

20^e édition

infoprodigital

Forum **TERATEC** 25

21 & 22 MAI 2025

Parc Floral, Paris

Performance through specialization – or not?

Axel Nackaerts, Scientific Director, imec

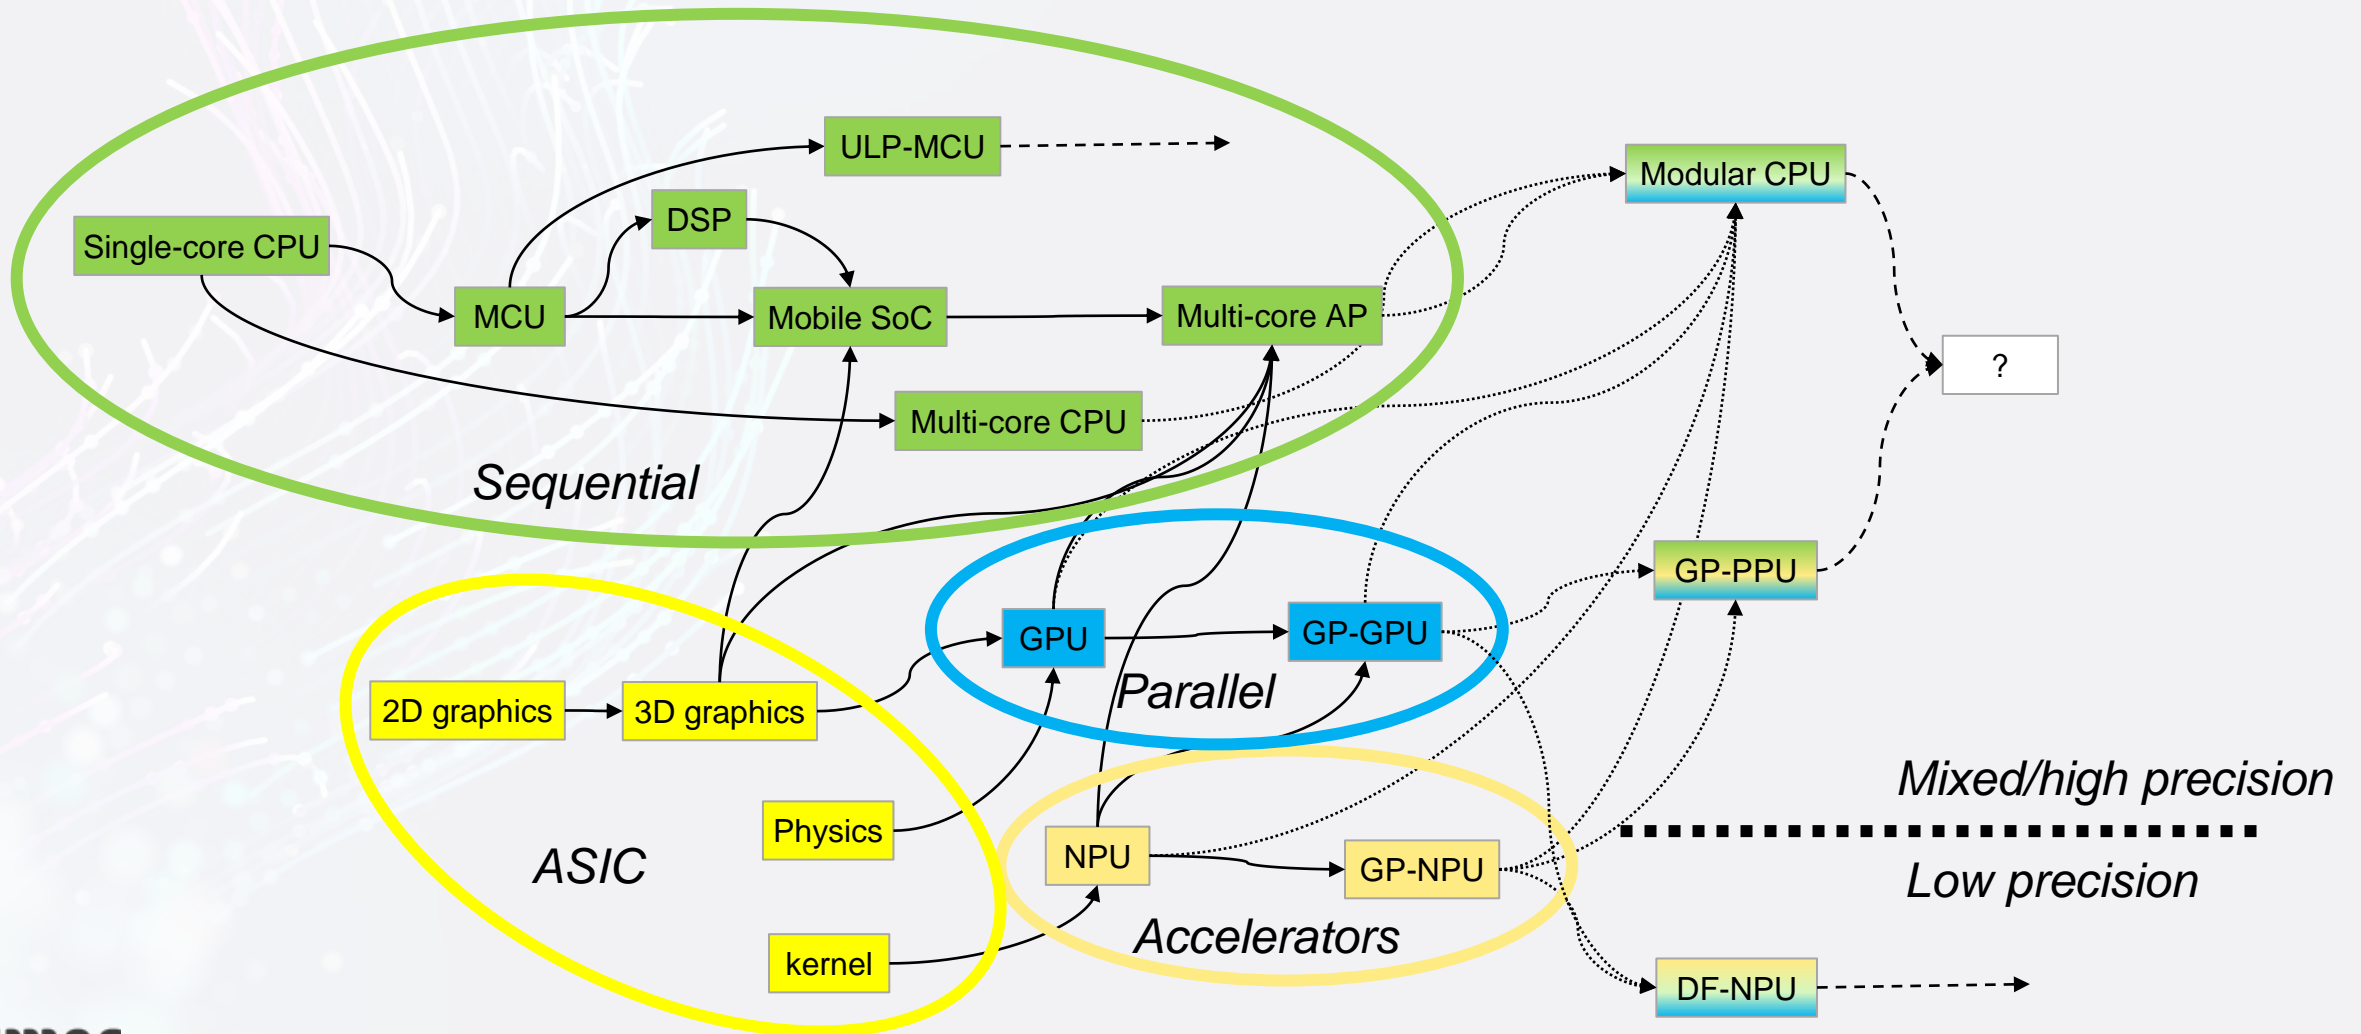
En partenariat avec

L'USINENOUVELLE

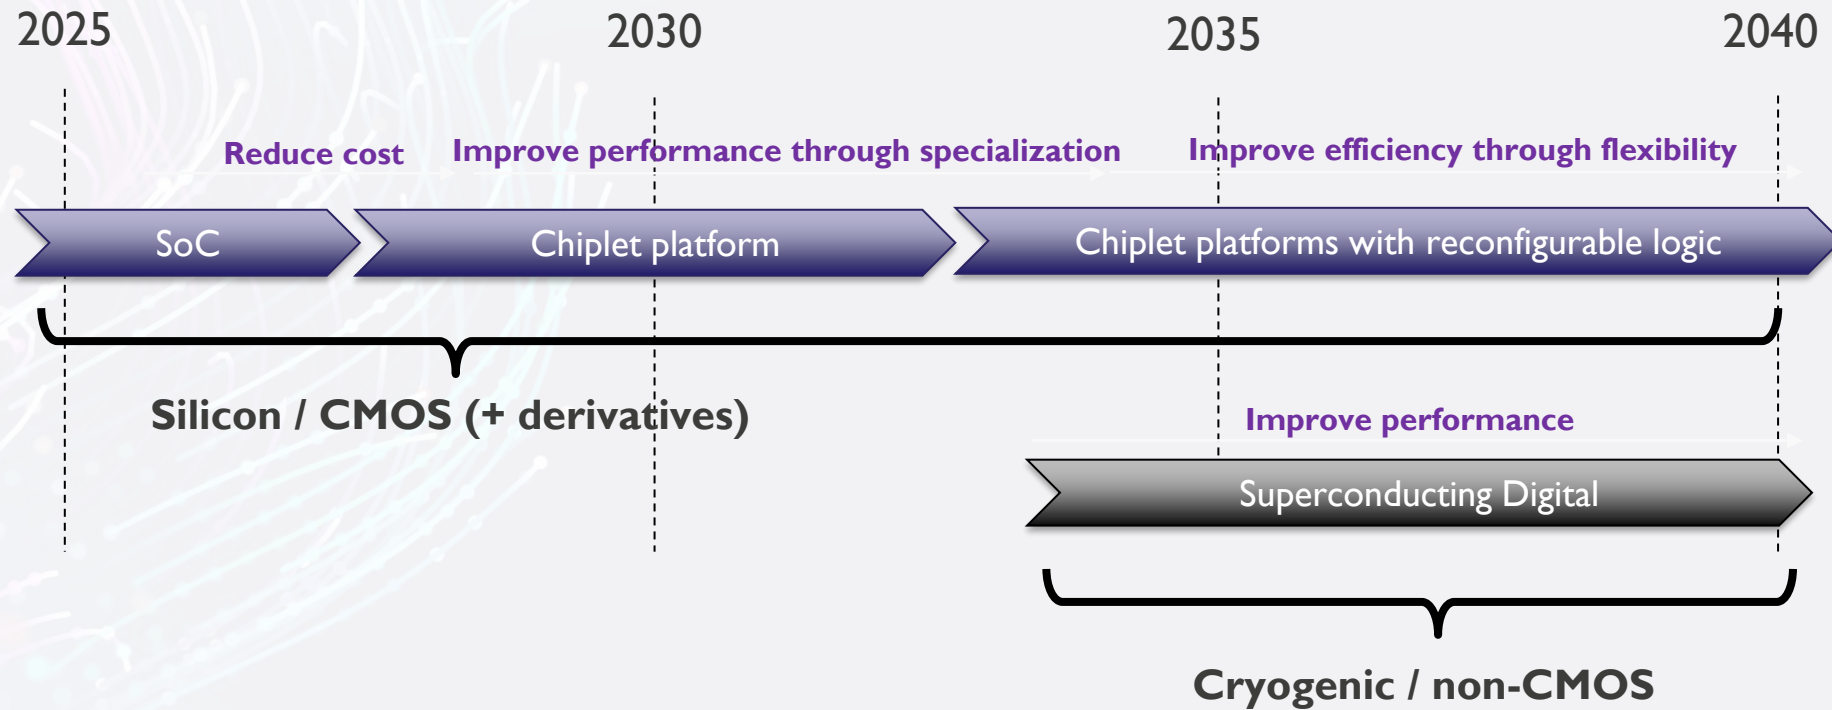
L'USINEDIGITALE



Evolution of Processors



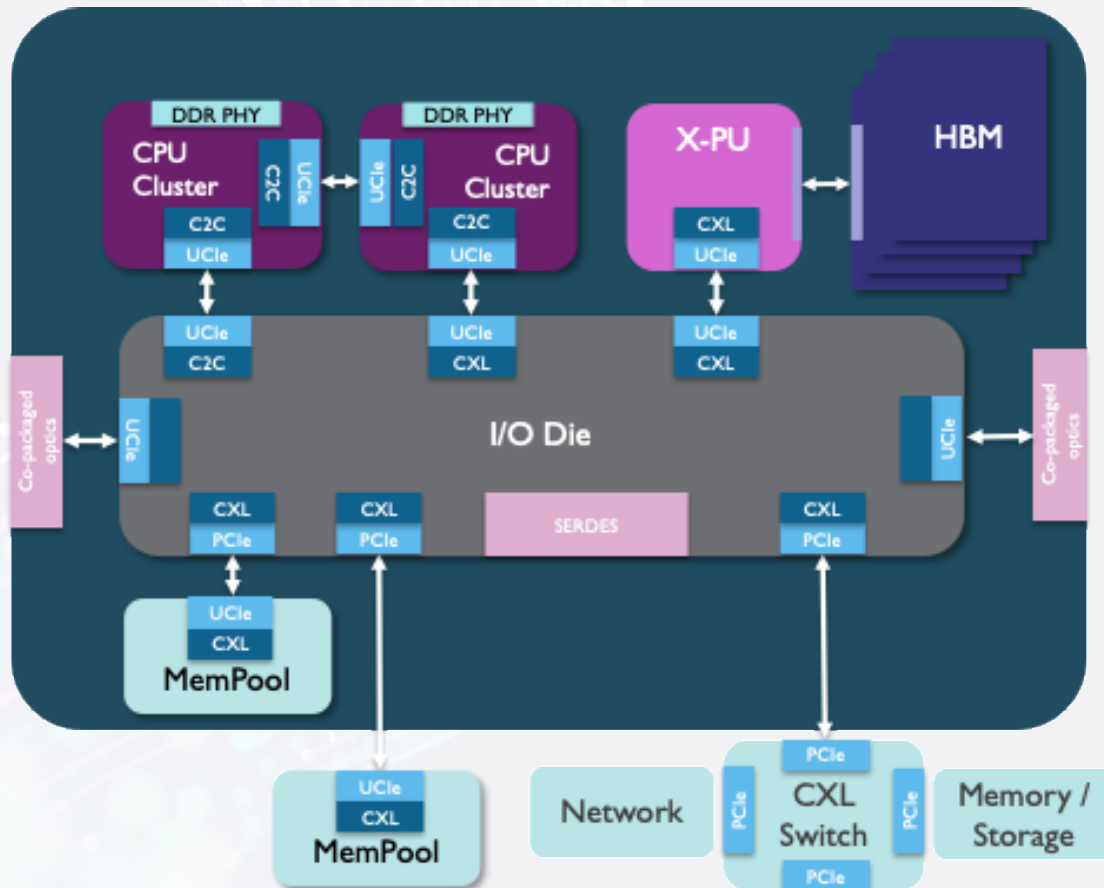
Evolution of implementation



Chiplet platforms



Chiplet platforms




- Combine different chiplets from different manufacturers to create an **application-**optimized solution
- Chiplets can be general-purpose or very specialized
- Combine different technologies for cost-effectiveness
- Increase the market for each chiplet by cross-domain re-use
- Requires adherence to standards
- Requires advanced integrators


Design Space

No 3D Stacking



Fewer Chiplets



More Chiplets



3D Stacking



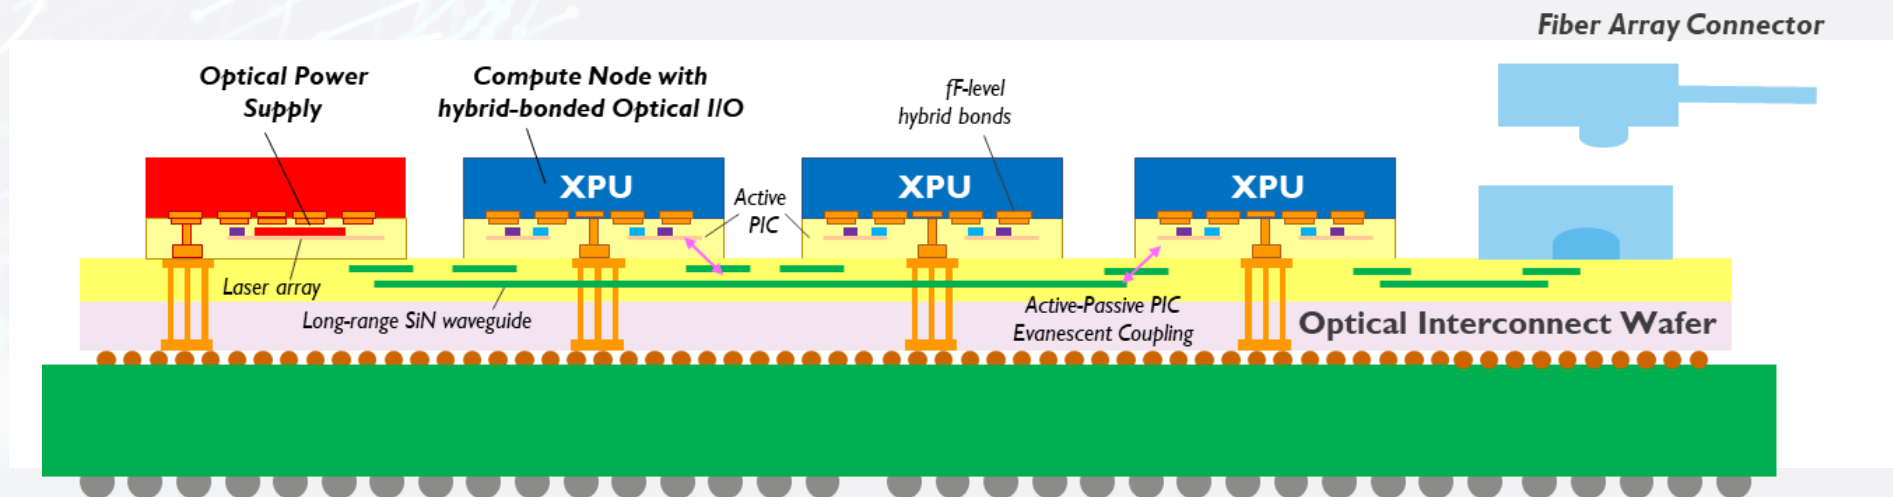
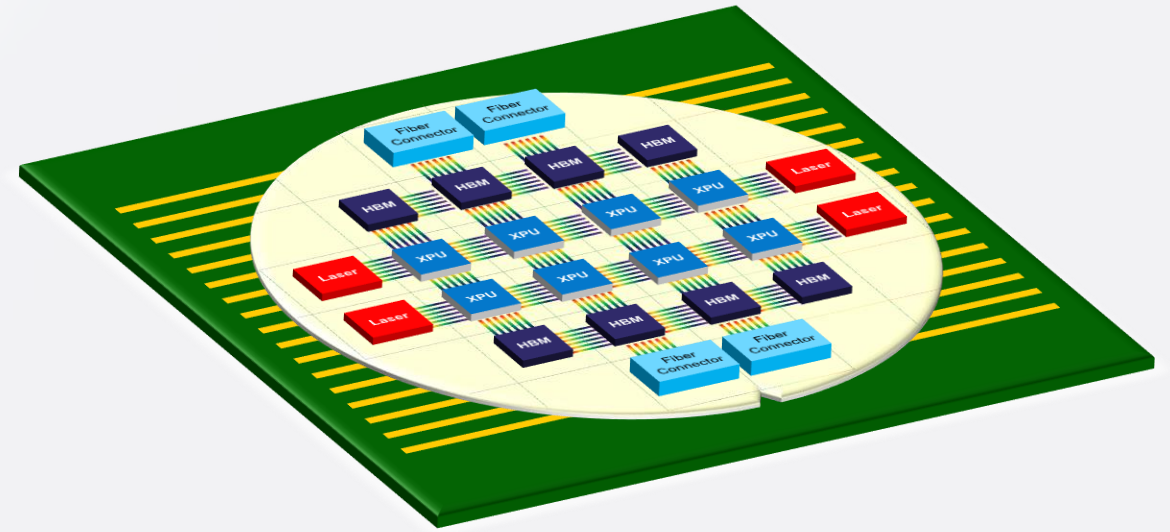
Packet-based communication

- Chiplet platforms are essentially miniature networks
- Processors already use Network-on-Chips that are packet / flit-based
- Migrate external interface from parallel to serial / packetized
- Global packet-based routing
 - networked communication from chip to chip
 - immense scalability thanks to packet switching and encapsulation
- *tasks* on chiplets communicate
 - Over multiple network protocols
 - With dynamically reconfigurable routes

Electrical to Optical

Packet-based communication fits optical platforms nicely

Power per transmitted bit is gaining on electrical



Side effects

- Chiplets likely run different “private” OS
- When combining multiple processors, how to control (shared) memory?
 - Software protocol?
 - Chiplet with memory controller?
 - Does this mean the “system” OS moves to this chiplet?
 - Distributed OS?
 - By the memory itself?
- Further de-coupling of hardware and software
 - Feature detection and enumeration
 - Virtualization of resources on the platform
 - Danger of fragmentation of code base!
 - More emphasis on JIT compilation

Reconfigurable Logic



Reconfigurable logic

- In time, yield, cost, thermals or mechanical constraints will limit the size of a system-in-package (or system-on-wafer), even when using chiplets
- Deciding which architecture to use is a risk due to the fast evolution of software
- High cost of designs will slow down architectural evolution
- Add reconfigurable logic
 - Allow to modify the architecture to fit the software
 - Reduce design time and cost – no tapeout of specialized circuits
 - Allow for post-shipment modifications, bug fixes, security fixes...
 - Start with adding e.g. an FPGA chiplet

Software-Defined Silicon (SDSi)

- Ultimately, offer a “canvas” that can implement arbitrary processors
- Canvas contains all the necessary building blocks or “functional grains”
 - Not bit-level, but intended for instantiating *processors*
 - Configuration of functional grains and flexible interconnection
- Canvas can dynamically reconfigure itself
 - JIT compilation of the architecture description
- Architectural description embedded in the application
 - Software developers choose the architecture – *code brings its own processor*
 - Processors become softcores, licensed to application developer
 - Dynamic adaptation of architecture during code execution

Software-Defined Silicon (SDSi)

- Reconfigurability leads to circuit-level overhead
 - Performance loss
 - Larger area
- Mitigate the losses by
 - Higher application performance due to architectural match
 - High-density design of the functional grains
 - Specializing the silicon technology to the properties of functional grains
- Technology is not intended for arbitrary digital circuits
 - Separate transistors properties according to the type of grain (dens, high-drive, SRAM, ...)_
 - Map a type of transistor to a separate 3D layer (CMOS 2.0)
 - Grains are 3D blocks spanning one or more layers of transistors



Conclusion

- Use chiplets to add more power-efficient, specialized functionality
- Migrate to global packet-based communication
- Move to optical communication between dies and packages
- Longer term, replace specialized chiplets by reconfigurable ones
- Ultimately, separate architecture from hardware implementation and move to specialized, layered silicon technology but fully flexible circuits
- Then repeat with Superconducting Logic

Thank you!

