SRC JUMP 2.0 Proposers Workshop
January 25, 2022

Theme 6: Advanced Monolithic and Heterogeneous Integration

Advanced Heterogeneous Integration

Prepared by Johanna Swan, Fellow, Intel Corporation

Adapted and presented by Henning Braunisch, Principal Engineer, Intel Corporation

The Future of Compute Performance

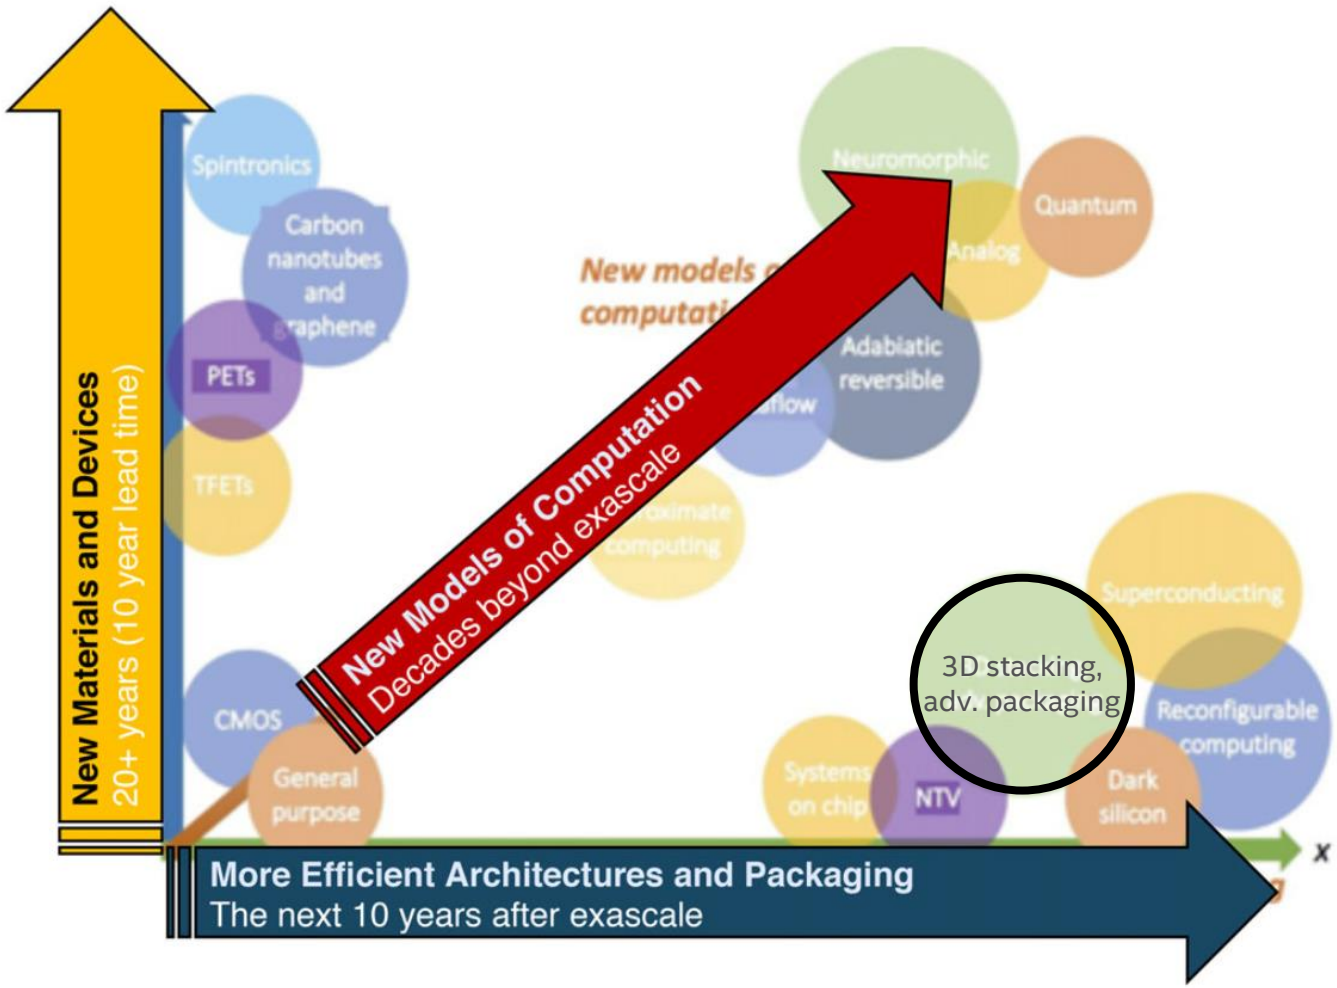

- Three Vectors
  - More efficient architectures and packaging leveraged for advanced heterogeneous integration
- 3D Stacking & Advanced Packaging
  - Focus on functional densification with performance, power, area, and cost

Ref.: J. Shalf, "The future of computing beyond Moore's Law," *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences,* 378 (2166), March 1, 2020.

Functional Densification: Intel's Ponte Vecchio

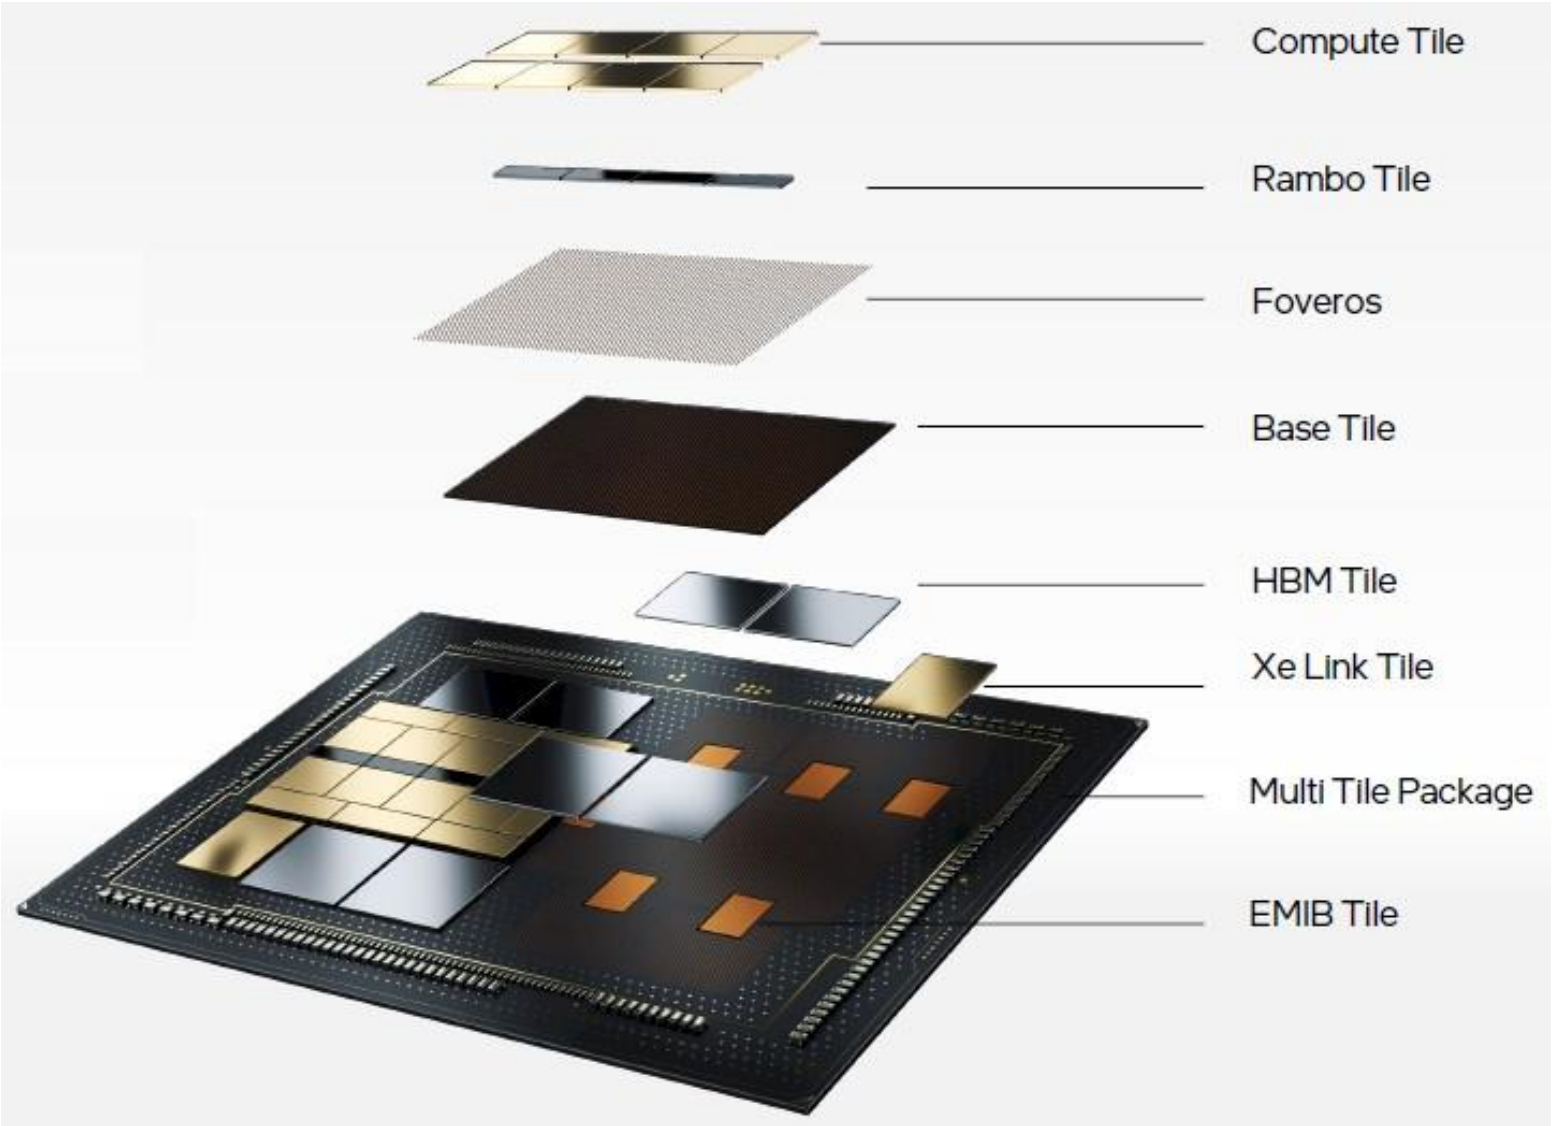

- 47 Tiles
- Five Process Nodes
- Foveros and EMIB Integration in the Same Package

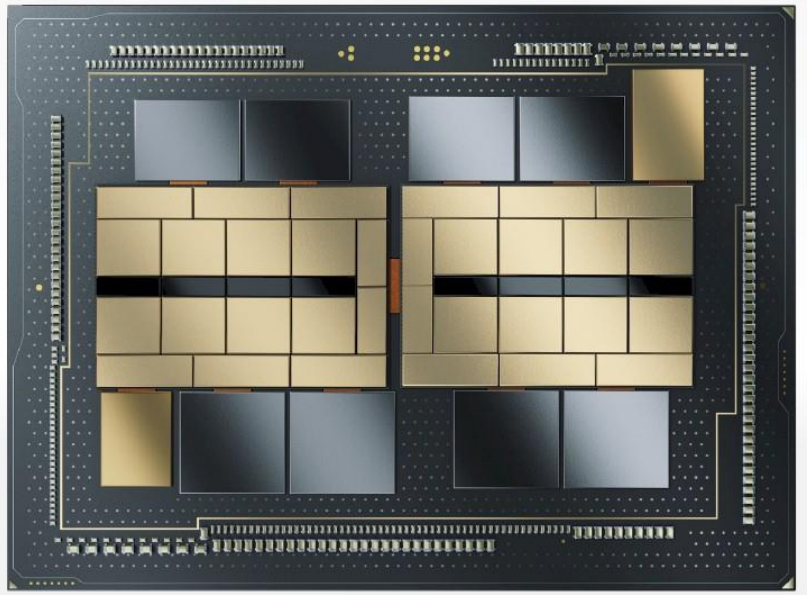

Challenge1: Harness & Exploit Increasing Design Complexity

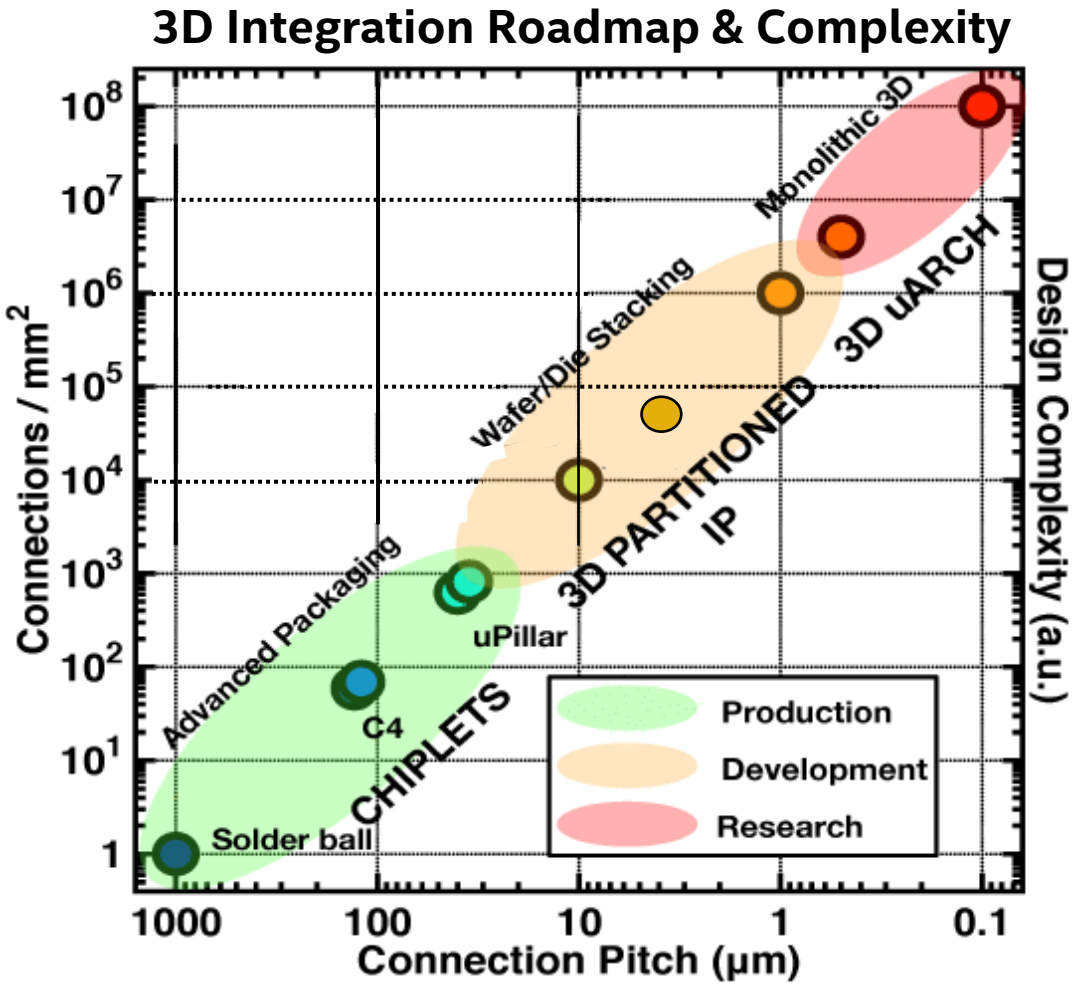

Ref.: ARM & GF, IEDM 2020 with adaptation

Solutions Needed

Full Stack Reference Designs With Performance Validation	At Scale i.e. mix-n-match ICs
Physical Implementation Flow with IEEE 1838 3DIC DFT Standard	Production Class Solutions Agnostic to Vendor
3D Design Space & System Level Co-Optimization	Opportunity to Drive Best In Class Solutions
Interface Standards	Proprietary & Open

- More tiles means even more complexity
  - Opportunity to leverage the Intel university shuttle program and driving the ecosystem with circuit, design & validation innovations

Imagine… AI Driven System Design

Article

A graph placement methodology for fast chip design

https://doi.org/10.1038/s41586-021-03544-w

Received: 3 November 2020

Accepted: 13 April 2021

Published online: 9 June 2021

Check for updates

Azalia Mirhoseini[1,4]✉, Anna Goldie[1,3,4]✉, Mustafa Yazgan[2], Joe Wenjie Jiang[1], Ebrahim Songhori[1], Shen Wang[1], Young-Joon Lee[2], Eric Johnson[1], Omkar Pathak[2], Azade Nazi[1], Jiwoo Pak[2], Andy Tong[2], Kavya Srinivasa[2], William Hang[3], Emre Tuncer[2], Quoc V. Le[1], James Laudon[1], Richard Ho[2], Roger Carpenter[2] & Jeff Dean[1]

Chip floorplanning is the engineering task of designing the physical layout of a computer chip. Despite five decades of research[1], chip floorplanning has defied automation, requiring months of intense effort by physical design engineers to produce manufacturable layouts. Here we present a deep reinforcement learning approach to chip floorplanning. In under six hours, our method automatically generates chip floorplans that are superior or comparable to those produced by humans in all key metrics, including power consumption, performance and chip area.

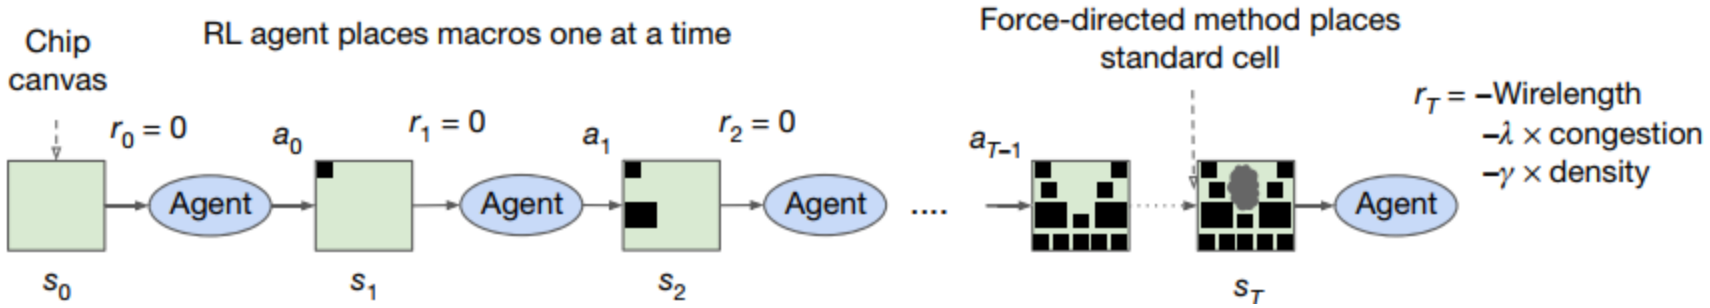

- Could a similar approach be leveraged to design & package heterogeneous systems? Considerations include…
  - Physical Design
  - Timing/Clocking Schemes
  - Thermal Management
  - Power Delivery
  - Mechanical Robustness through Process & Reliability Testing

Challenge 2: Find & Eliminate Impact of < 100 nm Defects

Hybrid Bonding Process

Dielectric Contact at T_{room} with van der Waals bond

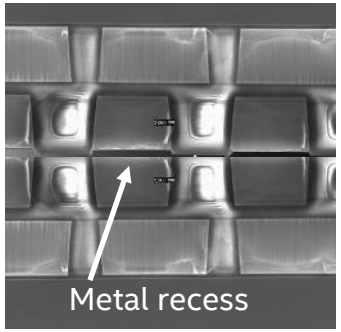

Dielectric Fusion at $T>T_{room}$ with Metal Gap

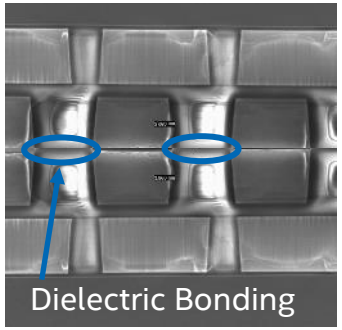

Metal Expansion & Connect at T_{hot}

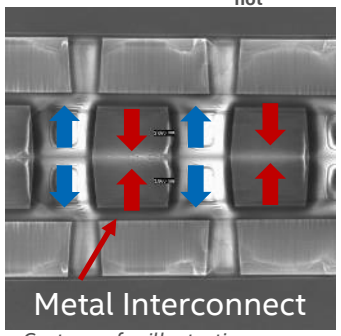

Cartoons for illustrative purposes

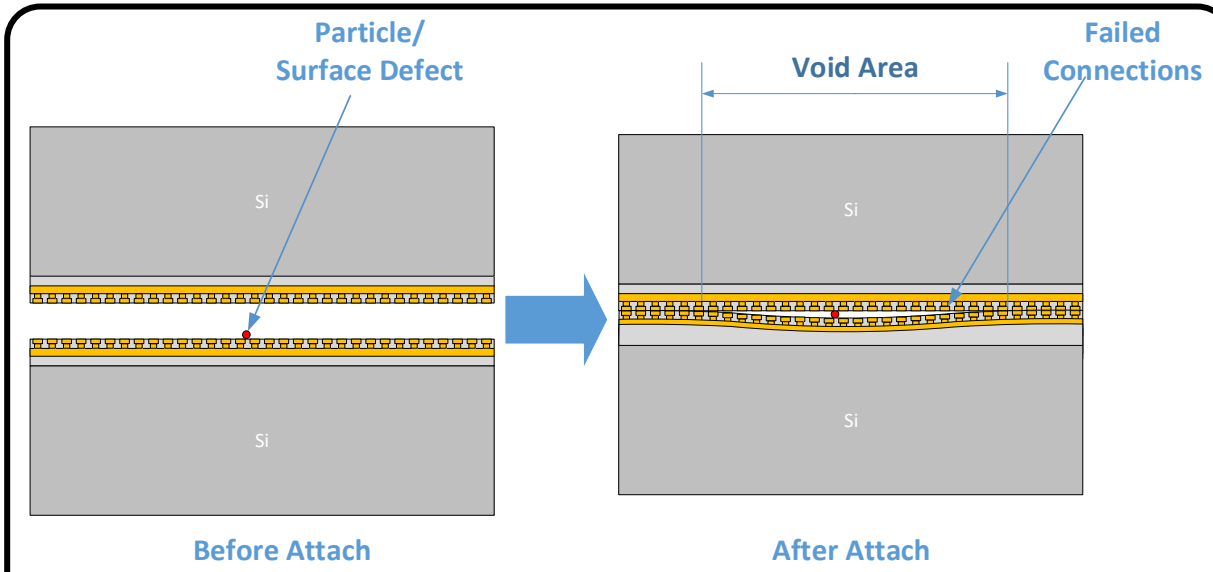

Voids due to particles/defects are 100–1000x larger

> 90 nm Debris Now Leads to Killer Defects at 9 µm Pitch

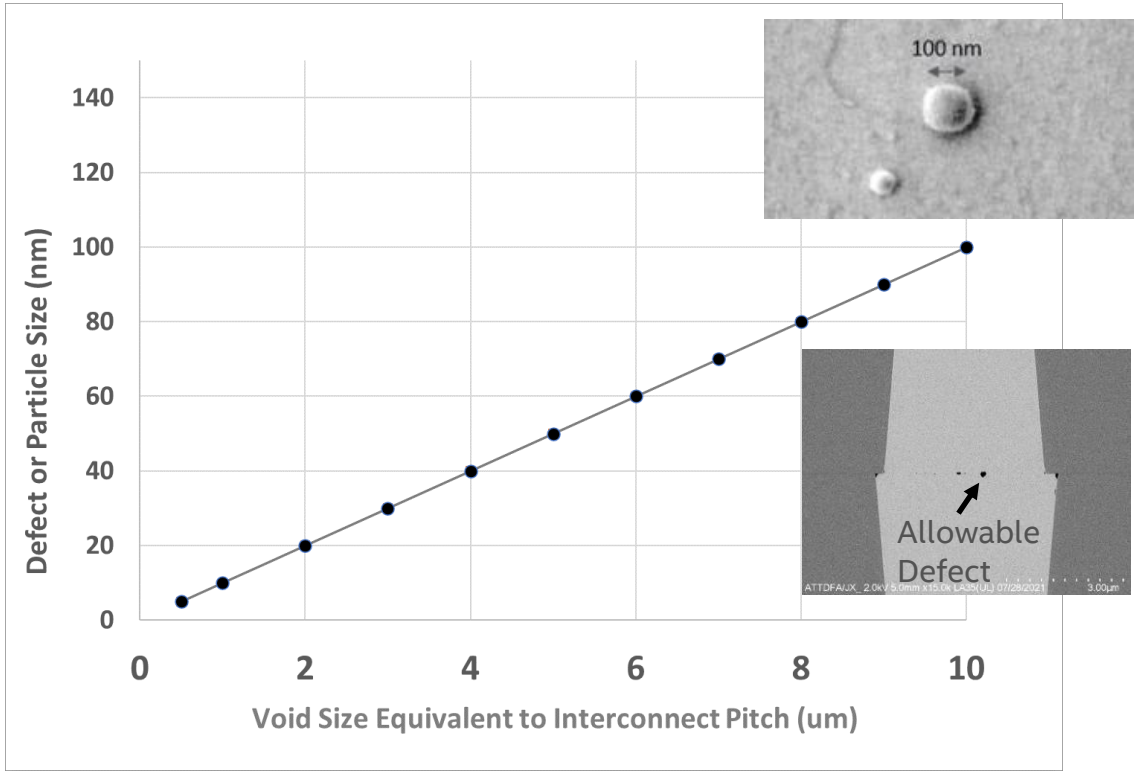

- Non-destructive, fast methods to find organic & inorganic defects before & after bonding to ensure yield and quality

Imagine… 3D Imaging Through Multi-Layer Metals

Articles
https://doi.org/10.1038/s41928-019-0309-z
nature electronics

Three-dimensional imaging of integrated circuits with macro- to nanoscale zoom

Mirko Holler[1*], Michal Odstrcil[1], Manuel Guizar-Sicairos[1], Maxime Lebugle[1], Elisabeth Müller[1], Simone Finizio[1], Gemma Tinti[1], Christian David[1], Joshua Zusman[2], Walter Unglaub[2], Oliver Bunk[1], Jörg Raabe[1], A. F. J. Levi[2] and Gabriel Aeppli[1,3,4]

The imaging of integrated circuits across different length scales is required for failure analysis, design validation and quality control. At present, such inspection is accomplished using a hierarchy of different probes, from optical microscopy on the millimetre length scale to electron microscopy on the nanometre scale. Here we show that ptychographic X-ray laminography can provide non-destructive, three-dimensional views of integrated circuits, yielding both images of an entire chip volume and high-resolution images of arbitrarily chosen subregions. We demonstrate the approach using chips produced with 16 nm fin field-effect transistor technology, achieving a reconstruction resolution of 18.9 nm, and compare our results with photolithographic mask layout files and more conventional imaging approaches such as scanning electron microscopy. The technique should also be applicable to other branches of science and engineering where three-dimensional X-ray images of planar samples are required.

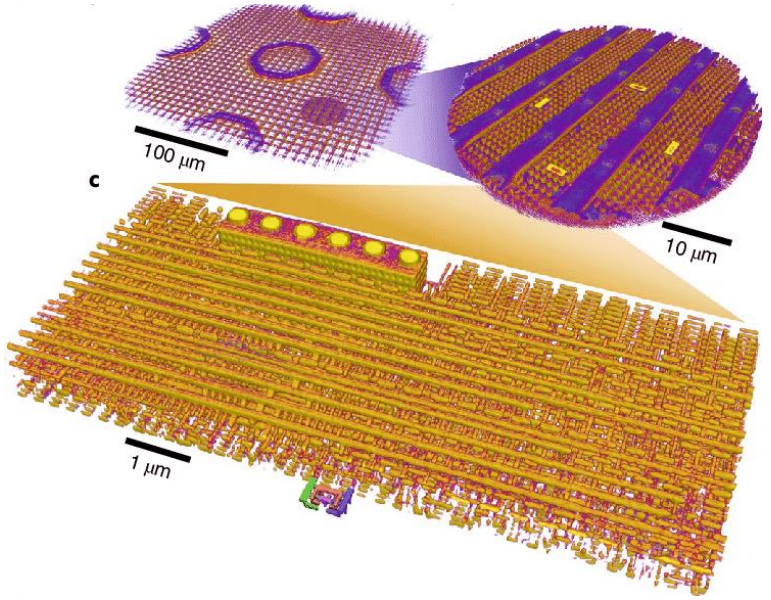

- Need large-area scanning at production ready speeds with nano-scale resolution through multiple metal layers without impacting device performance

Challenge 3: Place 20k+ Tiles per Hour at < 200 nm Accuracy

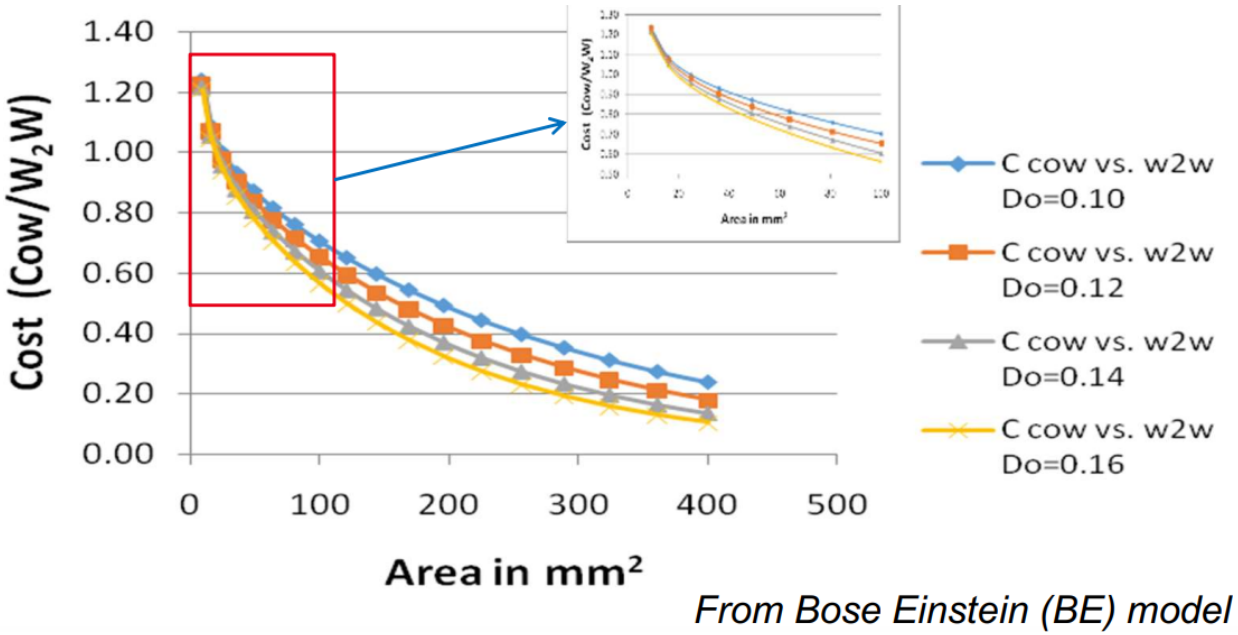

From Bose Einstein (BE) model

Ref.: D. Gitlin et al., "Generalized cost model for 3D systems," *2017 IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S)*

- Die to wafer bonding may be cost effective at large die sizes relative to wafer bonding especially at high defect densities
- As the die size reduces to less than ~30 mm² or the number of tiles per base tile increases, the cost of assembly increases unfavorably without further assembly advancements

Die-to-Wafer Equipment

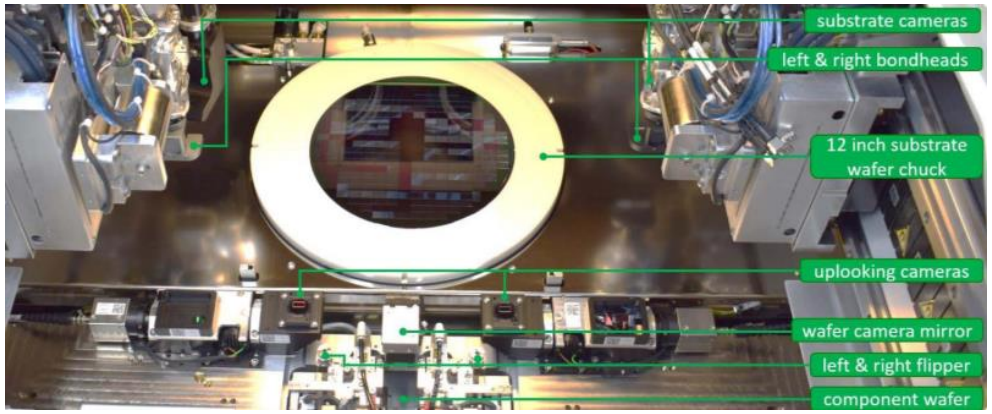

Nearly 200 nm Accuracy at 2000 Units Per Hour

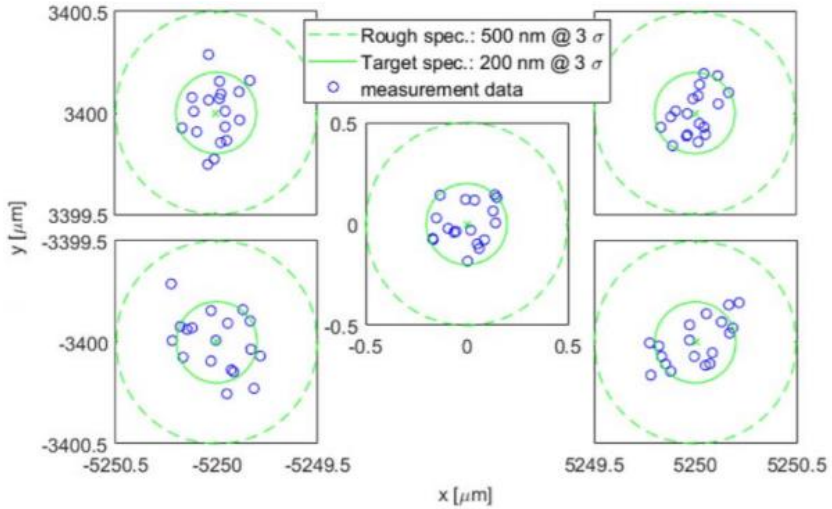

Ref.: B. Brandstätter et al., "High-speed ultra-accurate direct C2W bonding," *2020 IEEE 70th Electronic Components and Technology Conference (ECTC)*

Imagine… Nanoscale Self-Assembly of Dies in Batch

IEEE TRANSACTIONS ON COMPONENTS, PACKAGING AND MANUFACTURING TECHNOLOGY, VOL. 1, NO. 12, DECEMBER 2011 1873

Multichip Self-Assembly Technology for Advanced Die-to-Wafer 3-D Integration to Precisely Align Known Good Dies in Batch Processing

Takafumi Fukushima, *Member, IEEE*, Eiji Iwata, Yuki Ohara, Mariappan Murugesan, Jichoel Bea, Kangwook Lee, *Member, IEEE*, Tetsu Tanaka, *Member, IEEE*, and Mitsumasa Koyanagi, *Fellow, IEEE*

An advanced die-to-wafer 3-D integration using a surface-tension-driven multichip self-assembly technology was proposed to 3-D stack a large number of known good dies (KGDs) in batch processing. The parallel self-assembly with a unique multichip pick-up tool was newly applied to die-to-wafer 3-D integration to overcome throughput and yield problems in conventional 3-D integration approaches. In addition, novel batch transfer of chips self-assembled on a carrier wafer to the corresponding target wafer was demonstrated. By using the multichip self-assembly, many KGDs can be precisely aligned and temporarily placed on a carrier wafer all at once, and then, the self-assembled KGDs can be simultaneously transferred to another target wafer in a face-to-face bonding manner at the wafer level. Average alignment accuracy was found to be approximately 400 nm when a hundred 3-mm-square chips were self-assembled on carrier wafers with small droplets of an aqueous solution. The alignment accuracy was experimentally proven to be fairly dependent on liquid surface tension as a self-assembly parameter.

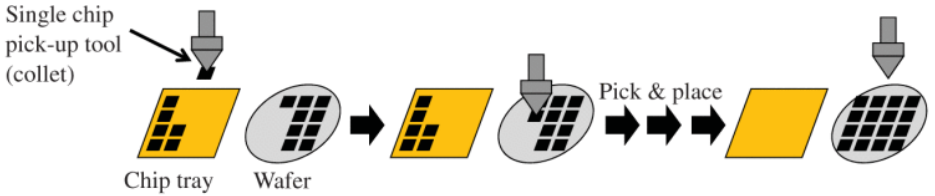

	Equipment A [18]	Equipment B [19]	Equipment C [20]	Equipment C [21]
Assembly throughput:	1.7 sec/chip* (~2 100 chips/hour)	0.72 sec/chip** (2 × 5000 chips/hour)	1.0 sec/chip (3.600 chips/hour)	0.29 sec/chip (~12 400 chips/hour)
Alignment accuracy:	± 2 μm	± 10 μm	± 10 μm	± 25 μm
Bonding style:	Flip-chip bonding	Flip-chip bonding	Die bonding (thermal compression)	Die bonding (adhesive bonding)

* This cycle time dose not include processing time such as loading, bonding, and vacuum release.
**This equipment having two bond heads can offer two machines in just one compact module.

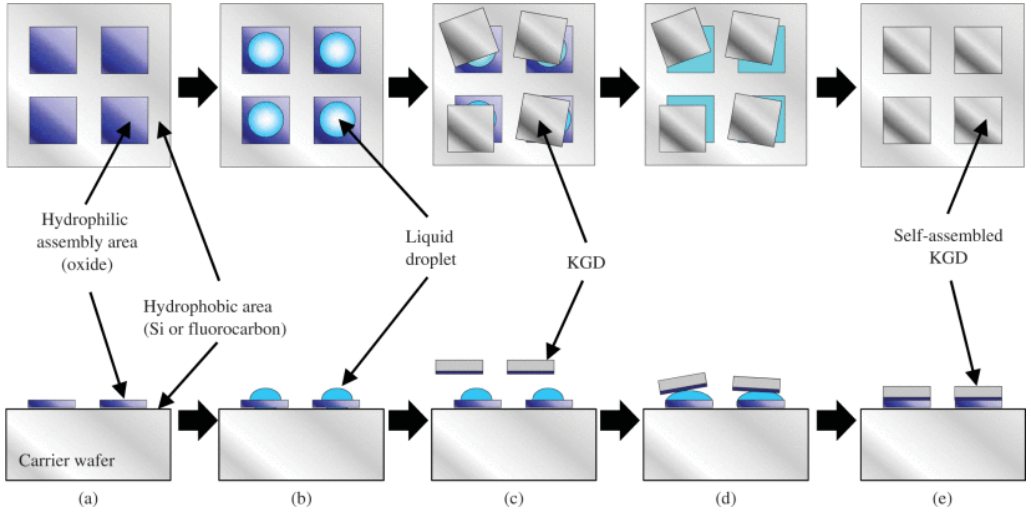

- Equipment & material advancements in combination with self-assembly are needed to achieve 10x faster die placements

Summary & Conclusions

- We are at the beginning of a period where the future of compute performance relies on efficient advanced 3D packaging of extremely heterogeneous systems
- Numerous challenges face the industry. A sampling includes:
  1. Harnessing & exploiting increasing design complexity
  2. Finding and eliminating the impact of < 100 nm defects non-destructively
  3. Placing 20k+ tiles per hour at < 200 nm accuracy
- This is an exciting time when fab processes & methodologies are colliding with traditional packaging know-how
  - **JUMP 2.0 Theme 6** provides an excellent opportunity for impactful research in this space

SRC JUMP 2.0 Proposers Workshop
January 25, 2022

Theme 6

Advanced Monolithic and Heterogeneous Integration

Marko Radosavljevic, Principal Engineer, Intel Corporation

intel®

Exploding design complexity of the future

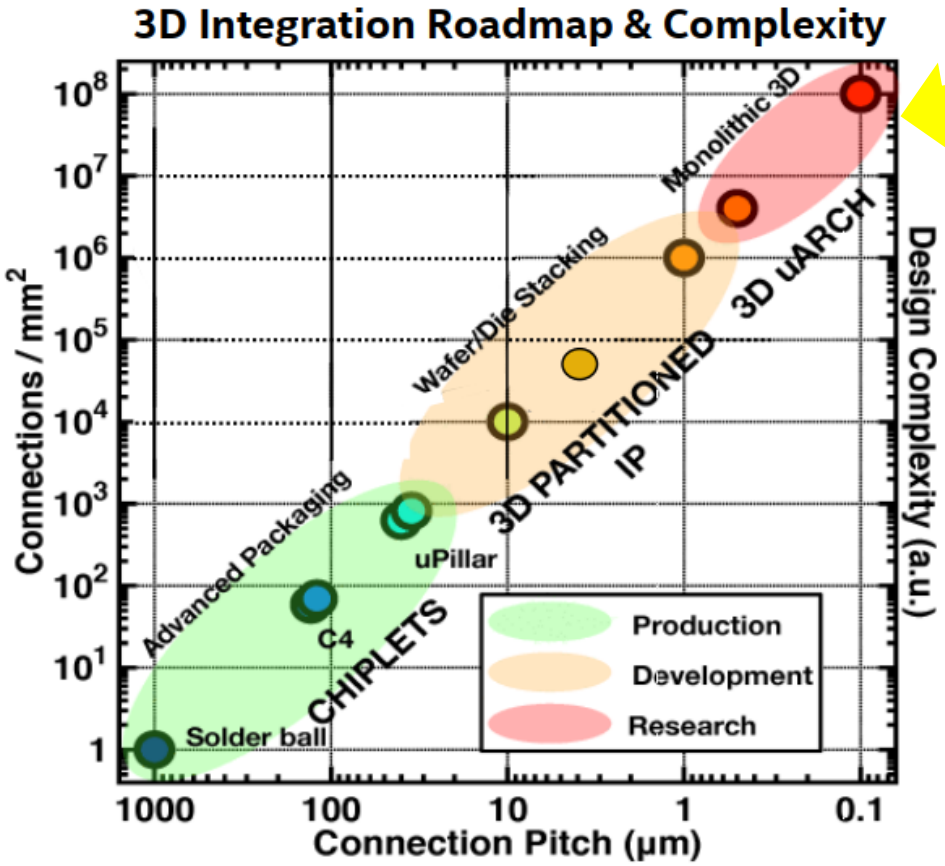

Ref.: ARM & GF, IEDM 2020 with adaptation

- Monolithic 3D provides most connections
  - Makes tiles more complex but perhaps reduces number of tiles
  - Correct choice for each application set?

Monolithic 3D with heterogeneous materials for heterogeneous applications

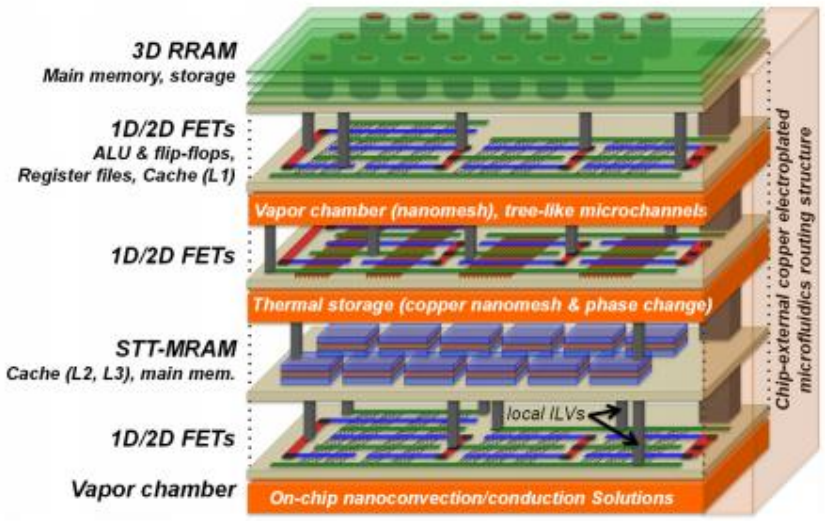

Ref.: Stanford SystemX

- Monolithic 3D but each layer virtually independent
  - Thermal budget considerations
  - Bottom up integration or layer transfer techniques
  - Is there tight pitch requirement → good inter-layer alignment
  - Infinite choice of applications
  - Differentiate vs. packaging techniques

Monolithic 3D

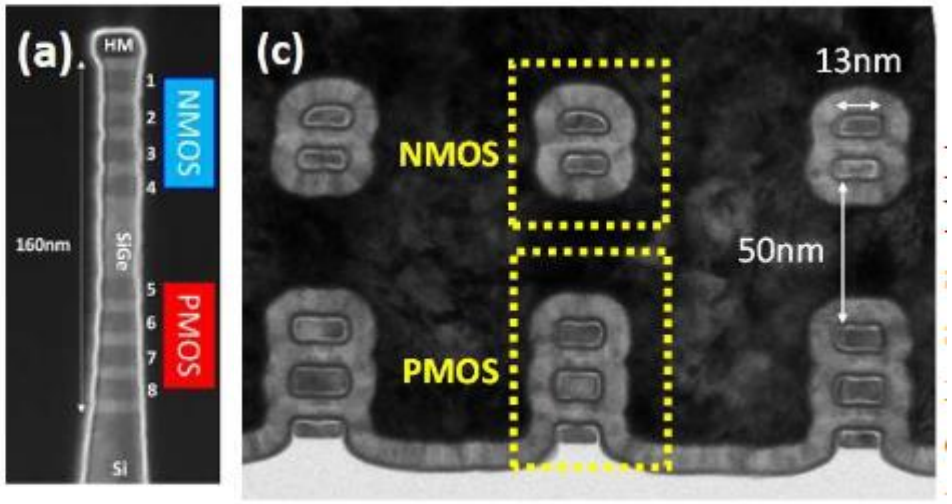

Ref.: Intel, 2021 IEDM

- Layers co-processed simultaneously
  - High aspect considerations (etching, filling,…)
  - Potentially self-aligned layers
  - Layer re-use/dual role
  - Combination with bottom up techniques and layer transfer
  - Compatible with multiple applications on same tile

Summary & Conclusions

- Exciting time to develop technologies that bring 3D integration to ultimate scaling limit
- Challenges abound:
  - Efficiently designing complex systems with novel functionalities
  - Utilizing new materials to enable multiple applications on same tile or in system
  - Utilizing novel techniques to enable high aspect ratio and inter-layer tight pitch
- JUMP 2.0 Theme 6 provides excellent opportunity for impactful research in this space


intel