutis
Gerardo Pelosi
Gerhard Kahmen
Germain Pham
Giacomo Lancellotti
Gian Domenico Licciardo
Giannis Kontogiannopoulos
Giannis Rokomos
Gianvito Matarrese
Gilberto Kreisler
Gilles Masson
Giorgio Bersani
Giorgio Melchiorre
Giorgio Rolandi
Giorgos Balis
Giorgos Dimitrakopoulos
Giovanni Agosta
Giovanni Alberto Sartorato
Giovanni Di Pietrantonio
Giulio Filippeschi
Giulio Galderisi
Giuseppe Iannaccone
Gordon Pipa
Gregory Worrell
Guangzhi Tang
Guilherme Corrêa
Guillaume Prenat
Günhan Dündar
Gunnar Tufte
Gunter Fischer
Gustavo Kanaiama
Gustavo Marcati
Guy Torfs
Hajer Abdelaziz
Halil Durmuş
Hamed Nasiri
Hampus Hoffman
Hana Maged
Hannes Simon Korn
Hans Ingelberts
Hans Vandierendonck
Hanspeter Schmid
Harald Pretl
Harris E. Michail
Harshit Krishna R
Hasan Uluşan
Hassan Rekabi Bana
Helen Clara George
Heloise Therien-Aubin
Hendrik Meersseman
Hendrik Wöhrle
Heng Yu
Henk Corporaal
Hiba Al Youssef
Hongliu Yang
Hong-Yi Huang
Hossein Esmailbeygi
Hung-Ming Kuo
Hwang Yoon
Hwann-Kaeo Chiou
Hyemin Ko
Hyunwoo Kim
Iacopo Nannipieri
Ilias Kalouptsoglou
Ilias Tsakiridis
Ioannis Bartsiokas
Ioannis K. Chatzipaschalis
Ioannis Kousparis
Ioannis Messaris
Ioannis Sofianidis
Ioannis Tompris
Ioannis Voyiatzis
Ion Vornicu
Iosif-Angelos Fyrigos
Iraklis Anagnostopoulos
Irene-Maria Tabakis
Ismael Seidel
Ivan Lisovyi
J Sai Maneesh Kumar
Jacopo Quercia
Jaekwan Kwon
Jafar Shojaii
Jaime Aguiar
James Lambert
Jan Genoe
Jaroslav Dzuba
Javad Ahmadi-Farsani
Javad Pourmohammad
Jelle Biesmans
Jens Rettkowski
Jens Trommer
Jens Wagner
Jhan Luke Okkabaz
Jianbin He
Jiayang Li
Jingjin Li
João Barreiros Rodrigues
Johan Holmstedt
Johannes Hellwig
Johannes Kuhn
Johannes Reibold
John Kauffman
John Liaperdos
John Paul Strachan
John Richard Hizon

Jonas Gillner
Jonathan Miquel
Jonathan Pitzschel
Jonathan Vrijsen
Jongwook Kim
Jorge Castro-Godínez
José Benjamin Zambrano Chavez
José Luis Güntzel
José Luís Güntzel
Jose M. de La Rosa
José M. de la Rosa
Jose Marie Acedillo
José Rodrigo F. Azambuja
José T. de Sousa
Josep Rius
Juan Martinez
Jules Dubois
Julia Nako
Juliano Maia
Junjie Zhao
Jussi Ryynänen
Kaijie Ding
Kamal Y. Kamal
Kamel Abouda
Kamel-Eddine Sadadou
Kandagatla Venkata Rohith Sai
Kareem Abdelraof
Kareem Atef
Kari Stadius
Karim Osama
Karol Bednarz
Karylle Shainne Cadalzo
Katia Samperi
Kazuki Monta
Kea-Tiong Tang
Kemal Ozanoglu
Kenny Vohl
Kerem Yurtseven
Kevin Valakuzhy
Khang Nguyen
Klaus Hofmann
Klaus Thelen
Konstantina Karagianni
Konstantinos Katsaros
Konstantinos Kontaxis
Konstantinos Liakos
Konstantinos P. Pagkalos
Konstantinos Spyridopoulos
Konstantinos Tatas
Kosmas Alexandridis
Kris Myny
Kunal Iiitd
Kuo-Hsing Cheng
Kyriaki Niotaki
Lahcen Attijani
Lakhdar Benhamimid
Lang Lin
Laurent Fesquet
Laurent Ouvry
Laurentiu Acasandrei
Laxmeesha Somappa
Lazar Ilić
Lazaros Karagiannidis
Lazaros Papadopoulos
Leila Hamadi
Leila Sharara
Léo Ribeiro
Leonard Schauer
Leonardo Antonietti
Leonardo Gasparini
Leonardo Miceli Lara
Leonardo R. Gobatto
Leonardo Soares
Leonidas Paratsokas
Lihong Zhang
Lirida Naviner
Lorena Anghel
Lorenzo Benatti
Luca Bertolini
Luca Frontini
Luca Sauer de Araujo
Luca Sterpone
Luca Tegazzini
Lucas Bublitz
Lucas Ott
Lucas Tinoco
Ludovica Deriu
Luigi Di Benedetto
Luigi Lombardi
Luis Arturo Martinez Alvarado
Luis Bravo
Luis Carlos Rodrigues Linares
Luís França
Luis G. Leon-Vega
Maarten Kuijk
Madhav Rao
Makoto Nagata
Maloth Vishnu Vardhan
Manil Dev Gomony
Manolis Katsaragakis
Marc Chevalier
Marc Dandin
Marc Rosales
Marc Tiebout
Marcela Zachariášová
Marcio Monteiro
Marco Bolognesi
Marco Fattori
Marco Jefferson Mamani Arapa
Marco Lanuzza

Marco Longo
Margherita Scognamiglio
Maria Pantazi
Maria Pantazi-Kypraiou
Maria Sophia Ralota
Maria Theresa De Leon
Mariam Ahmed
Marián Precner
Marius Enachescu
Mark Louie Pasturan
Marko Andjelkovic
Marko Kosunen
Marko Simicic
Markus Petri
Martin Lapeyre
Maryam Habibollahi
Masoud Askariraad
Massimo Gandola
Massimo Piotto
Massimo Ruo Roch
Matej Horský
Mateus Grellert
Mathieu Coustans
Matteo Petruzzi
Matteo Sommella
Matthew Kinsinger
Matthieu Couriol
Mattis Hasler
Mauricio Velazquez Lopez
Maurits Ortmanns
Maurizio Martina
Maurizio Valle
Max Uhlmann
Mayu Zhang
Md Sakibur Sajal
Megan Gunn
Mehdi Tahoori
Meng-Xuan Lin
Mert Tas
Meysam Akbari
Meysam Shayestehfar
Michael Grzybek
Michael Karagounis
Michael Tomlinson
Michał Choraś
Michele Caselli
Miguel Angel García Andrade
Mihalis Psarakis
Miltiadis Siavvas
Mina Nazarian Samani
Minas Dasygenis
Ming Cheng
Ming Zhang
Ming-Der Shieh
Mingi Jeong
Ming-Jay Yang
Minsoon Hwang
Min-Yuan Tsai
Mirko Brazzini
Mitra Mirhassani
Mohamad Raad
Mohamed Abd El Ghany
Mohamed Amine Yahiaoui
Mohamed Benbouzid
Mohamed El-Hadedy
Mohamed Salah Azzaz
Mohamed-Amine Hamoura
Mohammad Al Janaideh
Mohammad B. Rou
Mohammad Fahim Sultan Anoy
Mohammed Ismail
Mohit Dandekar
Morgana Macedo Azevedo Da Rosa
Morgana Rosa
Moslem Heidarpur
Mostafa Waseem
Mouaad Belguedri
Mouli Venkata Prakash Pittala
Muhammad Dawood Asghar
Muhammad Zia Hameed
Muh-Tian Shiue
Mykhailo Zarudniev
Nader Saffari
Nader Sehatbakhsh
Nathaniel Kaye Mellor
Nathaniel Orillaza
Nathaniel Renz Domingo
Nebojsa Maletic
Nefeli Maria Panagiota Kavouni
Neha Ashar
Nele Mentens
Nestor Evmorfopoulos
Nicolas Llaser
Nicolas Reyren
Nidhi Verma
Nikolaos Alexiou
Nikolaos A. Anagnostopoulos
Nikolaos Makris
Nikolaos Naskas
Nikolaos Papageorgiou
Nikolaos Passalis
Nikos Peppes
Niranjan Gopal
Nise O'Cuill
Nithin Sai Reddy
Nithya Kalyani Ramachandran
Nobukazu Takai
Noor Ul Amin
Norman Chang
Nuri Ince

Ojasweeta Singh
Okan Zafer Batur
Olivier Saint-Pe
Oran Hayes
Orazio Aiello
Otmane Ait Mohamed
Otto Simola
P V Phaneendra Kumar
Panagiota Fasili
Panagiotis Antoniou
Panagiotis Chaidos
Panagiotis Karathanasis
Panagiotis Vlachos
Panayiotis Pericleous
Pantazis Anagnostou
Pantelis Fraidakis
Paoline Coulson
Paolo Bruschi
Paolo Fioravanti
Paolo Stefano Crovetti
Paris Kitsos
Patricia Desgreys
Patrick Hickey
Paul Capgras
Paul Farago
Paul François
Paul Jason Co
Paul P. Sotiriadis
Paul-Catalin Medinceanu
Pavlos Stoikos
Pavlos Tosidis
Pawan Patel
Pelopidas Tsoumanis
Petr Bardonek
Petros Spachos
Philip Ostrovskyy
Philippe Descamps
Philippe Martin-Gonthier
Philippe Pouliquen
Phillip Pramberger
Piero Malcovati
Pierre Almairac
Pierre Courouve
Pierre-Emmanuel Gaillardon
Pieter Harpe
Pietro Patimisco
Pio Esposito
Qihan Wang
Qi-Xian Wu
Quanwen Zheng
Rafael Neves de Mello Oliveira
Rafael Rodrigues
Rafael Schivittz
Rafael Soares
Rahimullah Shaik
Rainier van Dommele
Ralf Wunderlich
Ramanathan Muthiah
Ramiro Taco
Ramy Raafat
Rana Walied Ahmad
Raphael Maier
Razan Khalifeh
Reda Mohellebi
Regina Dittmann
Remco Schalk
Richard Niu
Richard Prajwal
Richard Schroedter
Rikuu Hasegawa
Rita Louro Barbosa
Robert Bogdan Staszewski
Robert Krämer
Robert van Veldhoven
Roberto Miorelli
Rodrigo Lopes
Rodrigo S. Cenço
Roel Wuyts
Rohan Purkait
Rolf Paulus
Roman Obermaisser
Ron Louis Nierva
Ronald Tetzlaff
Rosa Rodríguez-Montañés
Rosalba Liguori
Rosana Rodriguez
Roy Taylor
Rozhin Salehi Saheb
Rúben Teixeira
Sabrin Alla
Sagar Bhivsane
Sahibia Kaur Vohra
Said Hamdioui
Sakshi Arora
Salvador Manich
Salvatore Pennisi
Samit Hasan
Sandip Halder
Sandro Rossi
Sanguk Lee
Sara Alessa
Sarah Azimi
Sarthak Dharmendra Nandanwar
Sasya Merugu Sri
Sathish N
Savvas Karipidis
Savvina Nastou
Sebastian Fieldhouse
Sebastian Fischer
Sebastian Haesler

Sebastian Wittlich
Sebastiano Strangio
Seerat Sekhon
Sefa Çetinkol
Sejin Park
Sergei Krylov
Sergio Lecce
Sergio Saponara
Seung-Ki Lee
Seungyoung Ahn
Shamma Alblooshi
Shamoil Ejaz
Shiuan-Han Huang
Shiva Nejati
Shreejith Shanker
Shubham Sharma
Sima A. Alidokht
Sima Ghandi
Simone Contardi
Simone Perriello
Socratis Nazlidis
Sorin Cotofana
Sorin Hintea
Sotirios Aias Karioris
Sotirios Xydis
Sounak Samanta
Souris Sahu
Spyridon Athanasiadis
Spyridon Nikolaidis
Spyridon Vlassis
Spyros Tragoudas
Stan Oosterlynck
Stavros Simoglou
Stavros Stathoudakis
Stefaan Decoutere
Stefan Heinen
Stefan Slesazeck
Stefano Cherubin
Stefano D'Amico
Stefano Gregori
Stefano Quer
Stefanos Horianopoulos
Steffen Seitz
Stoyan Kirilov
Stylianos Kotsobolos
Sudeb Dasgupta
Sudip Bandyopadhyay
Sumanta Chaudhuri
Sumit Datta
Sven Baerten
Sven Finke
Sven Krause
Sylvain Engels
Takayoshi Namura
Takuya Wadatsumi
Tales Pimenta
Tao Sun
Tarsio Onofrio
Thanasis Douklias
Theodora Sanida
Theodoros Panagiotis Chatzinikolaou
Theofilos Spyrou
Thierry Kohl
Thierry Lagutere
Thomas Bedecarrats
Thomas Lapauw
Thomas Mikolajick
Thomas Noulis
Thomas Van den Dries
Thuy Pham Trong
Tibor Izsák
Tim Lauber
Ting-Chang Chang
Tobias Krawutschke
Tobias Schwanke
Tobias Stelzer
Tom Herzog
Tomas A. Lopez
Tomáš Ščepka
Tommaso Zanotti
Tong-Da Lee
Trishya Angela Babs
Tristan Strahler
Tristan Thierfelder
Tsu-Hsiung Huang
Tsung-Han Tsai
Tze Hin Cheung
Ulf Kulau
Umer Amer Khan
Uri Stroh
Urmimala Chatterjee
Uygar Yildiz
Valentin Stadtlander
Valentina Davidoiu
Valentino Peluso
Valeri Mladenov
Varatheepan Paramanayakam
Vasileios Kalenteridis
Vasileios Konstantakos
Vasileios Ntinas
Vasileios Tenentes
Vasilis Pavlidis
Vassilis Alimisis
Vassilis Paliouras
Vassilis Vassios
Verneri Hirvonen
Vikramjeet Singh
Viktor Weinelt
Vincent Cros
Vincent Goiffon

Vincenzina Zecchino
Vincenzo Spagnolo
Vinicius Correa
Vishesh Kiran Kumar Bafna
Vishnu P S
Vladimir Milovanović
Vojkan Vidojkovic
Wannes Nevens
Wen-Hsuan Hsieh
Wenshan Hu
William Fornaciari
Wim Dehaene
Xavier Lesage
Xenofon Vasilakos
Xiao-Chun Li
Xiaonan Xing
Xin Xu
Xingwei Zhou
Xinming Shi
Xu Wang
Yağız Kırtaş
Yajun Ha
Yanbing Icey Chen
Yanxia Lin
Yash Sengupta
Yasmin Halawani
Yike Li
Yintang Yang
Yiorgos Tsiatouhas
Yiqin Hou
Yoosang Na
Youngjin Park
Yu Shang-Yin
Yu-Jun Chen
Yun-Ting Hsu
Yusra Rachidi
Yusuf Serhat Özkan
Yuto Moriguchi
Yuying Zhuo
Zane Gunn
Zhaoxi Li
Zhongcheng Lei
Zidu Li
Zipeng Mi
Ziyi Zhao
Zoya Ahmed
Zuo Liu

Track Chairs

Track Chairs

George Athanasiou
Georgios Keramidas
Panagiotis Mousouliotis
Vasilis Ntinas
Debjit Pal
Georgios Panagopoulos
Maria Papadopoulou
Emre Salman
Nikolaos Sketopoulos
Dimitrios Soudris
Zoran Stamenkovic
Boris Vaisband



Reviewers

Reviewers

Aathreya Bhat
Abdelati Reha
Abderrahmane Ghimouz
Abdoul Rjoub
Abolfazl Roshanpanah
Abu Baker
Achilles Boursianis
Ahmad Yarahmadi
Ahmed Elwakil
Ahmed Madian
Ahmet Samil Demirkol
Aikaterini Tsilingiri
Alejandro Medina
Alessandro Catania
Alex James
Alexandre Serres
Alfio Dario Grasso
Ali Dogus Gungordu
Ali Ibrahim
Alireza Ahrar
Alkis Hatzopoulos
Amin Safaei
Anca Molnos
Andrea Ballo
Andrea Bandiziol
Andrea Bonfanti
Andreas Tsiougkos
Angela Slavova
Angelos Kyriakos
Anthony Kolar
Antonello Florio
Antonio Carlos Schneider Beck
Antonio Lopez-Martin
Antonio Strollo
Argyrios T. Hatzopoulos
Argyris Kokkinis
Aristotelis Tsekouras
Arnaldo Oliveira
Arvin Delavari
Asma Ben Letaifa
Athanasios Kakarountas
Athanasios Papadimitriou
Atilla Uygur
Ayoub Fentis
B. Bhuvan
Bandan Bhoi
Bappaditya Dey
Barak Hoffer
Baudouin Martineau
Bertrand Granado
Bhaskar Choubey
Boon Chong Ang
Boris Vaisband
Bruno Miranda
Carl James Debono
Carol de Benito
Catalin Visan
Ch Chan
Chao-Lieh Chen
Christoforos Kachris
Christoforos Theodorou
Christophe Jego
Christos Baloukas
Christos Georgakidis
Christos Spandonidis
Costas Psychalinos
Cristian Grava
Cristina Meinhardt
Dalibor Biolek
Daniel Chillet
Danijel Dankovic
Danilo Santoro
Debjit Pal
Deepak Dasalukunte
Devesh Bhaskaran
Dietmar Fey
Dimitri Galayko
Dimitrios Garyfallou
Dimitrios Kampelopoulos
Dimitrios Papakostas
Dimitrios Prousalis
Dimitrios Soudris
Dimitris Bakalis
Dimitris Karampatzakis
Dimosthenis Masouros
Dionysios Reisis
Dolatshahi Mehdi
Dominique Morche
Dr. Joseph Anthony Prathap
Edna Barros
Edson Lisboa
Edward David Moreno
Efstathios Kyriakis-Bitzaros
Eftichios Koutroulis
Elisabetta Moisello
Elissaios Alexios Papatheofanous
Emanuele Gemo
Emeric de Foucauld
Emmanouil Kavvousanos
Emmanuel Amankrah
Erdem Köse
Erik Wehr

Erivelton Nepomuceno
Esteban Tlelo-Cuautle
Ettore Napoli
Evangelos Tzouvaras
Evi Keramida
Fábio Coutinho
Fabio Frustaci
Fabio Ramos
Faik Baskaya
Fanny Spagnolo
Farshad Moradi
Federico Alimenti
Fernando Corinto
Fernando Moraes
Filippo Carloni
Florent de Dinechin
Foad Beheshti
Fotis Giannopoulos
Fotis Plessas
Francesco Centurelli
Frank K. Gürkaynak
Fu-Qi Chua
Gabriel Galeote Checa
Gaetano Palumbo
Gage Hills
Gennaro Di Meo
George Athanasiou
George Lentaris
George Souliotis
George Theodoridis
Georgia Tsirimokou
Georgios - Panagiotis Kousiopoulos
Georgios Chararas
Georgios Dimitriou
Georgios Karakonstantis
Georgios Keramidas
Georgios Selimis
Georgios-Ioannis Paliaroutis
Gianfranco Avitabile
Gianluca Giustolisi
Gianluca Traversi
Gildas Leger
Giorgos Dimitrakopoulos
Giovanni Agosta
Giuseppe Esposito
Giuseppe Ferri
Giuseppe Scotti
Guilherme Corrêa
Hamidreza Rezaee-Dehsorkh
Haodong Chang
Hassan Al-Lawati
Hassan Rabah
Hassen Aziza
Henrique Freitas
Hesam Abbasi
Hieu Ngo
Himshekhar Das
Hitesh Shrimali
Hongliu Yang
Hsin-Liang Chen
Hugo García-Vázquez
Hüseyin Canbolat
Ian O'Connor
Ignacio Santamaria
Ihsan Cicek
Ikhwana Elfitri
Ioannis Antoniades
Ioannis Kouretas
Ioannis Messaris
Ioannis Sofianidis
Iosif-Angelos Fyrigos
Iraklis Anagnostopoulos
Janaina Guimaraes
Javier Beloso Legarra
Jayant Nandwalkar
Jean-Baptiste Begueret
Jelena Popovic-Bozovic
Jesus Manuel Muñoz-Pacheco
Jia Zhao
Jiaxuan Cai
Joao Casaleiro
Joao Goes
Johan Arbustini
John Kalomiros
John Vourvoulakis
Jose de Jesus Morales-Romero
Jose Delgado-Frias
Josep Maria Sánchez Chiva
Josie Esteban Rodriguez Condia
Juan Guerrero
Juan Oliver
Juan Pedro Dominguez-Morales
Julien Lintignat
Kamil Mielcarek
Kawori Sekine
Keertana Alugoti
Kemal Ozanoglu
Khalid Javeed
Khalil Yousef
Khoirom Johnson Singh
Kiichi Niitsu
Klaus Tittelbach-Helmrich
Konstantinos A. Tsintotas
Konstantinos Laoudias
Konstantinos Manolopoulos
Konstantinos Nakos
Konstantinos Tatas
Korkut Tokgoz
Lan-Da Van
Lazaros Moysis

Lazaros Papadopoulos
Lennart Bamberg
Leonel Sousa
Li Honggui
Liangliang Chang
Liangtao Dai
Lilian Bossuet
Linlin Cai
Lirida Naviner
Loukas Michalas
Luca Pancioni
Luca Sterpone
Lucas Cambuim
Luigi Giuffrida
Luis Lolis
M.B.Hemanth Kumar
Manar Abdelatty
Manel Ben Romdhane
Manolis Katsaragakis
Marcela Zachariasova
Marco Carminati
Marco Privitera
Marco Storace
Maria Papadopoulou
Marius Enachescu
Mark Zwolinski
Marko Andjelkovic
Marko Kosunen
Markus Petri
Massimo Macucci
Mathieu Coustans
Mathieu Thevenin
Matias Miguez
Matteo Sonza Reorda
Maurizio Martina
Maurizio Valle
Mervat Mahmoud
Mete Celik
Metin Yazgi
Michael Birbas
Michael Cave
Michael Peter Kennedy
Mihalis Psarakis
Minas Dasygenis
Ming Zhang
Mirhojjat Seyedi
Mitar Simić
Mohamed Abouelatta
Mohammed Alselwi
Mohd Shahrizal Rusli
Mustafa Yelten
Neelesh Biswas
Nele Mentens
Neslihan Serap Şengör
Nestor Evmorfopoulos
Nicola Massari
Nicolas Dehaese
Nicolas Sklavos
Nicole McFarlane
Nihan Kahraman
Nikolaos Athanasios Anagnostopoulos
Nikolaos Georgoulopoulos
Nikolaos Sketopoulos
Nikolaos Tampouratzis
Nikolaos Vlassopoulos
Nikolaos Zazatis
Nikos Petrelis
Nishan Chettri
Nithin Kumar
Odysseas Bouzos
Oguzhan Cicekoglu
Okan Zafer Batur
Olympia Axelou
Onur Yusuf Çinar
Orazio Aiello
Otmane Ait Mohamed
Panagiotis Bertsias
Panagiotis Karydopoulos
Panagiotis Mousouliotis
Pankaj Siri Bharath Bairu
Paolo Crovetti
Paolo Marconcini
Paris Kitsos
Pavankumar Bikki
Pelopidas Tsoumanis
Peter Szolgay
Piergiulio Mannocci
Piero Malcovati
Pritam Bhattacharjee
Rafael Castro Lopez
Rafael Schivittz
Rahul Chaurasia
Rashed Al Amin
Remi Queheille
Remy Vauche
Revna A. Vural
Ricardo Martins
Richard Schroedter
Robert Henderson
Rodrigo Picos
Rodrigo Possamai Bastos
Roman Sotner
Ronald Tetzlaff
Ruizhe Zhang
Sahbi Baccar
Saikat Chatterjee
Salvatore Pennisi
Sana Ullah
Sanchari Das
Satyadev Ahlawat

Sedat Kilinc
Sercan Aygun
Sergio Callegari
Sergio Saponara
Shahaboddin Ghajari
Sherif Abouzeid
Shi Bu
Shifu Wu
Shruti Jain
Siddik Yarman
Silvia Mella
Sohaib Majzoub
Sorin Cotofana
Sotiria Galata
Sotirios Goudos
Spyridon Nikolaidis
Spyridon Vlassis
Srikar Bhagavatula
Stavroula Kapoulea
Stefan Popa
Stefano Concadoro
Stefano Gregori
Suprio Bhattacharya
Sylvain Feruglio
Sylvain Guilley
Tales Pimenta
Tamas Zsedrovits
Tamoghno Das
Tayfun Nesimoglu
Taylor Barton
Thomas Noulis
Tiago Almeida
Tobias Zekorn
Tommaso Rizzo
Tuba Ayhan
Tyler Archer
Ugur Cini
Vagner Rosa
Valeri Mladenov
Varun Jeoti
Vasileios Kalenteridis
Vasileios Konstantakos
Vasileios Tenentes
Vasiliki Gogolou
Vasilis Ntinas
Vasilis Pavlidis
Vassilis Alimisis
Vassilis Paliouras
Vianney Lapôtre
Victor R. Gonzalez-Diaz
Vikas Rana
Vincent Rijmen
Vladimir Milovanović
Wanderson Roger Azevedo Dias
Yi Chen
Yianna Danidou
Yihang Cheng
Yinan Wang
Yingjie Chen
Yiorgos Tsiatouhas
Yue Jiet Chong
Zoltan Nagy
Zoran Stamenkovic



