



TECHNOLOGY CHALLENGES, CHIPLET OPPORTUNITIES

GABRIEL LOH

UCLA - CHIPS - November 2019

THE WORKLOADS OF THE FUTURE REQUIRE INCREDIBLE AMOUNTS OF COMPUTE POWER

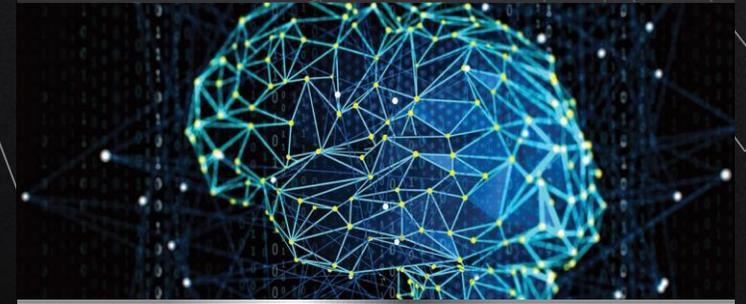
HIGH PERFORMANCE COMPUTING



CLOUD, HYPERSCALE & VIRTUALIZATION



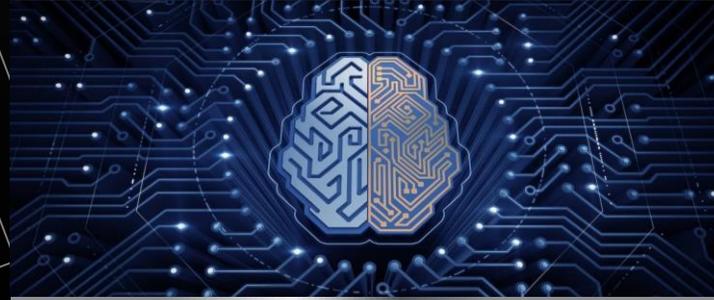
MACHINE INTELLIGENCE



BIG DATA ANALYTICS



IMMERSIVE & INSTINCTIVE COMPUTING



SOFTWARE-DEFINED STORAGE



GPU AND CPU PERFORMANCE TRENDS

GPU Single Precision Floating Point Operations Per Second Trend



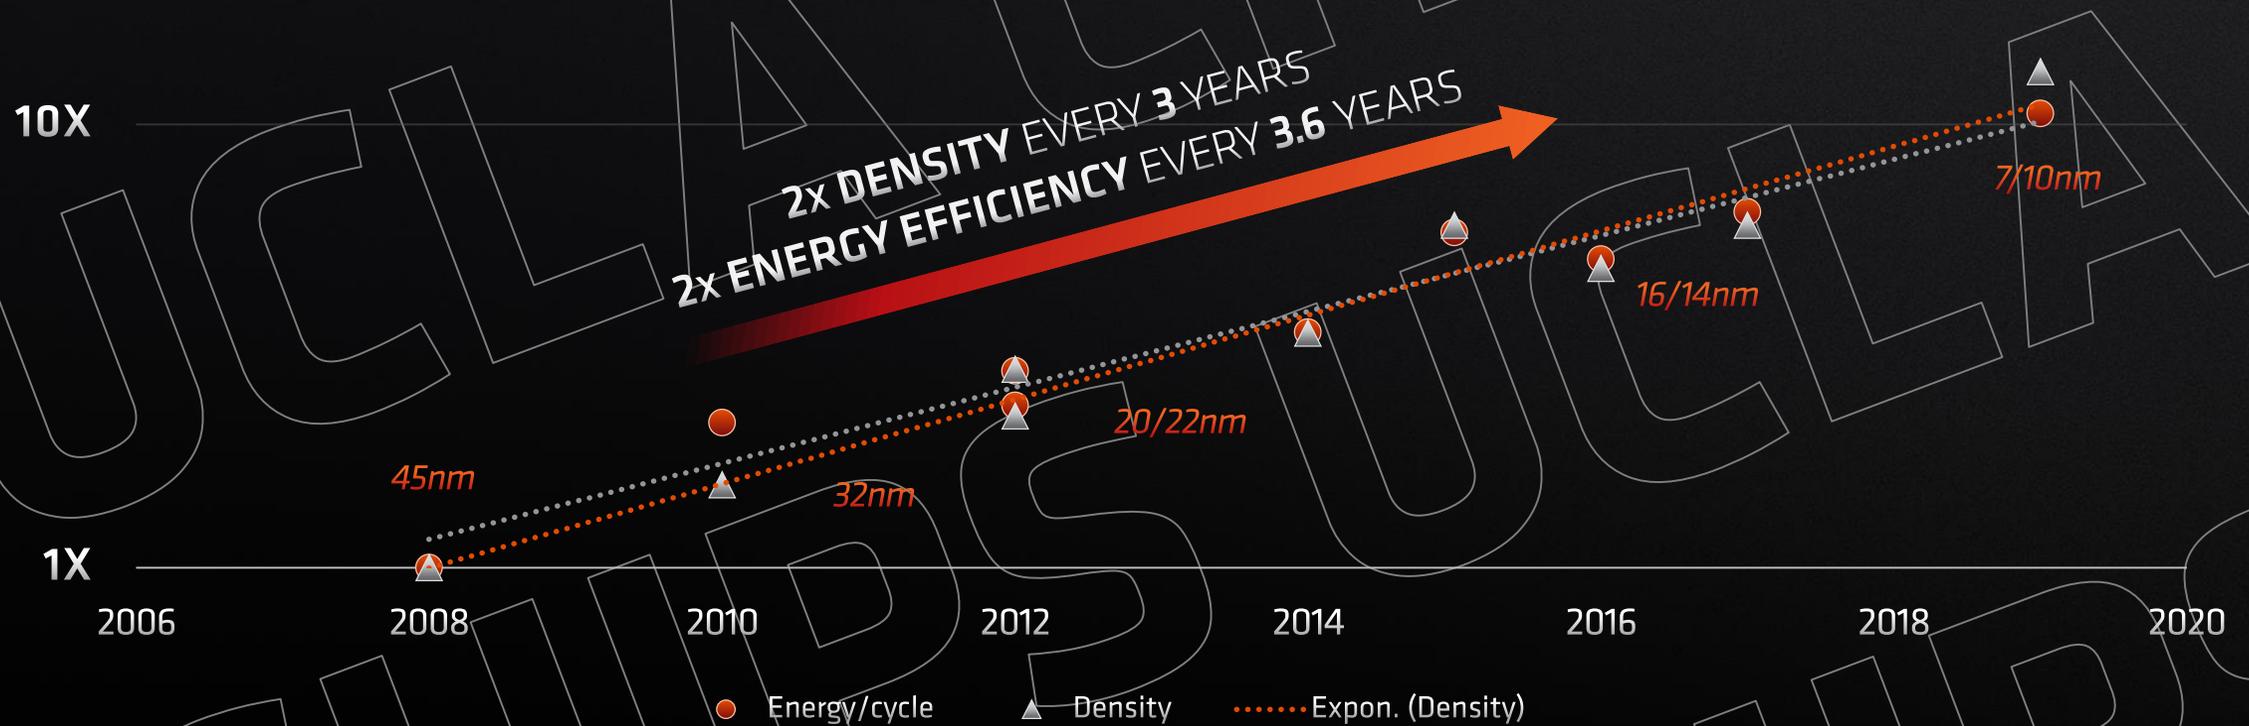
Specint[®]_rate2006 2P Server Performance Trend Over Time



CONSISTENT AND EXPONENTIAL CPU/GPU PERFORMANCE GAINS

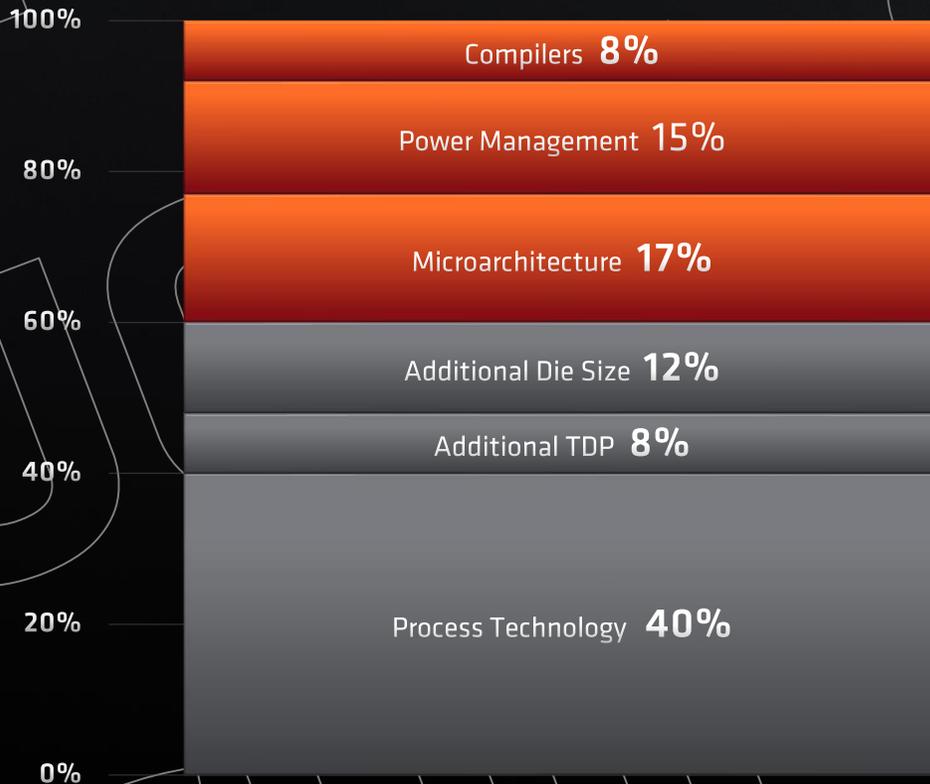
PROCESS TECHNOLOGY DELIVERED SIGNIFICANT GAINS

Technology Energy Efficiency and Density Across Process Nodes



SIGNIFICANT PART OF ENERGY EFFICIENCY DERIVES FROM MOORE'S LAW

PERFORMANCE GAINS OVER THE PAST DECADE



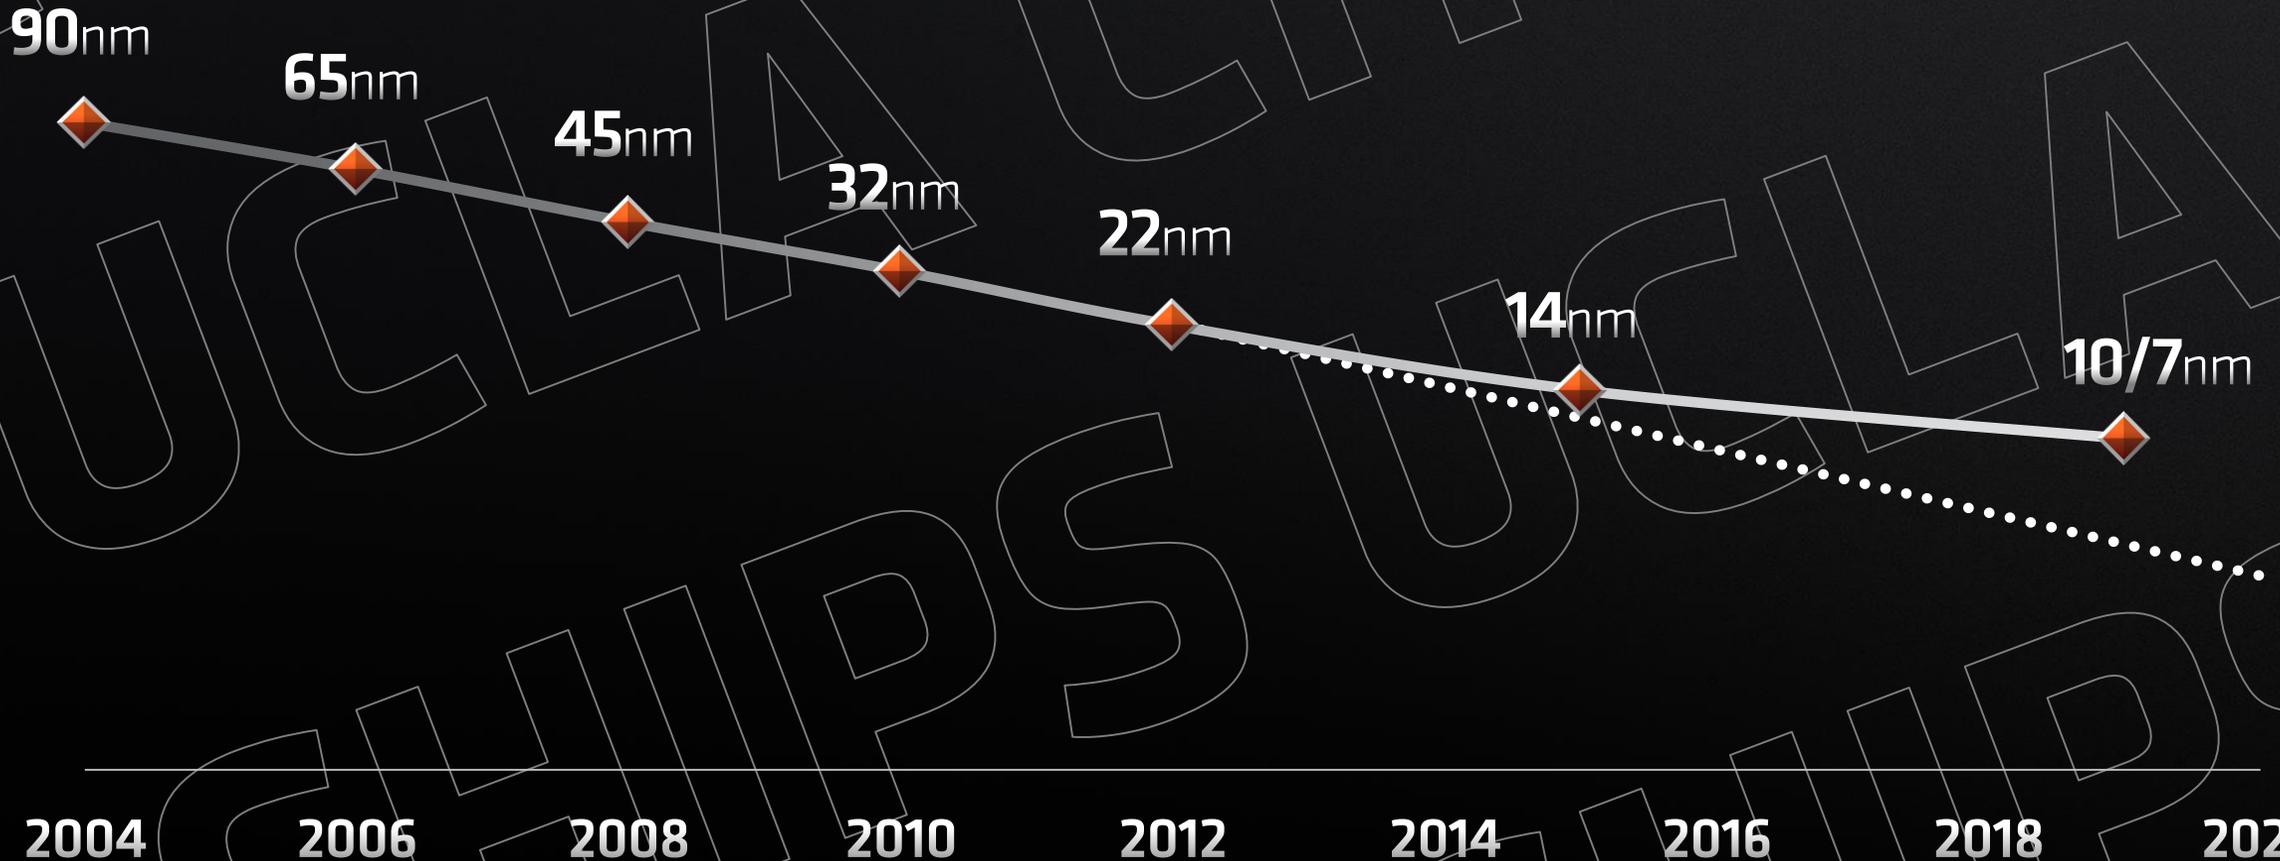
- Integration of System Components
- Micro-Architectural Efficiency
- Power Management
- Software
- More Silicon Power
- Bigger Die

HIGHER PERFORMANCE, DENSER, LOWER POWER TRANSISTORS

ELEMENTS OF 2x IN 2.5 YEAR PERFORMANCE GAIN OVER THE PAST DECADE

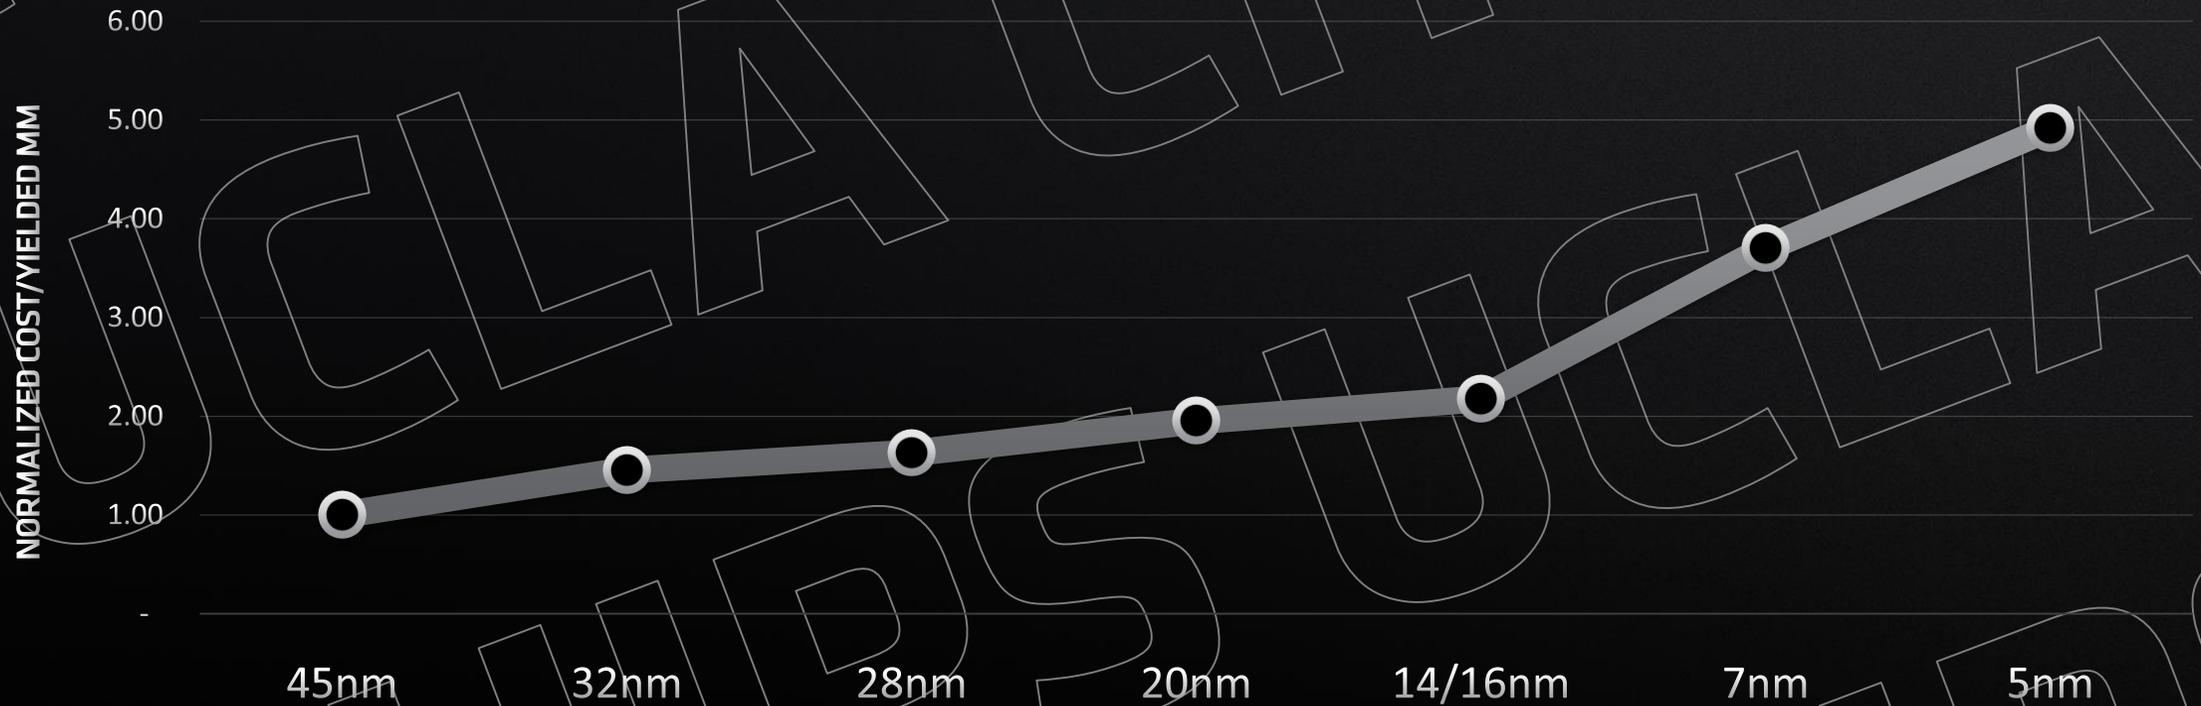
BARRIERS TO CONTINUED PERFORMANCE IMPROVEMENT

MOORE'S LAW KEEPS SLOWING



WHILE COSTS CONTINUE TO INCREASE

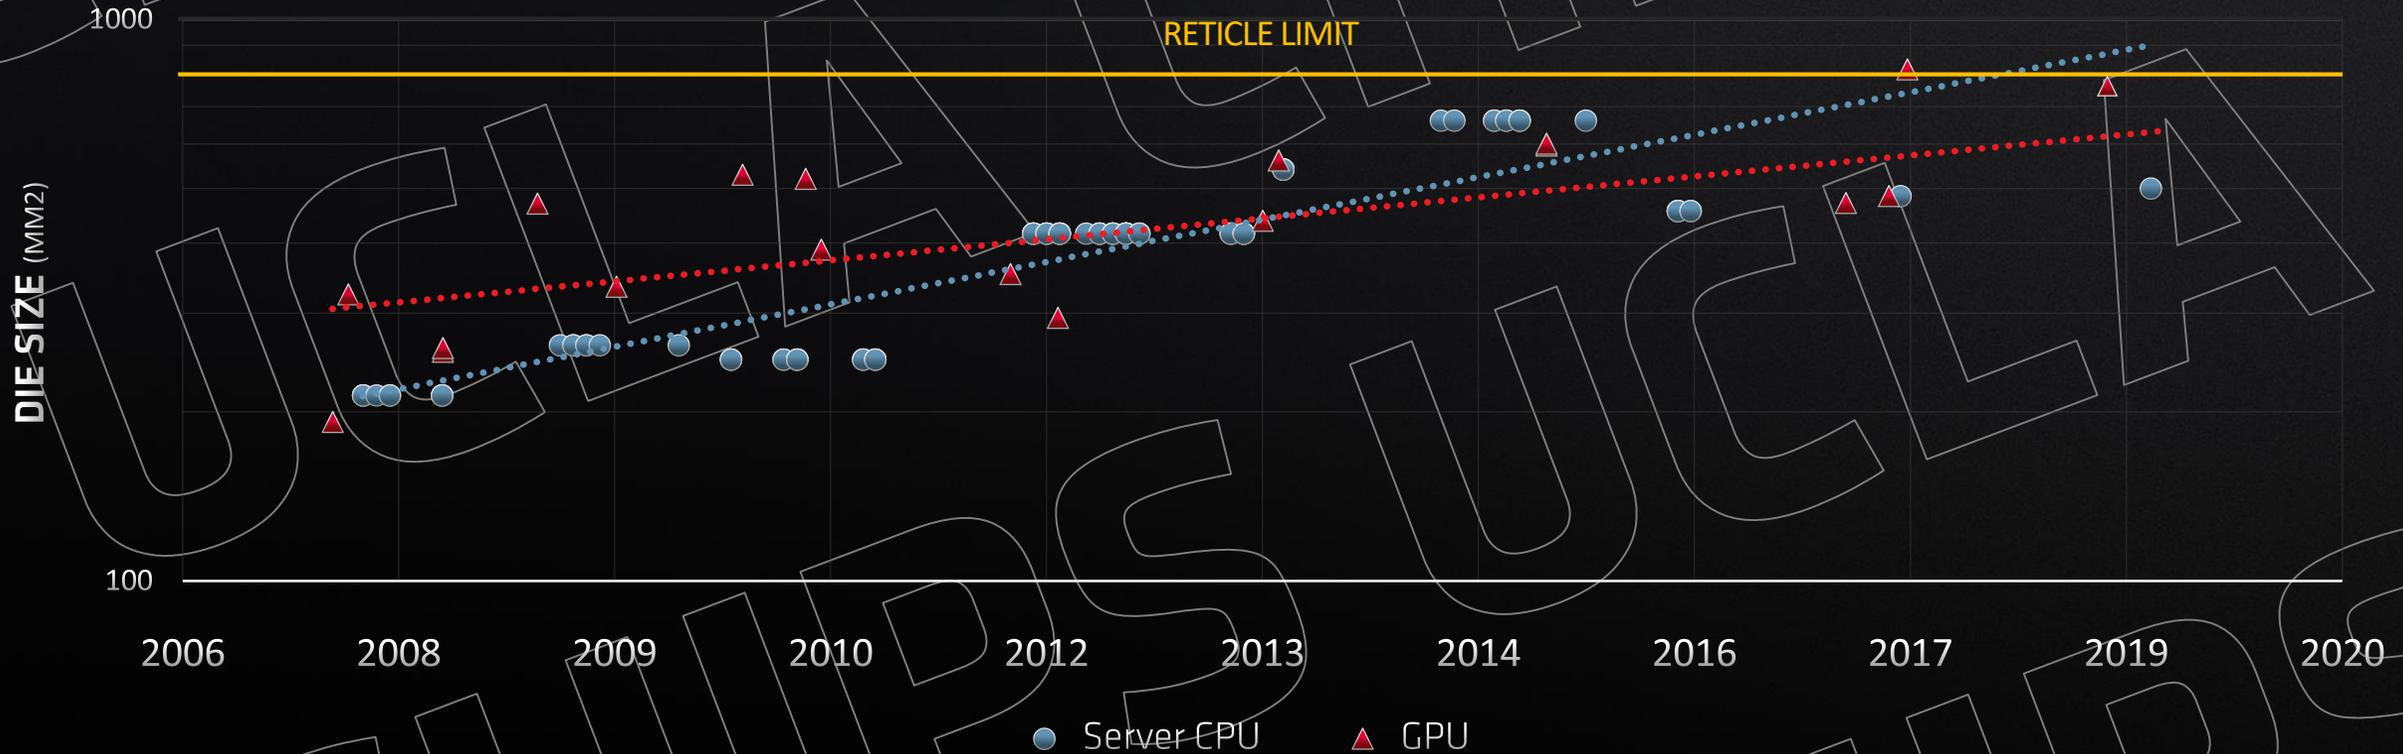
Cost Per Yielded mm² for a 250mm² Die



INCREASING DIE SIZES ARE ECONOMICALLY PROBLEMATIC

DIE SIZE TREND

Die Size Increases Over Time in Server CPUs and GPUs

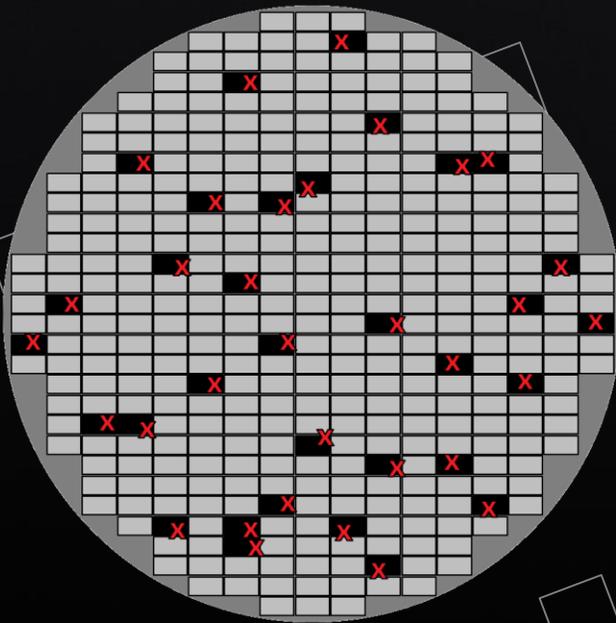


DIE SIZES INCREASING AT AN UNSUSTAINABLE RATE

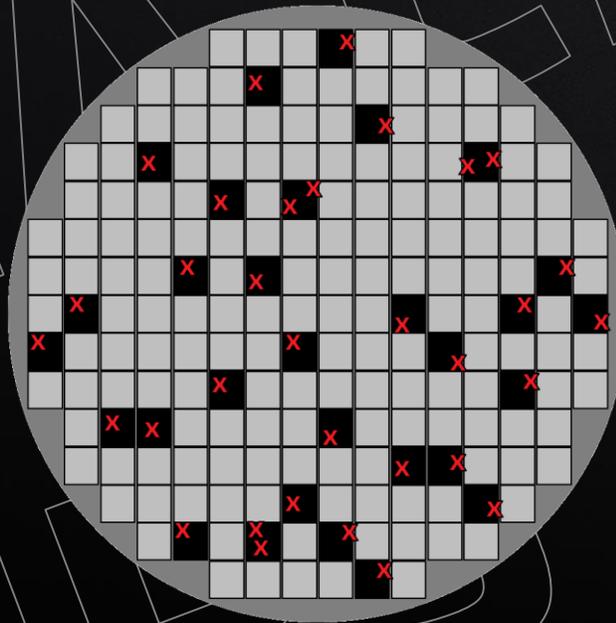
CHIPLLET ARCHITECTURES TO EXTEND PERFORMANCE GAINS

BIGGER CHIPS TO OFFSET TECHNOLOGY SLOWDOWN

- If technology scaling only gives you (say) 1.5x more devices per 24 months, why not just make chips 1.33x bigger to get 2x transistors?

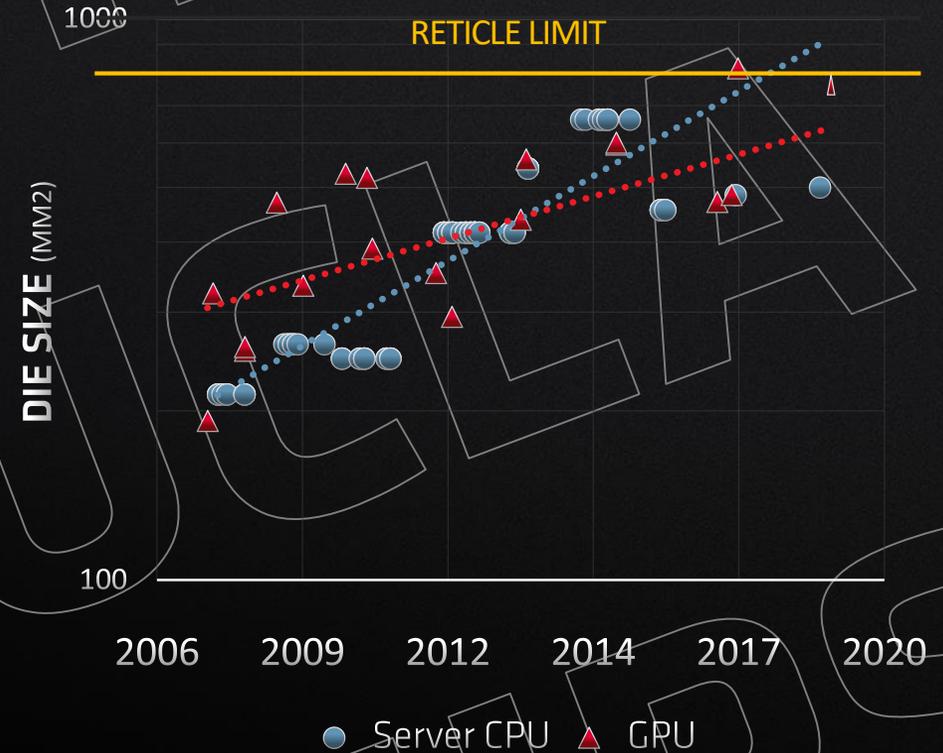


395 chips → 362 good die
(8% yield loss)



192 chips → 162 good die
(16% yield loss)

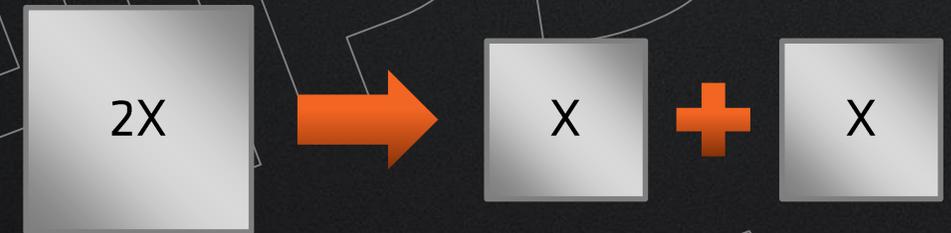
(hypothetical/academic example [1], not real yield rates)



[1] Kannan, Enright Jerger, Loh, "Enabling Interposer-based Disintegration of Multi-core Processors", International Symposium on Microarchitecture (MICRO), 2015

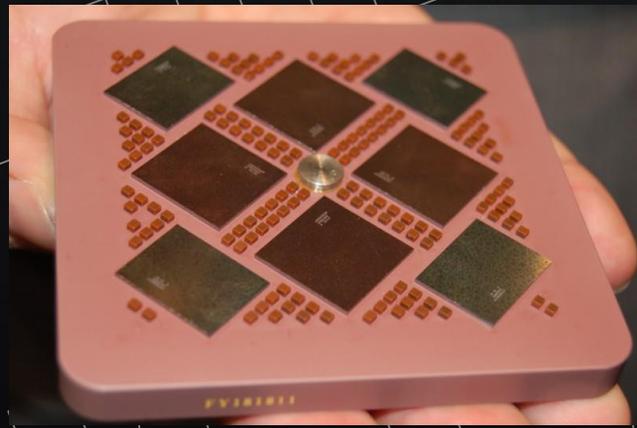
CHIPLETS BACKGROUND

- Alternative: build multiple smaller chips
- Historically not needed for most markets
 - Except for the largest systems, Moore's Law was sufficient to meet compute needs
- Chiplets not free
 - Additional area for interfaces, replicated logic
 - Higher packaging costs
 - Additional design effort, complexity
 - Past methodologies less suited for chiplets

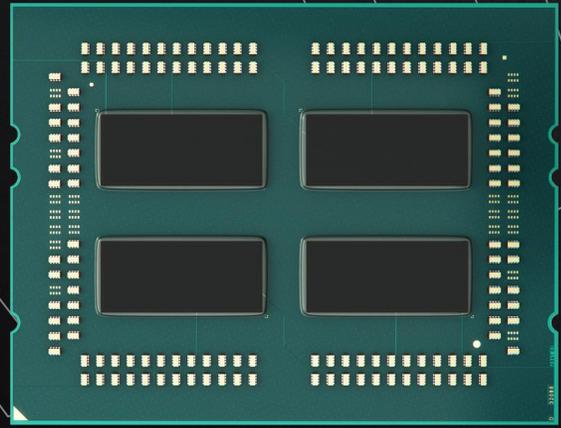


2X device functionality
costs > 2X silicon area

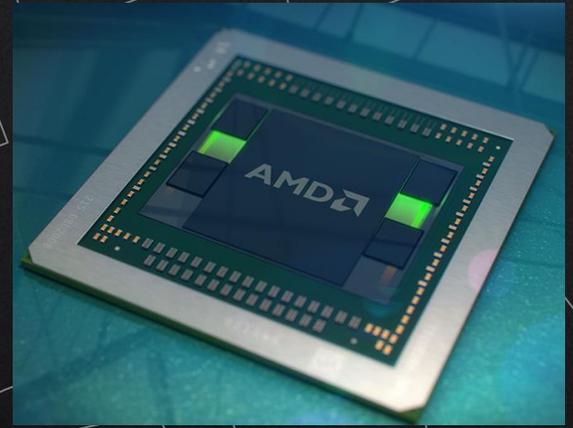
PROGRESSION OF ADVANCED INTEGRATION/PACKAGING



Ceramic Substrate MCMs [1]



Organic Substrate MCMs



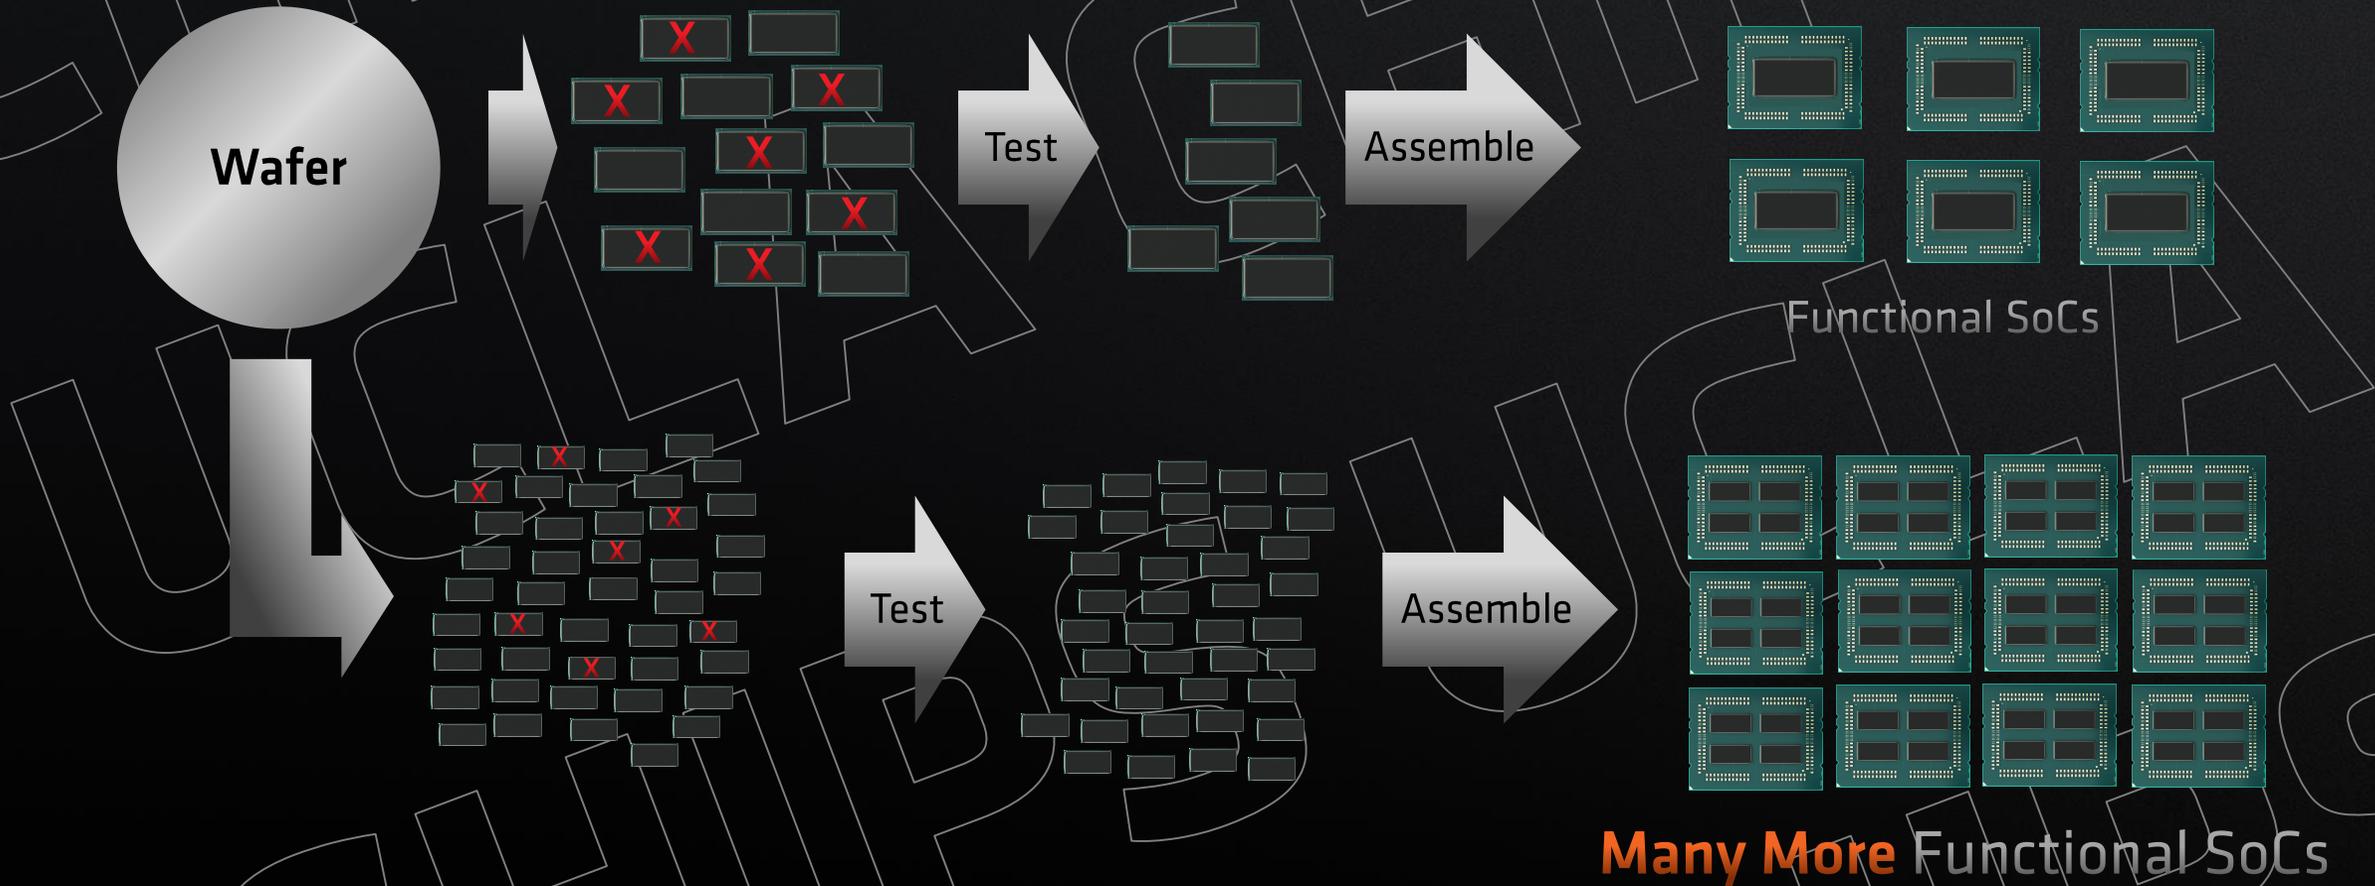
Silicon Interposer

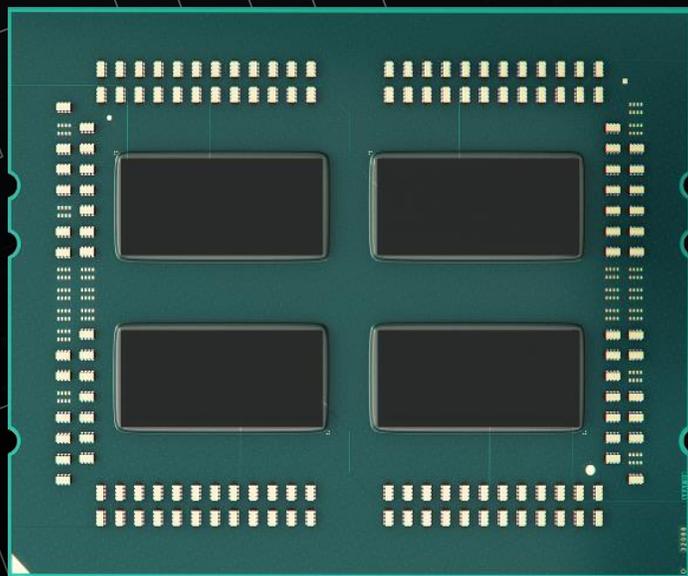
- The concept of partitioning systems into multiple chips is not new
- Evolution of packaging technology has changed the trade-offs in terms of cost, bandwidth, latency, energy, etc.

End of Moore's Law
 + Advanced Packaging Advancements
 → New era for multi-chip/chiplet approaches



HIGH-LEVEL APPROACH TO CHIPLETS





AMD EPYC 2017

MULTICHIP MODULE

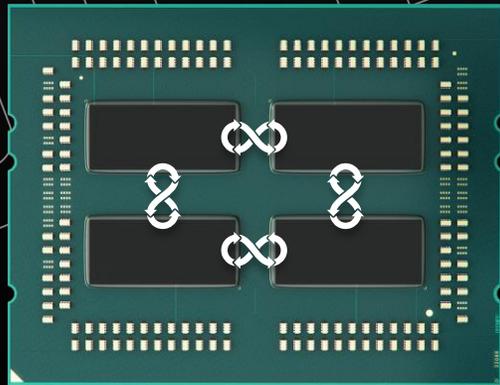
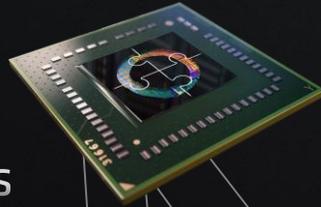
- Die Size > Reticle Limit
- Increased Peak Compute
- Highly Configurable
- Manufacturing Cost Reduced 41% vs. Monolithic

1ST GENERATION
CHIPLLET DESIGN
UP TO 32 CORES PER PACKAGE

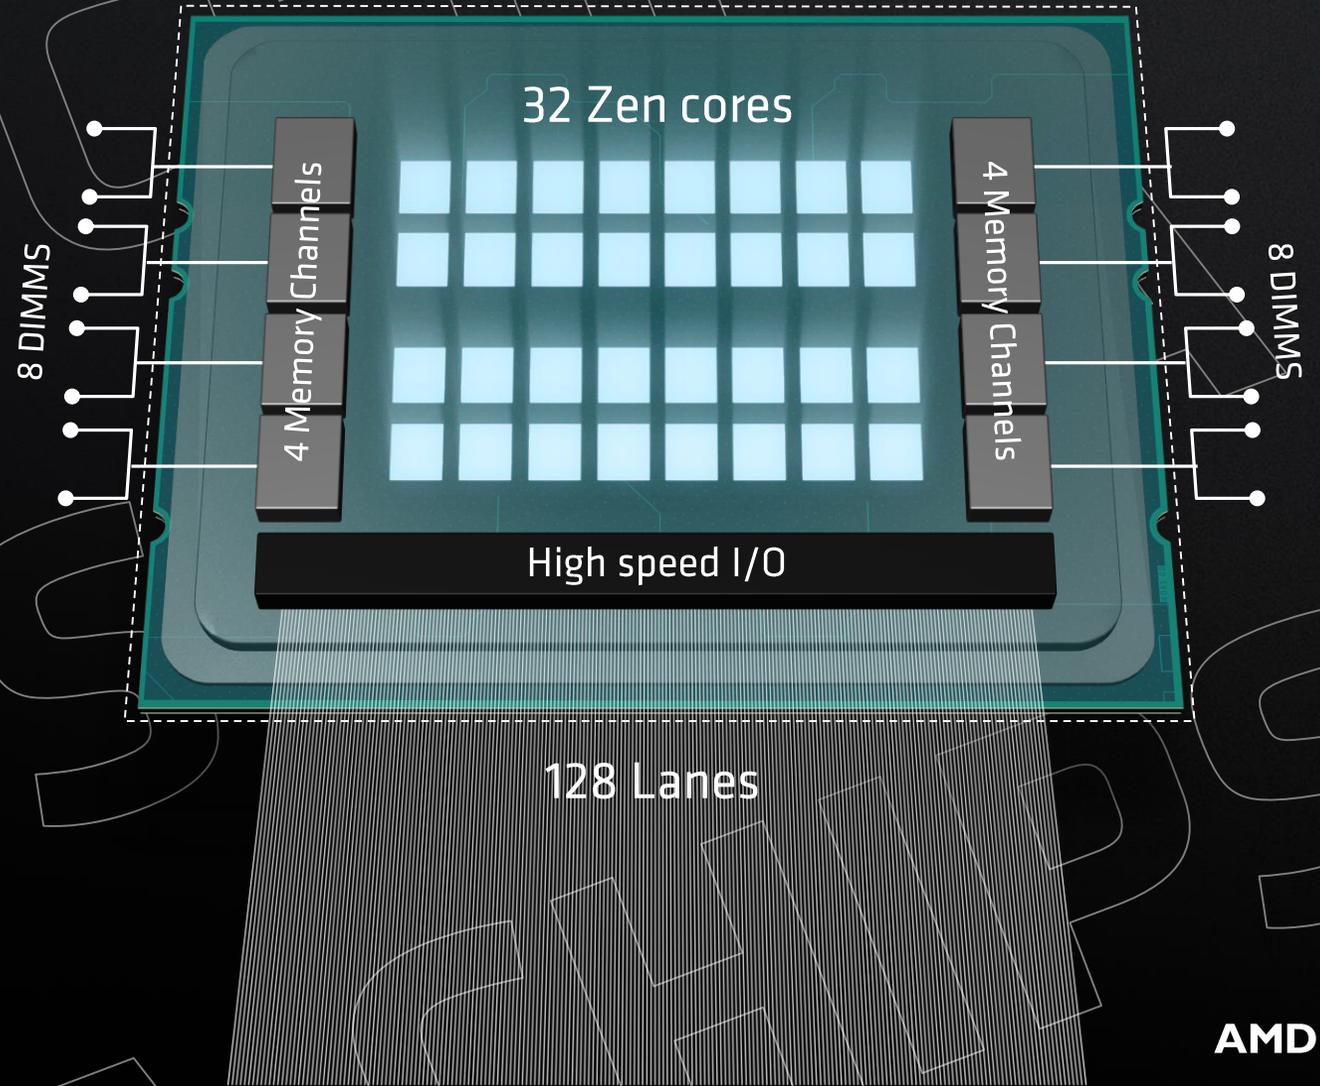
1ST GENERATION EPYC DETAILS

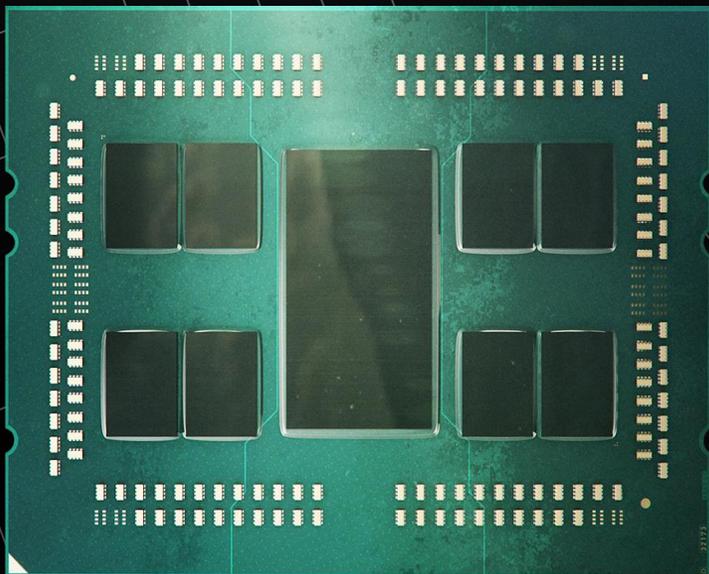
Per Chip

- Eight Zen cores
- 2 DDR4 memory channels
- Infinity Fabric interconnect to other chips



Per Socket





AMD EPYC 2019

MULTICHIP MODULE

- Significantly More Compute in Same Package
- I/O and CCD in Optimal Process Technologies
- Even Greater Configurability
- Optimized Memory Access

2ND GENERATION CHIPLLET DESIGN

UP TO 64 CORES PER PACKAGE

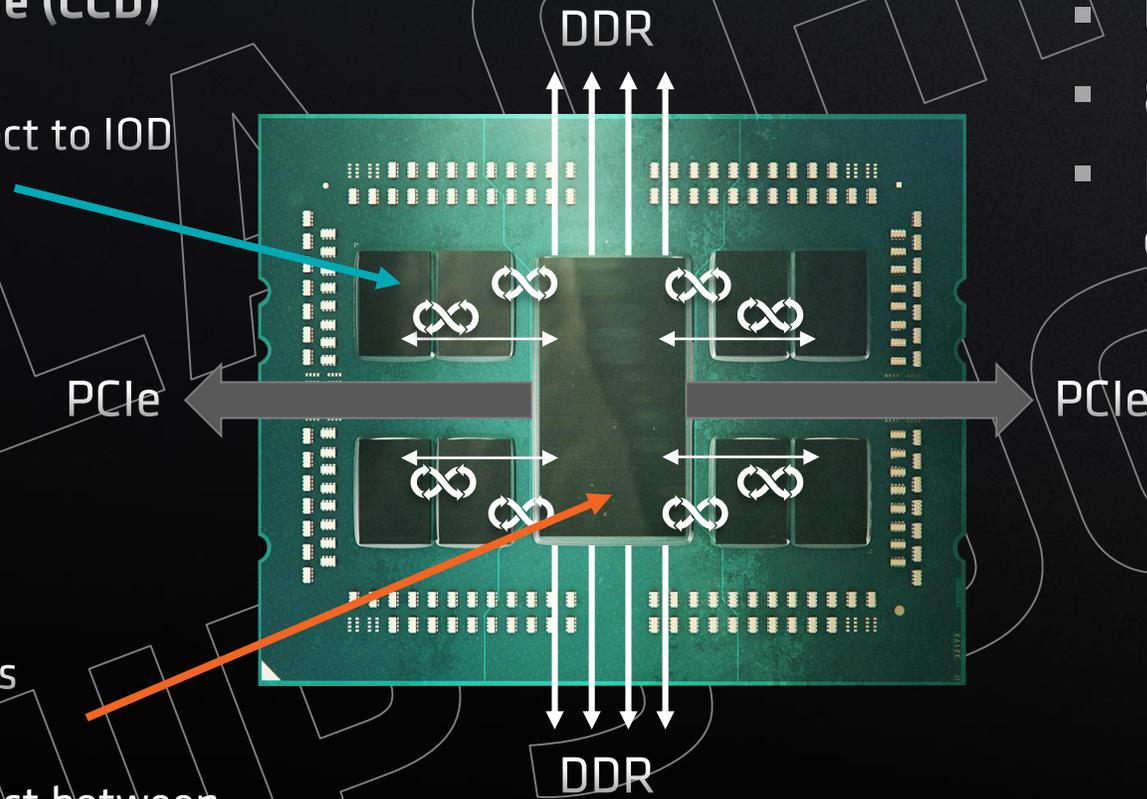
2ND GENERATION EPYC DETAILS

- **Per 7nm Core-Cache Die (CCD)**

- Eight Zen cores
- Infinity Fabric interconnect to IOD

- **12nm I/O Die (IOD)**

- 8 DDR4 memory channels
- 128 lanes PCIe
- Infinity Fabric interconnect between cores, memory, IO



Per Socket

- Up to 64 cores
- 8 DDR4 memory channels
- Infinity Fabric interconnect to other sockets

REVOLUTIONARY CHIPLLET DESIGN



Each IP in its
Optimal Technology

Infinity Fabric™ Enables
Modularity (MCM),
Scaling (CCD Count)

Optimized I/O Die enables Common
Latency to All Cores/Caches

CHIPLLET RESEARCH AREAS

DESIGN FOR REUSE AND FLEXIBILITY

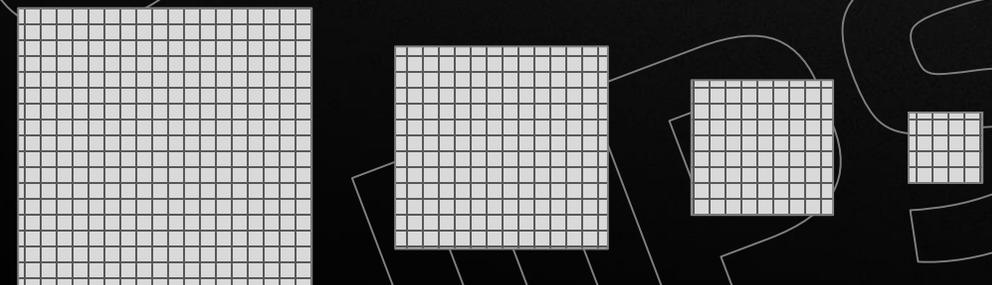


■ Generality vs. Optimization

- Interface widths and speeds, supported functionality/protocol(s), fixed pinouts
- Memory options (e.g., not all systems need HBM)
- Form factor/chiplet size (hard to support too many sizes, unnecessarily large silicon increases costs)
- Power delivery: pinouts, supported voltage(s), voltage regulation, current draw, required decap
- Thermal budgeting/allowance, cooling solutions

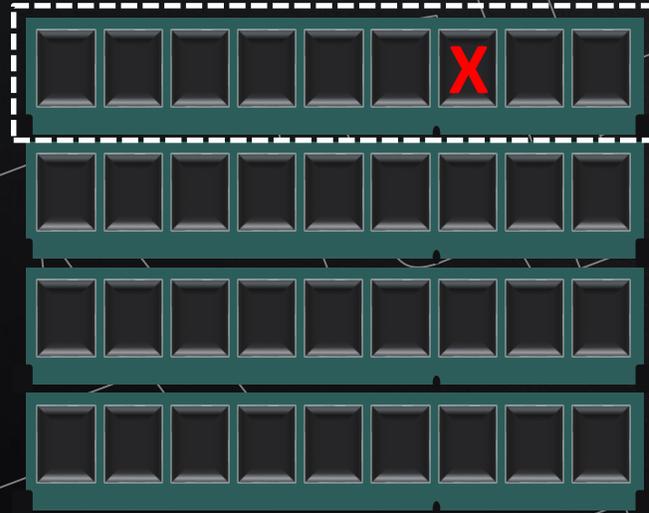
■ Design for reuse in large-scale design

- How to leverage wafer-scale systems across smaller scales?

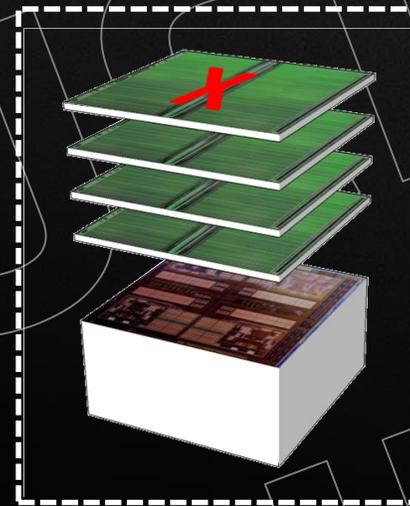


DESIGN FOR "RA" (=RAS~~X~~)

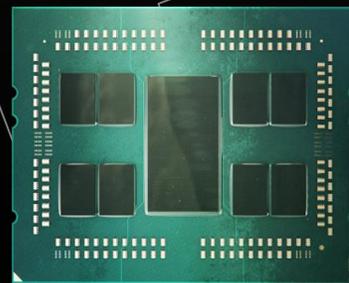
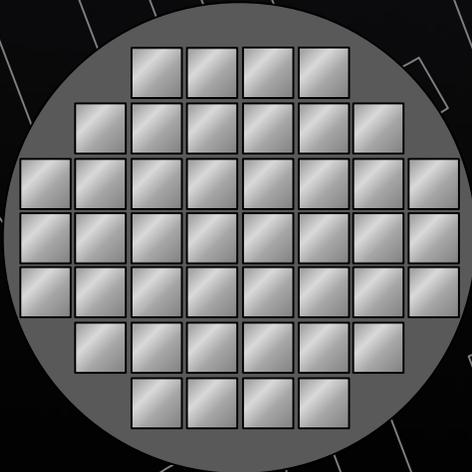
- Highly-integrated solutions may present serviceability challenges



Stacked DRAM + Processor:

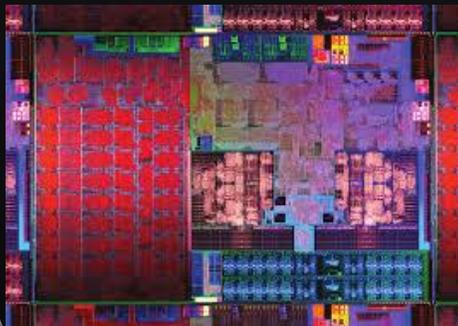


- Potential replacement cost increases with scale of integration



ARCHITECTURAL PARTITIONING

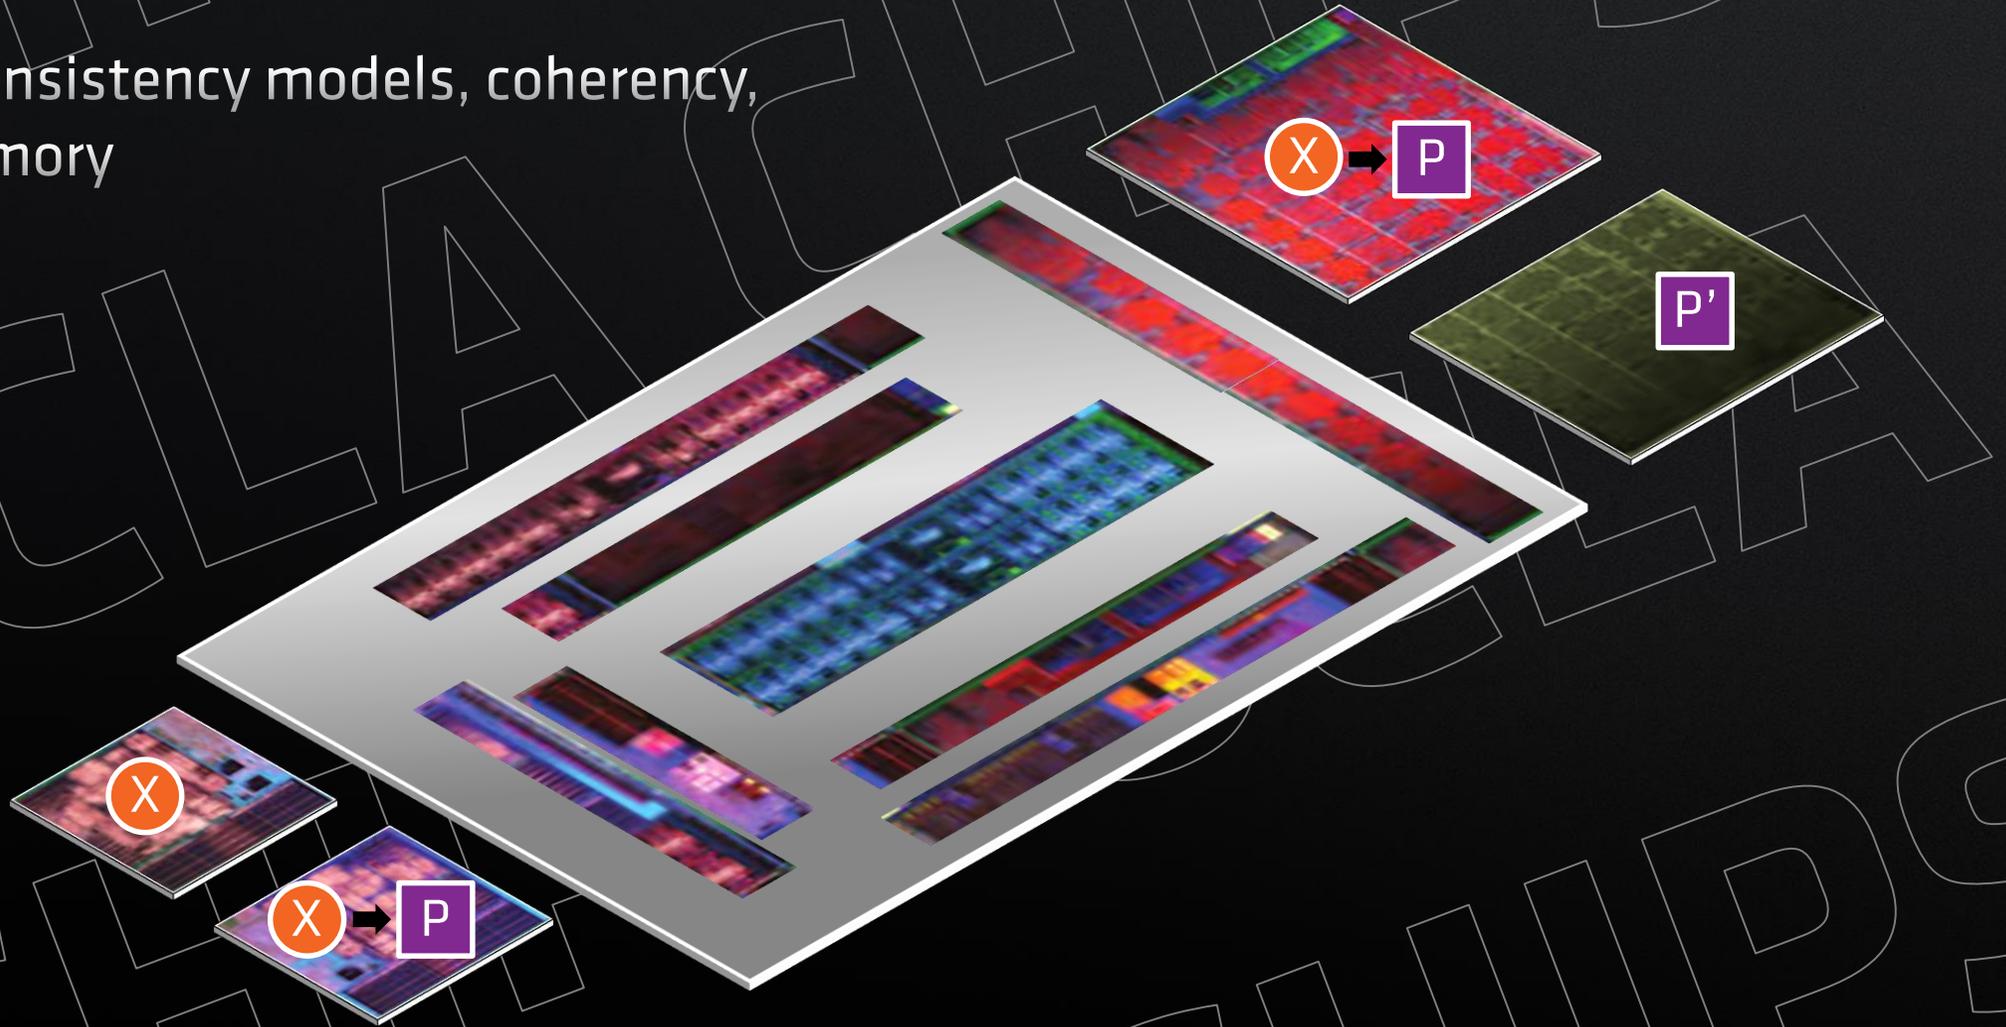
- Decomposition for a system



- Multi-objective technology selection and optimization
- Decomposition for N arbitrary, undesigned, unimagined systems
 - Interfaces:
 - Datapath interfaces: physical, protocol
 - Control interfaces: power management, debug, profiling, security, system alerts (e.g., thermal emergency)

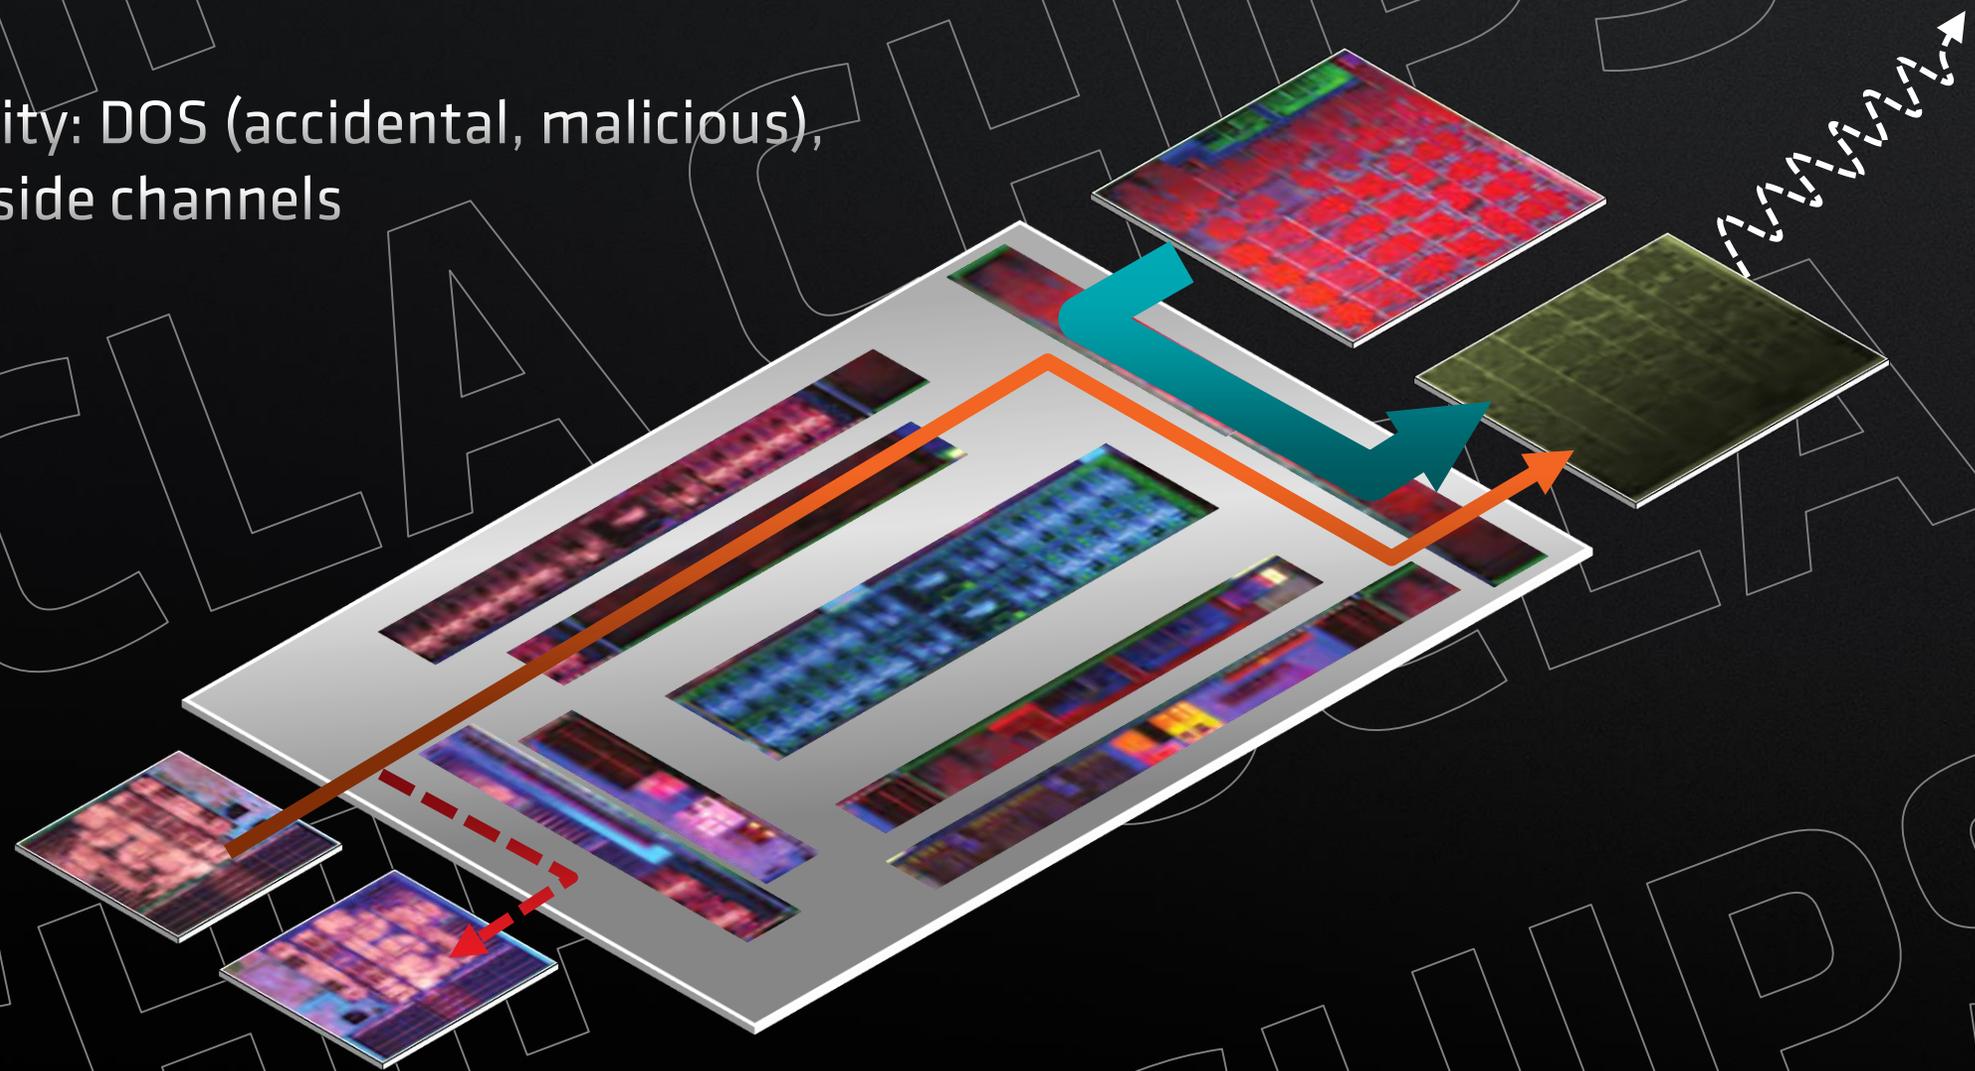
MEMORY CONSIDERATIONS IN PARTITIONED ARCHITECTURES

- Memory consistency models, coherency, virtual memory



SECURITY, UNTRUSTED CHIPLIETS, ETC.

- QoS, Security: DOS (accidental, malicious), snooping, side channels



CHIPLLETS ARE CHANGING HOW WE DESIGN

Fighting Moore's Law

Silicon Reuse

Product Flexibility

Process Optimization

Architecture Innovation

DISCLAIMER & ATTRIBUTION

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED "AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Vega, Polaris, and Zen are codenames for AMD architectures and are not product names.

ATTRIBUTION

© 2019 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, Ryzen, Radeon, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions.

Microsoft and DirectX are registered trademarks of Microsoft Corporation in the US and other jurisdictions. Xbox is a registered trademark of Microsoft Corporation.

"PlayStation", the "PS" Family logo, and "PSP" are registered trademarks of Sony Interactive Entertainment Inc.

Other names are for informational purposes only and may be trademarks of their respective owners.

AMD 