EURO SERVER

EUROPEAN COMMISSION

SEVENTH FRAMEWORK PROGRAMME

Riding the perfect storm

John Goodacre

Director Technology and Systems
ARM Ltd Cambridge

ARM Low-Power Leadership

MPSoC'14

Abstract

EURO SERVER

- EUROSERVER is a European commission FP7 part-funded project which is combining the technology trends of nanotechnology 3D integration, low-power SoC processor integration and the impossible requirements from next generation compute to investigate and build a solution for scalable, cost effective and flexible ARM-based micro-server system architecture suitable across multiple markets. This talk will introduce the vision and the goals for the project and the approach the consortium is taking to realize a ground breaking solution out of this perfect storm.

ARM

BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación

cea

CHALMERS

EUROTECH Imagine. Build. Succeed.

FORTH Foundation for Research & Technology - Hellas

onapp™

ST life.augmented

TECHNISCHE UNIVERSITÄT DRESDEN

The Perfect Storm

EURO SERVER

- The ability to inexpensively extract heat from chips of any size maxed out.
- The ability to lower voltages with decreasing feature size slowed dramatically.
- The design complexity of single core microprocessors hits point of diminishing returns where more transistors could add little to the core's per cycle performance .
- We are approaching a limit on economically viable off-chip interconnect with the available technologies because of electrical issues.
- The cost of increasing wire-based signalling rates begins to grow considerably, especially in power and complexity of the interface circuits.
- The economics of memory chip production stops the growth in size of memory die.
- The electrical and power issues associated with driving off of a memory chip at high rates through inexpensive commodity packaging (such as found on commodity DIMMs) to a microprocessor chip more than a few inches away reached a point where further improvements become fairly difficult and/or expensive.

SANDIA REPORT

SAND2013-9229
Unlimited Release
Printed October 2013

EUROSERVER Approach: Next Generation Compute

EURO SERVER

- Market
  - Competition is good – "one size doesn't fit all" – Make it flexible
  - Critical mass is good – " a problem shared" – Keep it compatible
  - Innovation is expensive and time consuming….
    - Exploiting it even more so – Collaborate appropriately
- Technology
  - Towards the end of the road…
    - Utilize and extend the advancements in nanotechnology integration
  - Increasingly high costs, NRE, Recurring
    - Make the best economically accessible SoC, 3DIC, FDSOI/FINFET
  - Use "Embedded" processors now feature capable for new markets
- The Common Requirement
  - More for the same, more for less, specialized, specific, accepted by the critical mass software community

Headline Objectives

EUROSERVER

- EUROSERVER will design and prototype technology, architecture, and systems software for the next generation of "Micro-Servers" to be used in building data centers.

- We will progress on the current state of the art in the following domains:
  - **Reduced Energy consumption** by: (i) using ARM (64-bit) cores, which are the world-leading low-power processors, (ii) drastically reducing the core-to-memory distance by using novel silicon interposer packaging technology, and (iii) improving on the "energy proportionality“

  - **Reduced Cost** to build and operate each microserver, owing to (i) improved manufacturing yield when multiple smaller "chiplets" are placed on a larger silicon interposer(2.5D) , and (ii) reduced physical volume of the packaged interposer module and (iii) and energy efficient semiconductor process (FDSOI)

  - **Better Software efficiency**: Next Generation system SW that (i) manages the resources in a server that consists of multiple coherence-islands, and (ii) isolates and protects the multiple workloads from each other when they use the shared server resources of I/O, storage, memory, and interconnects.

EUROSERVER Vision

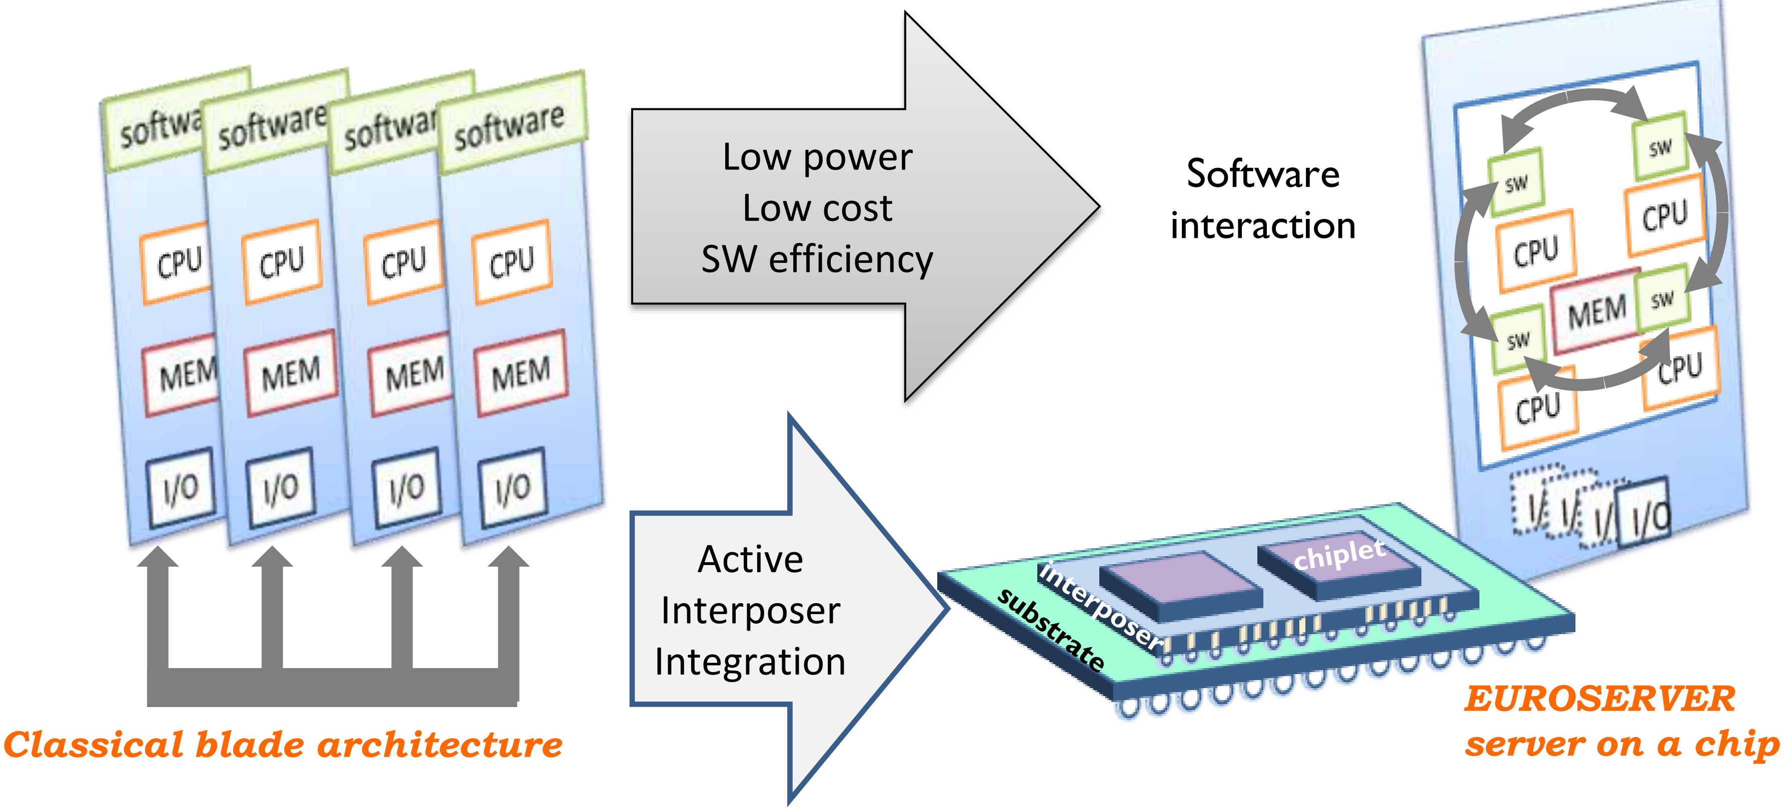

- Euroserver redesigns the entreprise server
  - Lower cost through Improved Chiplet Yielding and Flexible System Integration
  - Energy efficiency : low-power 64 bit processor, "Everything local",
  - Software Mutualization and sharing of I/O and common resources

EUROSERVER Applications

EURO SERVER

Application Area

- Data Centres & Cloud
- Telecom infrastructures
- High-end Embedded

Typical workloads

- Web-server hosting (LAMP/WAMP)
- Distributed databases (HADOOP)
- OLAP, OLTP workloads Traditional, relational databases (MySQL)
- Network communications
- Vehicle on-board computer
- Automatic vehicle location tracking
- Advanced security and surveillance

Focus of Developements

EURO SERVER

1. Next Generation Compute System Architecture
   - Scale-out and flexible heterogeneous compute
   - Define the Unimem memory model of a virtual-mappable shared GAS
2. Nanoscale, silicon-on-silicon 3D integration and IP
   - “Chiplet” concept in 3DIC integration and test
   - Heterogeneous silicon interface bridging
   - HMC “memory compressing” controller
3. Software Architecture and Frameworks
   - Utilize the “Unimem” memory hierarchy
   - Resource sharing and system wide reallocation
4. Applicability of Euroserver solution
   a) Device PCB Realization: Development System
   b) Embedded Mircoservers: Wireless Basestation
   c) Scale-out Servers: Cloud Services

Focus Area 1 / 4
System and Memory Architecture

EURO SERVER

System Architecture

- Multiple independent nodes
  - Heterogeneous IO and compute within and between nodes
  - All share access to a common GAS
  - Provide consistent/coherent access to "all" of its local memory space
- Topology Agnostic
  - Address based routing across GAS
  - Hierarchical address ownership
- Globally aware virtualization layer
  - Sharing common IO/Peripherals
  - System-wide power management

"Unimem" Unified Memory

- Multiple coherent Islands
  - Heterogeneous Compute
  - Coherent Local Memory
  - Essential IO and Peripherals
- Access to shared GAS
  - Additional PA
  - Shared IO and Peripheral
- Memory Model
  - Any node can locally VA map any GAS location
  - Single node can cache VA map of shared GAS

Focus Area 2/5: Nanoscale Integration

Actual configurations to be confirmed

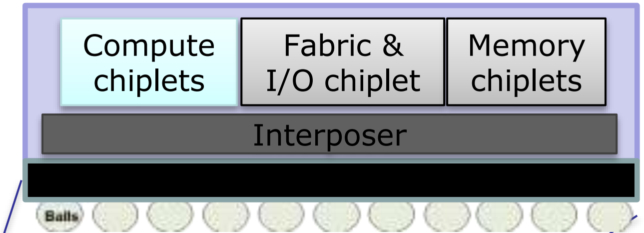

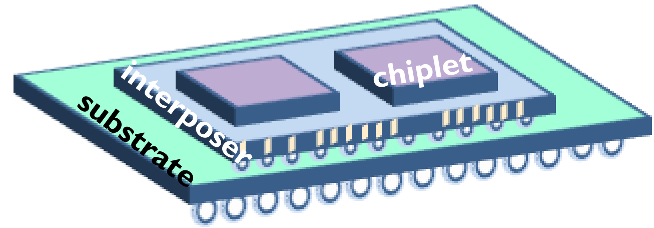

-Structure of "chiplets"
- IP to bridge silicon
-Technology to realize

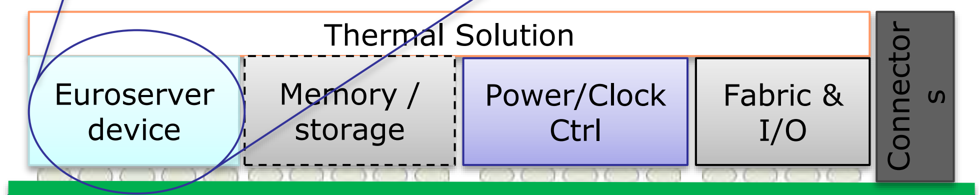

Micro-server board view
-How to reuse across different markets
- Market specific requirements

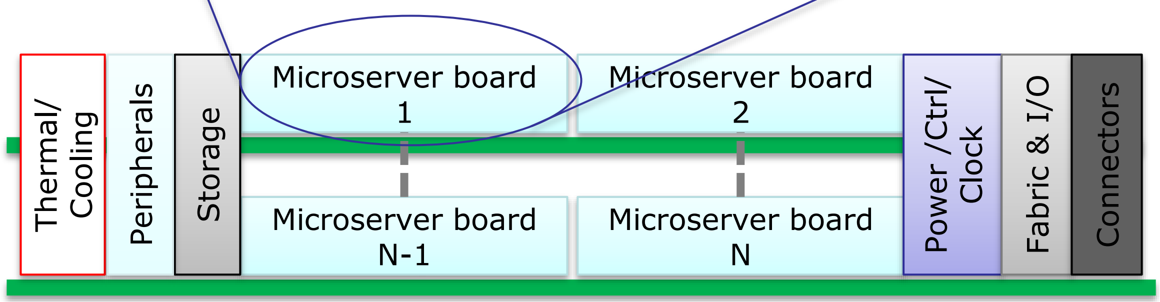

Modeling Device Cost Impact

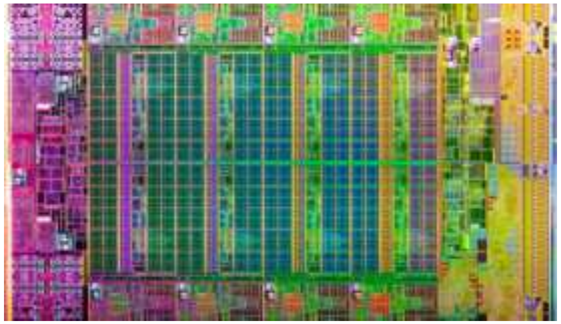

Consider a traditional 300mm2 2D-SoC

Currently yielding means $400-800 per device
In < 20nm, that cost nearly doubles!

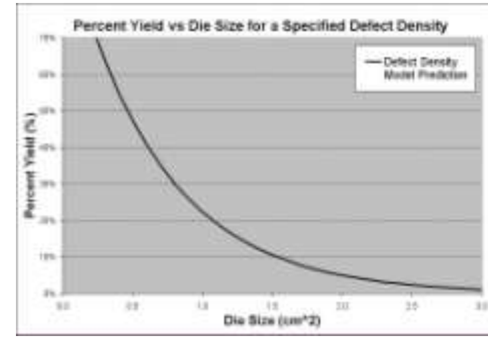

Assume a 60:40 split between logic best in < 20nm and that in >= 20nm

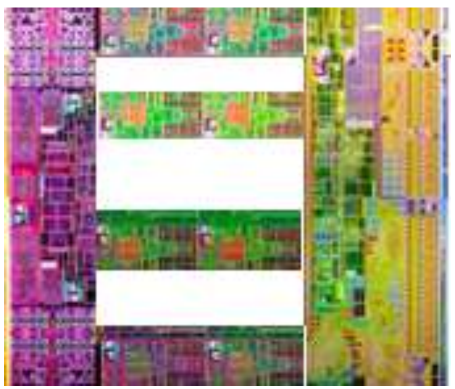

+

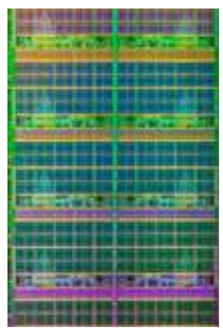

=

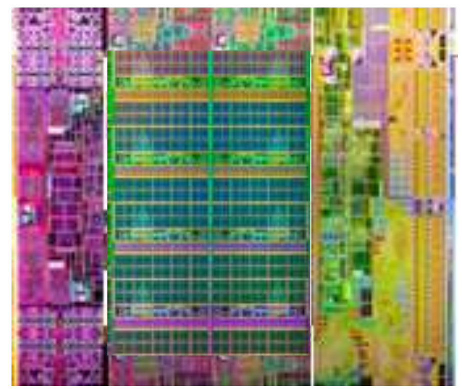

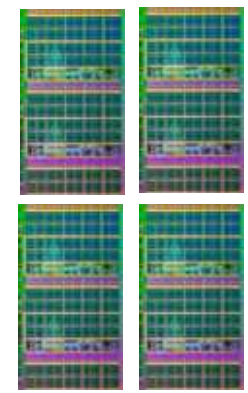

Rest of SoC in
low cost > 20nm

Key IP/processors
in <=20nm

And 3D stack

Replicated chiplets
increases yield
yet again.

Both die are smaller so yield higher and hence ~20x lower cost, even when adding *today*'s 3D cost

Focus Area 3/5: Software Architecture

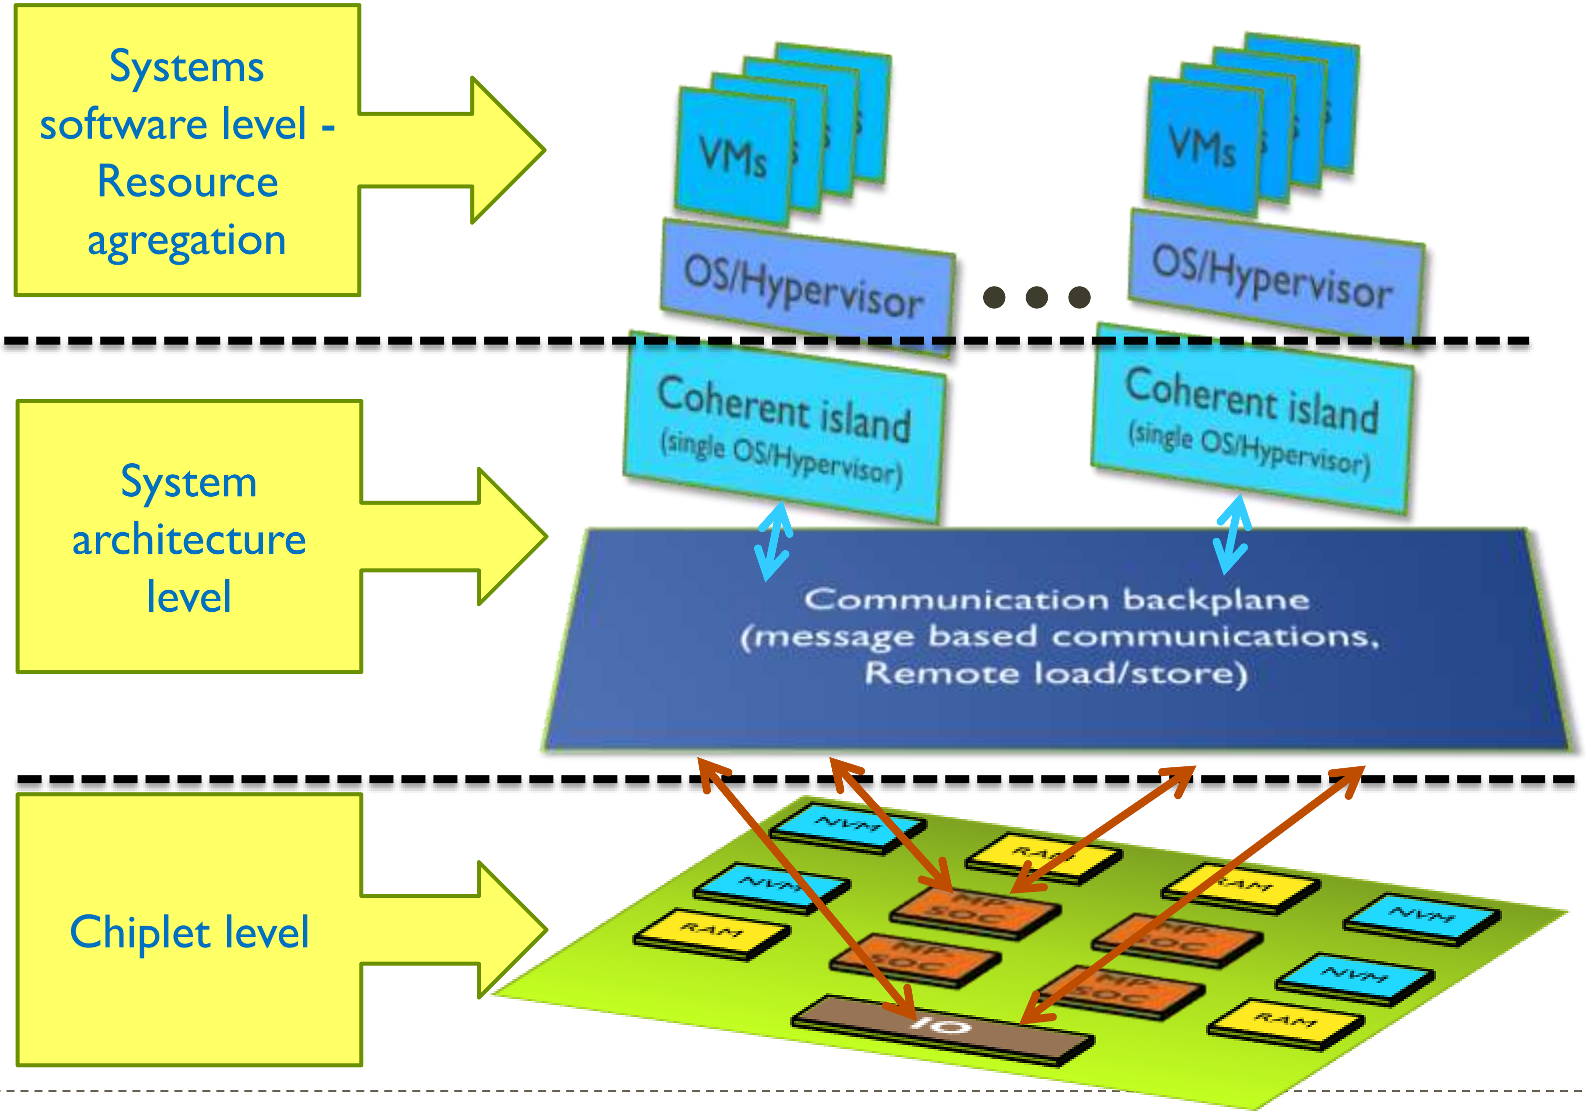

Access anything, anywhere

Flexible over specific "compute"
Heterogeneity and configuration

DMA to support current RDMA paradigms

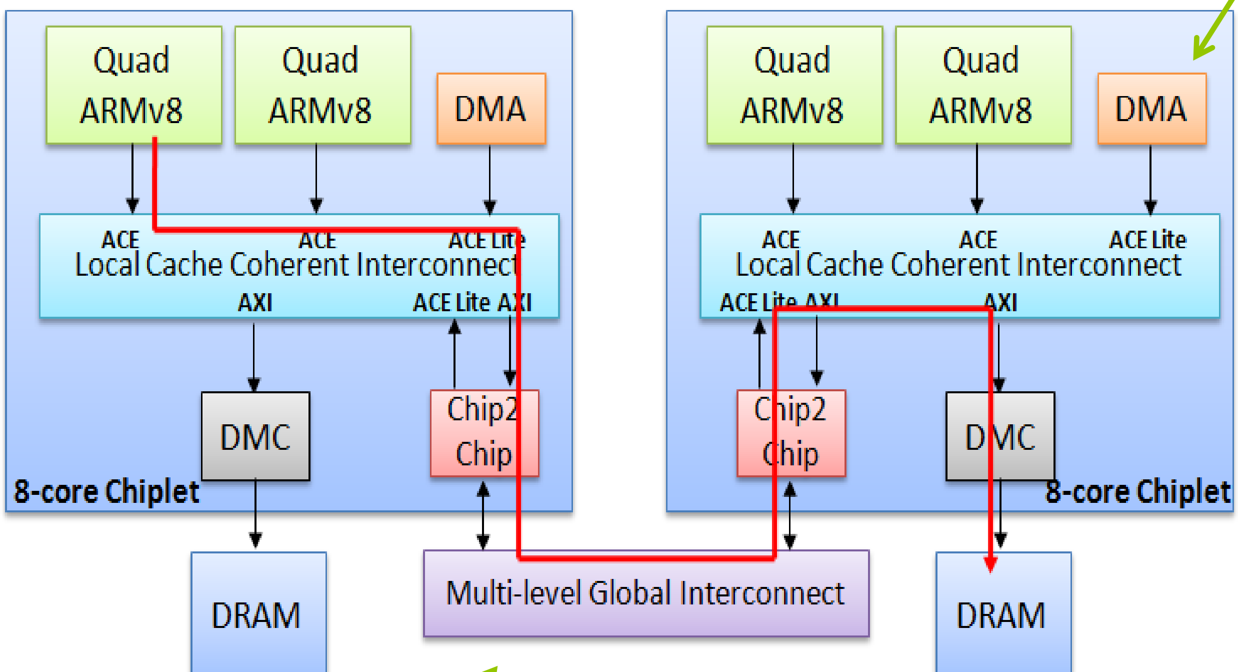

directly VA map
between nodes over address based
routing topology

Local DRAM partitions

Architecture enables duplicating only the resources the app needs

Scaling using current commodity architecture

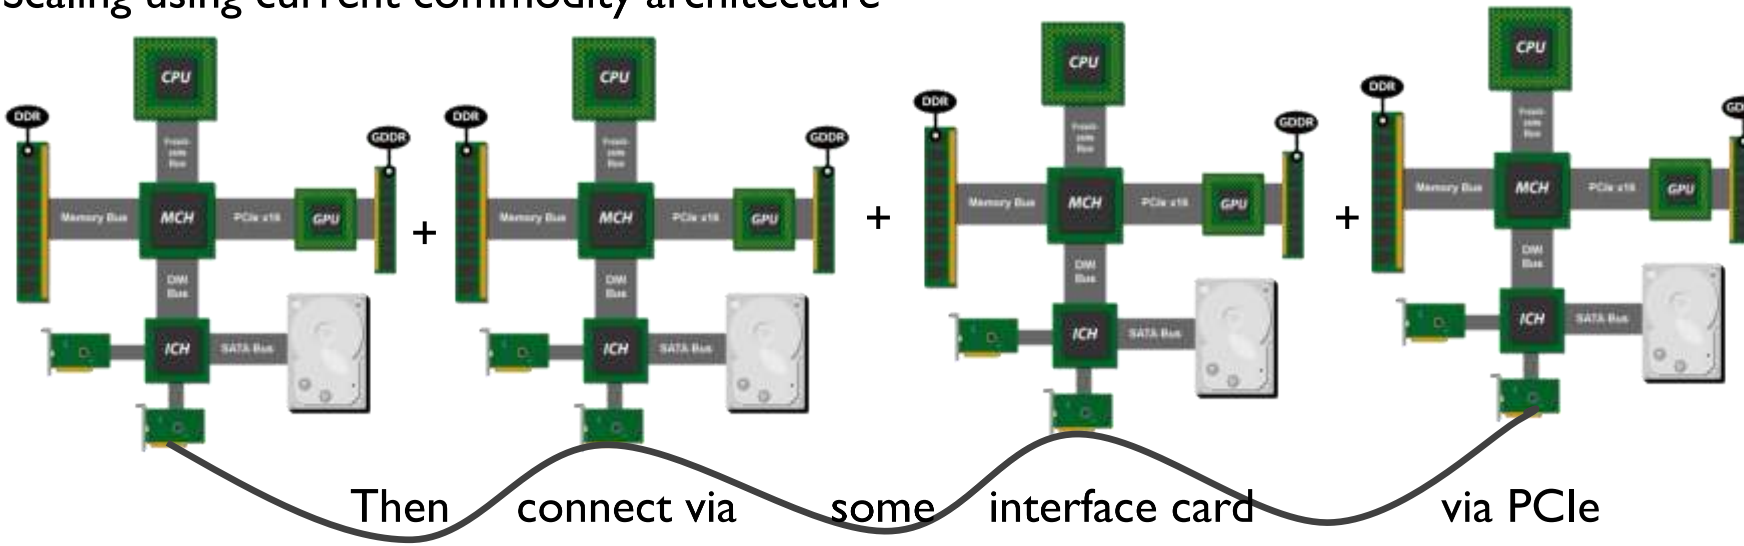

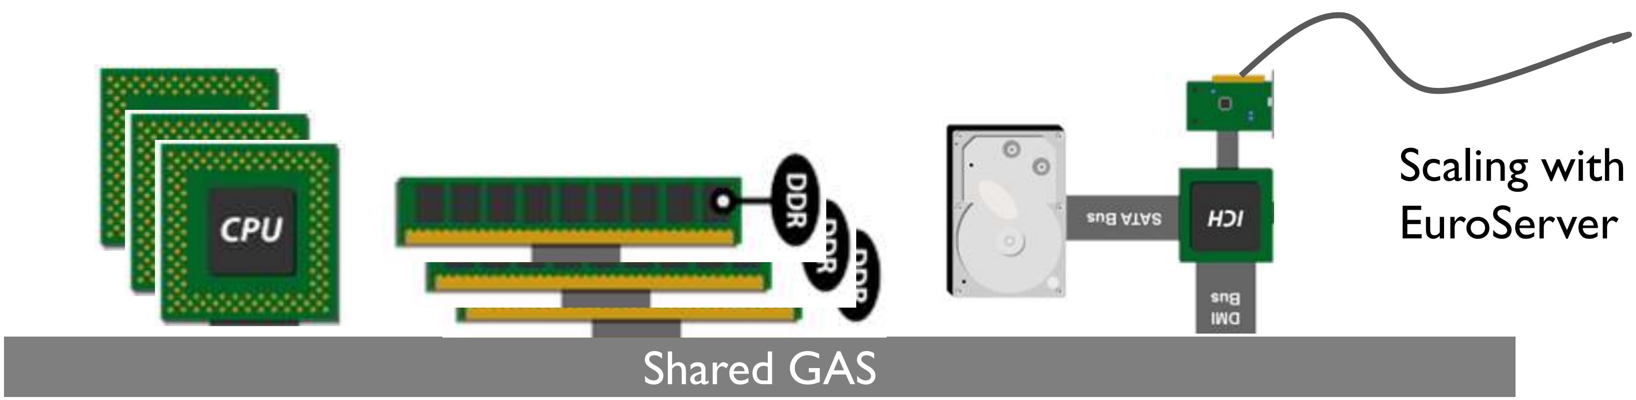

Constructed using chiplets then direct virtually map remote memory and devices

Focus Area 4/4 part-a: Demonstration of a Micro-server

- Tentative specs
  - Form factor: standard ComExpress or Custom
  - Size: ~ 8 x 12 cm
  - Components:
    - 1 Euroserver Device (compute node)
    - Single board-to-carrier connector
    - Local storage, local memory
    - Network Fabric Interconnect
    - Board management and control logic and sensors
    - Local test, debug and peripherals to consider
  - The Micro-server board will populate (in different number) both the Embedded Server and the Enterprise Server

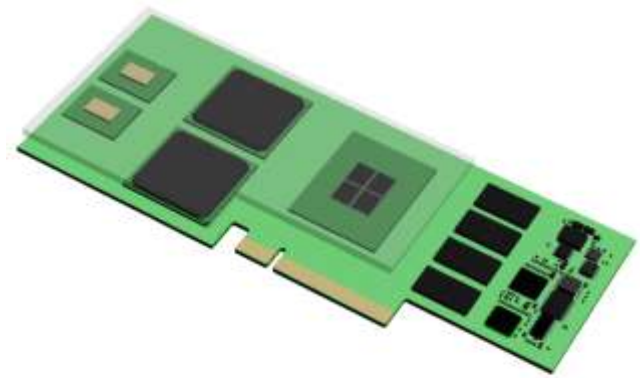

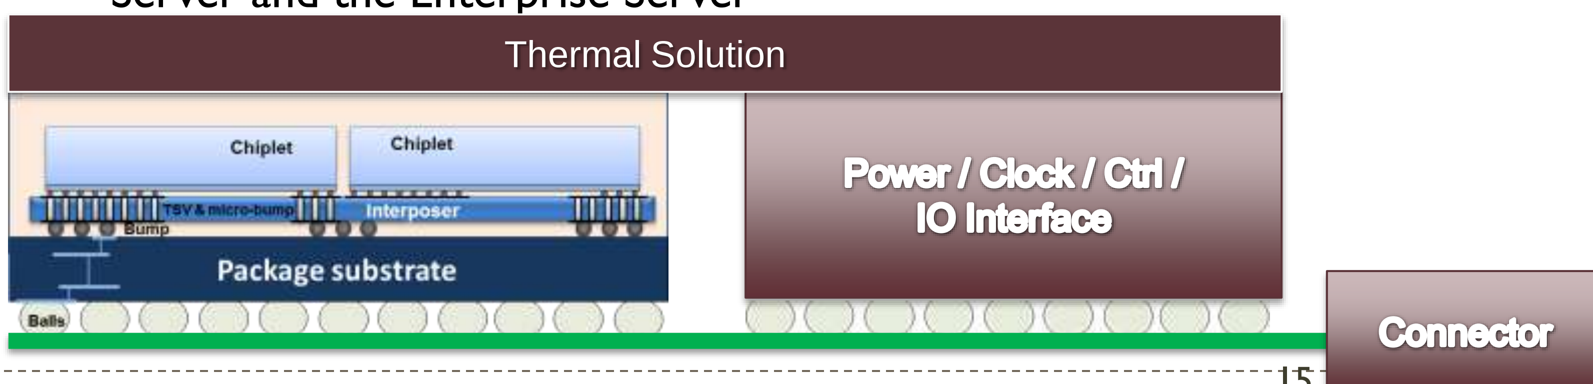

Key Focus 4/4 part-b:
Demo in an Embedded Server

EURO SERVER

- Tentative specs
  - Target form factor: Eurotech Mounted Mobile Computer
  - Size: 13 x 25 x 8 cm
  - Prototype targeted for rugged sealed enclosure and extended temperature range
  - Modular system design
    - Shared design for micro-server board and carrier board with Enterprise Server version
    - Up to 2 micro-server slots
    - Support for optional general purpose card (using 1 slot) for specific function or peripheral
  - 9 to 36 Vdc in; 30W max;
  - Passive conduction cooling
  - Support for Backup battery pack to investigate

Key focus 4/4 part-c: Demo as an Enterprise Server

EURO SERVER

- Tentative Specs:
  - Compatible form factor with standard server rack design

  - Minimum density design
    - 1" node pitch (16 nodes per rack width)
    - It enables 64 micro-server cards (16 x 4) in a standard cabinet and 12 kW compute power in a 42U rack
  - Carrier board design shared with Embedded Server version
  - Population options: any combination of micro-server node cards and general purpose cards (each card uses one slot)
  - Support for multiple thermal and cooling options

Vision for the Future

- Target Different Application Area - HPC:
  - Realize the next level of system architecture for Next Generation Compute. (adding stacked dram and integrated IO)
  - Apply the System and Memory Architecture along with more advanced Nanotechnology and Software Framework and Runtimes

- Increased "compute-chiplet" heterogeneity
  - Core selection, combinations and integration granularity
  - Semantically Aware, Generic Runtime Support

- Investigate application of Photonic Technology
  - On silicon for active interposer's interconnect
  - On substraight for interconnection of Silicon in MCM
  - Extending across PCB