



OPEN
Compute Project

BoW: A Die-to-Die Interface Solutions Specification Update


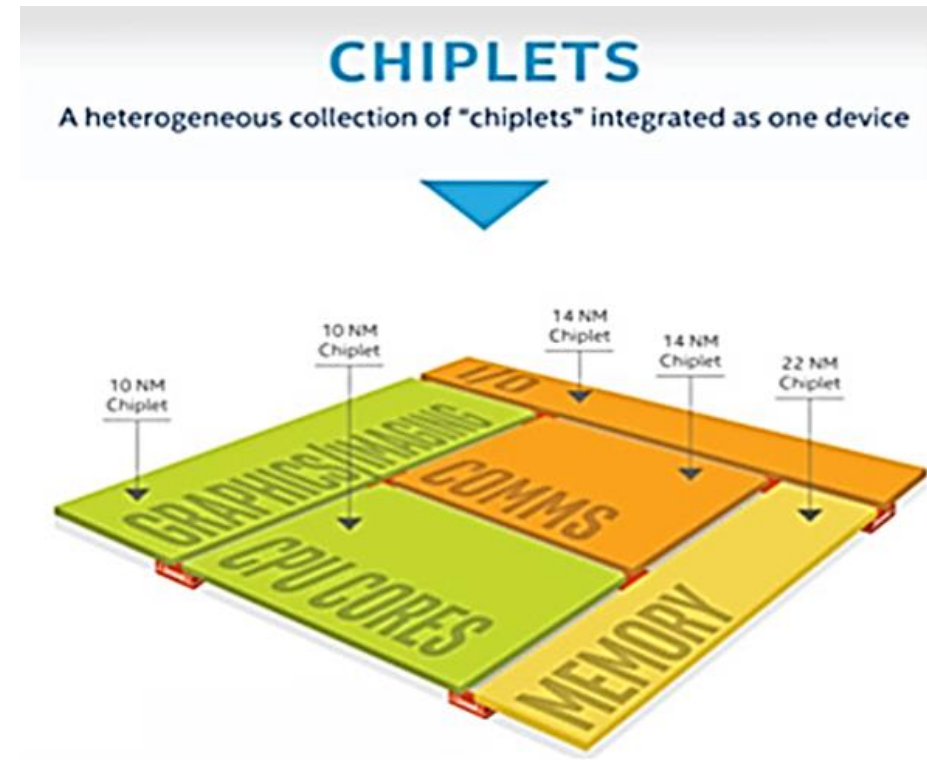
ODSA Regional Summit
September 27, 2019
Amsterdam, Holland

Ramin Farjadrad
Mark Kuemerle
Bapi Vinnakota

Why Heterogeneous

- Die disaggregation
 - Partition large die to improve yield
 - Reduce Time-to-Market
 - Use optimized process technology per functions
 - Reduce development cost in advanced nodes
- Package aggregation
 - Smaller footprint
 - Lower power
 - Lower Cost

→ Need new on-package inter-die connectivity



High-level Targets for Inter-Die Connectivity IP in MCM Packages

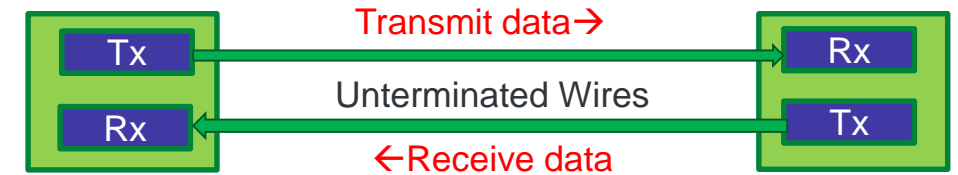
- Performance Targets:
 - Throughput Efficiency → 100Gbps/mm to beyond 1Tbps/mm (die edge)
 - Energy Efficiency → 1pJ/bit - 0.5pJ/bit for typical distance
- Small silicon area per IO port for dense integration
 - Goal: IP silicon area not to limit IO density
- Minimal analog/complex circuitry to offer easy/fast process porting
 - Limit the maximum baud rate of the interface → ~10G-16Gbaud
 - Common clock to replace CDR with DLL, Digital PLL to generate clock, etc
- Single supply IP supporting wide V_{dd} range: 0.70V – 0.9V
 - To be compatible with most existing SoC/ASIC in popular/available process nodes
- Logical compatibility with other simple interfaces

What's new with BoW?

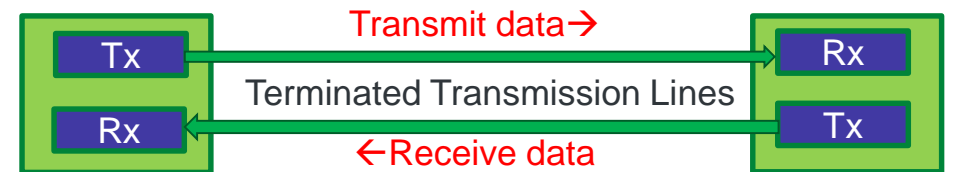
- Releasing v0.7 of specification to solicit input from the industry
 - <https://www.opencompute.org/wiki/Server/ODSA>
 - Early suggestions already pointing to beneficial updates
- Redefining 'modes'
 - BoW Base is now "BoW"
 - BoW Fast becomes termination mode of BoW
 - More modes are possible, as long as base "BoW" mode is supported
- Allow BoW to run at 8Gbps with extremely short interconnect
 - Compatibility of "BoW" does not require supporting such high data rates
 - Longer distances would require termination mode to be enabled
- Added Channel Specifications

BoW (ODSA Proposal): Single-ended Signaling

- BOW
 - Unterminated lanes → up to 5 Gbps/wire up to 10mm
 - Higher data-rate support up to 8Gbps for very short interconnect
 - DDR Source Synchronous with clock alignment

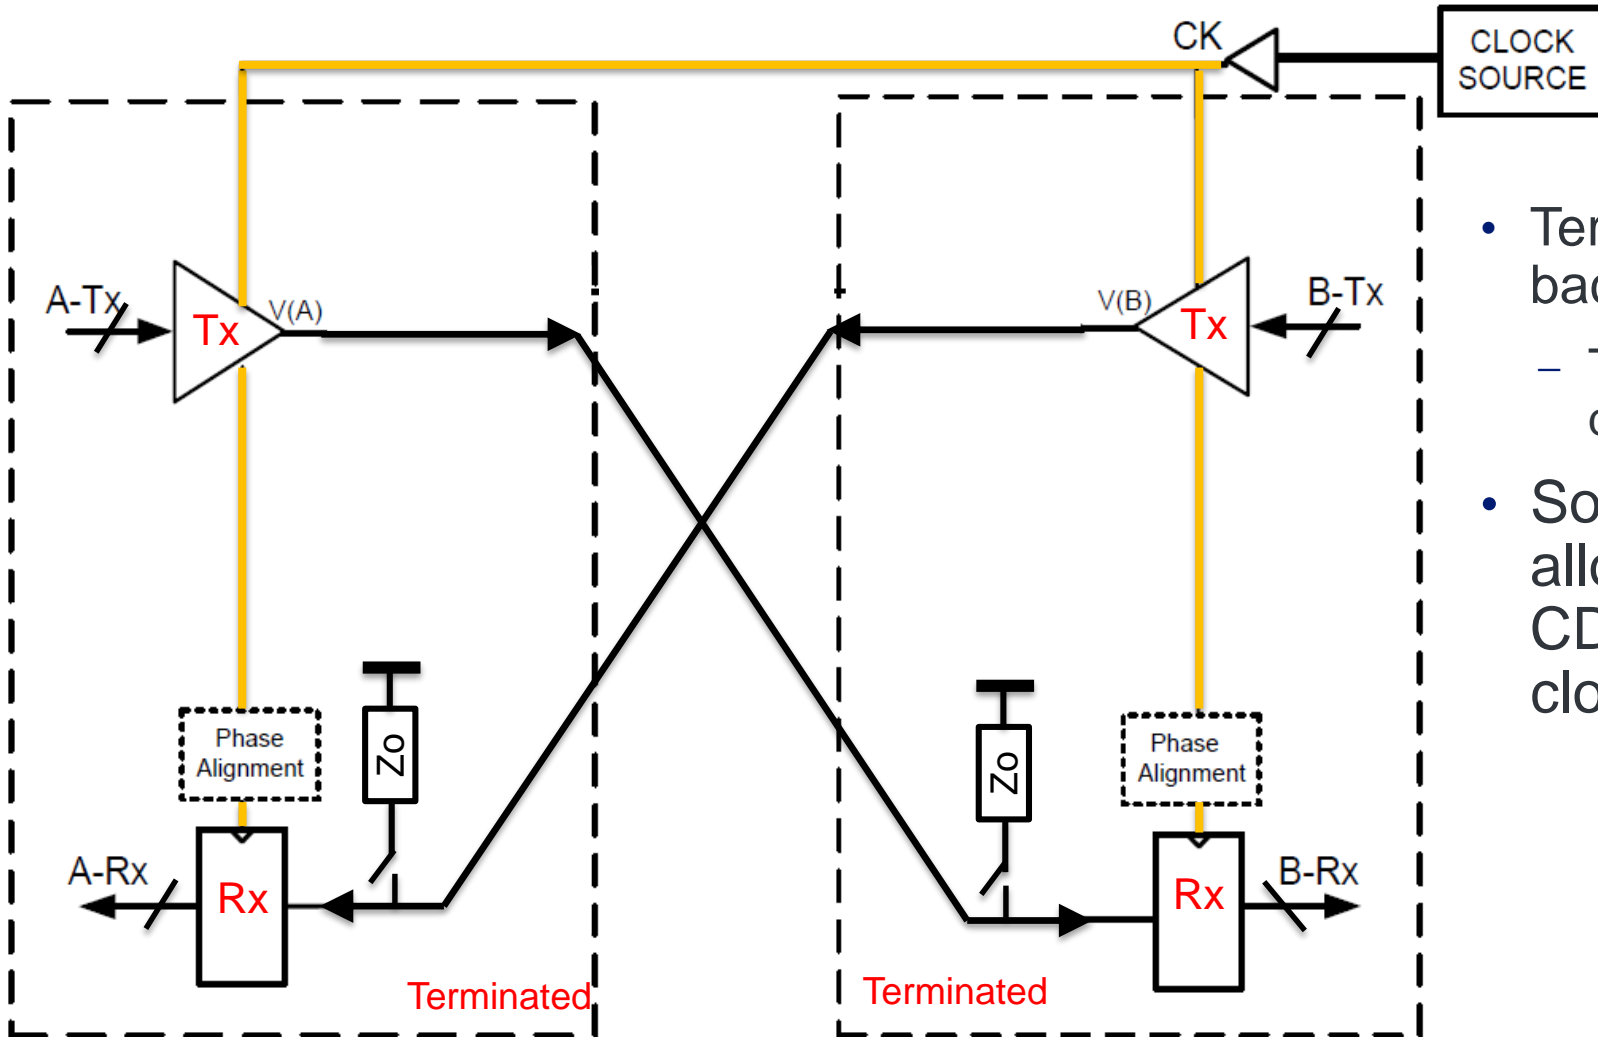


- Additional modes supported as long as BoW is supported
 - Terminated mode → up to 16Gbps/wire* up to 50mm
 - DDR Source Synchronous with clock alignment
 - Future enhancements possible
 - Reference BoW Turbo from prior discussions



Note: Baud rate limited to 16Gbaud for simplicity of design and ease of port

BoW : IO Block Diagram

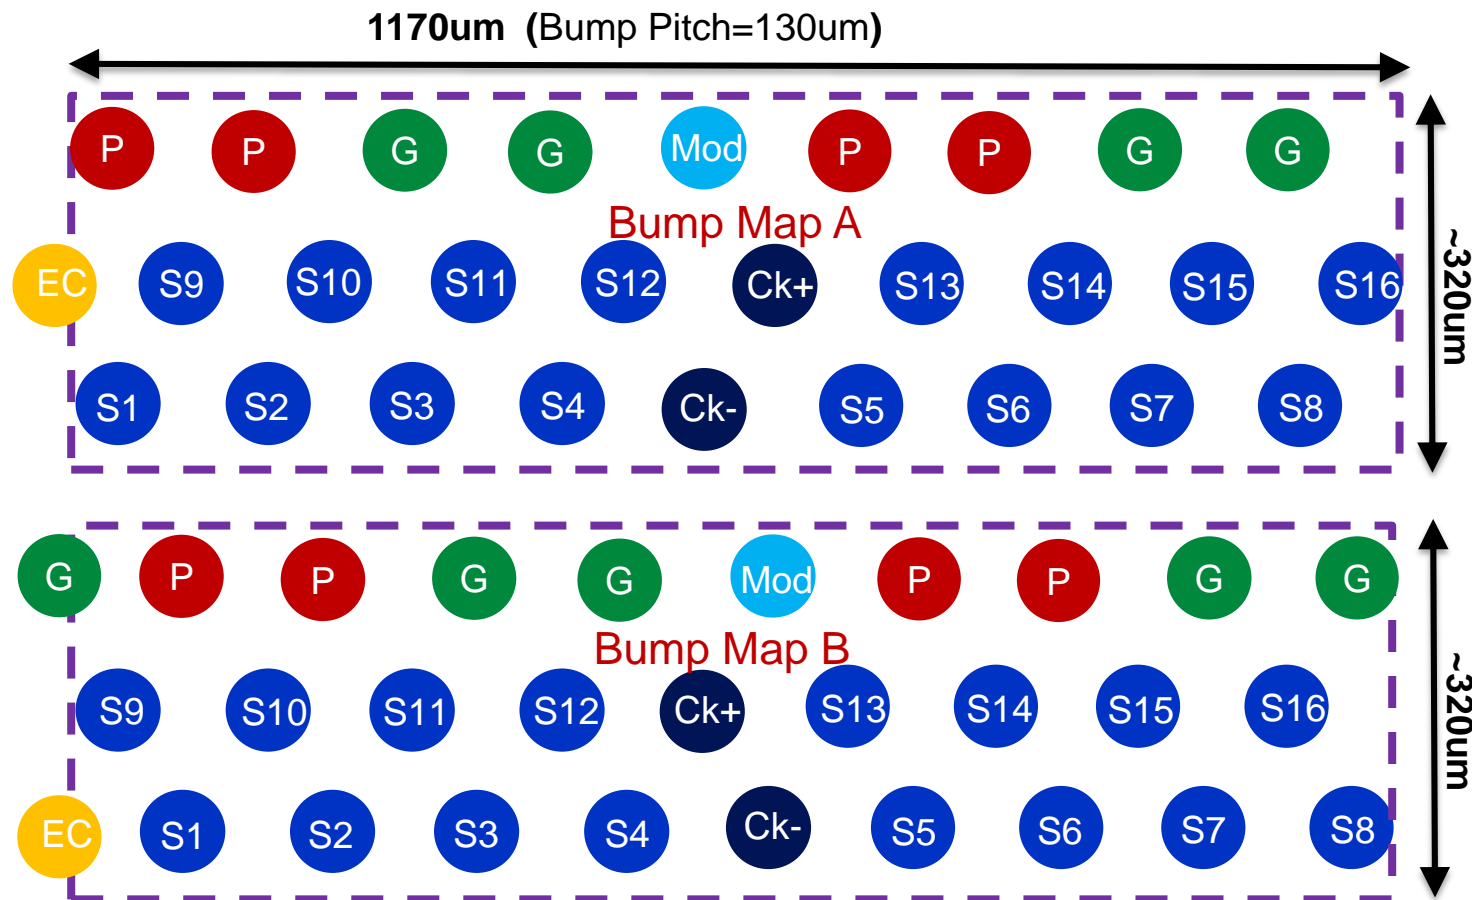


- Termination mode configurable to be backward compatible to Bow
 - This is achieved by adding the option for disconnecting the line terminations
- Source-synchronous clocking allows replacement of the complex CDR with a clock phase alignment clock

BoW Bump Map Flexibility

- BoW specification does not specify a specific bump pitch
 - It does enforce signal ordering to allow fan-in and fan out to finer bump pitches
- Examples shown are based on 130 um bump pitch
 - More coarse or fine bump pitches are possible and are the choice of the integrator
 - Enables support for older technologies with low datarates, and more advanced packaging with very high datarates

BoW Slice (Building Block) Bump Map



BoW Slice with same circuitry/layout but different RDL comes in 2 bump maps to create efficiency in building larger BoW modules

- A common bump map to be used for BoW Slice in all operation modes
 - 16 high-speed data bumps
 - 1 differential clock bumps
 - 1-2 Error correction bump (Optional)
 - To keep integer multiples of clock
- Aggregate BoW Throughputs (Tx+Rx)
 - BoW :
16x5Gbps/pad = 80Gbps
 - Terminated mode:
16x16Gbps/pad= 256Gbps
- Vertical stacking of Bow Slices increases throughput per die edge

BoW “Light ECC” Proposal

- Latency overhead is of extreme concern in die-to-die USR connectivity
- Short-reach NRZ links inherently have very low bit error rate, typically $BER < 1E-15$
 - An ECC that corrects at least 1 error per frame improves $BER < 1E-30$
- Reed-Solomon FEC is preferred over other ECC codes as it correct a group of bits for minimal extra overhead rather than single bits
 - Good for burst errors due to supply glitches
- Proposed ECC per BoW Slice → Reed-Solomon FEC = RS(34, 32, 8)
 - 16bits of ECC overhead → Corrects 8 bits in every RS_Frame (~6% ECC overhead)
 - **Total decoder latency per BoW Slice @16Gbps = ~10ns**

Test & Calibration Options

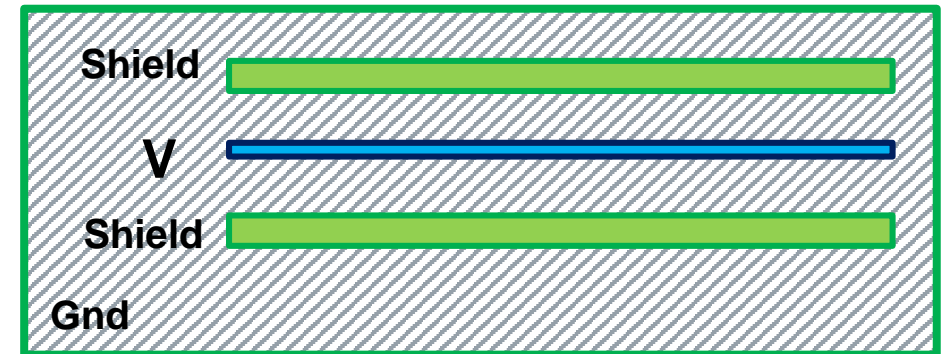

- 1149.1 Legacy/1149.6 High-speed (JTAG scan)
- IEEE 1500 (HBM type systems)
- At-Speed Self Test
 - Serial At-Speed PRBS Self Loopback: Tx(Port-N_Chiplet-A) → Rx(Port-N_Chiplet-A)
 - High-Coverage Wafer Test Screen (pre-package) → loopback traces on ATE load board
- At-Speed System Test/Compliance
 - External PRBS Loopback: Tx(Port-N_Chiplet-A) → Rx(Port-M_Chiplet-B)
 - Eye Monitor for At-Speed Test/Compliance: Tx(Port-N_Chiplet-A) → Rx(Port-M_Chiplet-B)
 - Per data bump: Measure Errors Rate for each phase and voltage threshold
- Calibrations
 - PLL/DLL Lock
 - Rx Phase Lock
 - DCC Calibration (optional)

BoW Slice Specifications

Parameter	Parameter
Single Supply Voltage	0.70V-0.90V (+/-5%)
Baud rate/bump	BoW: 1-5Gbaud, Terminated: 1-16Gbaud
Max Throughput/mm (stackable up to 4x)	BoW: 80Gbps, Terminated: 256Gbps BoW: 320Gbps, Terminated: 880Gbps
Tx Output Amplitude	400mV (+/-80mV)
Energy Efficiency	< 0.7pJ/bit
Trace length	Un-terminated: 10mm, Terminated: 50mm
Latency	< 30/Gbaud (<3ns @10Gbaud)
Slice Bump Allocation (excludes global AUX)	16 Data, 2 Clock, Opt. Ctrl
BER Target	<1E-15 (No ECC) <1E-30 (No ECC)
ESD / CDM protection	250V/50V



Intra-Package Signaling: Differential or Single-Ended

- Reasons for Differential Signaling
 - Immunity to common-mode noise coupling
 - Proper return current Path
 - Double effective amplitude with a given supply level
- Intra-Package Routing Advantages
 - Multi-layer substrate → Plane shielding between different layers
 - Trace pitch >> Bump pitch → Enables shielding between routes
 - The conducting plane/shields → Return current path
 - Ultra short routes → Ultra low loss → No need for double amplitude
- **Single-ended Signaling inside Package**
 - Can have same benefits of differential signaling
 - Saves half of the precious high-speed signaling bumps on MCM
 - Offers same throughput per bumps at half baud rate
 - Reduces interface design complexity → Faster/Easier Design

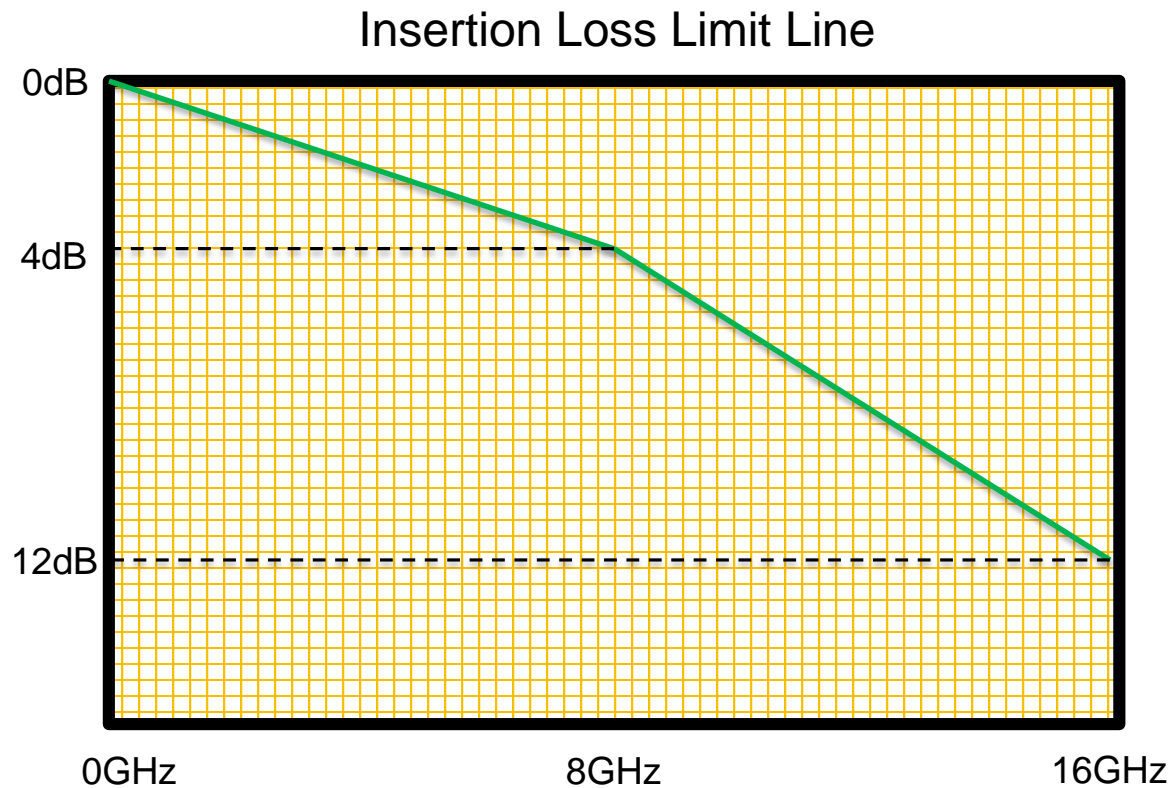


Intra-Package Signaling: Differential or Single-Ended

- In case of no side conductor shields, the differential signaling does not necessarily provide immunity to crosstalk from adjacent lanes
- Each conductor of the differential pair is exposed to a different aggressor signals on either side
- Single-ended signaling at lower baud rate provide better isolation to adjacent aggressors
- The conducting planes on top/bottom provide the Return current path




Channel Specs: Insertion Loss Limit Line




- Underterminated channel is limited by the round-trip reflection delay and limited to <10mm for a 5Gbaud with proper slew-rate control
- Termination mode channel can be long regardless of the baud rate, but to minimize the equalization requirement, it shall meet the provided Insertion Loss limit line

Channel Specs: Return Loss Limit Line



- BoW channel is expected to be unterminated and does not need to follow any Return Loss spec
- Termination mode channel shall meet the provided Return Loss limit line for proper operation and meeting target BER

Channel Specs: Power-Sum Crosstalk Limit Line



- Crosstalk is defined in form of the sum of crosstalk power of all aggressors on a target trace.
- The proposed power-sum crosstalk spec:
PS Xtalk Limit = -10dB – 37dB.e^(-f/8GHz)
- A victim trace in between two aggressor traces on the same package substrate layer with air gap spacing of 50um (or more) with at least a reference plane under or above will meet the proposed power-sum crosstalk limit

Summary

- Concept proven in 14nm Silicon (easy to port to other nodes)
- Throughput close to 1Tbps/mm edge over 50mm organic substrate (pitch= 130um)
 - Full-duplex Throughput/bump up to 16Gbps
 - Higher bandwidth possible with fine bump pitches
- Small area per port
- Less than 0.7pJ/bit in 14nm, lower with small trace lengths
- Single power supply 0.7V-0.9V: Compatible with synthesized logic circuits
- Easy and quick to port into other process nodes (vs 112G PAM4)
- Backward compatibility
 - A Chiplet using termination interoperates with basic BoW interfaces
- High Wafer-level test coverage per Chiplet to improve final product yield

How You Can Help

Seeking volunteer(s) for

- Foundries to support implementation of BoW PHY IP in their high-end process nodes
- Development of the BoW PHY IP in multiple process nodes (28nm to 7nm and even 5nm)
 - Availability of the BoW in wide range of process nodes will encourage larger and faster adoption
 - Open implementation would further benefit ecosystem
- Help us develop more detailed test implementations
 - New sub-group forming up within ODSA
- Many items not yet covered in proposal, including calibration, initialization

