





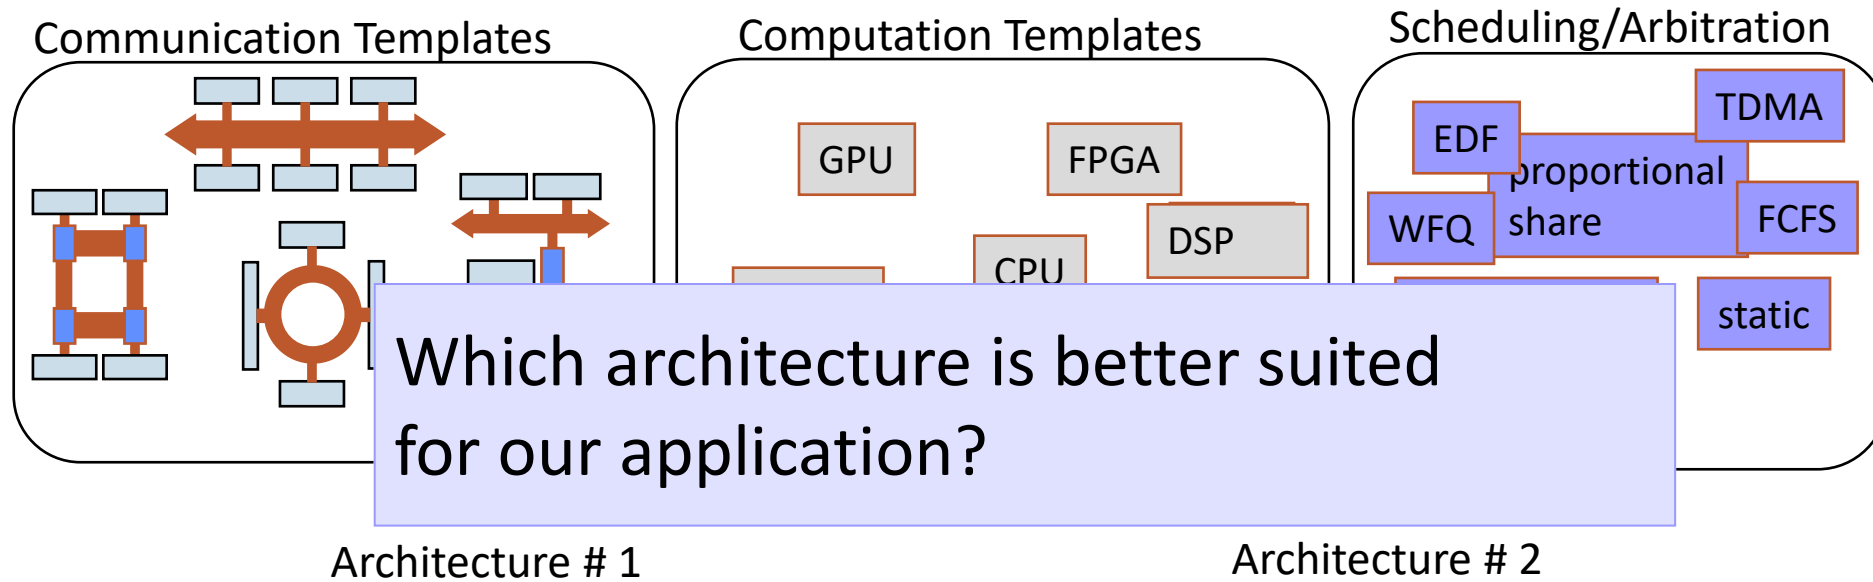
Title: Complex application mapping to heterogeneous compute resources

AUTHOR 1: WESLEY SKEFFINGTON, AMD, ALBANY, NY, USA, WESLEY.SKEFFINGTON@AMD.COM

AUTHOR 2: SURYA CHONGALA, AMD, HYDERABAD, INDIA, SURYA.CHONGALA@AMD.COM

AUTHOR 3: DEEPAK SHANKAR, MIRABILIS DESIGN INC, CHENNAI, TN, INDIA, DSHANKAR@MIRABILISDESIGN.COM

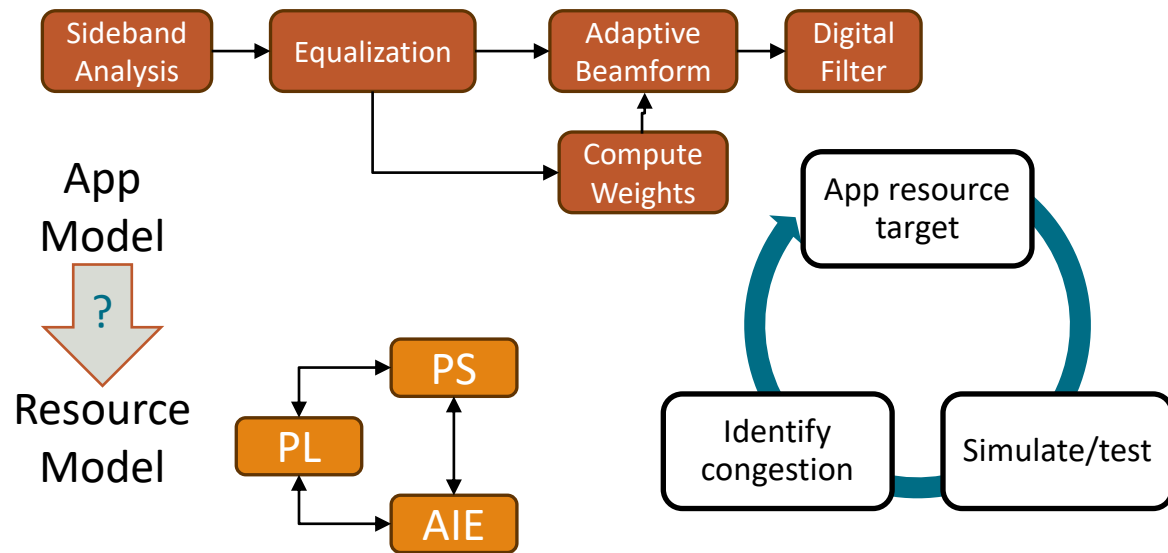
What is System Architecture Exploration?



System-Level Application Exploration Tool

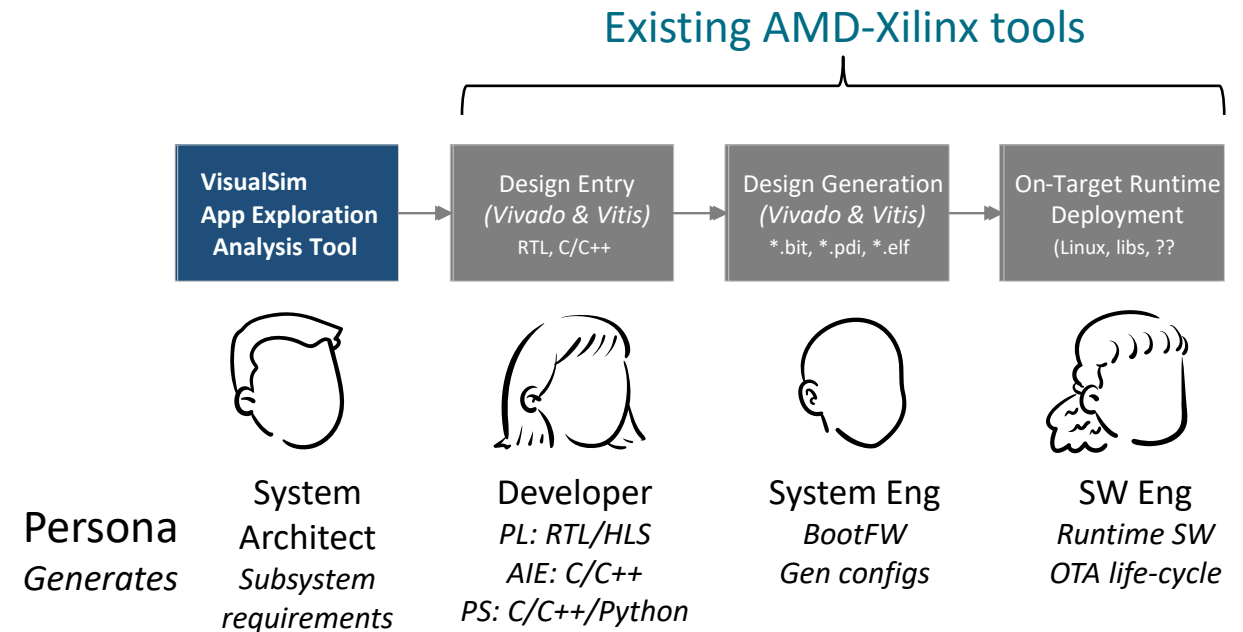
What problem does it solve?

- Stochastic models focused on **early application mapping** exploration to heterogeneous SoC compute resources.
- **Users are system architect** responsible for high-level complex system mapping, not design entry
- **Rapid trade-off** and **design iteration** prior to detailed algorithm design entry



How does it fit with existing tools?

- Prototype that extends AMD-Xilinx general toolset with **pre-design-entry** focused application analysis
- Lower fidelity **stochastic based model** vs. full design entry simulation tools
- Provides guidance to down-stream sub-system (AIE, PL, PS) design entry development teams



Introduction to VisualSim System-Level- IP Architecture Platform for AMD Versal SoC

VisualSim is system-level modelling and simulation SW

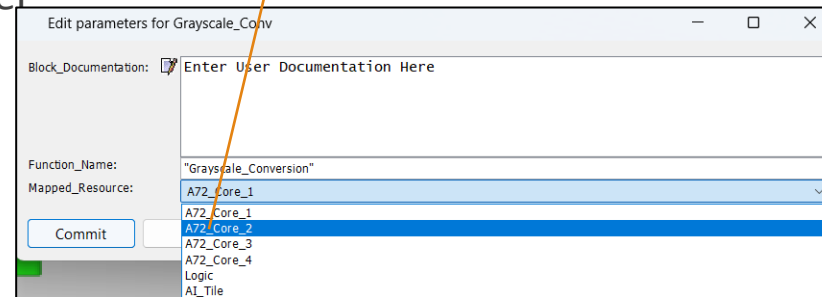
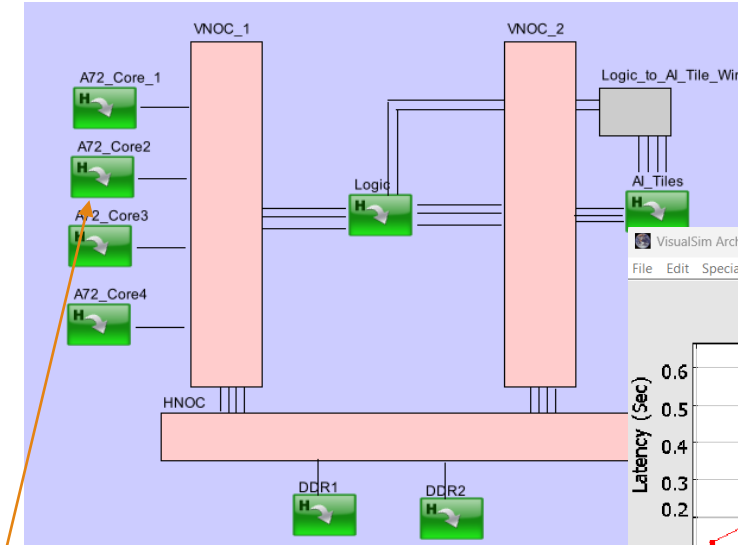
Platform for rapid trade-off

- Performance/power/area during planning
- Study speed, power, failure and bottlenecks



Optimize implementation, resource and timing constraints of algorithm tasks

VisualSim Versal SoC is a stochastic model containing

- Heterogeneous compute resources
- DDR and HBM memory interfaces
- Statistics for latency, throughput, utilization and power
- User expandable resource usage table
- External Interfaces
- Task block with mapping function
- Traffic generator for workloads




App_Name	Num_DSP	Num_LUT	Num_FF	BRAM_Size_Bytes	Num_AI_Tiles	PS_Delay_Sec	Logic_Delay_Sec	AI_Tile_Delay_Dec
"FIR"	2	200	250	1024	2	1.0e-3	100.0e-6	500.0e-6
"IIR"	4	400	300	1524	3	1.0e-3	100.0e-6	500.0e-6
"Grayscale_Conversion"	10	3000	7000	15240	2	100.0e-6	10.0e-6	50.0e-6
"FFT"	50	30000	70000	552400	60	10.0e-3	1.0e-3	5.0e-3
"Edge_Image"	2	3000	5000	15240	2	10.0e-3	1.0e-3	5.0e-3
"iFFT"	50	32000	75000	552400	5	10.0e-3	1.0e-3	5.0e-3
"Edge_Image_Enhancement"	1	300	700	5240	2	10.0e-3	1.0e-3	5.0e-3
"Segmentation"	100	50000	80000	952400	50	100.0e-3	10.0e-3	50.0e-3




Motivation for Analyze and Validate with VisualSim

- Complex behavior
 - input stream
 - data dependent behavior
- Contention
 - limited resources
 - scheduling/arbitration
- Interference of multiple applications
 - limited resources
 - scheduling/arbitration
 - anomalies

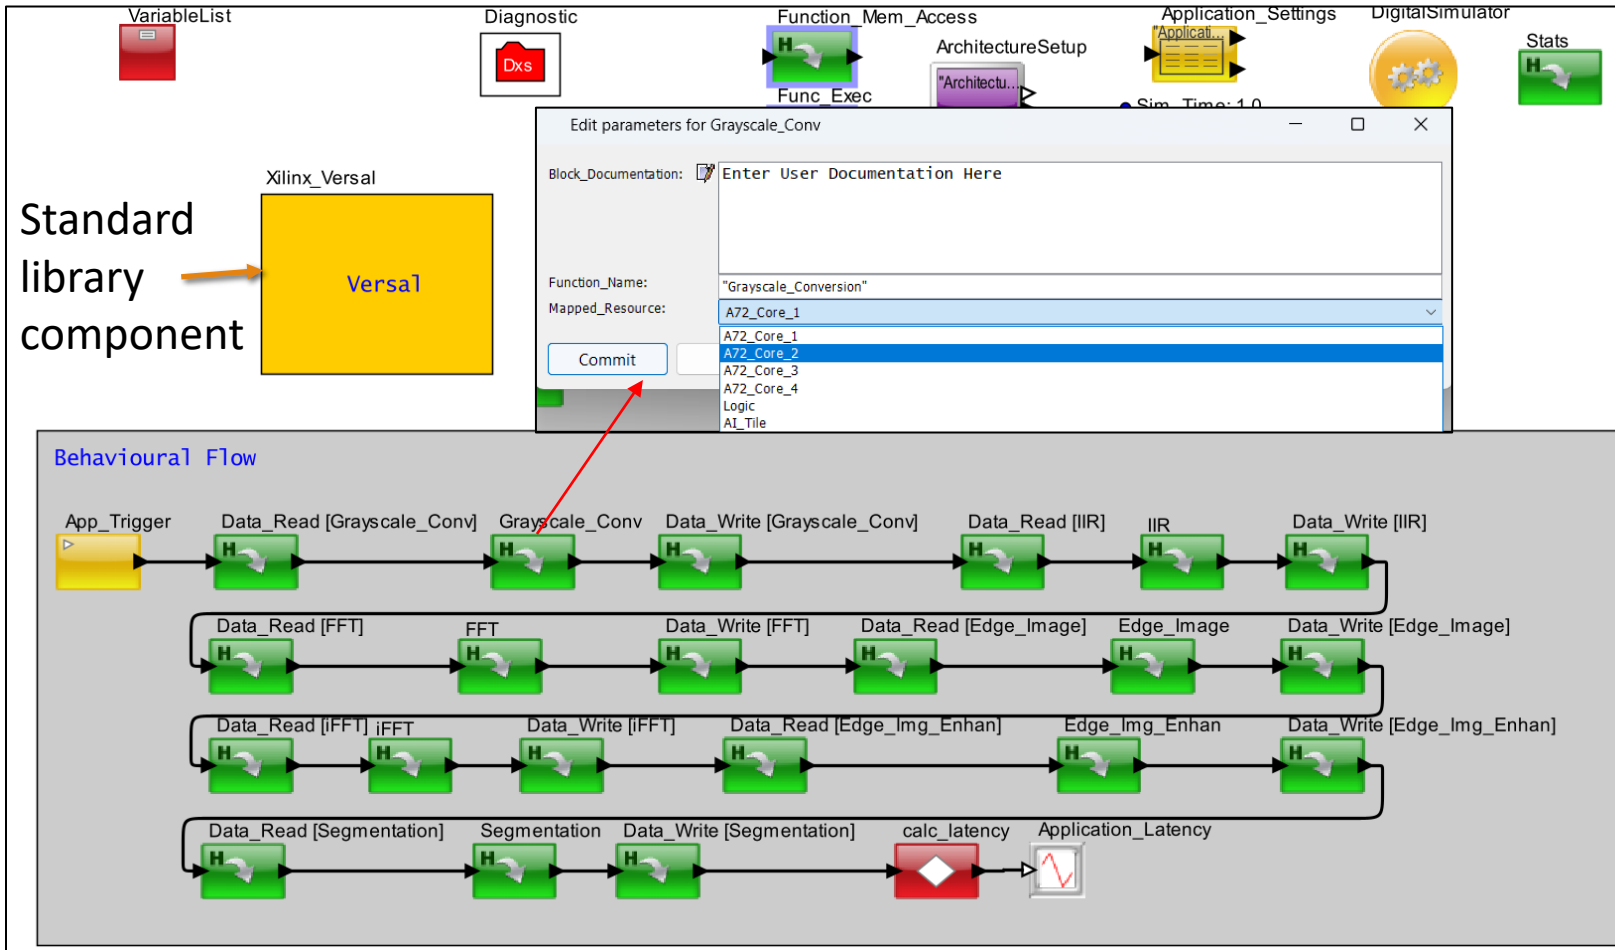


Performance: System Sizing and Range Prediction



Unpredictable System Response

Mapping Algorithm to Versal SoC-FPGA



Implementing an image processing algorithm on AMD-Xilinx Versal FPGA.

Each task is mapped to a resources.

Basic/Starting Configuration

Grayscale_Conversion - PS [A72 Core 1]

IIR – Logic (PL)

FFT – AI Engine Tile


Edge_Image - Logic (PL)

iFFT – AI Engine Tile

Edge_Image_Enhancement – Logic (PL)

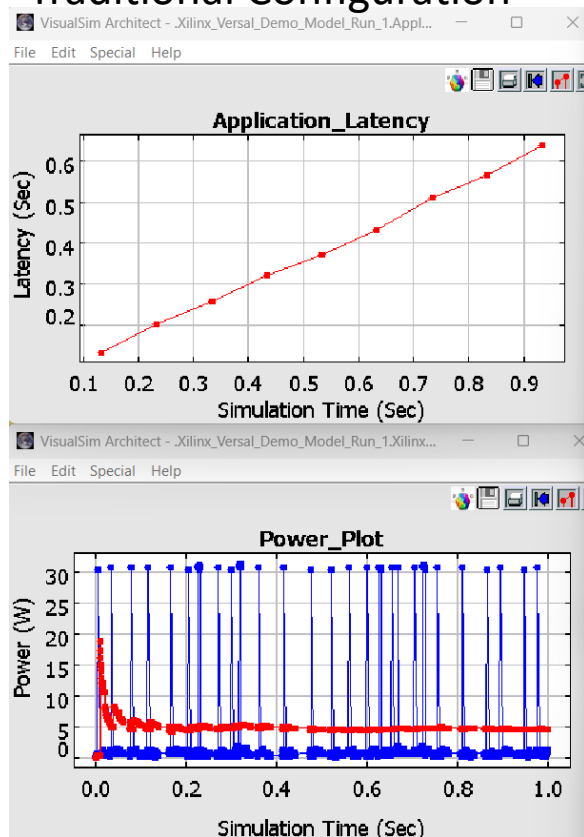
Segmentation – PS [A72 Core 2]

Image processing algorithm



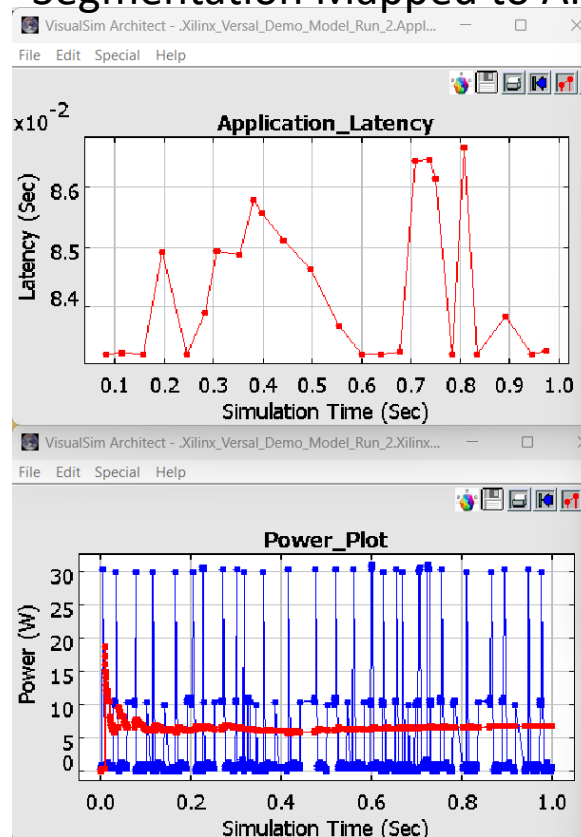
Experiments with Different Implementations

Run 1 – Traditional Configuration



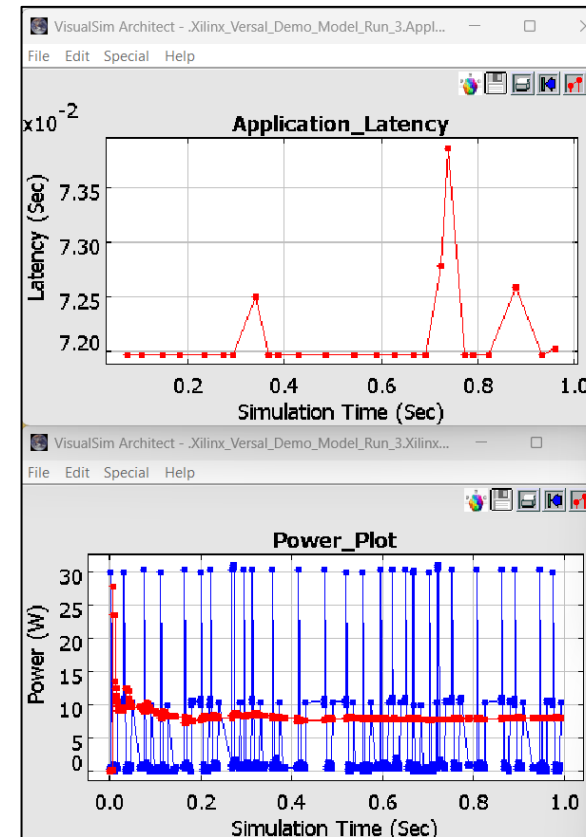
Application latency increasing over time.
Latency increases due to Segmentation.
Remap segmentation task AI Tiles

Run 2 – Using App Explorer Segmentation Mapped to AI Engine



Application latency in bounded range.
NoC Utilization is high.
Changed interconnect for Segmentation
from NoC to Direct

Run 3 – Using App Explorer Using Direct Path between PS and AI



Latency is deterministic
Latency requirement (App latency
< 80 msec) is met.
Utilization across NoC is acceptable

AIE + PL Modeling Prototype

Current PoC methodology:

- PoC is modelled using the pre-characterized libraries of basic functions for an application domain. Function represented as a grid of pre-characterized basic functions.
- Versal AIE and PL mapping and performance estimation is performed by an integer linear programming (ILP) solver approach.
- Application requirements and device capabilities are represented as a set of variables and constraints in ILP.
- Identify the number of AIE tiles/DSPs/LUTs/FFs and orientation, streaming channels, memory requirements for each function and availability of routing resources between multiple functions.
- Versal PL mapping is done by abstract placer and router that utilizes the pre-characterized libraries.
- Compared to a spreadsheet-based tool, the solver-based approach brings in the data-flow centric early application mapping and performance estimation for multiple compute resource target types.

Conclusion & Next Steps

Summary

App Exploration PoC tool has been demonstrated for a signal processing use-case using static models and on target function characterization.

User was able to iterate mapping design choices and establish an application mapping trade-space analysis to iterate to their design requirements.

Demonstrate concept of system architect being able to rapidly map their application to SoC for a faster compute subsystem requirements definition.

Next steps

Continued of AIE+PL solver-based functional mapping to increase tool flexibility and scalability.

Expand functional block libraries to support a broader set of application use cases including AI/ML inference workloads.

Development of user defined functions for both PL and AIE based targets.

Questions and feedback?

THANK YOU!

Schedule Notes - [Link](#)

Wednesday, June 26

PDT

[expand all](#) · [collapse all](#)

TIME	TYPE	SESSION / PRESENTATION	CONTRIBUTORS	LOCATION	TAG	PLAN
10:30am - 12:00pm PDT	Embedded Systems and Software	Edge Intelligence & GenAI: Exploring Challenges and Ethics ▼		2012, 2nd Floor	AI Embedded Systems Engineering Tracks	
10:30am - 12:00pm PDT	Back-End Design	Electron Signatures for Predicting the Diagnosis! ▼	Badhri Uppiliappan	2008, 2nd Floor	Back-End Design Design Engineering Tracks	
10:30am - 12:00pm PDT	Front-End Design	Industry Trends in the Front End ▼	Vikas Sachdeva	2010, 2nd Floor	AI Design Engineering Tracks Front-End Design	
1:30pm - 3:00pm PDT	IP	Advances in mixed signal IP design ▼		2012, 2nd Floor	Engineering Tracks IP	
1:30pm - 3:00pm PDT	Back-End Design	Design Automation Advancement in the Analog Domain		2008, 2nd Floor	Back-End Design Design Engineering Tracks	
1:30pm - 3:00pm PDT	Embedded Systems and Software	Embedded Systems and Software ▼	Natraj Ekambaram	2010, 2nd Floor	AI Embedded Systems Engineering Tracks	
Session Chair: Natraj Ekambaram ,						
Description: Embedded Software is a critical differentiator in today's designs, and its close relationship with hardware aspects specifically impacts hardware/software flows and development methodologies. This session will discuss aspects like accuracy enhancements for virtualization to enable early software development, system design aspects of power, firmware verification, reinforcement learning for test optimization, automotive AUTOSAR software aspects, and mapping challenges of software to computing resources.						
Presentations:						
1:30pm - 1:45pm PDT	No-Code Power and Clock System Design Presenters: Hoyeon Jeon , Ahchan Kim , Ingyu Kim , Jongbae Lee					
1:45pm - 2:00pm PDT	Functional Accuracy Enhancement of In-House Virtual Platform using Commercial IP Model Presenters: Jooho Wang , Dongyoung Lee , Myeongjin Kim , Sangwoo Han , Seungik Ha , Jinbeom Kim , Jaeyeong Jeon , Songyi Park , Jongseong Park , Kyungsu Kang , Jaewoo Im					
2:00pm - 2:15pm PDT	Complex application mapping to heterogeneous compute resources Presenters: Wesley Skeffington , Surya Chongala , Deepak Shankar					
2:15pm - 2:30pm PDT	Automated Generation of SSD Stress Tests Using Offline Reinforcement Learning Presenter: Sunghee Lee					
2:30pm - 2:45pm PDT	Development of SystemC-based Security VP for In-House SED SSD Firmware Verification and Application of libFuzzer Presenter: CHANGWON KIM					
2:45pm - 3:00pm PDT	Open Source AUTOSAR Classic Platform Presenter: Moisés Urbina Fuentes					



Wes Skeffington

AMD Fellow – Systems Architecture
AMD Inc.



Wes is an AMD Fellow within the Engineering Architecture team focused on the technical definition and roadmap of platform level products including AMD Kria and Embedded+. His application space focus includes industrial, healthcare, robotics, and automotive applications using heterogeneous Arm SoCs and Arm SoC + x86 multi-chip platforms. He has 20+ years of experience in developing embedded systems with multiple application requirements that include hard real-time functions, general-purpose OS, heterogeneous HW targets, virtualization, and deploying these in regulated product domains.



Deepak Shankar

Founder
Mirabilis Design Inc.

Deepak Shankar is the Founder and Chief Technologist at Mirabilis Design. Deepak has experience in sales and marketing of software to semiconductor, aerospace, defense and automotive companies across the world. He has published and spoken at over 50 Technology and Business conferences on topics ranging from autonomous driving to space travel and semiconductors to future of electronic design. He has won numerous awards for his work in systems engineering and product development optimization. Deepak Shankar has an MBA from University of California Berkeley, MS from Clemson University and BE from Coimbatore Institute of Technology.