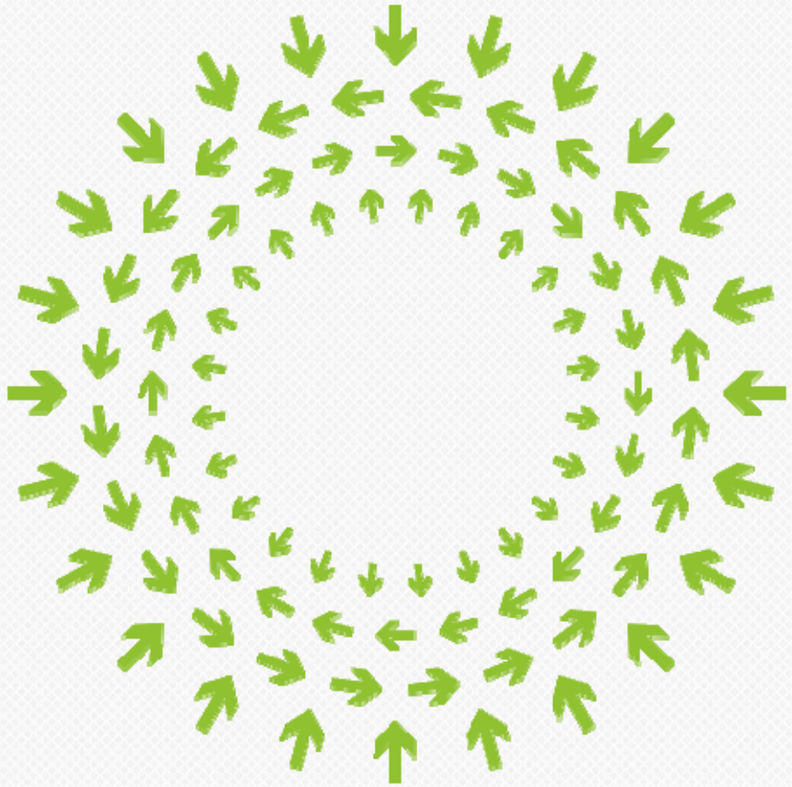

OPEN

Compute Engineering Workshop

March 9, 2015

San Jose

OPEN
Compute Engineering Workshop

OCP Mezzanine Card 2.0 Update

Jacob Na
Facebook Thermal Engineer
Jia Ning
Facebook Hardware Engineer

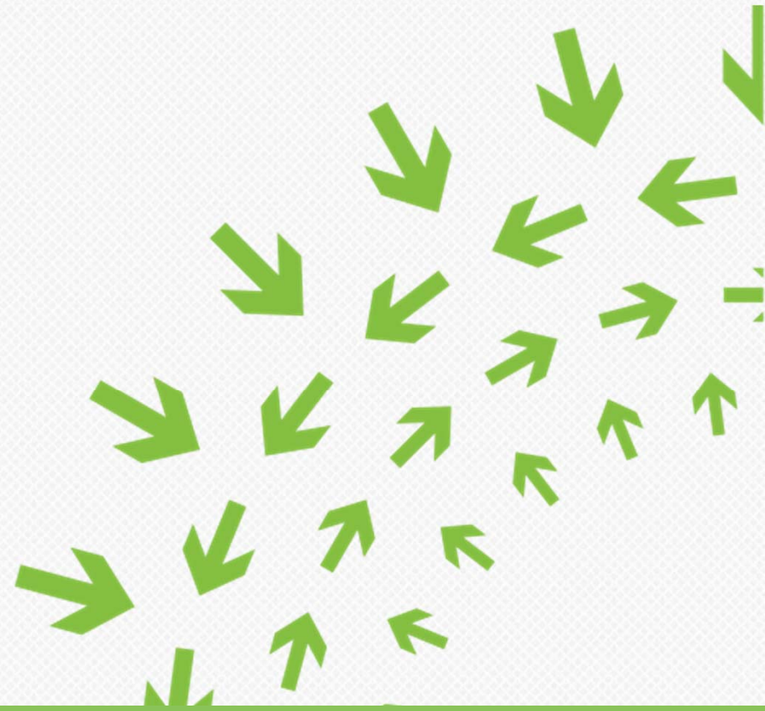

Background

OCP Mezz 1.0

- 2x SFP+
- PCIe x8
- OOB on I2C
- 8MM stacking
- OCP Intel MB V2.0/V3.0

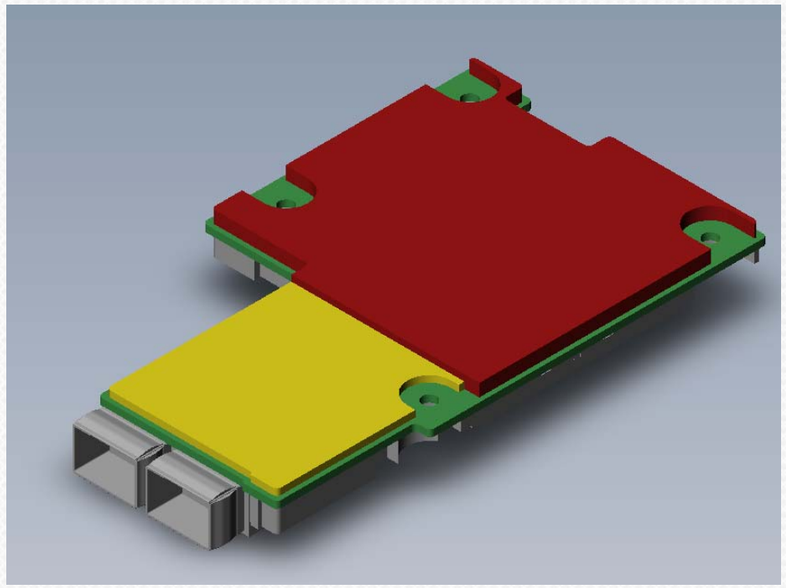

- Developed in 2011
- A.k.a “Mezz 0.5”

Overview

OCP Mezz 2.0

- 2x SFP+/2x QSFP/2x QSFP28
- 4x SFP+/4x 10GBaseT
- PCIe x8 or x16
- OOB on I2c or NC-SI
- KR to Host
- Optional 12mm stacking

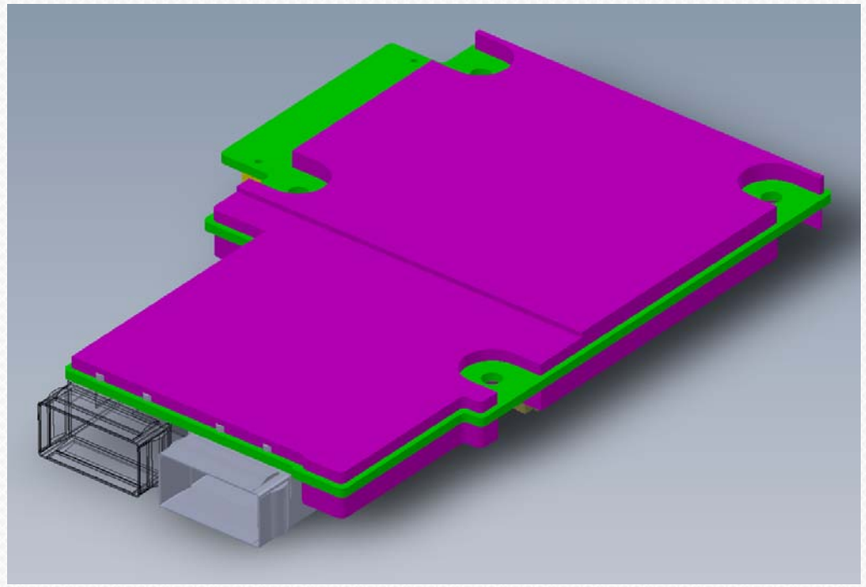

- Spec released and approved in 2014 Aug with Rev0.40

  Rev 0.40 Release
- Update draft to Rev0.45

  Rev 0.45 Draft
- Future Plan

  Future Plan

Form Factor Change in Horizontal Plane

- Add Connector B
- Enlarge board space
- Optional I/O area

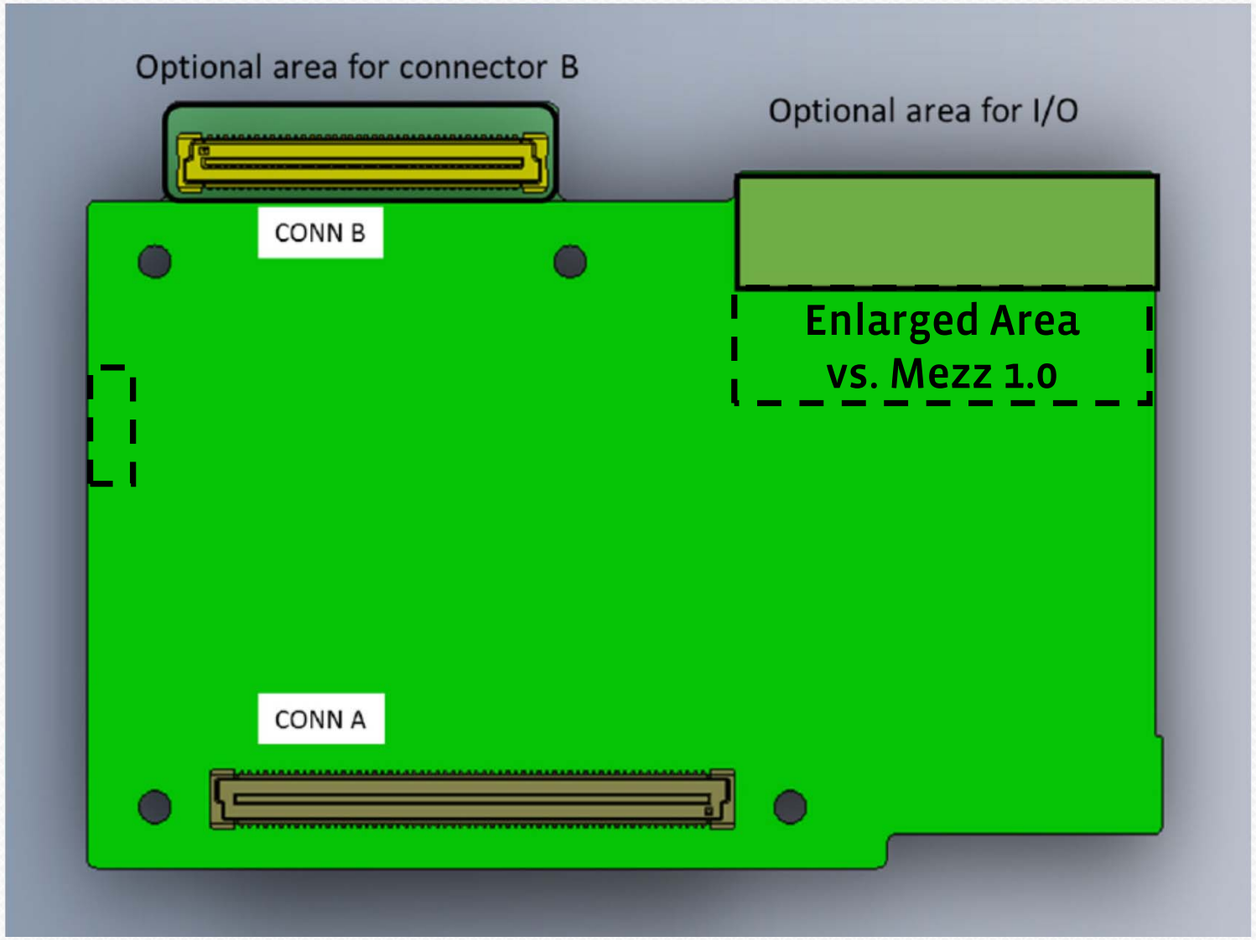

Form Factor Change in Vertical Plane

- Type 1
  - Similar to Mezz 1.0
  - Optional I/O at secondary side
- Type 2
  - 12mm stacking for 11.5mm heatsink
- Type 3
  - Controller on secondary side

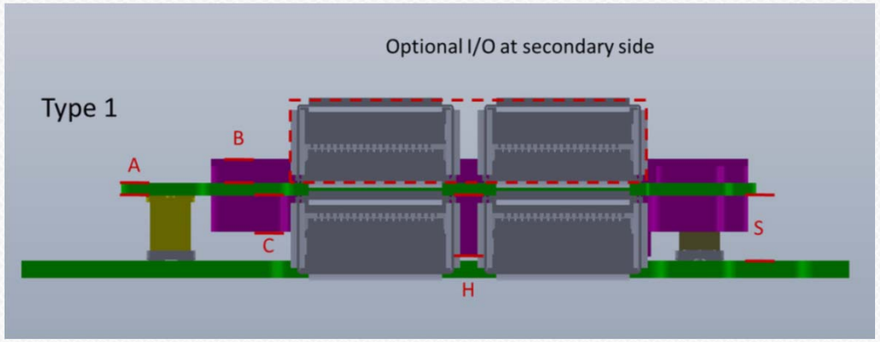

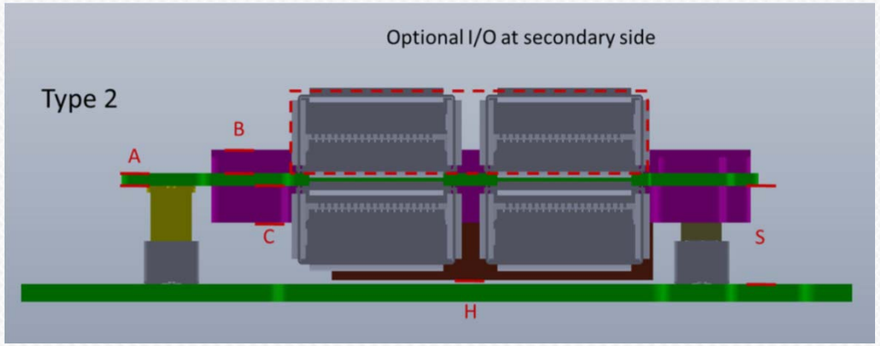

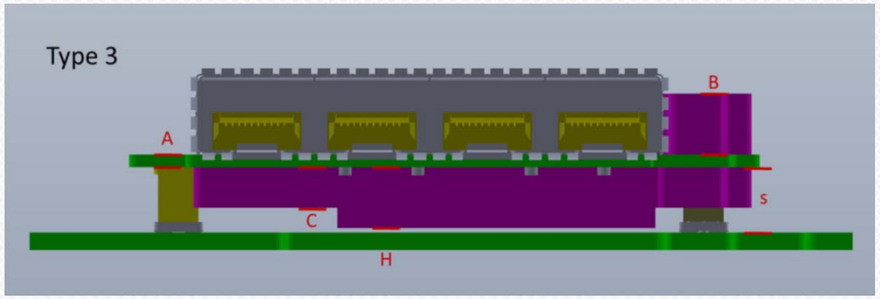

Network Interface

- 2x QSFP
- 4x SFP+
- 4x RJ45/10GBASE-T

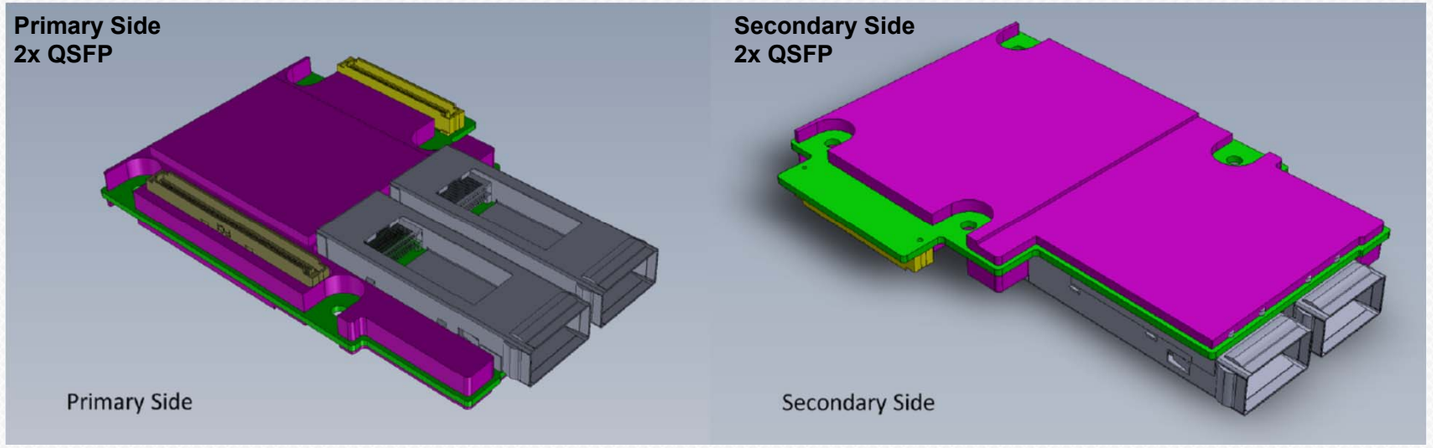

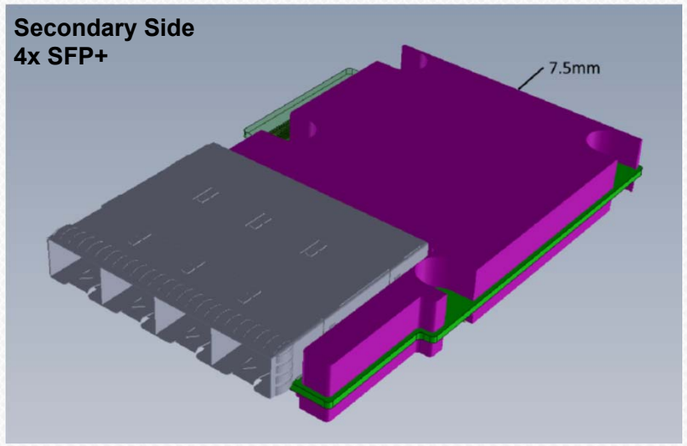

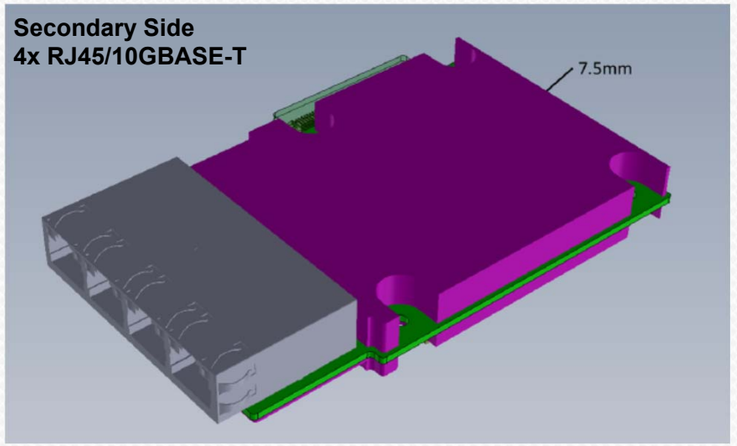

Rev 0.45 Draft

Network Interface

- 50G/100G QSFP28
  - Keep rear surface unchanged for backward compatibility to baseboard
  - Shift mating plane forward by 2.19mm for SFF-8662 Connector/SFF-8663 Cage

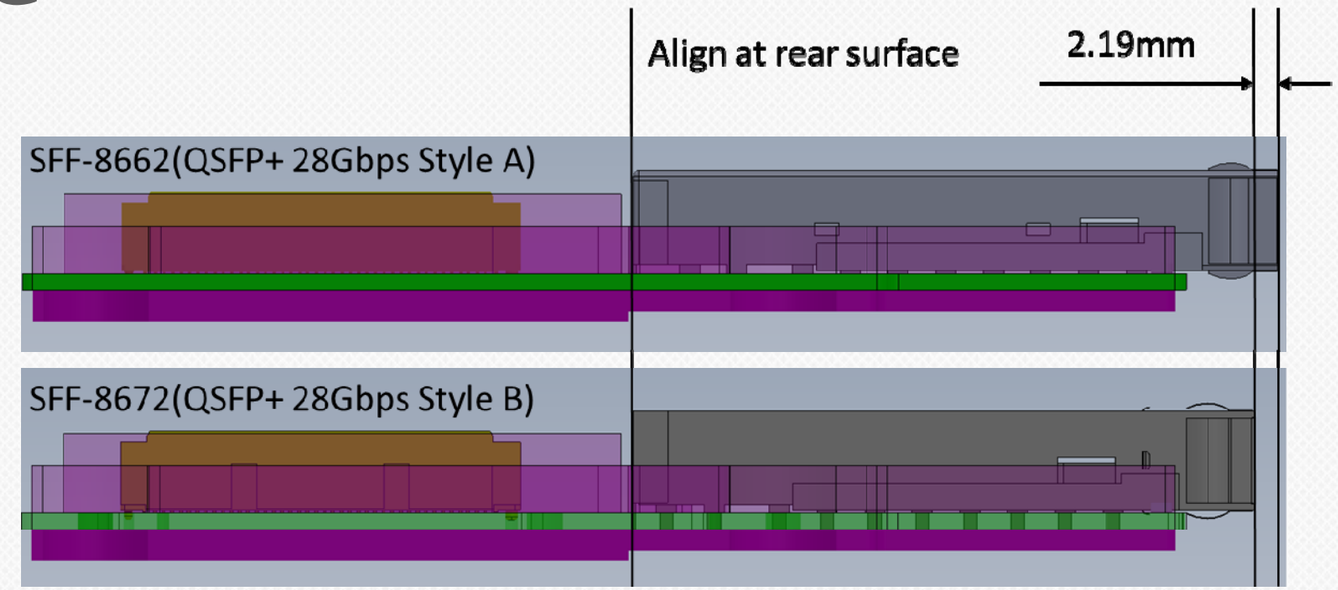

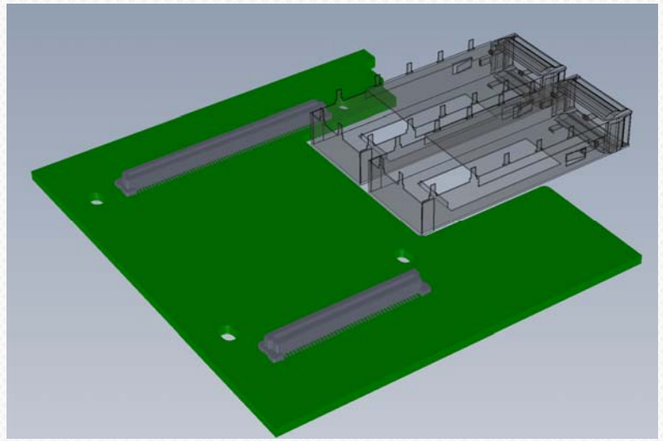

Host Interface Connector A (120pin)

- Add NC-SI
- Thermal reporting
- 4x PERST#
- EEPROM(Card ID)
- Baseboard ID

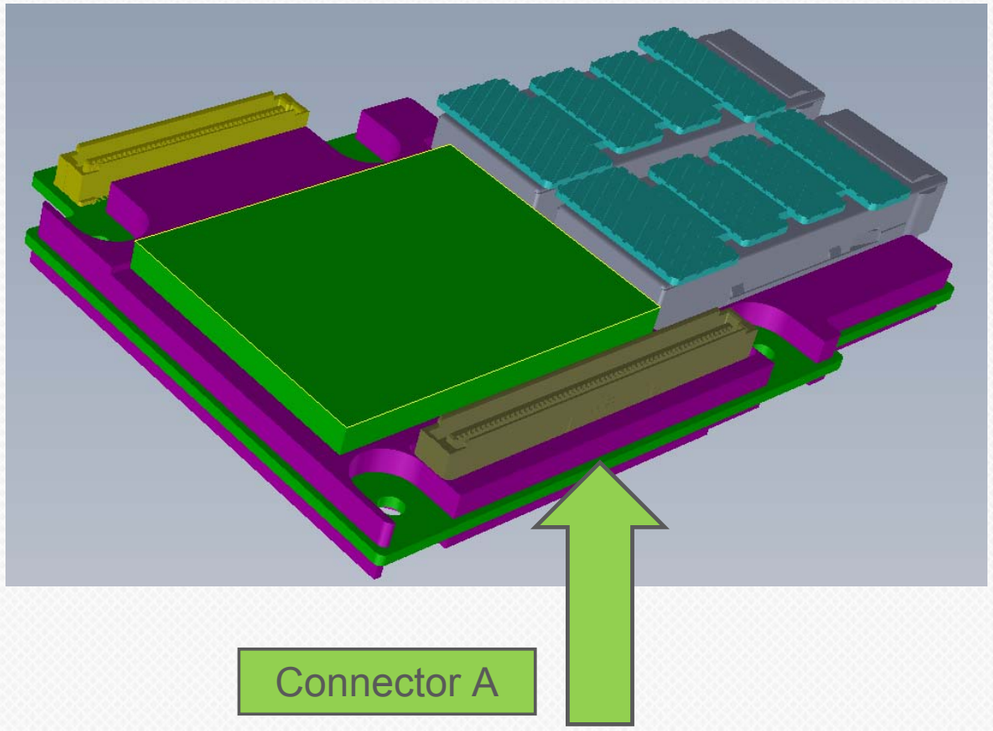

Host Interface Connector B (80pin)

- Connector B (80pin)
  - New Added
    - ->Enable up to x16 PCIe Gen3
  - Use on as needed base
    - -> Impacts backward compatibility

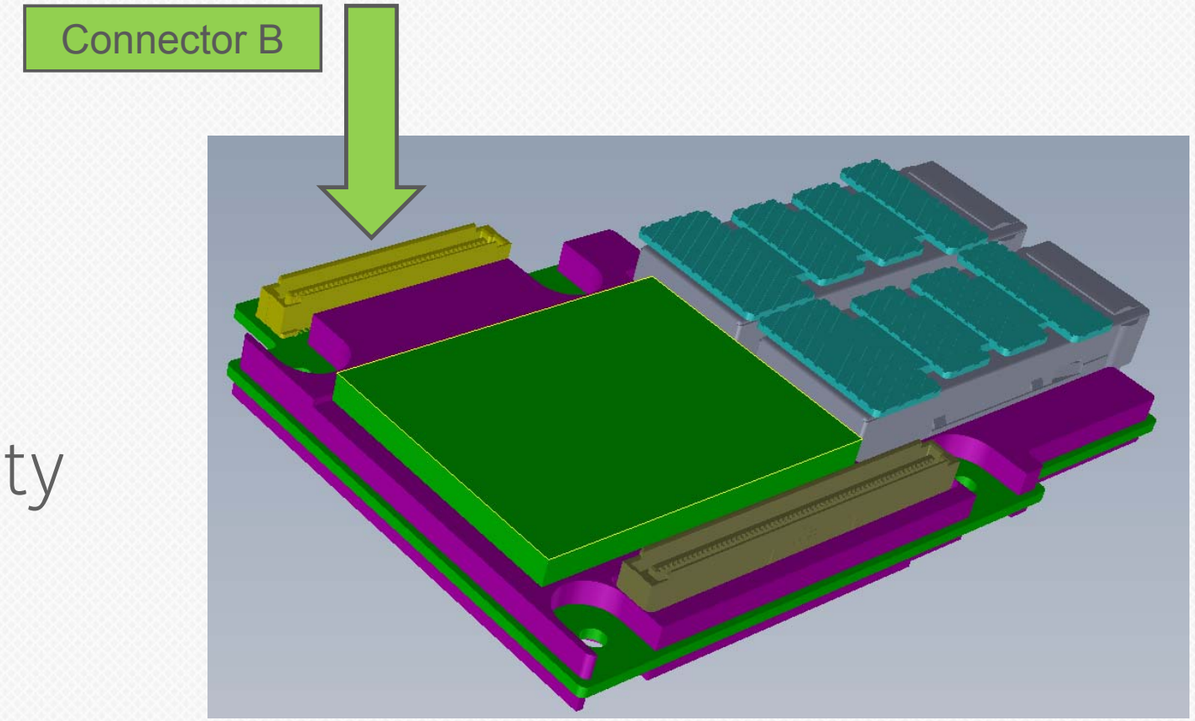

Host Interface KR Mezz

- Up to 16x KR10 channels
- 4x KR10 example
- Pin assignment was changed in draft

Rev 0.40 Release

KR/Repeater Numbering Sequence																
# of KR	*Pin 1		Connector A						*Pin 1		Connector B					
4 KR			0	1			2	3								
8 KR			0	1			2	3			4	5			6	7
16 KR	8	9	0	1	10	11	2	3	12	13	4	5	14	15	6	7

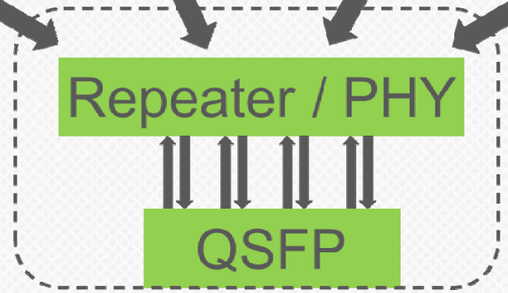

Rev 0.45 Draft

KR/Repeater Numbering Sequence																
# of KR	*Pin 1				Connector A				*Pin 1				Connector B			
2 KR									0	1						
4 KR									0	1	2	3				
8 KR									0	1	2	3	4	5	6	7
16 KR	8	9	10	11	12	13	14	15	0	1	2	3	4	5	6	7

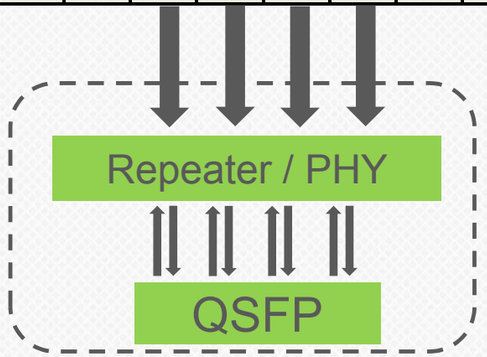

Baseboard ID

- Add weak Pull low for Connector B
- Add more details to Baseboard ID Table

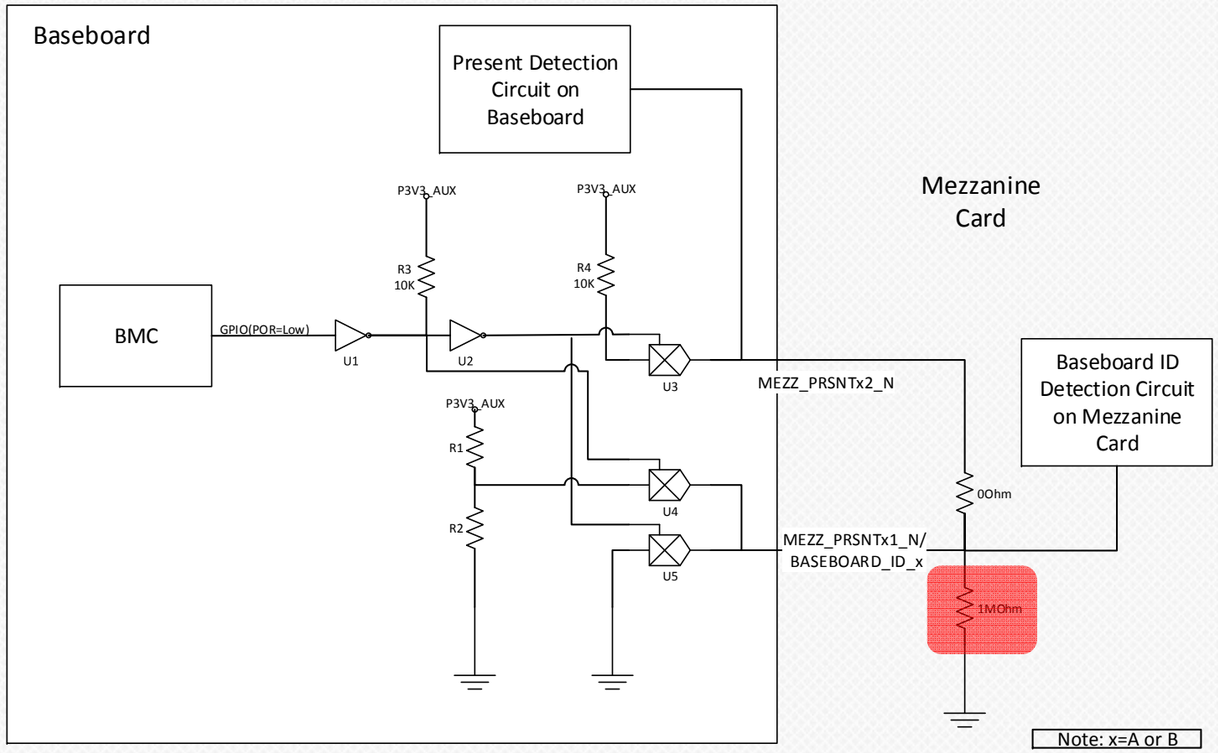

Rev 0.40 Release

Table 8: Baseboard ID definiton

ConnA R1	ConnA R2	ConnB R1	ConnB R2	Baseboard type
NC	0Ohm	Connecter B NC		PCIe x8
NC	0Ohm	NC	0Ohm	PCIe x16
Other				Reserved

Rev 0.45 Draft

Table 8: Baseboard ID definiton

ConnA R1	ConnA R2	Baseboard type on Connector A
NC	0 Ω	One x8 PCIe Root Port on baseboard Connector A; No Connector B on Baseboard
10 KΩ	887 Ω	One x16 PCIe Root Ports on Baseboard Connector A and B
10 KΩ	2.10 KΩ	One x8 PCIe Root Port on baseboard Connector A; Connector B presents on Baseboard
10 KΩ	3.83 KΩ	Two x4 PCIe Root Ports on baseboard Connector A
10 KΩ	6.49 KΩ	Four x2 PCIe Root Ports on baseboard Connector A
10 KΩ	11 KΩ	Eight x1 PCIe Root Ports on baseboard Connector A
10 KΩ	20.5 KΩ	RFU
10 KΩ	48.7 KΩ	RFU
10 KΩ	NC	Up to 8x KR on baseboard Connector A

ConnB R1	ConnB R2	Baseboard type on Connector B
NC	NC	No Connector B on baseboard; Mezzanine card samples Baseboard_ID_B as 0V with weak pull low on Mezzanine card side
10 KΩ	887 Ω	One x16 PCIe Root Ports on Baseboard Connector A and B
10 KΩ	2.10 KΩ	One x8 PCIe Root Port on baseboard Connector B
10 KΩ	3.83 KΩ	Two x4 PCIe Root Ports on baseboard Connector B
10 KΩ	6.49 KΩ	Four x2 PCIe Root Ports on baseboard Connector B
10 KΩ	11 KΩ	Eight x1 PCIe Root Ports on baseboard Connector B
10 KΩ	20.5 KΩ	RFU
10 KΩ	48.7 KΩ	RFU
10 KΩ	NC	Up to 8x KR on baseboard Connector B

Power

- Allow P12V_Aux/P12V to accommodate system without sufficient P12V_Aux
- Baseboard is preferred to use P12V_Aux
- Mezzanine card preferred to design for both

Spec Version	P12V Definition
Mezz 0.5	P12V or P12V_AUX
Rev 0.40 Release	P12V_AUX
Rev 0.45 Draft	P12V or P12V_AUX

Thermal

-Temperature Reporting

- Thermal reporting interface is added for Mezz 2.0

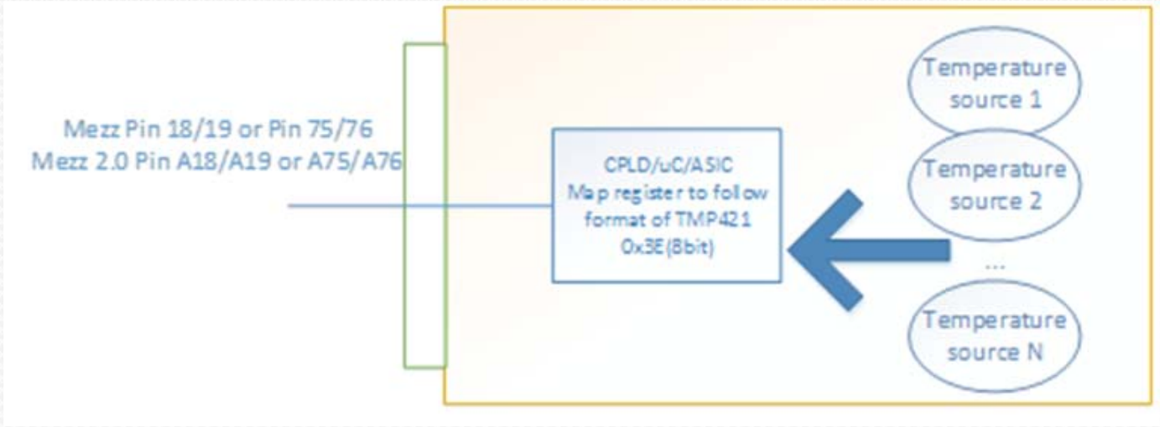

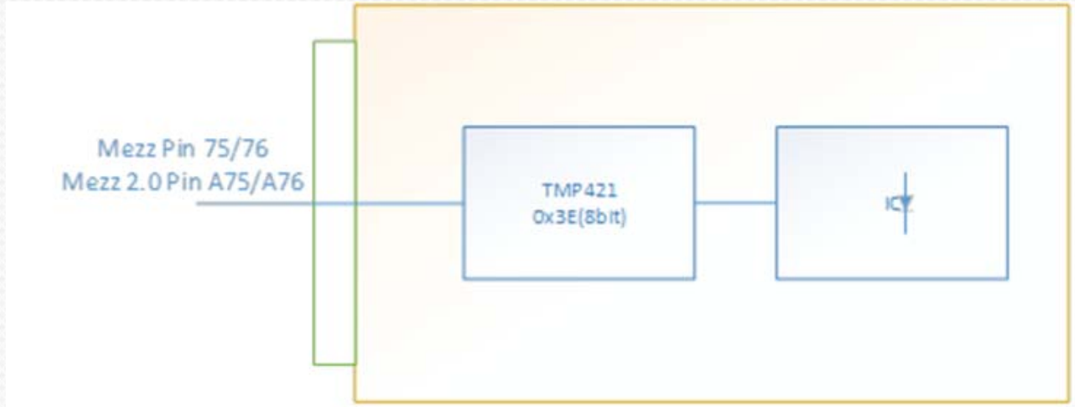

Thermal

-Environment Requirement

- Plan to Change Air Flow Requirements
  - Mezz 2.0 Rev0.40->170LFM at heatsink area
    - Hard to setup simulation assumption
    - Not ideal to correlate to system simulation
  - Future plan:
    - Specify a simulation assumption/testing configuration
    - Criteria correlates to Ambient Temperature/link activity
    - Leave key Criteria to User

Thermal

-Simulation assumption example

- 8mm stacking(Type 1)
  - Chassis Inner dimension: 84mm x 12mm x 300mm
- Assumptions/model defined for chassis
  - Open chassis in both x+/x- axis
  - Non-conducting material
  - Automatic algebraic turbulence
  - Inflow airflow

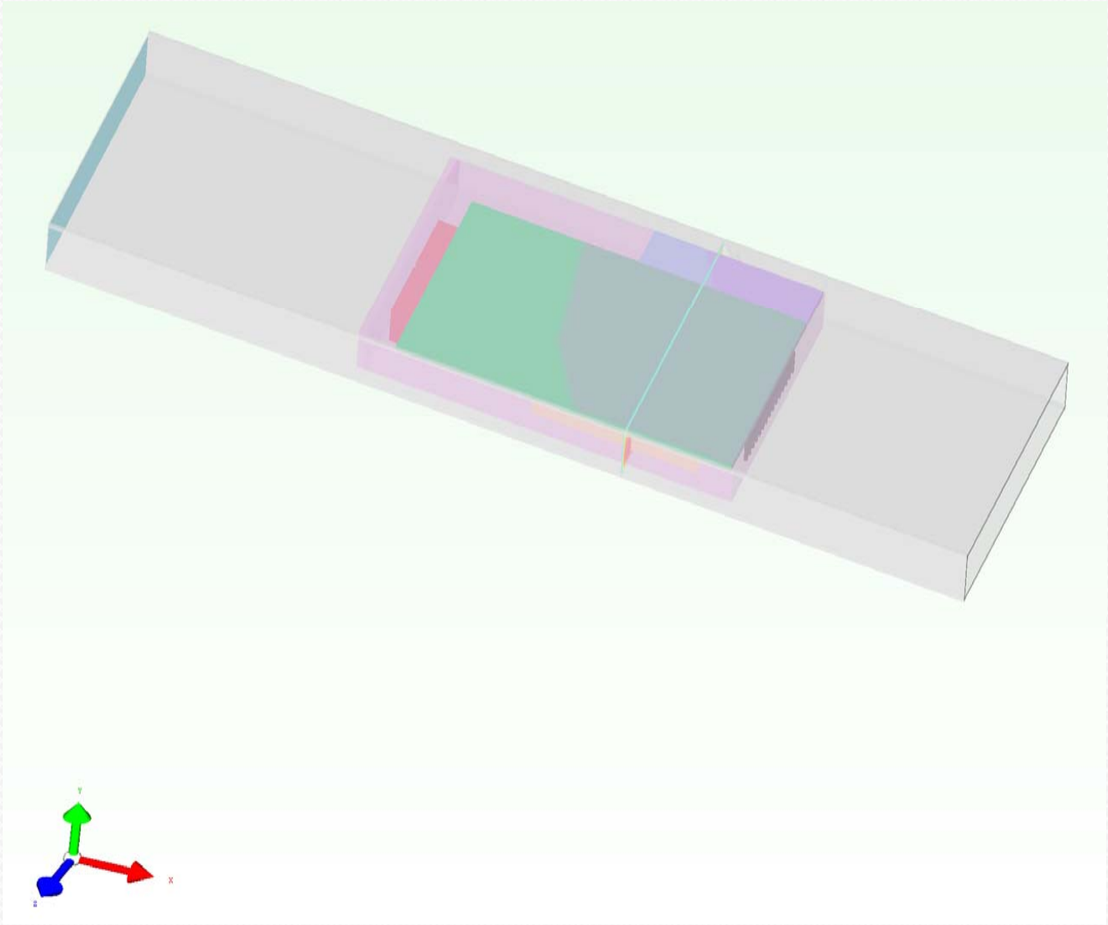

Thermal

-Simulation assumption example

- Location of Mezz card is specified

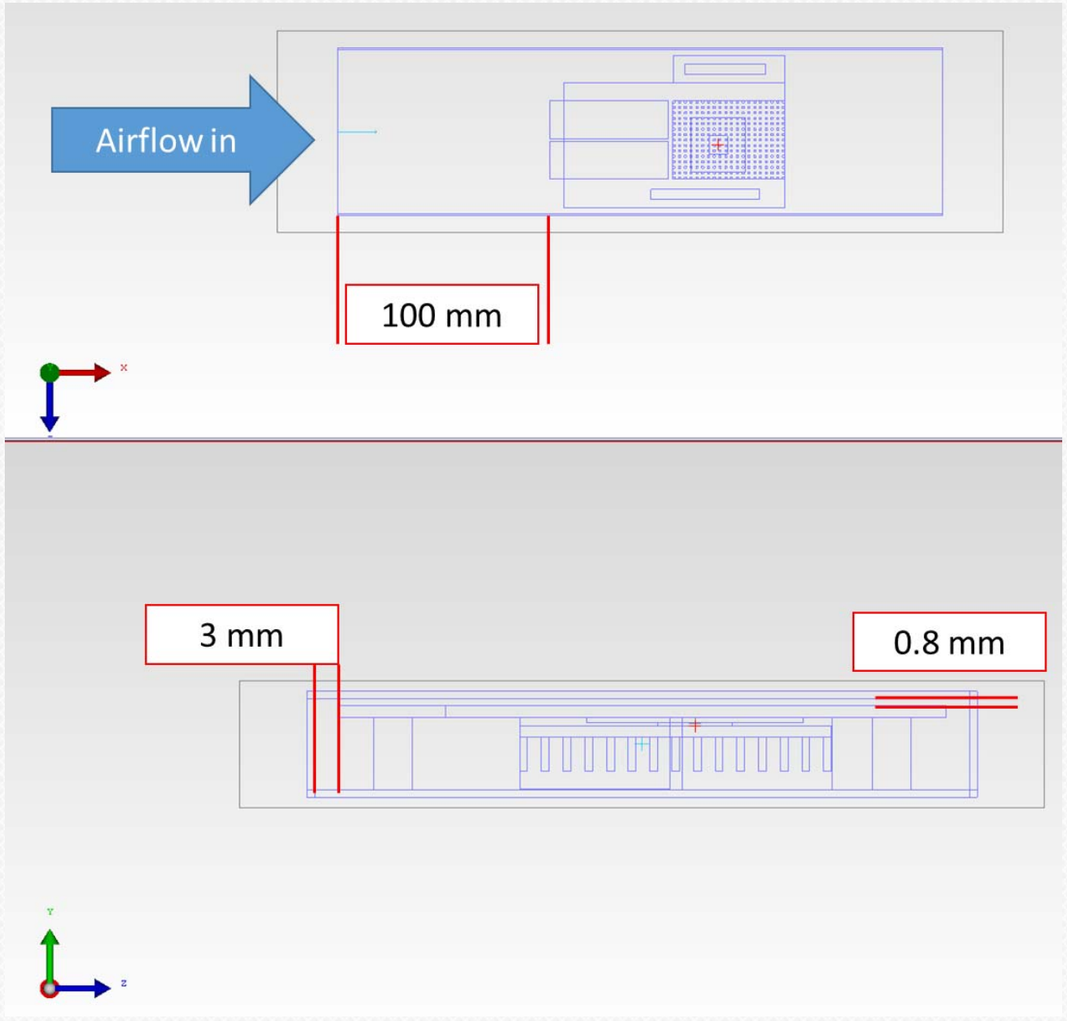

To be worked on:

- KR loss budget
- KR sideband(LED) signals
- Thermal specification to fit broader use cases

Q&A

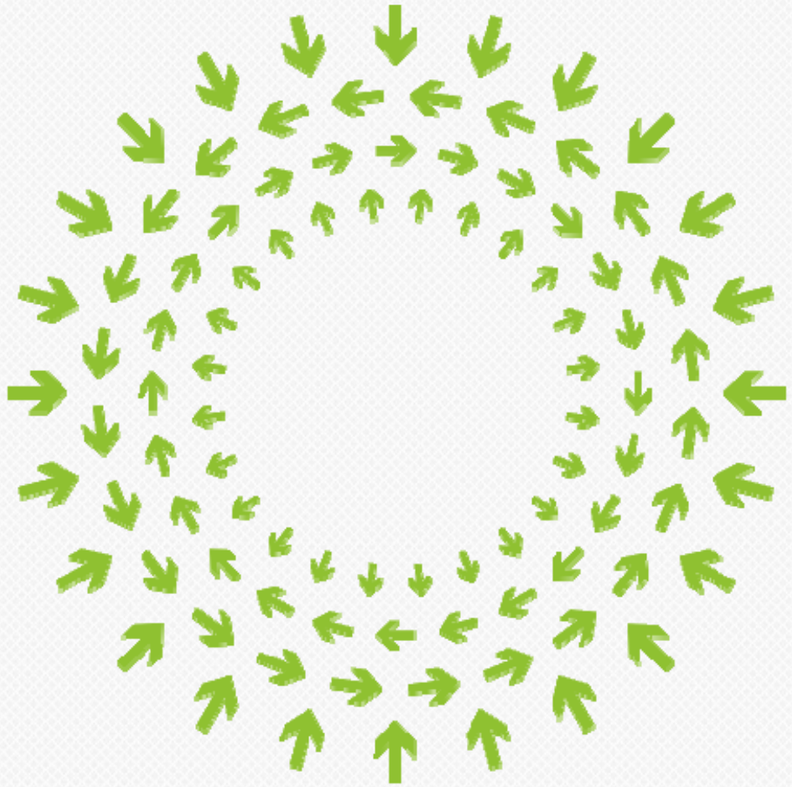

OPEN

Compute Engineering Workshop

March 9, 2015

San Jose