MANCHESTER 1824
The University of Manchester

Should the semiconductor industry embrace silicon modularity ?

John Goodacre
Professor Compute Architectures
School of Computer Science
University of Manchester

17th INTERNATIONAL FORUM ON MPSoC
for software-defined hardware
July 2-7, 2017
Les Trésoms Hotel Annecy, France

Abstract

After years of research and talk of stacking memory on processors, 2.5D interposers, 3DIC, TSV, advanced substrates, what is driving the commercial design and assembly structures for tomorrow's digital chip? What level and logically structures can be used to modularise a chip? Will one technology and approach address all needs? How can die from different company exist in a single package? This talk will look at some of the research that's occurred in the EU, is still happening, and might happen, to look at realizing some of the benefits promised by such integrations. What business situation could be a tipping point in defining a new market and business for building processor devices through interconnected silicon modules?

In the beginning…

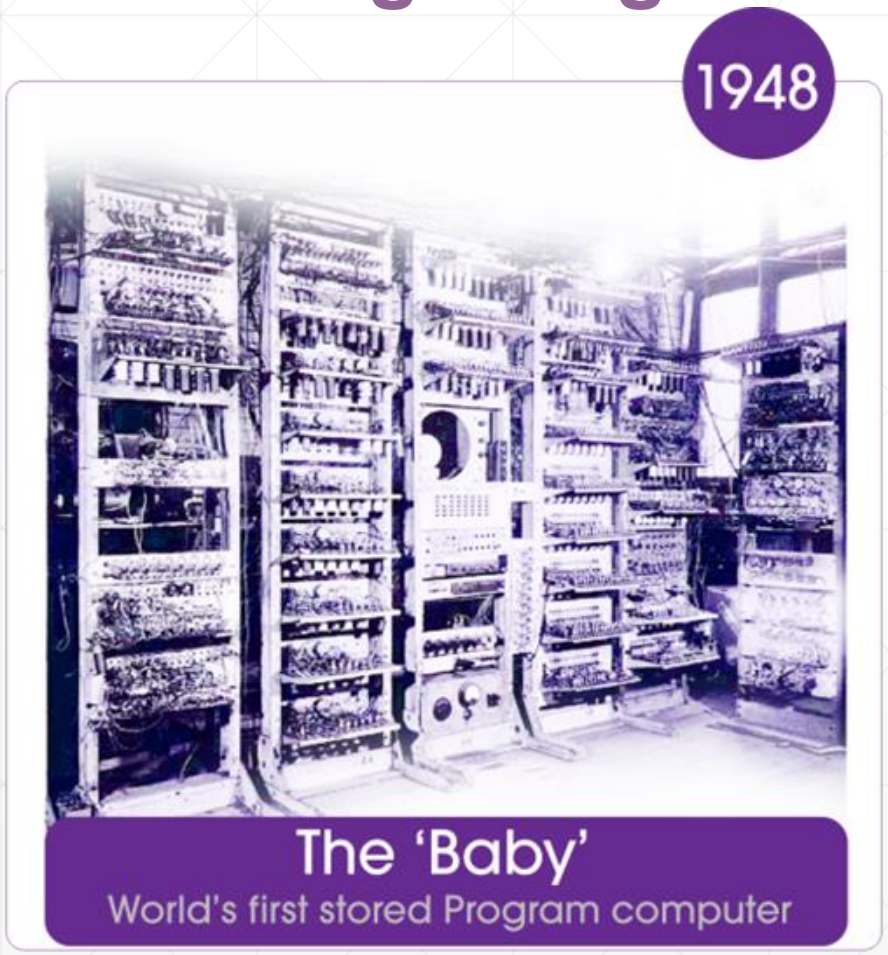

- The World's first stored program computer:
- Invented and constructed in the University of Manchester by Tom Kilburn and Freddie Williams;
- Functions implemented across modular "racks"

….and things haven't really changed since!

Modularity of the PC architecture

- Processor / Northbridge / Southbridge
- Created an entrenched ecosystem
  - (for quite a while) North/South chipset
  - Plugin graphic adaptors
  - Plugin network adaptors
  - 3rd party innovation limited to plugins
- Market differentiation through the PCB and the specific modules used.

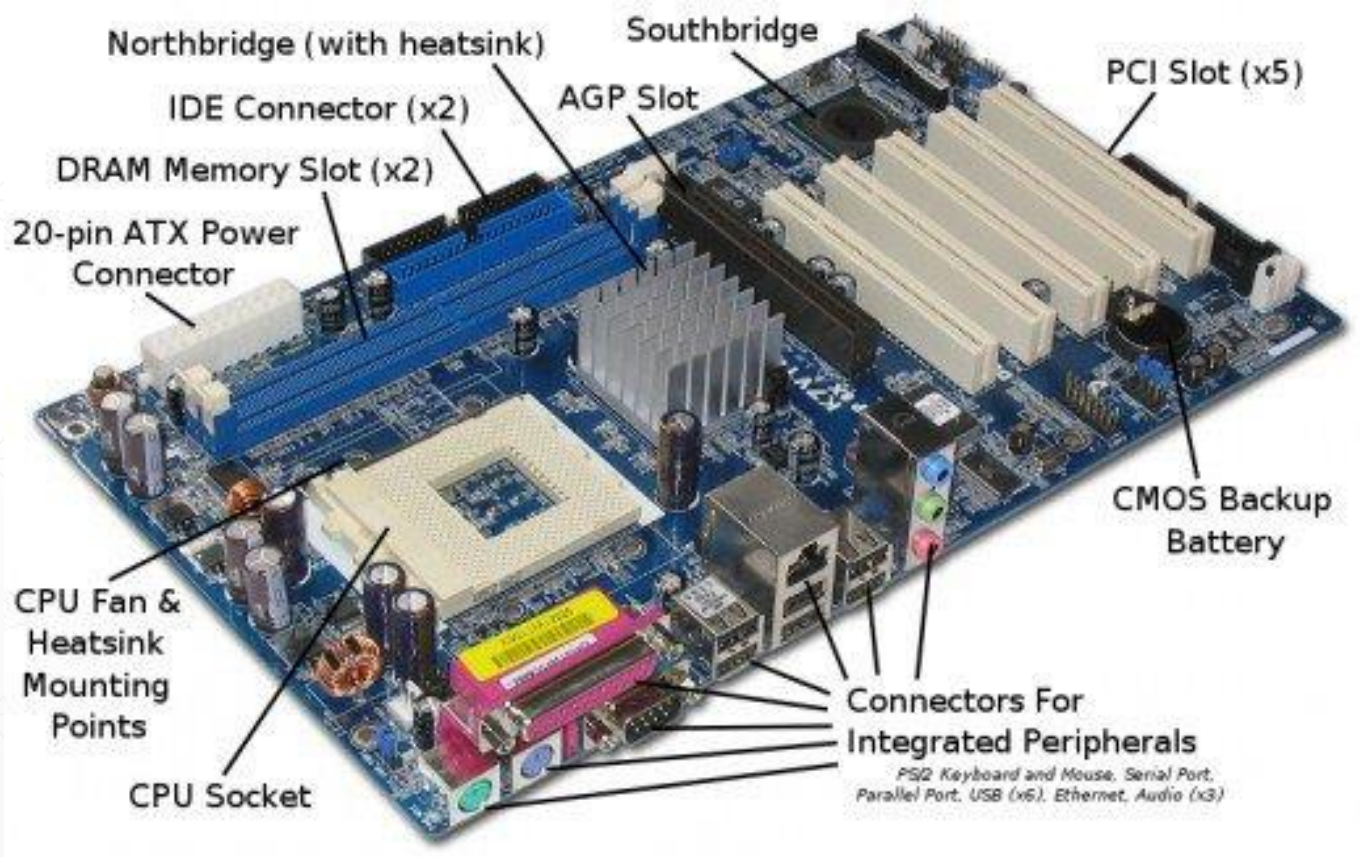

Some multi-chip modularity – even 15 years ago

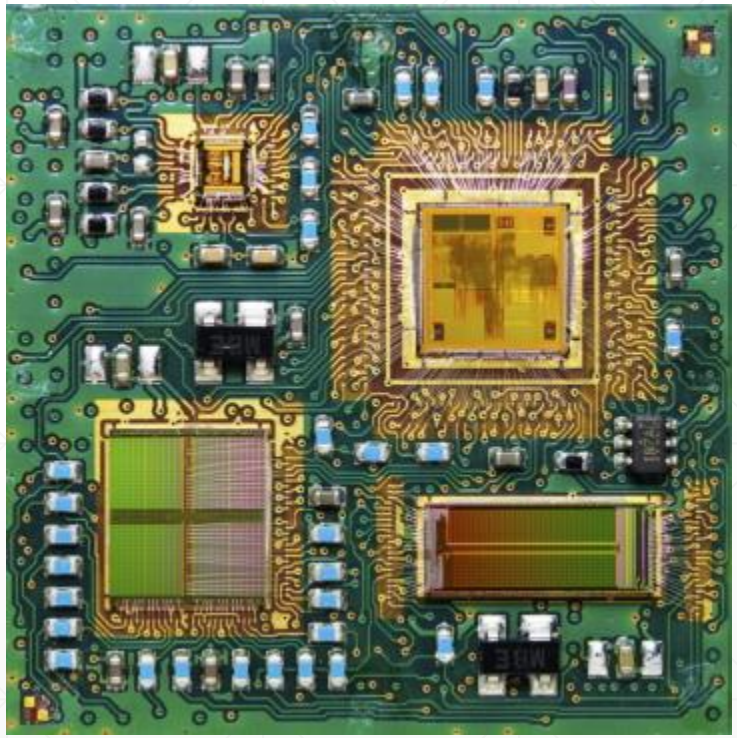

http://developer.axis.com/old/products/mcm/

- **Reduced production costs** - There are a lot less components to mount and this will result in a smaller PCB with fewer layers. In most cases you can use a 2- or 4- layer PCB instead of a 4- or 6-layer PCB.
- **Reduced BOM (bill-of-materials)** - All components included in the MCM are sourced by Axis from the suppliers in very high volumes and at better rates than can be achieved for smaller volumes. Buying the MCM is normally cheaper than buying all the components separately - at least for volumes less than 10K per year. After that, the cost advantage wears off.
- **Increased production yield** - There are less components that can fail during production and the hardware is much simpler in design.
- **Lower inventory costs** - Instead of keeping stock and sourcing some 50 different components you only need one component which will also reduce logistical problems.
- **Decreased footprint** - The ETRAX 100LX measures only 27x27 mm, which will makes it possible to design products with a very small footprint. In practice, the connectors of your product will be the thing that decides the size of it.

…and then into System on Chip

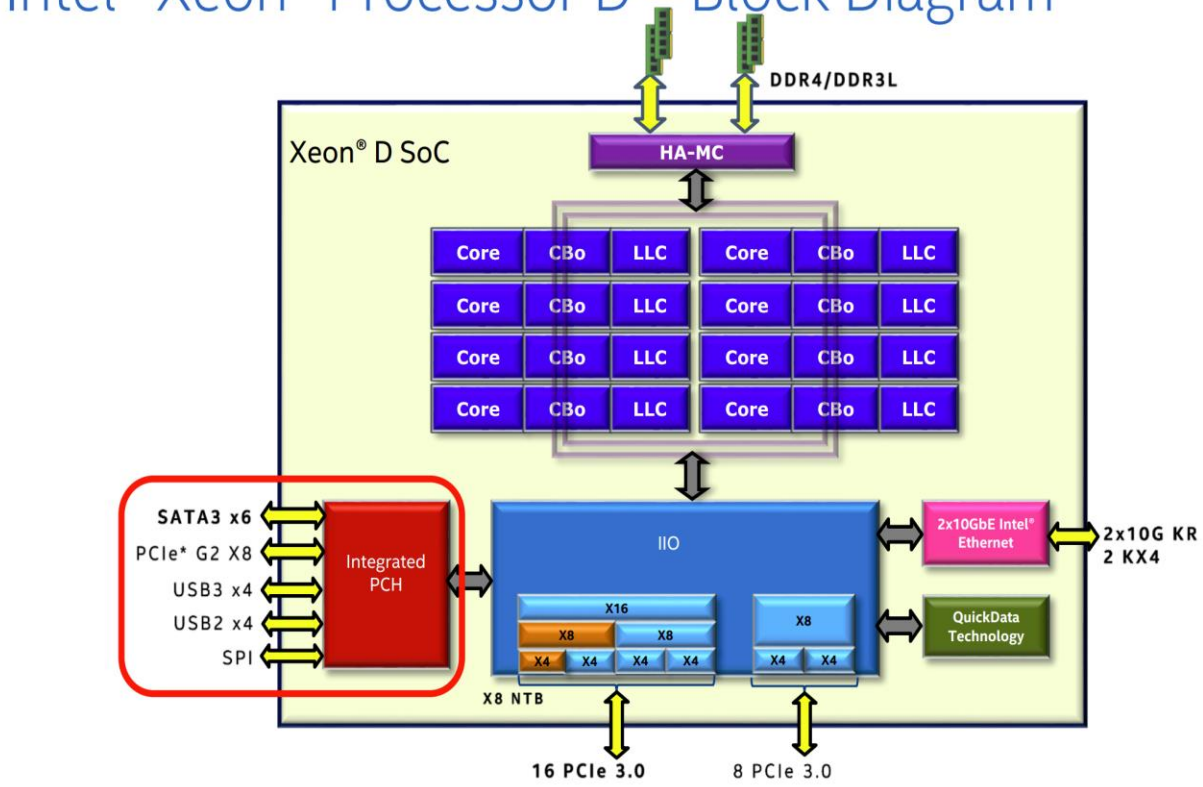

- Enabled by IP business models (aka ARM)
- Much more open to innovation, and ability to create application specific devices
- Limited x86 SoC offerings

More expensive transistors, slower interconnect

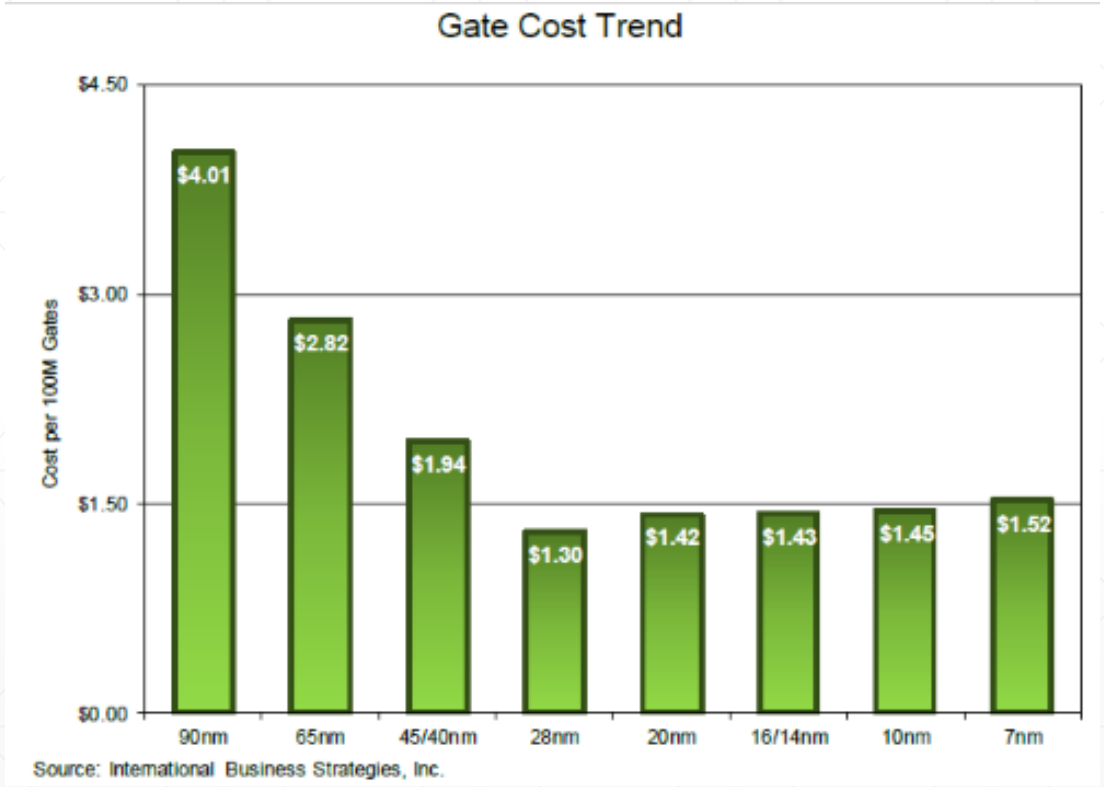

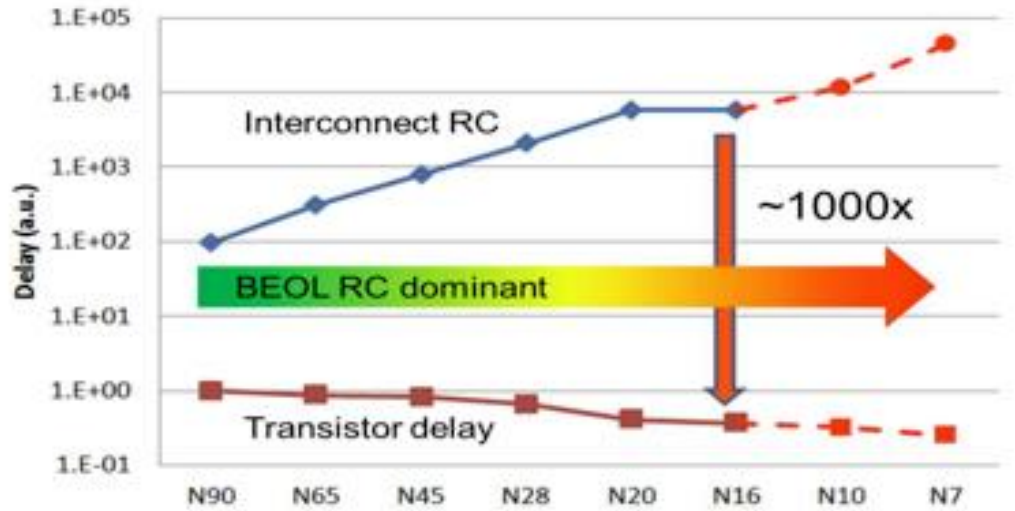

Fig. 17: BEOL performance/area/cost scaling is the foremost issue for 10nm/7nm nodes.

A perfect storm of technological limits

For PCB design

- Pitch of pins on device
- Noise from the high-swing voltages needed to communicate over PCB
- Power density and cooling
- Design complexity to innovate beyond device reference design
- Limits of bandwidth/latency and energy cost of interconnect

For SoC design

- Design and fabrication costs (NRE)
- Thermal density (Dark silicon)
- Voltage reductions
- Bits per pin / rc -> speed / parallelism
- Pins per device / device fanout
- Die size / yielding

Idealized Roadmap: SiP/MCM - 2.5D - 3D - Monolithic

- Driven by technologists
- Early adaptors push each generation hard
  - Eg interposer challenges to support FPGA
- No heterogeneous die business model
  - Known good die – who to blame
- No “open” standard between components
  - Connect nvidia GPU to Intel CPU ????
- Some adoption within a single company on a single design

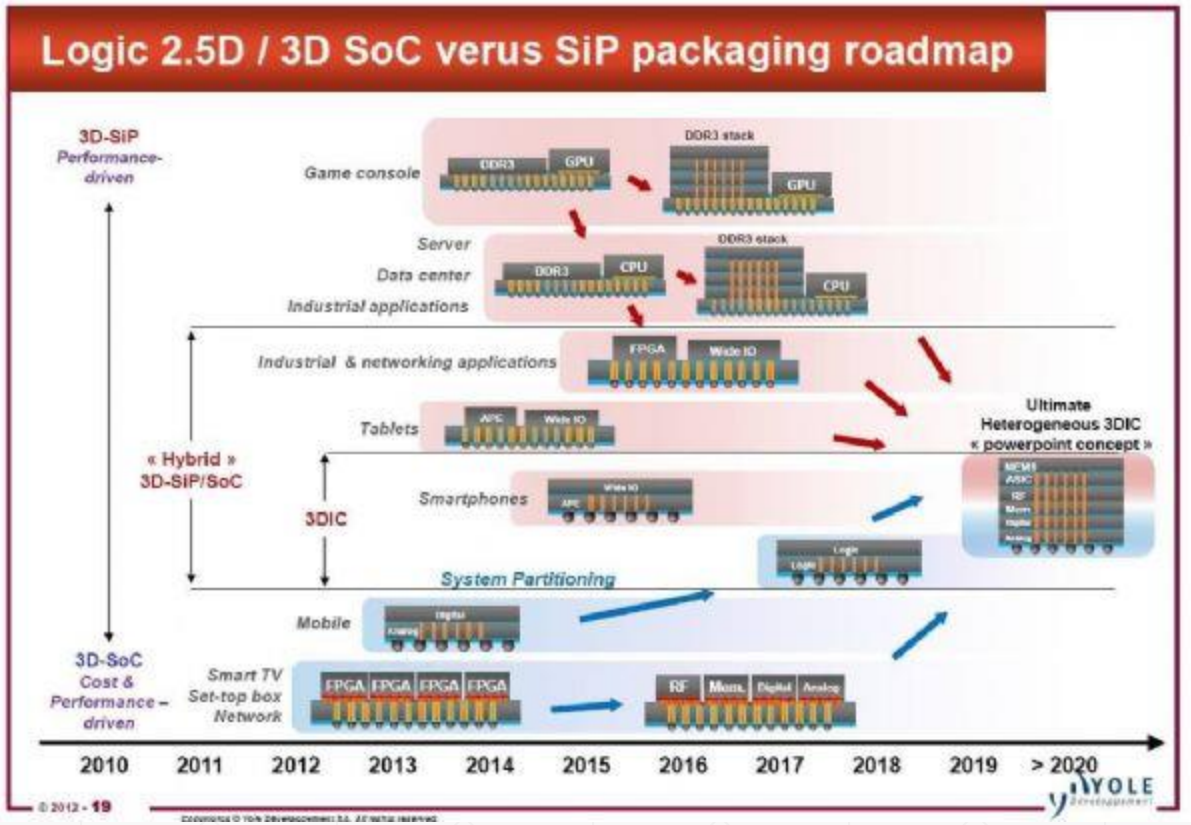

Commercial hopes and roadblocks

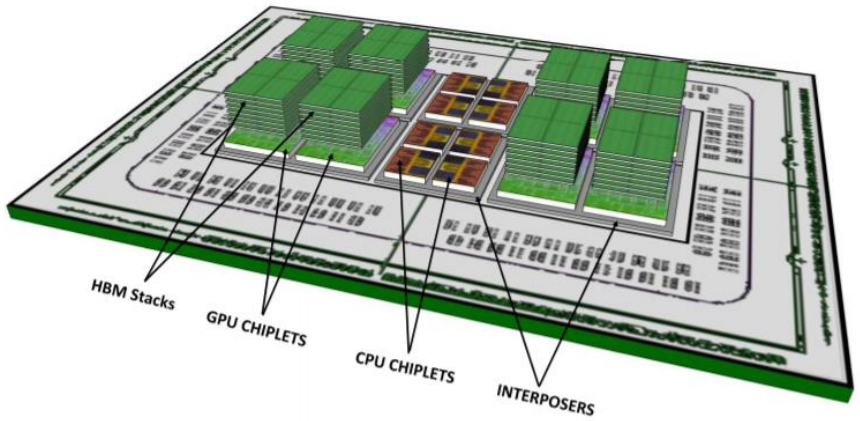

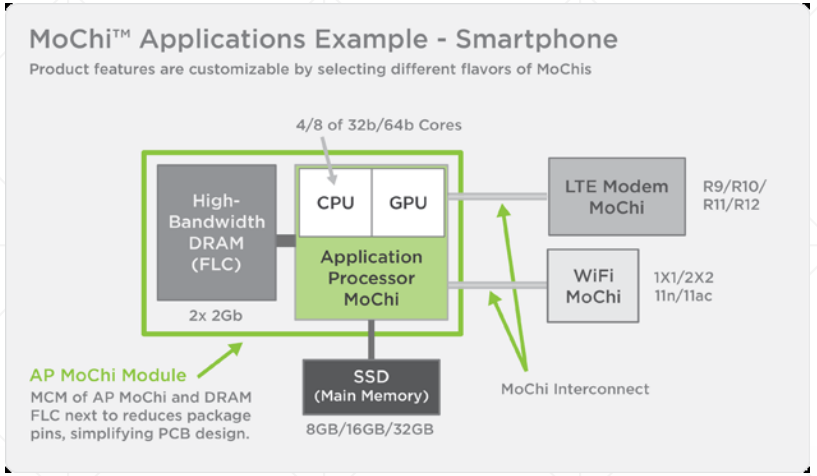

- Scaling device larger that silicon yield allows
- Short reach lowers power and latency
- Simplifies customer use

- Traditional design partitions
- Interconnect capable of PCB-level drive
- Modularity composed at PCB level

http://www.computermachines.org/joe/publications/pdfs/hpca2017_exascale_apu.pdf

http://www.marvell.com/architecture/mochi/

Missing steps:

Silicon modules as reusable IP

- Silicon Modules
  - Accepted design partitioning and associated system architecture
  - Proof of higher per-die yields, lower device cost, incremental development, design reuse
  - Ability to delivery of innovation
  - Business model to share silicon fabrication costs
  - IP licencing restrictions
- Open in-package interconnect
  - Ecosystem of module and modular device vendors
  - Availability of the “PCIe” for in-package interconnect

Investigations and research towards silicon modularity

- EuroCloud — 3D Stacking of DRAM, TSV approaches
- EUROSERVER — Silicon chiplets, design partitioning, system architecture
- ExaNODE — Active interposer, heterogeneous MCM, run-time use of architecture
- EuroEXA — ARM silicon module, MCM device, HPC applications
- (EuroPROC) — Ecosystem of modules, open-interconnect IP, EuroProcessor

MANCHESTER 1824
The University of Manchester

EUROSERVER: Design and Silicon Partitioning

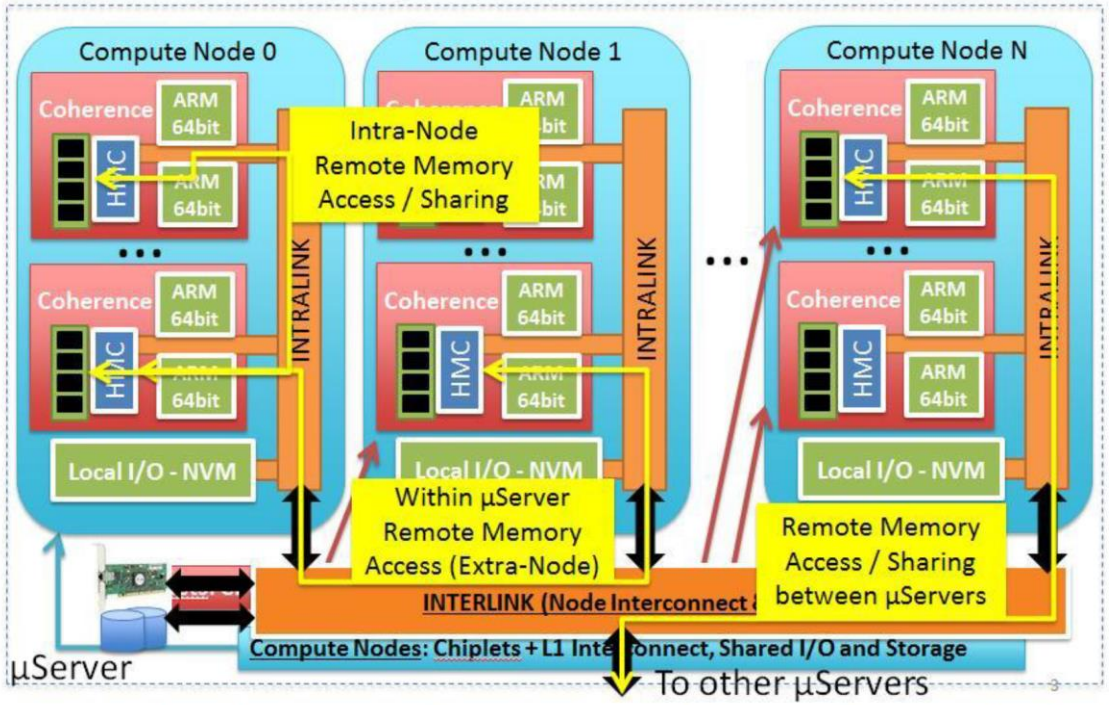

Figure 3: Proposed EUROSERVER architecture showing compute nodes (light blue), cache-coherent areas (light red) and 3 types of remote memories (yellow).

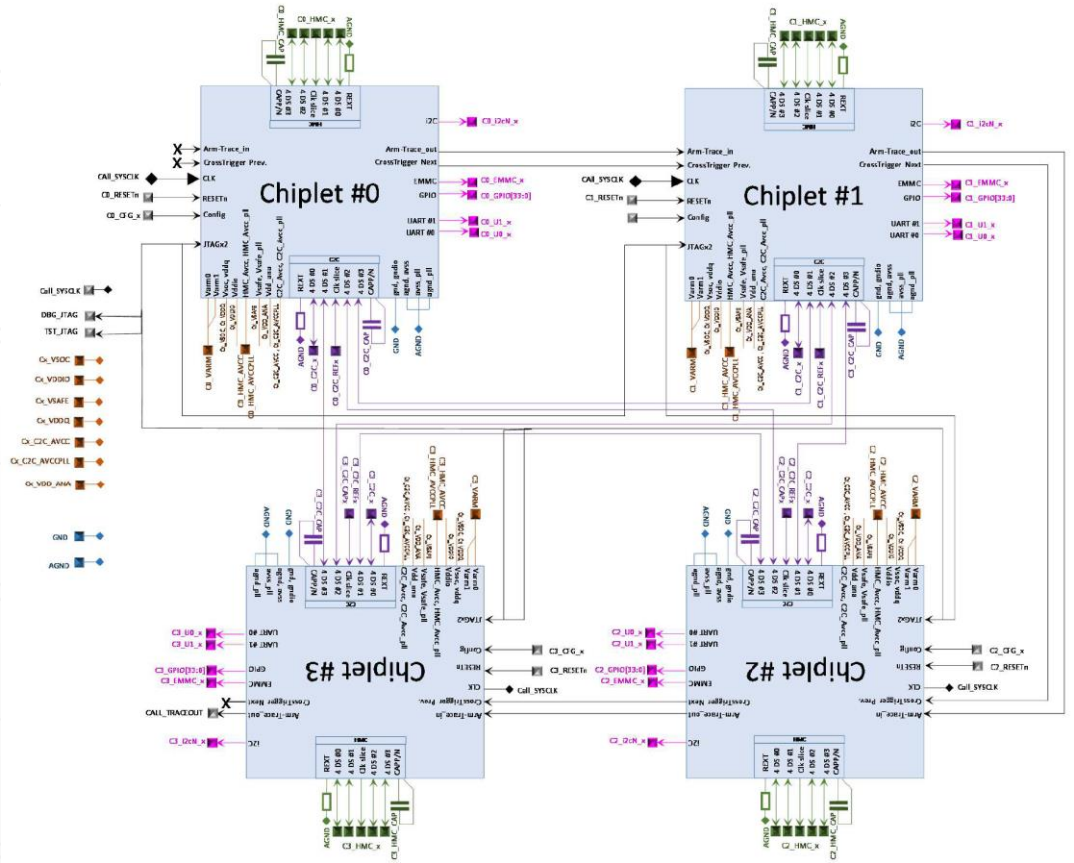

EUROSERVER: D4.3 Kernel-level memory and I/O resource sharing across chiplets
https://projectnetboard.absiskey.com/viewdocument/a64f37-7caac9-8ce525-226202-000020

EUROSERVER: D5.2a Specification of the interposer
https://projectnetboard.absiskey.com/viewdocument/290655-ff63c7-4327c4-3bc061-000003

ExaNODE: Progression on Integration and modularity

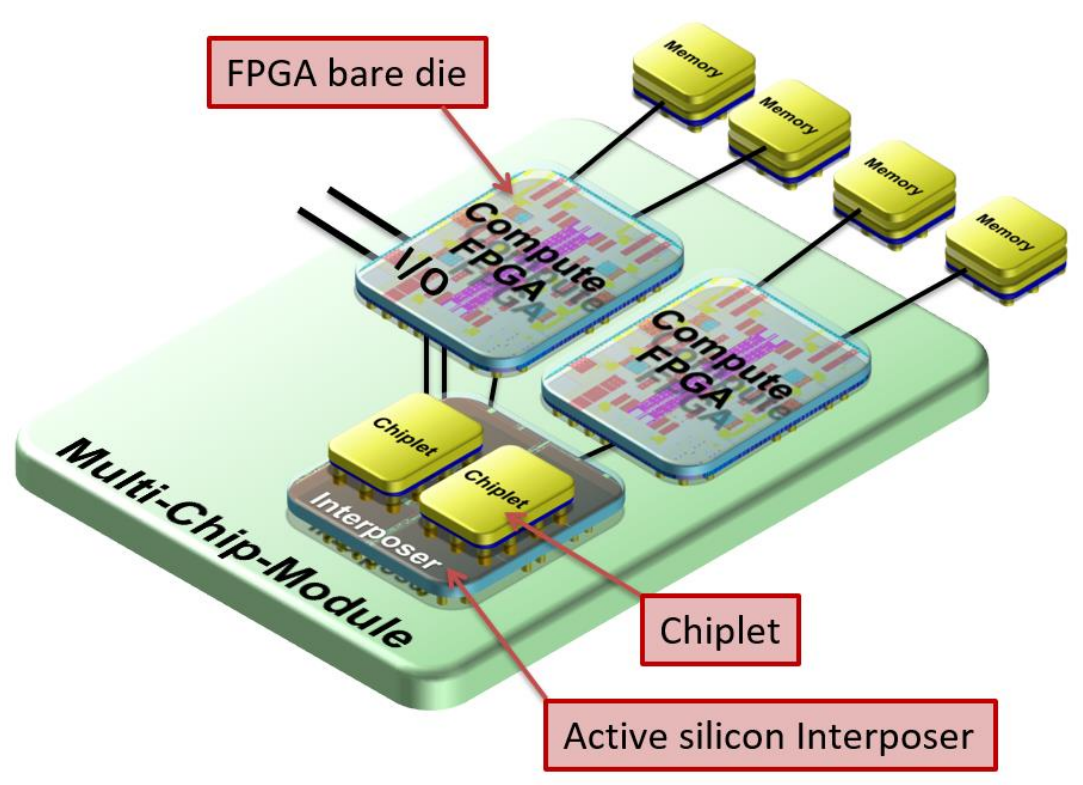

- 3rd party (FPGA) bare-die integration
- Advanced chiplet-chiplet interconnect
- Active interposer (wires + logic)
- Unimem system architecture support
- Software and run-times support

ExaNODE: Path to open-standard interconnect

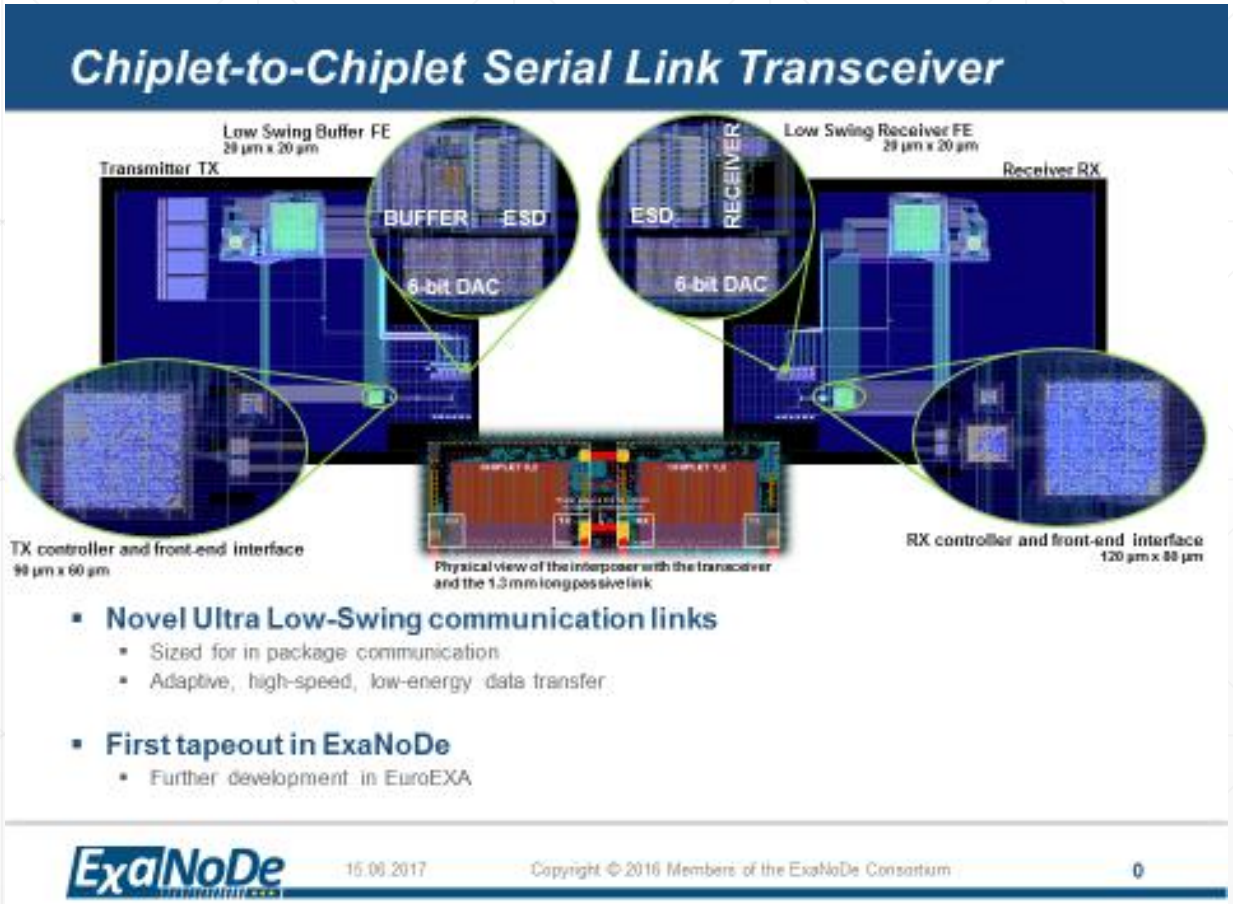

- Designed for in-package interconnect
- Adaptive to specific substrate
- ExaNoDe tapeout on silicon substrate
- EuroEXA tapeout on organic substrate

- Goal to create a open ecosystem of modules around a open interconnect

EuroEXA: (starting Sept'17)

- Creation of a native Unimem capable ARM silicon module
- HPC co-design of a mix of modules
- Deployment of device and system architecture at scale
- Furthering capabilities and maturity of the inter-module interconnect
  - Align to substrate physical layer
  - Increased bridging capabilities

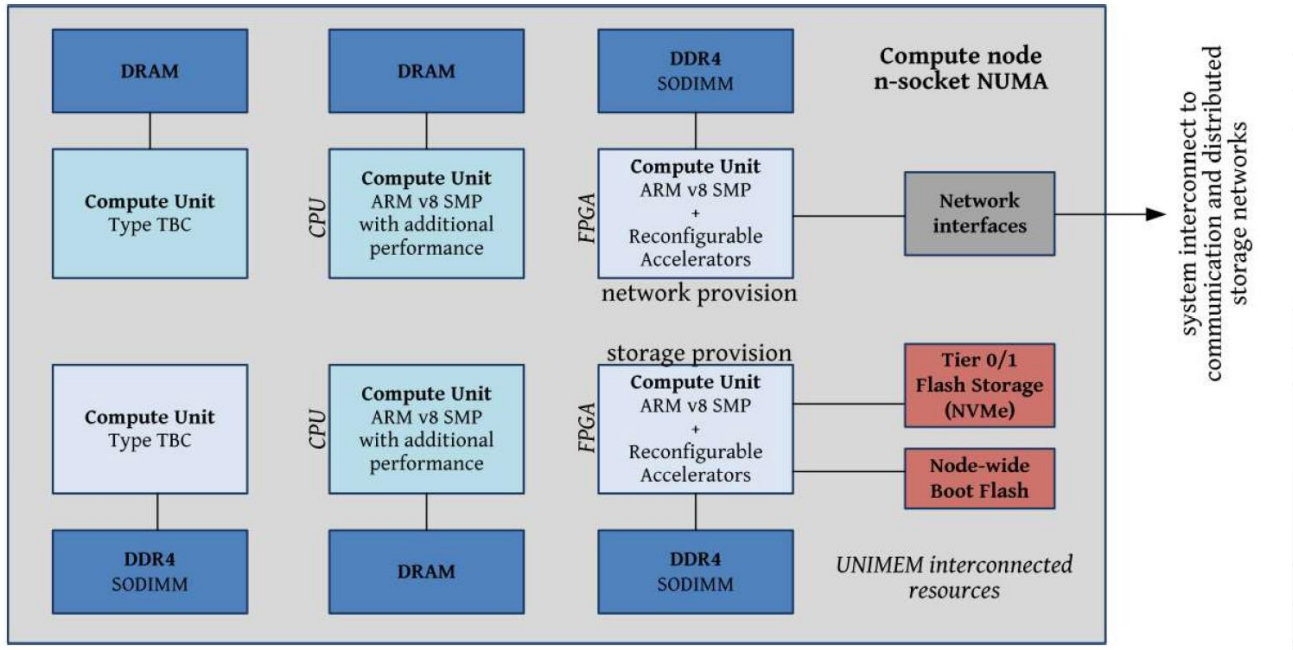

FIGURE 1.4 POTENTIAL EUROEXA NODE CO-DESIGN FOR DEPLOYMENT IN TESTBED 3

Proposed progression

- EuroPROC:
  - Creation of various silicon modules
    - Smart memory module
    - Programmable IO accelerator module
    - Integration of security
    - Advancement of general purpose compute module
  - Open-Hardware release of module interconnect PHY and bridge
    - Kick start ecosystem with multiple implementations
    - Free to use specification and existing implementation macros
    - New implementation to be shared into ecosystem

Conclusion

- Fabrication and assemble technology offers multiple approaches for silicon modularity
- There is no industry agreed design partitioning scheme – other than those used at PCB level
- There is no generally accepted interconnect – other than those at PCB level
- These is no simple business module to integrated multiple vendor die – unless packaged

- There are strong reasons the for these issues to be solved
- There are a number of EU projects driving the approach