What is the right Die-to-Die Interface? A Comparison Study

Shahab Ardalan (LMNS)
Bapi Vinnikota (BRCM)
Tawfik Arabi (AMD)
Elad Alon (BCA)

Outline

- Updated results to 2020
- First analyzed D2D interfaces in 2019
- Precise definition of all parameters
- Dropped some categories

Die to Die Interface

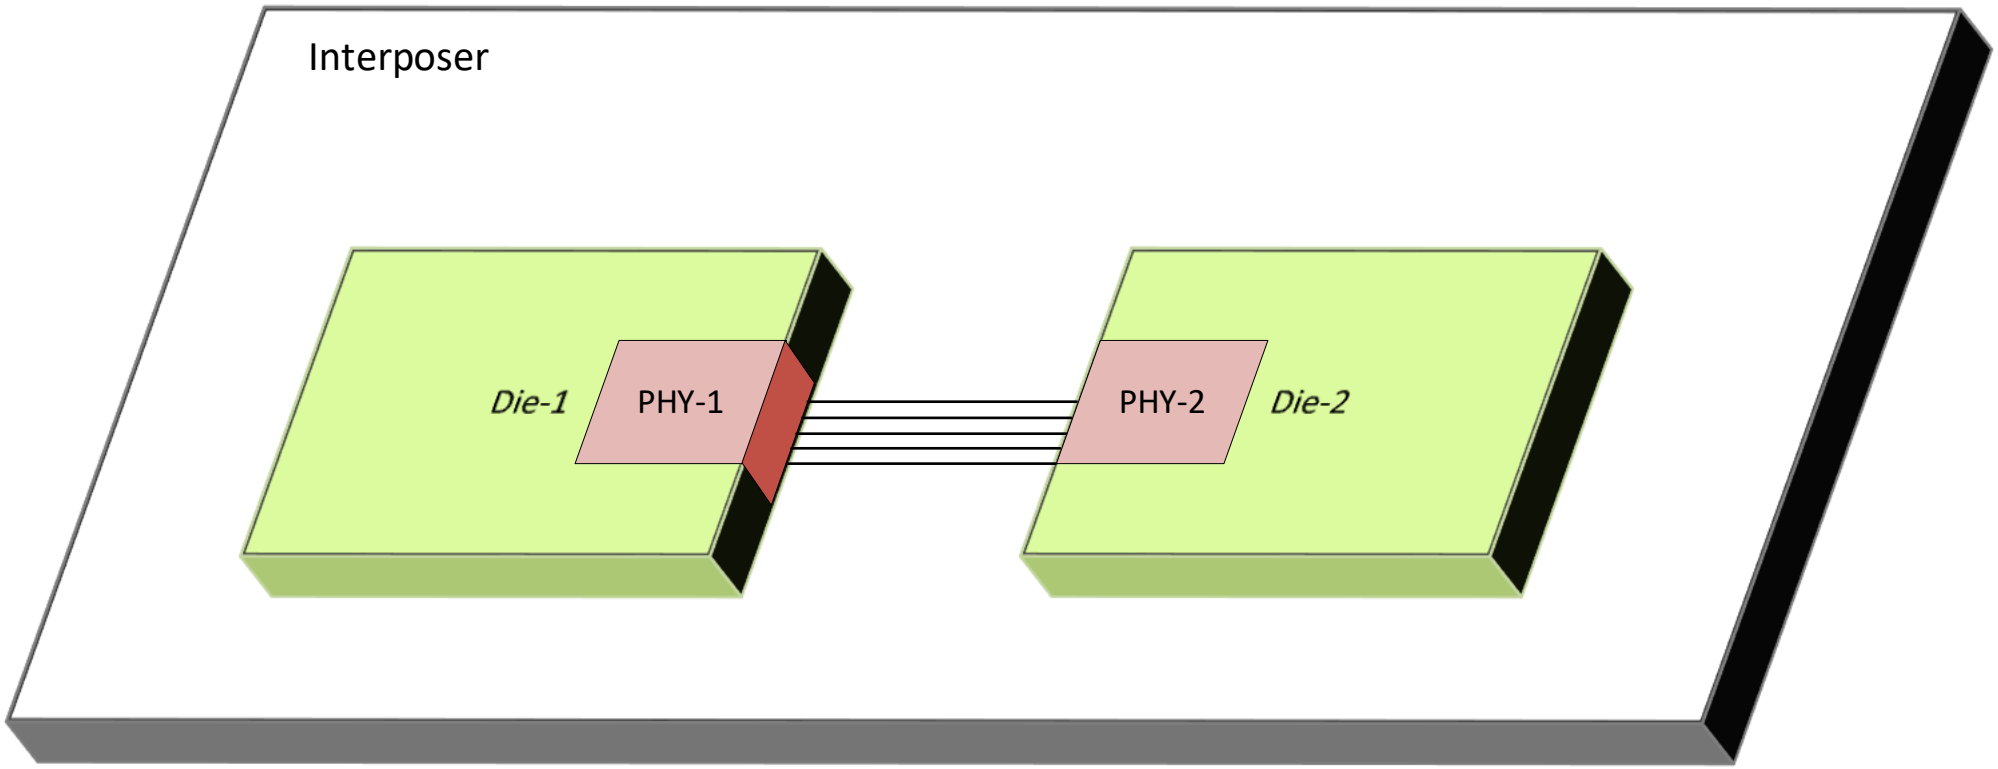

Process

- Open call to participate
- Active participation from: Alphawave, AnalogX, BoW workstream, Kandou, Marvell, Open HBI workstream, SiFive (later)
- From public information: Intel AIB, PCIe (reference, thanks to Toru Takaishi (Socionext))
- Used last year’s spreadsheet as a starting point.
- Lots of meetings. It’s important! Everyone cares deeply.

Metrics

Metric	Impact	2019 Components	2020	2022
Piece Cost per Unit	Manufacturing cost/complexity	Pad-limited area Substrate		
Operating Cost	Power-performance at rated throughput	Figure of Merit (Tb/mm)/(pJ/bit)	FOM2 = FOM/reach Latency at 1GHz logic at 1E-15 BER	Updated FoM: including Cost, BER, & reach
~~Product Risk~~ Design Impact	Chiplet/product design NRE/schedule risk	Routing Freedom ~~Low power states~~ Process Node Diversity IP Integration Complexity Production test/assembly		
~~Interface Risk~~ Product Integration	Interface technology NRE/schedule risk	Licensing Fee ~~Multi-sourcing~~ Interface test, assembly IP dev/port complexity	Open standard	IP Maturity and availability
Packaging Cost	Manufacturing Cost			Differentiation between 2D, 2.XD and 3D integration

Unidirectional vs. Bidirectional Links

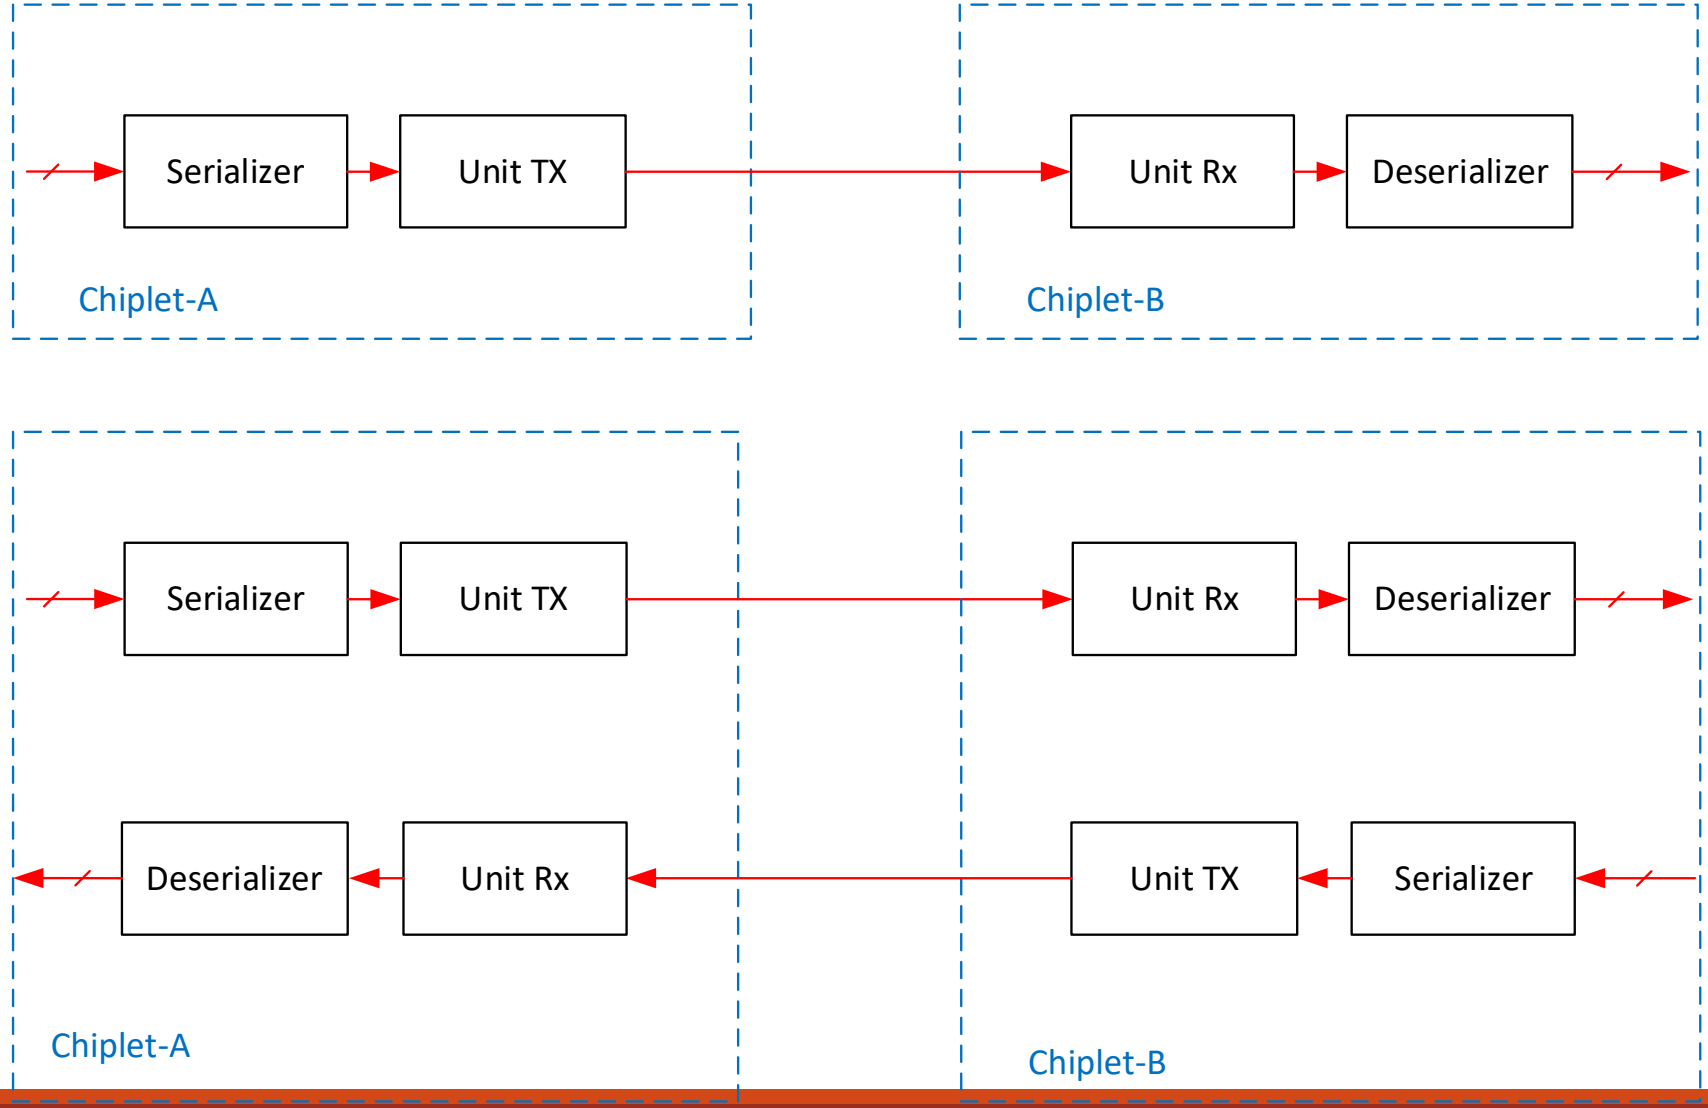

Section 1: Unit Unidirectional simplex link

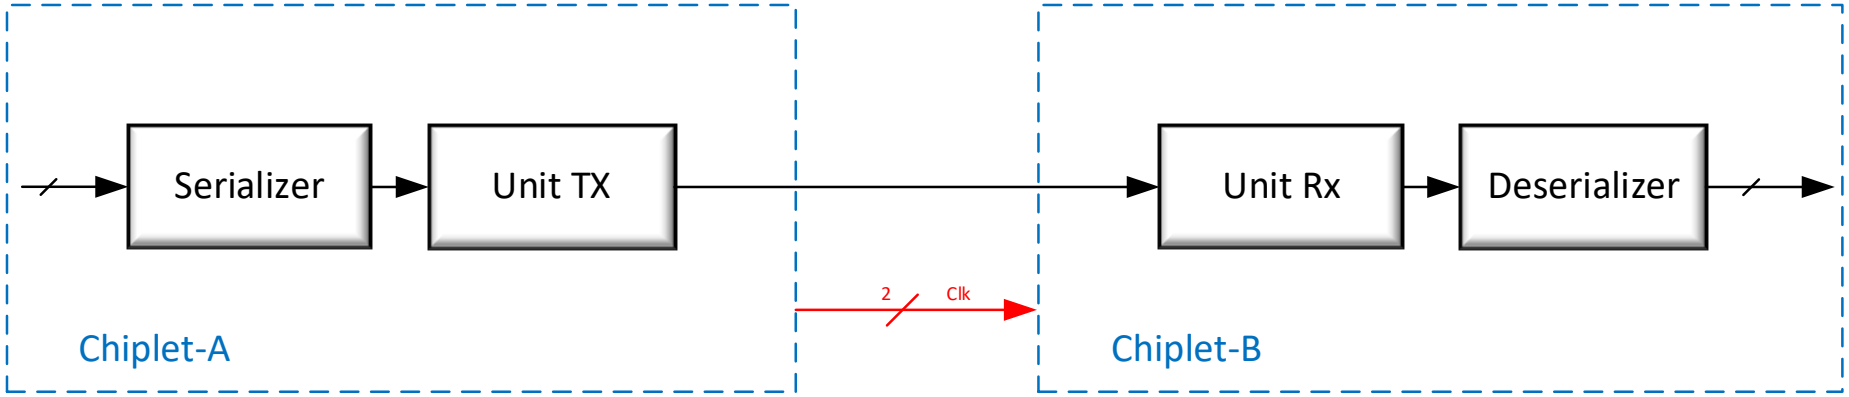

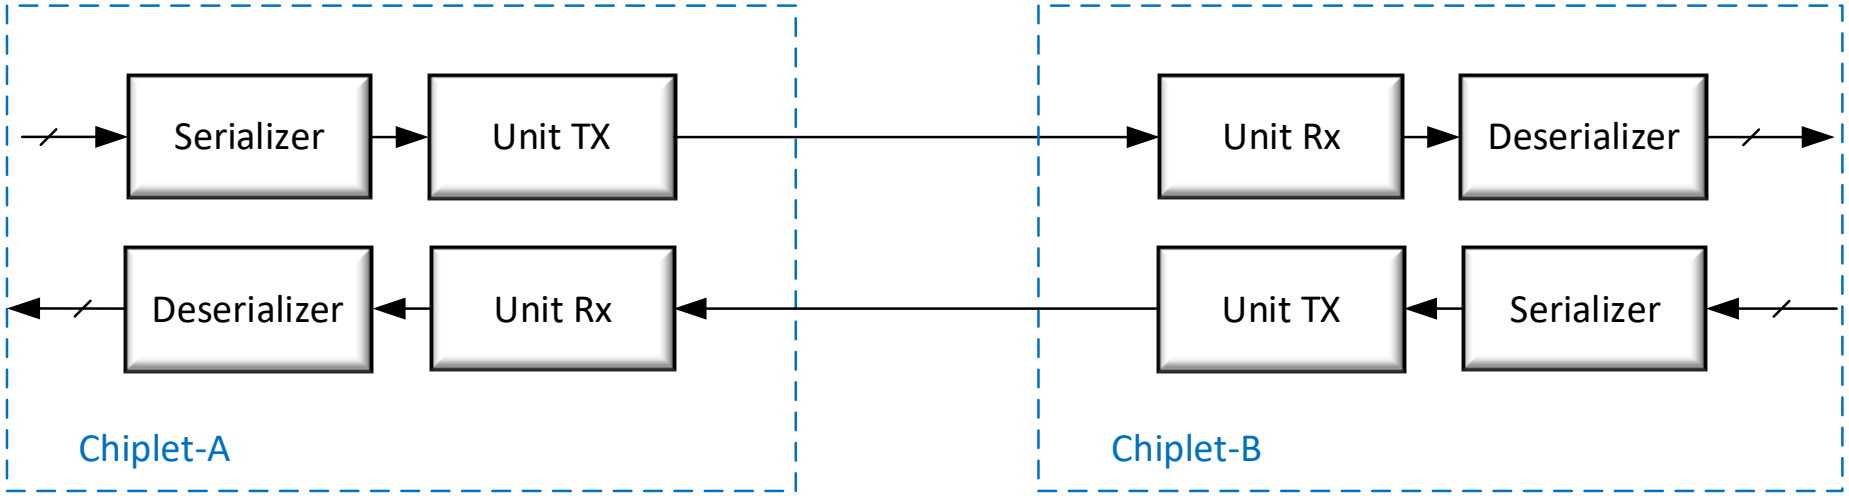

Reporting ½ of Area

Bow Link

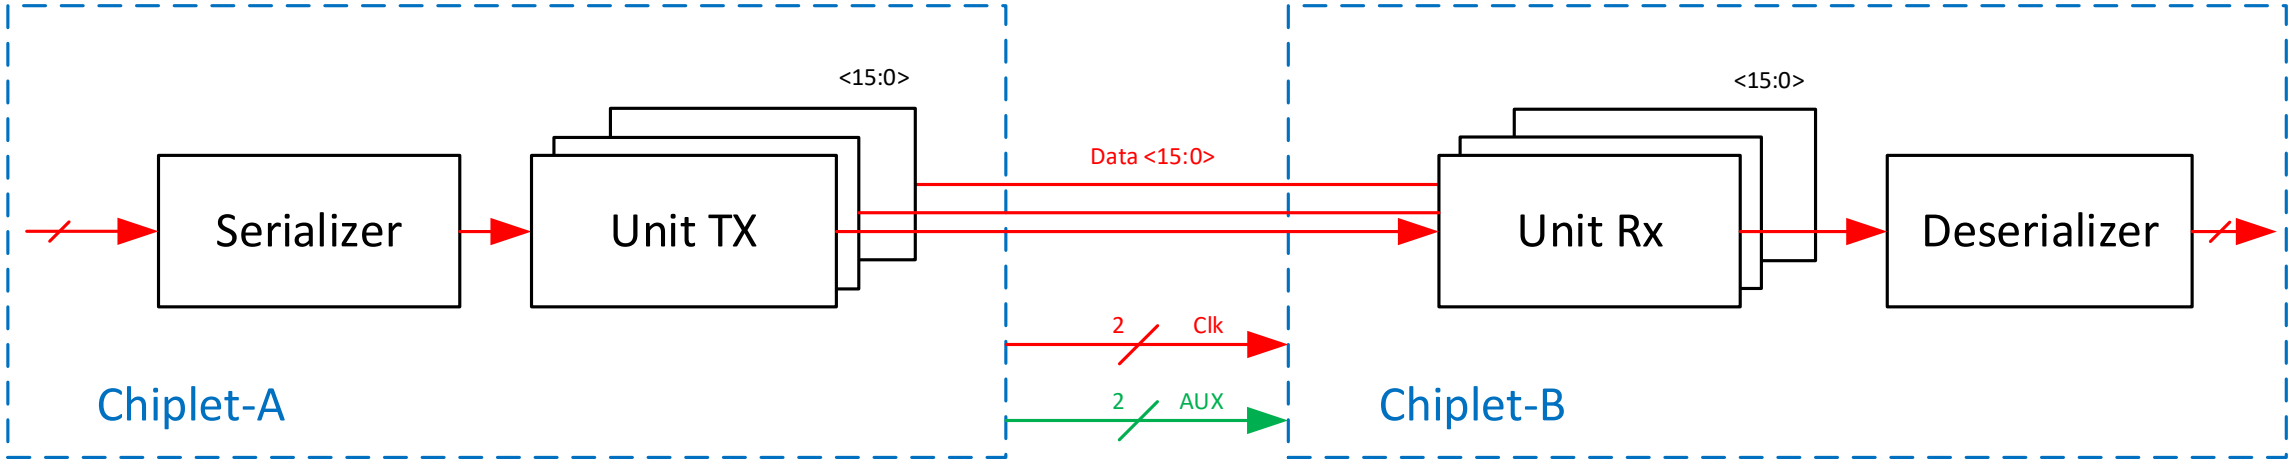
<15:0>
<15:0>
Serializer
Unit TX
Data <15:0>
Unit Rx
Deserializer
2
Clk
2
AUX
Chiplet-A
Chiplet-B

Section 2a: Unit simplex link parameters

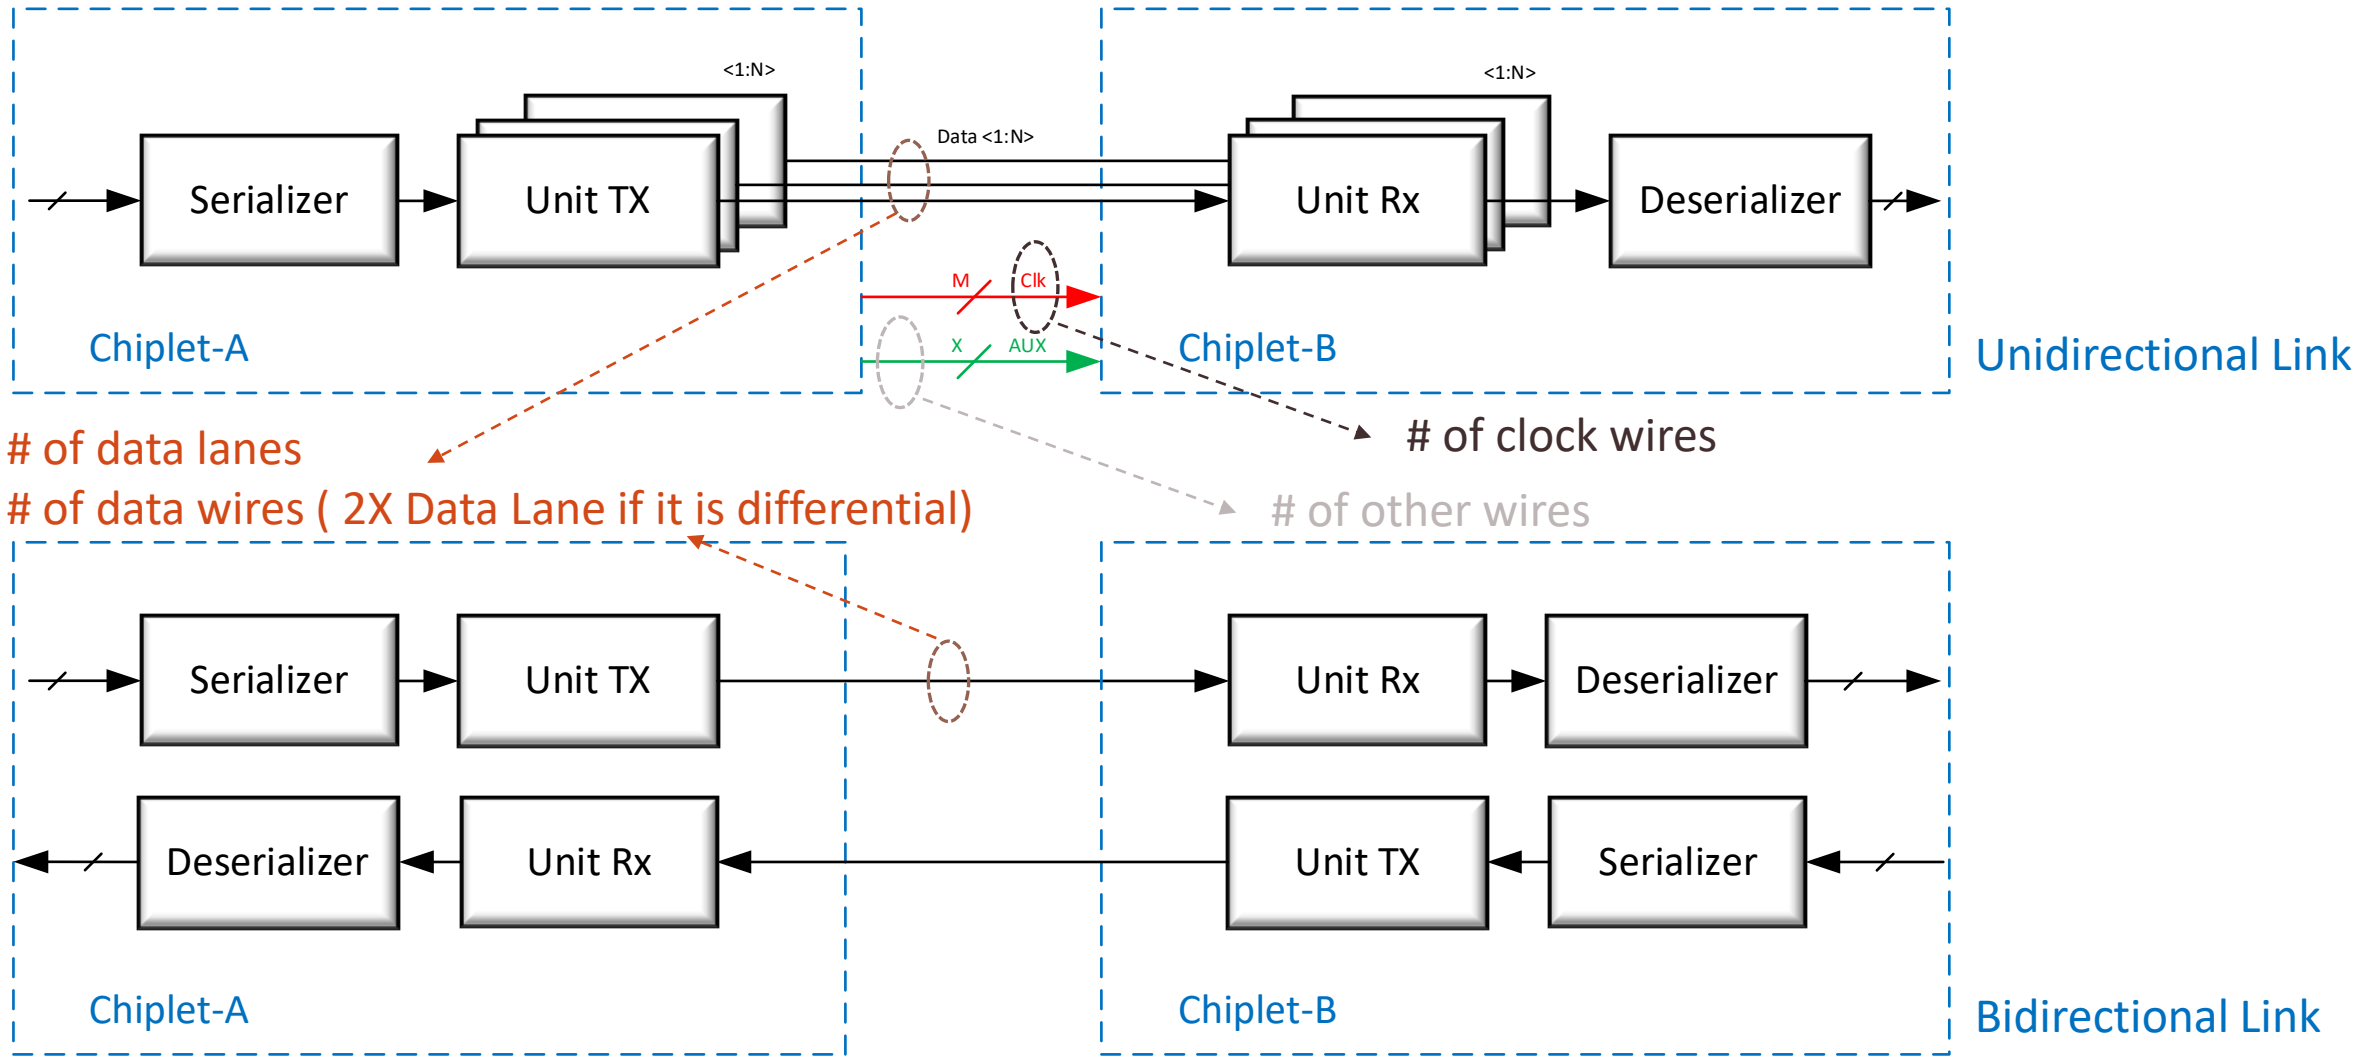

Section 2b: On-die parameters*

Performance:

Unit bandwidth (unit_bw) in each direction	Gbps	(speed_per_data_lane) *(num_data_bits) Examples: BoW: 16*5 = 80Gbps XSR: 112 X 1 = 112Gbps
Unit Power (Tx+Rx+Clock)	mW	Tx, Rx and Clocking power [mW/Gbps X unit_BW in Gbps] Example BoW: 80Gbps X 0.5mW/Gbps = 40mW XSR: 112Gbps X 1mW/Gbps = 112mW

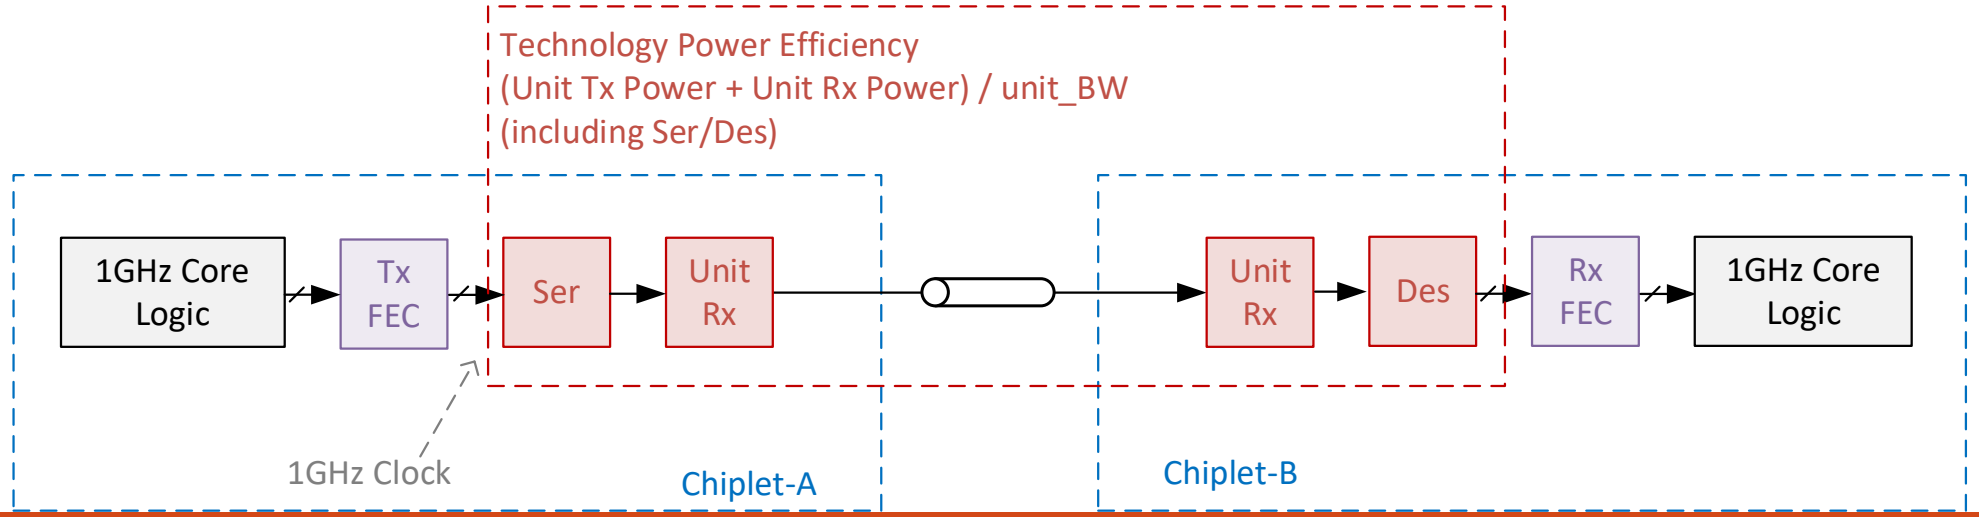

*estimated at 7nm

Section 2b: On-die parameters*

Delay / Latency:

Latency Without-FEC	In ns, from Tx MAC output to Rx MAC Input @1GHz core logic clock
BER Without FEC	Expected BER without error correction
Latency at 1E-15 BER	In ns, from Tx MAC output to Rx MAC Input @1GHz core logic clock

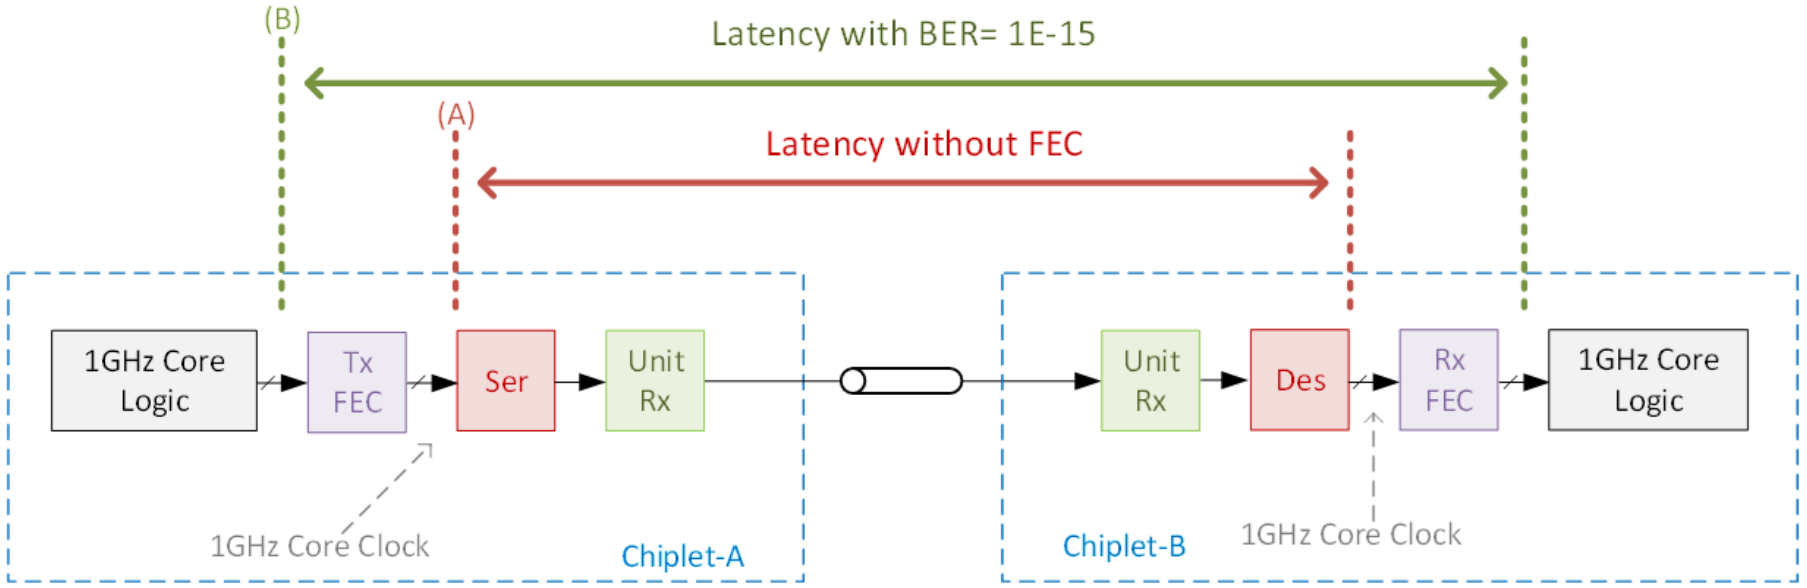

*estimated at 7nm

Area & reach definition

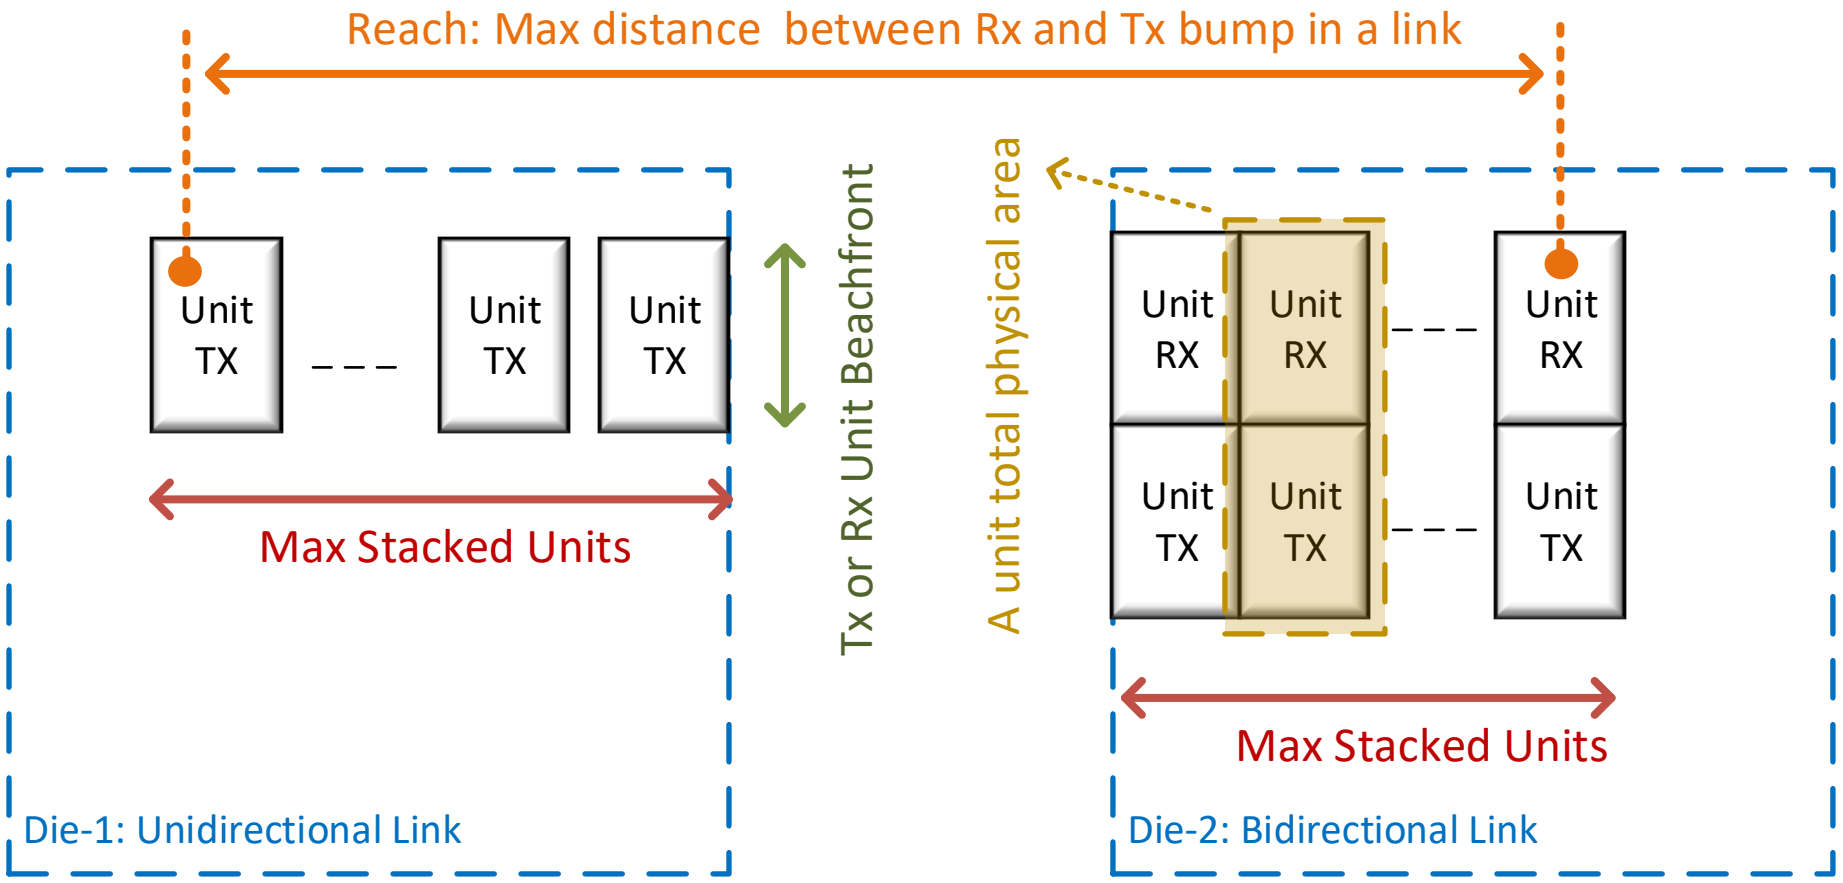

Performance Comparison

	Pump Space [um]	Power Efficiency [pJ/bit]	Edge Density [Tbps/mm]	Area Density [Tbps/mm^2]	FOM-1: Edge_Density/Power_Eff. [Tbps/mm / pJ/bit] Larger is better	FOM-2: Power Efficiency per Reach [pJ/bit / mm] Smaller is better
AIB 2.0	55	0.5	1.64	-	3.28	0.1
Open HBI 1.0	40	0.4	2.29	2.04	5.71	0.1
BoW- 256	130	0.55	0.88	0.69	1.59	0.01
BoW- 128 / Micro	40	0.5	1.78	1.07	3.56	0.1
AQ LinkP	130	0.55	1.91	1.46	4.48	0.01
AX-C2C	130	1.5	1.09	0.27	0.73	0.01
AX-DieIO	130	0.8	2.19	0.4	2.74	0.01
Kandou / CNRZ-5	130	1	1.33	0.67	1.33	0.02
XSR /Alphawave	130	1	3.98	0.51	3.98	0.02
PCIe Gen5	150	7	0.22	0.1	0.03	0.05

Learnings

- Power is important (pJ/bit): Numbers in the spreadsheet represent typical projected power. Expect variation across PVT (Process, Voltage, Temperature).
- Defining latency is very hard, we worked hard to minimize room for specmanship. For a target latency of 1E-15, FEC may dominate latency. Actual FEC impact is PHY specific.

Call to Action

- We have a workable set of unambiguous definitions. You're welcome to add your PHY to this list.
- Tuesdays @ 9:00, please check OCP event calendar