






# High Performance Processing Systems Enabled by 3D Integration

Bob Conn

May 5, 2011



# 3D HPC - Big Issues



Early 2009 – first SiCB @ siXis

Each edge:  
 \* Chiplet-to-external I/O ≈ 400 signals  
 \* Power & Ground ≈ 400  
 \* on 0.5mm pitch  
 \* Pads double as probe points

February 2011  
 BWRC  
 32 dual core low power processors  
 4 Tezzaron memory chips

- Design and Architecture
- Fab
- Interconnect
- Bandwidth
- **Power and Cooling**
- Assembly and Rework
- Testing
- Reliability
- **Cost**
- Supply Chain

# 3D HPC Issues – some detail

## Design and Architecture

- 3D silicon circuit boards are a new paradigm - not quite our friendly FR-4 but easier access than ICs. Needs new tools.
- **Ease of design**
- End of packaged ASICs, rise of MOSIS and silicon circuit boards
- IC designers or PCB designers?
- **Large area silicon circuit boards – greater than 5cm x 5cm**

## Fab

- **The TSV bandwidth problem**

## Interconnect

- **How to get off of silicon**
- Mechanical connectors and silicon
- Wire bonding
- Bumps, pillars, posts, etc.
- Conductive elastomer
- Nanotube forests!
- Flexible cable
- Optical connectors
- Free-space laser optical – ideal , no connector at all!
- Near field RF

## Bandwidth

- TSVs must work at 10+GHz
- Devices must move from 1000 I/O to 100,000 I/O – or equivalent bandwidth
- Integrated optical grading
- **Integrated optical**

A closer look at 3D HPC issues

## Power and Cooling

- How to get 200 amps into a silicon stack
- **Not much sense in exotic cooling – it's too big.**
- pJ/signal
- Integrated power
  - **Better energy storage – high Q inductors and transformers**
- Next gen will have commercial bare die with TSVs. One side is for power and the other for signals.
- **100x reduction for core power and I/O power**

## Assembly and Rework

- Cheap MOSIS
- Top and bottom die attach
- Bare die sources
- Indium, self-reflow
- Conductive elastomers
- **Nanotube forests**

## Testing

- Bare die testing

## Reliability

- What does one do with lost computational blocks or partial memory

## Cost



- Comparable to exotic PCB prototypes today

## Supply Chain

- Where are the bare die?
- Where are the second sources?
- Where are stacked die?

# System Level HPC SiCB stack/module

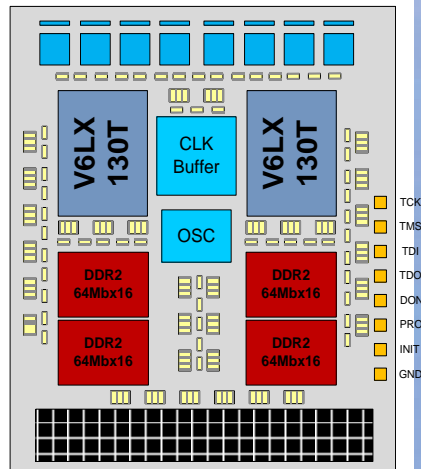

Typical High Performance Computing SiCB Modules



**Silicon always needs FR-4** (or some other substrate)  
 - how do we get power and signaling between SiCB and FR-4?

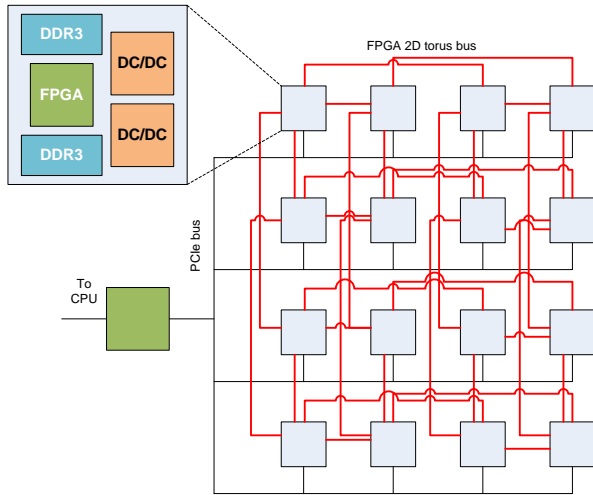
4" x 4" x 4"  
 32 FPGAs  
 64 memory chips  
 1920 DSP blocks  
 600w

# High Density DSP Design Concept




- Scalable DSP system using Silicon Circuit Boards
- System design objectives:
  - Equivalent volume of a softball ~ 10x10x10 cm<sup>3</sup>
    - 8 PCBs with SiCBs and DC conversion
  - Power consumption 600 watts
    - Air cooled - @ limit
- SiCB based DSP module
  - 4 Xilinx V6LX 130T DSP bare die
    - 1920 DSP blocks total, 500 MHz clock
    - 2X DSP density with V6SX part
  - 8 GB of DDR2
  - External IO via PCIe
- Controller: COTS CPU with Linux OS

# HPC Design and Architecture



We need a design system that allows 3-4 PhD students to design specialty silicon that connects to a custom silicon circuit board for less than \$100k and less than 3 – 4 months of effort.



Pluggable fiber optics allow re-configuration of the interconnect architecture

- New systems level architectures
  - **Increased I/O – how do we get on and off the SiCB?**
  - **The Cooling and Power Problem**
  - Torus, mesh, hypercube, linear string, etc.
  - Fast turn ASICs with TSVs and 1000s of I/O
- Ease of Design
  - 3D silicon is a new paradigm - not quite our friendly FR-4 but easier access than ICs. Needs new tools.
  - End of packaged ASICs, rise of MOSIS and silicon circuit boards
  - **IC designers or PCB designers?**
- Large area silicon circuit boards – greater than 5cm x 5cm – are needed

## Fab

- We need some uniformity across the set of 8" fabs
- Lower cost TSVs
- Large area silicon circuit boards (>4" on a side)
- Embedded passives
- n-layers of copper
- 4 microns extending to 10 microns thick copper
- Futures – embedded optical and transistors

- I've avoided Fab details as many folks are working in this area.

	Today	5 years	10 years				
Silicon Size	5cm x 5cm	10cm x 10cm	full wafer				
Silicon thickness	100 $\mu$ to 500 $\mu$	60 $\mu$ to 500 $\mu$					
Copper layers	5	8	10	how many do we need?			
Copper thickness	1 $\mu$ to 4 $\mu$	1 $\mu$ to 6 $\mu$	1 $\mu$ to 10 $\mu$	thicker is better for power			
Carrier substrate	FR-4, teflon, ceramic	FR-4, teflon, organics	FR-4, teflon, organics				
TSVs	8:1	12:1	16:1				
alignment	1 $\mu$	0.7 $\mu$	0.2 $\mu$				
min line/space	5 $\mu$ /4 $\mu$	4 $\mu$ /3 $\mu$	3 $\mu$ /2 $\mu$	no need for smaller - too much I*R loss			
Internal pad pitch Cu-Cu, Cu-SnCu	20 $\mu$	10 $\mu$	10 $\mu$	no need for I/O pad to be smaller than transistors			
External ball/bump pad pitch	0.5mm	0.3mm	0.1mm				
via diameter	8 $\mu$	4 $\mu$	2 $\mu$				
embedded passives	barely	R, L, C	2x improvement	sooner is better			
embedded power regulators	none	local POL	fully embedded regulators				
embedded optical	none	some	fully embedded optical				

# The Cooling and Power Problem

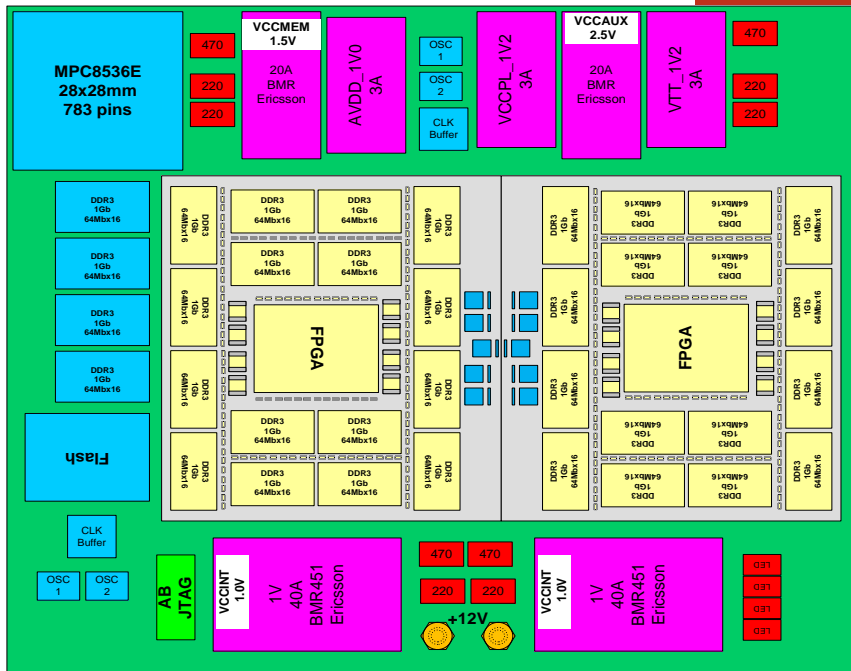
- Traditional HPC Systems can be considered as 3 separate volumes: Power generation, cooling, electronics
- Each of these takes approximately 1/3 of the volume inside a 19" rack:
- 1/3 Power Generation
  - Power generation from AC power in to distributed DC
  - Local power regulation on individual PCBs
- 1/3 Cooling
  - Air flow volume
  - Heat sinks
  - Fans
  - Ducts
  - Plumbing
- 1/3 Electronics
  - PCBs
  - Connectors
  - Components

**If we shrink the electronics to zero our system is  
2/3 of it's original size!**

- Cooling and power need to shrink!
  - we must have lower power components!
  - 5 years – cooling and power volume less than half of today's for a given wattage.
  - 10 years – cooling and power volumes less than 10% of today's.
- How to get 200 amps at 1.0v +/- 50mV into a silicon stack
- **Not much sense in exotic cooling – it's too big.**
- pJ/signal
- Integrated power
  - **Better energy storage – high Q inductors and transformers**
- Next gen will have commercial bare die with TSVs. One side is for power and the other for signals.
- **100x reduction for core power**
- **100x reduction for I/O power**
- No need for termination resistors in a lossy system



# The Cooling and Power Problem



FR-4 118mm x 140mm  
 SiCBs – each 56mm x 61mm

Power regulators matched to SiCB needs  
 require more area than the silicon circuit boards!

### Module Features

- 1 CPU - MPC8536E with memory
- 4 FPGAs – V6VSX475T
- 8 GB DDR3 Memory
- Oscillators and clock buffers
- Optical I/O

Each FPGA:

- 100 GFLOPS
- (9) 20Gb/s Highways:
  - (2) Intra-module
  - (1) PCIe to CPU
  - (6) general use
- (4) memory ports 2Gb each, BW= 128 Gb/s

Each Module:

- 0.4 TFLOP (dp)
- (24) General Highways = 480 Gb/s external I/O
- 1) CPU Highway = 20 Gb/s
- (16) memory ports 2Gb each, BW= 512 Gb/s

Cold plate cooling  
 150W maximum  
 All on-board power regulation  
 5" x 6" x 1"




SiCB bumps = vertical copper signals:

JTAG	4+4
CPU PCIe	4+4
Status Bus	12-40
Clocks	8+8
Power bumps	1000
Outtrigger bumps	100

Unused FPGA I/O 360

# Interconnect

- What about re-work?
- What about reliability?
- What about mechanical connectors on silicon?
- Need uniformity of interconnect solutions across fabs




## How to get off the silicon

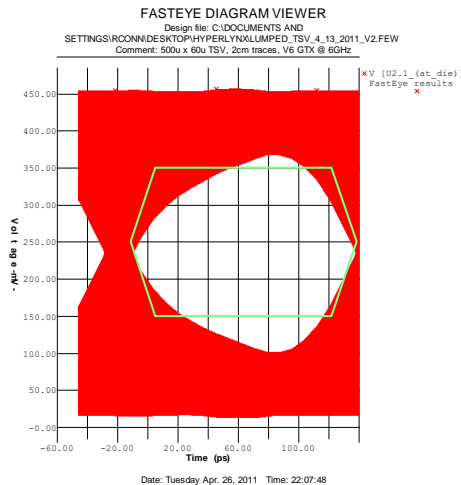
- Mechanical connectors and silicon \*minimize
  - Silicon can't take mechanical connector stress
  - Limited bump area for reliability reasons
- Flexible miniature ribbon cable \*now
- Wire bonding – fading fast \*history
- Bumps, pillars, posts, etc. \*now
  - Plenty of work done in this area
- Conductive elastomer \*now
  - Works well to 300 micron pad sizes
- **Nanotube forests!** \*10 years
  - nanovelcro
- Indium solder \*5 years
  - Different reflow temperatures for different parts
- Tiny flexible cable \*now
  - Careful handling
- Optical connectors \*now
  - Careful handling
- Free-space laser optical – ideal , no connector at all!  
\*5 years
- Near field RF \*5 years
  - Clock distribution within a module of several SiCBs
- Coupled capacitor \*5 years

# Bandwidth

- Devices must move from 1000 I/O to 100,000 I/O – or equivalent bandwidth
- Eye diagrams of a 2cm trace @ 6GHz – one with no TSV and one with a TSV. Xilinx V6 LVCMOS 1.5v, 16mA drive, 500 micron thick wafer and TSV.



2cm trace with no TSV



2cm trace with TSV

TSVs do poorly at 6GHz and up

## How to improve Bandwidth

- Fix TSVs!
  - Higher background resistivity
  - Smaller diameter
  - Thinner wafers
- More I/O pins
- Internal Optical
  - Optical grading
  - Silicon diodes and detectors
- External Optical
- Inter-module RF
- Is there an Amdahl's Law or Rent's rule for silicon circuit board based systems?

# Electrical Testing

- SiCB testing – do-able today
  - Flying probe - minimum pad size = 19 $\mu$ . Preferred minimum is 23 $\mu$ .
  - Visual inspection
  - Similar to FR-4 PCB testing
- Bare die testing
  - Big Problem today
  - More sources of tested bare die are available today
  - Probe card for JTAG/BIST
  - Flying probe for I/O continuity
- Module testing
  - JTAG/BIST
- Integrated system monitoring
  - Temperature
  - Voltage
  - Traffic


## We need work in testing

- Bare die need to be built with probe-able JTAG
- Better Flying Probe accuracy
- 10 years - ideal bare die would have probe-able power pads, optically testable I/O pads, optically accessible JTAG port.

## Don't worry Issues

- None

Cold plate  
Stiffener  
Alignment  
SiCB with bare die  
FR-4 with on-board regulation  
SiCB with bare die  
Alignment  
Stiffener  
Cold plate




- As far as I can tell there is nothing that we can ignore now.
- There are too many interconnected issues to eliminate any of them.

- Design
  - PCB tools are used for 3D silicon and silicon circuit boards
- Fab
  - The TSV bandwidth problem
- Bandwidth
  - 40 GHz copper all day
  - Integrated optical
- Power and Cooling
  - pJ/signal
  - Integrated power
    - Better energy storage – high Q inductors and transformers
  - Next gen will have commercial bare die with TSVs. One side is for power and the other for signals.
  - 100x reduction for core power
  - 100x reduction for I/O power
- Assembly and Rework
  - Cheap Mosis
  - Bare die sources
  - Indium, self-reflow
  - Conductive elastomers
  - Nanotube forests
- Testing
  - Bare die testing
- Reliability
- Cost
  - Comparable to exotic PCBs today
  - Less expensive in volume
- Supply Chain
  - Lots of change as processes become standardized

# 4 to 7 years out 2015 to 2017

Cold plate  
 Stiffener  
 Alignment  
 SiCB with bare die  
 FR-4 with on-board regulation  
 SiCB with bare die  
 Alignment  
 Stiffener  
 Cold plate



# 8 to 12 years out 2019 to 2023

## Design

- Low cost design capture tools spanning PCB to SiCB
- Integrated with layout tools
- Partitioner included
- Multi-medium simulations
- Integrated embedded optical
- Chips are built with TSVs
- Chips are built with self- I/O test
- Chips are built specifically for silicon circuit boards
  - Self -test
  - Self I/O test
- 10,000 I/O
- Memory chips with no multiplexing

## Fab

- Composite structures – mechanical structure for strength, silicon for active devices, glass for routing all in one fusion bonded stack
- Integrated optical gradings
- Integrated POL regulators
- Integrated RF dust
- Direct write patterning – e.g. laser ablation
- Organic dielectrics for better signaling, better reliability, lower cost

## Bandwidth

- 100GHz copper
- ??? optical
- Integrated optical grading
- **Integrated optical**

## Power and Cooling

- Single voltage power of 12v or 48v
- pJ/signal
- Integrated power
  - **Better energy storage – high Q inductors and transformers**
- Next gen will have commercial bare die with TSVs. One side is for power and the other for signals.
- **100x reduction for core power**
- **100x reduction for I/O power**

## Assembly and Rework

- Cheap MOSIS
- Bare die sources
- Indium, self-reflow
- Conductive elastomers
- **Nanotube forests**

## Testing

- BIST
- Bare die self-testing

## Reliability

- Better than packaged parts

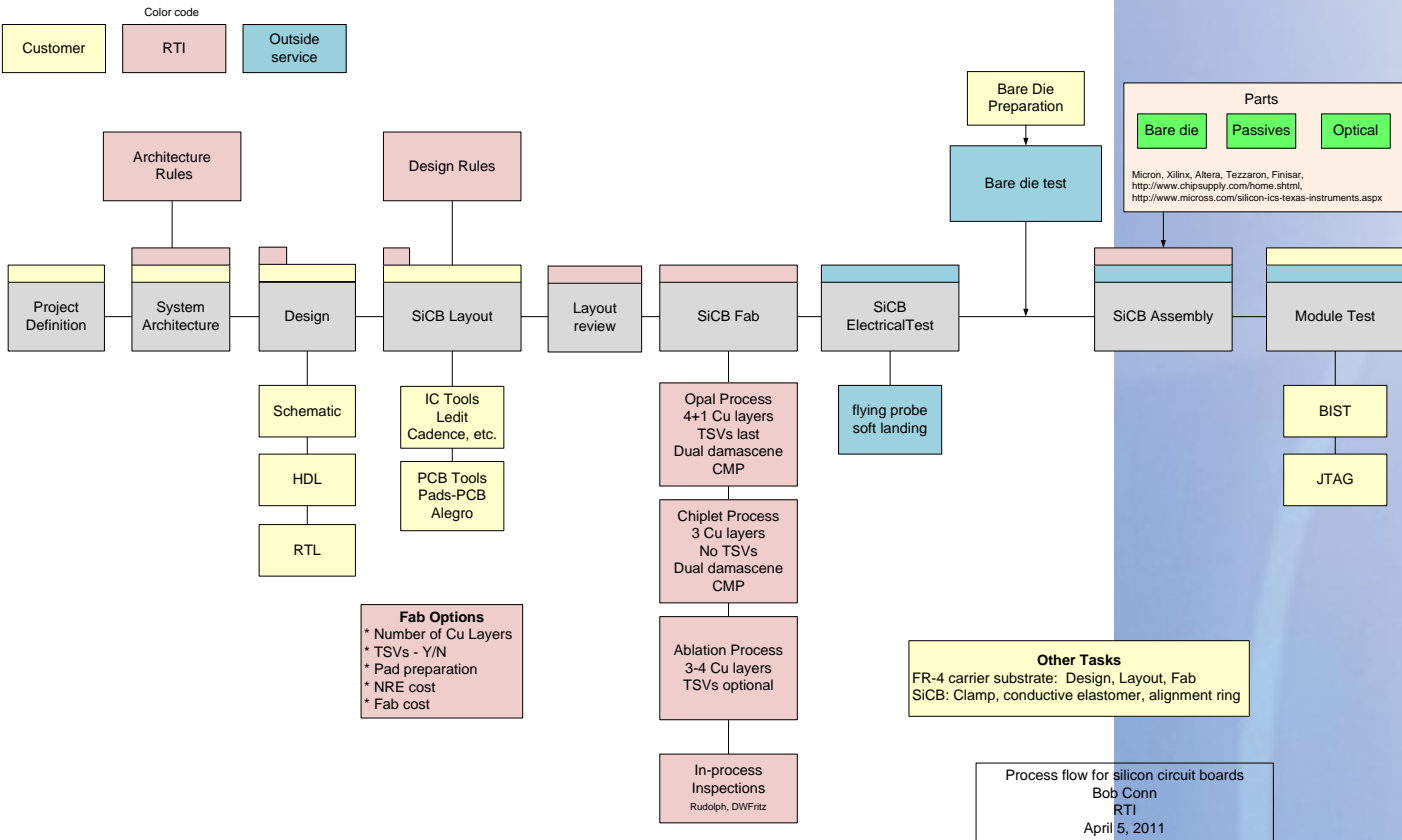
## Cost

- Less than commodity FR-4

## Supply Chain

- Normal way of

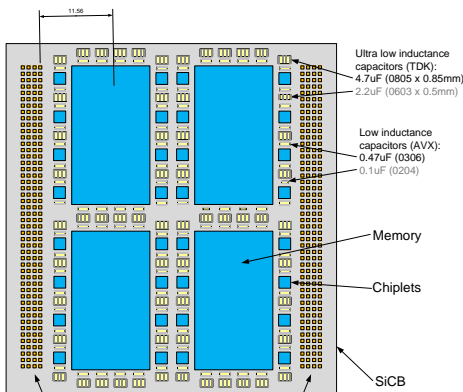
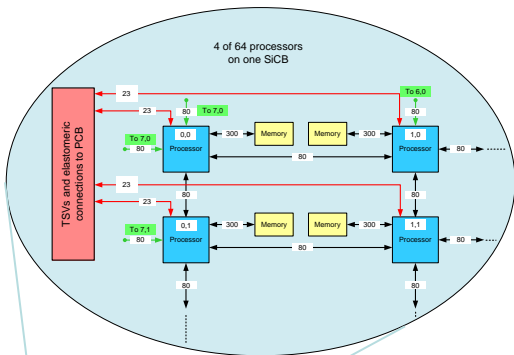
# Supply Chain



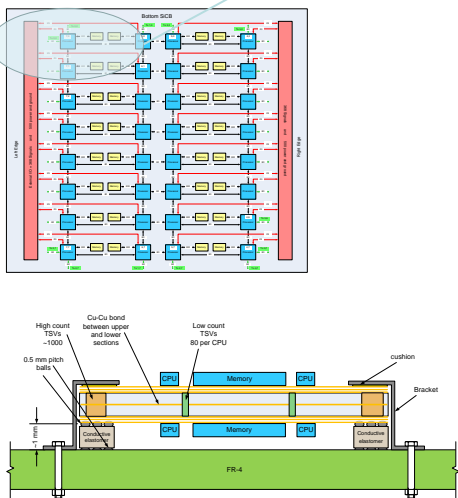

- No coherent supply chain today
- Most 8" fabs should be able to build silicon circuit boards
- Requires silicon circuit board + bare die + FR-4
- Weak points are
  - Non-uniform understanding of 3D silicon
  - Incompatible processes – my bare die can't attach to your silicon circuit board
  - Bare die testing



# UC Berkeley-BWRC



- Each edge:
- \* Chiplet-to-external I/O  $\approx$  400 signals
  - \* Power & Ground  $\approx$  400
  - \* on 0.5mm pitch
  - \* Pads double as probe points



## 32 Chiplets

- Low power custom die
- 2 CPUs, router, memory controller
- 256-bit wide DDR2 memory ports
- I/O power is approximately 1pJ

## 4 stacked memory die

- Tezzaron Octopus




## SiCB is 52mm x 55mm

- Seven copper layers for power and interconnect
- Wafer bonded silicon circuit boards
- Power : 1.0v @ 2.6A; 1.5v @ 16A
- Connection to external FR-4 PCB through conductive elastomer with a 0.4mm pitch

## Goals

- Demonstrate energy optimized multiprocessor design methodology for DOE high-performance computing applications
- Provide proof-of -concept for rapid prototyping through the use of standard substrates, standard components, and shuttle runs
- Explore low-volume deployment.

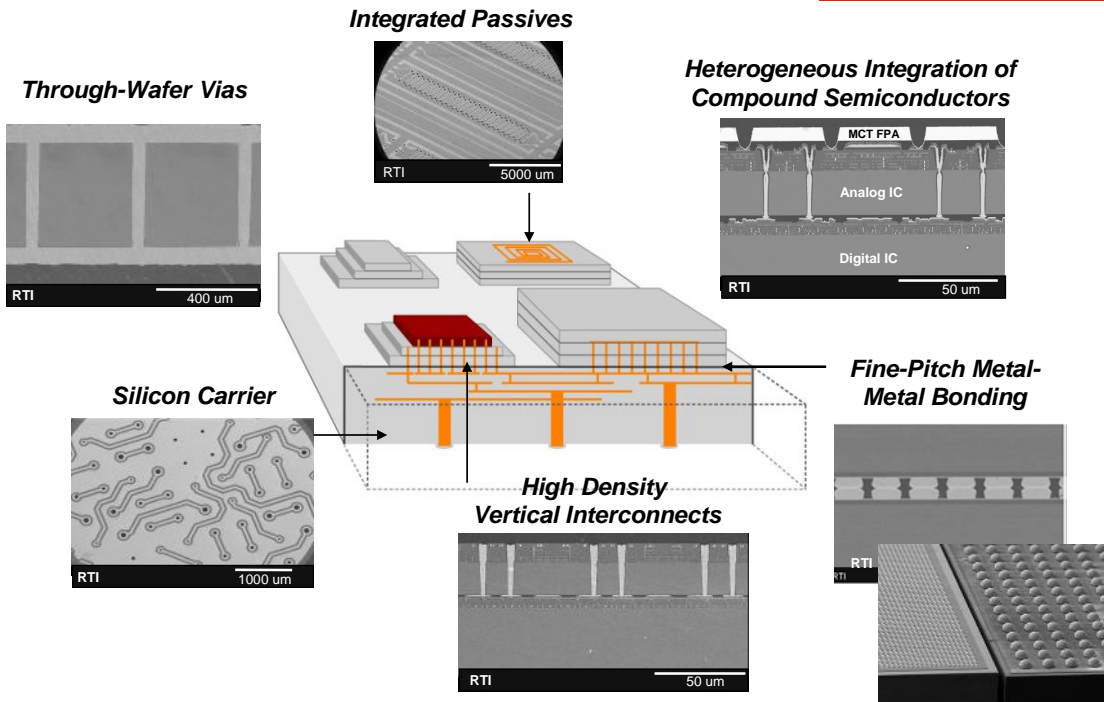
## CNT Anisotropic Conductive Films





- Anisotropic conductive films (ACFs) provide electrical connectivity in the vertical direction, and electrical insulation between adjacent pads
  - Current applications include low density, low performance applications such as displays and RFID
  - Current ACF technology does not simultaneously provide high conductivity and high pad density
- Higher performance, higher pitch solutions open new opportunities in 3-D chip stacking
  - High density, high performance ZIF connectors
  - Removable modules
  - CTE mismatch isolation at Si – FR4 interface

# 3D Fab Technology Tool Box

# 3D Microsystems Solutions



- These are the fab capabilities needed for complete 3D capability



“IC Foundry Agnostic”  
3-D Heterogeneous Integration

## Conclusions

- 3D silicon is here
- Lots of work to get to a smooth running prototype stage – 5 years
- Lots more work to get to routine production – 10 years
- Today we can do it – it's just engineering