




# Building Specialized HPC Systems with Composable **Chipllets**

Salishan Conference  
April 22 -25, 2024

Andreas Olofsson  
Zero ASIC  
[andreas@zeroasic.com](mailto:andreas@zeroasic.com)

# About me



1998-2006

Analog Devices

VLIW DSPs  
High Complexity  
Custom Logic  
**100 eng/chip**

2006-2008

Analog Devices

Camera ASICs  
Custom RISC  
**3 eng/chip  
(+support)**

2008-2017

Adapteva

Parallel DSPs  
Custom RISC  
Parallella,  
Epiphany  
**1-3 eng/chip**

2017-2020

DARPA



EDA/Design  
Chiplets  
**Goal: 1 eng/chip**

2020-()

Zero ASIC

Chiplets  
**Goal: 0 eng/chip**

# From my Salishan 2018 Silicon Compiler Presentation



	General Purpose (ref Dally, Kogge)	Scientific Machine (ref. Anton)	Machine Learning (trust me*)
OPS (Exa)	1	1	1
Memory (TB)	3,600	100	10
Power (MW)	68	2	0.2
Cost	\$500M	\$10M	\$2M

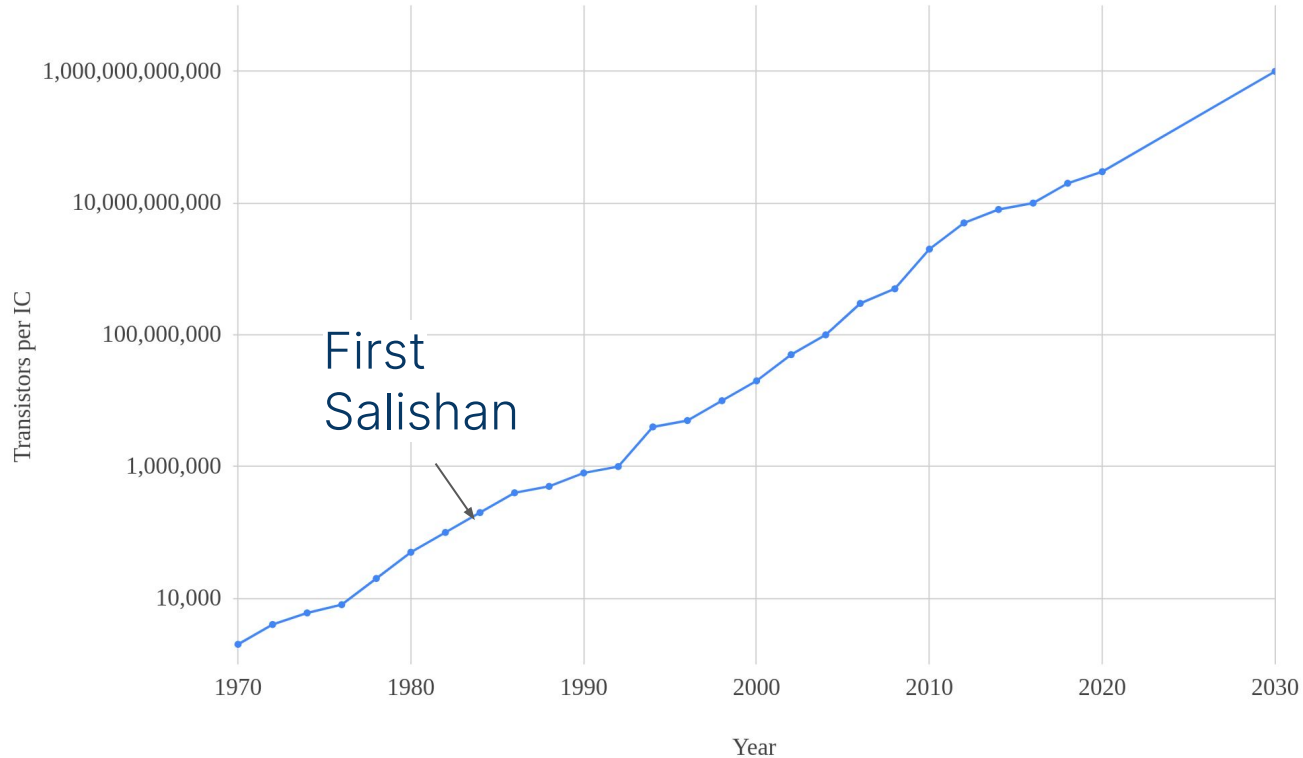
\*\$4/GB DRAM, 5 TOPS/W, shared MPW, IDEA/POSH working, market silicon pricing at 14nm

# Talk Outline

1. Troubling trends
2. Semiconductor economics 101
3. Chiplets to the rescue
4. Composable chiplets
5. Chiplet infrastructure
6. Chiplet examples
7. Chiplets for HPC



# Troubling Trends

# Moore's Law




The existential question....where will the next 1,000,000,000X come from?


# “Business-as-Usual Will Not Be Adequate”




1,000,000X



<10X




We all know what happens for  $\lim_{x \rightarrow 0^+} (1/x)$ , right?



*Optimist*


(or)



*Pessimist*




# Derivative: Escalating Fab Costs



Fab cost is a derivative of scaling complexity. As x goes to zero....


# Derivative: Escalating Mask Set Costs



Process complexity manifests itself as mask cost + unit cost increases.

# Derivative: Escalating Wafer Pricing


Cost per Wafer vs. Node



**\$250K - \$500K  
per FOUP!**

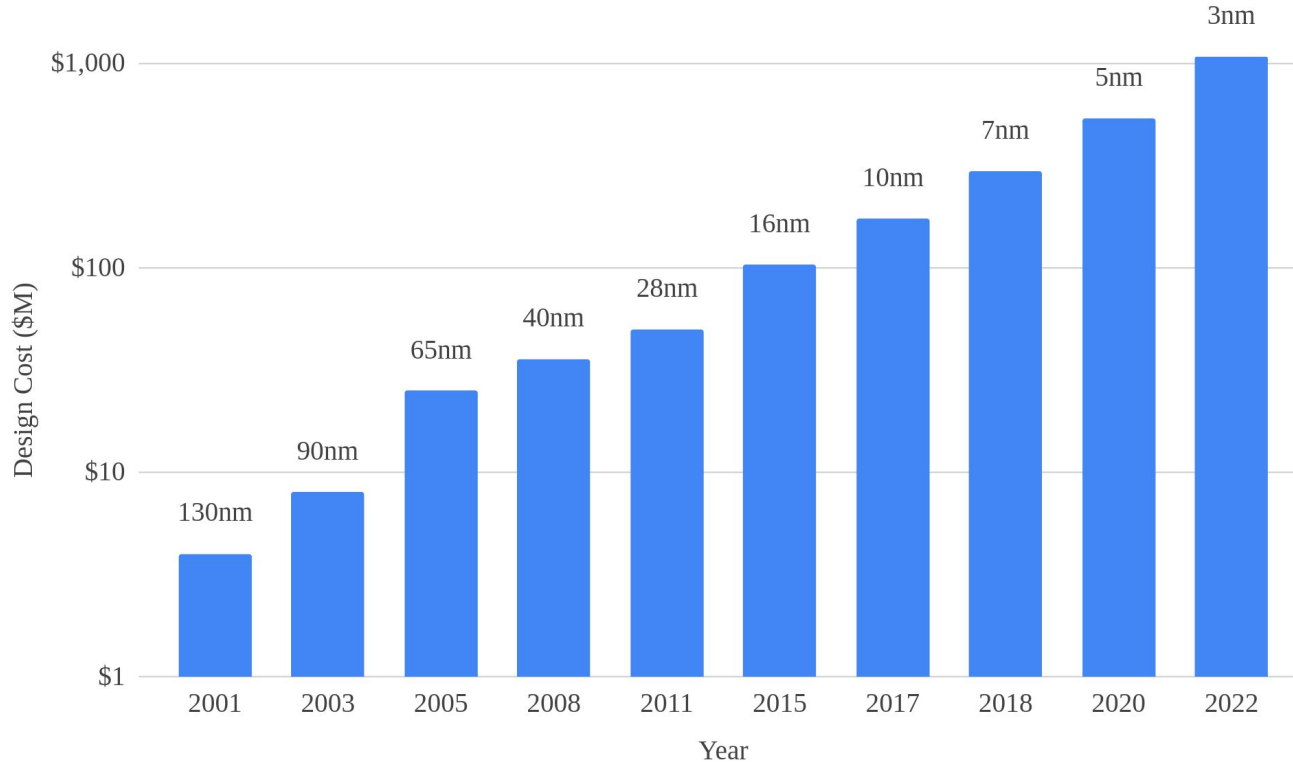


# Derivative: Transistor Costs are Saturating



Porting to the next node no longer automatic...

# Derivative: Escalating Design Costs



Design efficiency improvements have not kept pace with Moore's Law!

$$\text{Cost} = \sim mx + b$$


x = functions  
m = \$/function  
b = minimum cost (masks sets, qual,..)

# Important Existential Questions

- How much more computing performance do we actually need?
- Is the current cost of ASIC design actually a problem?
- What happens to the global economy when Moore's Law dies?
- What happens when chip design cost/time approaches zero?

# Semiconductor Economics 101

# Breaking Down Chip Design Cost

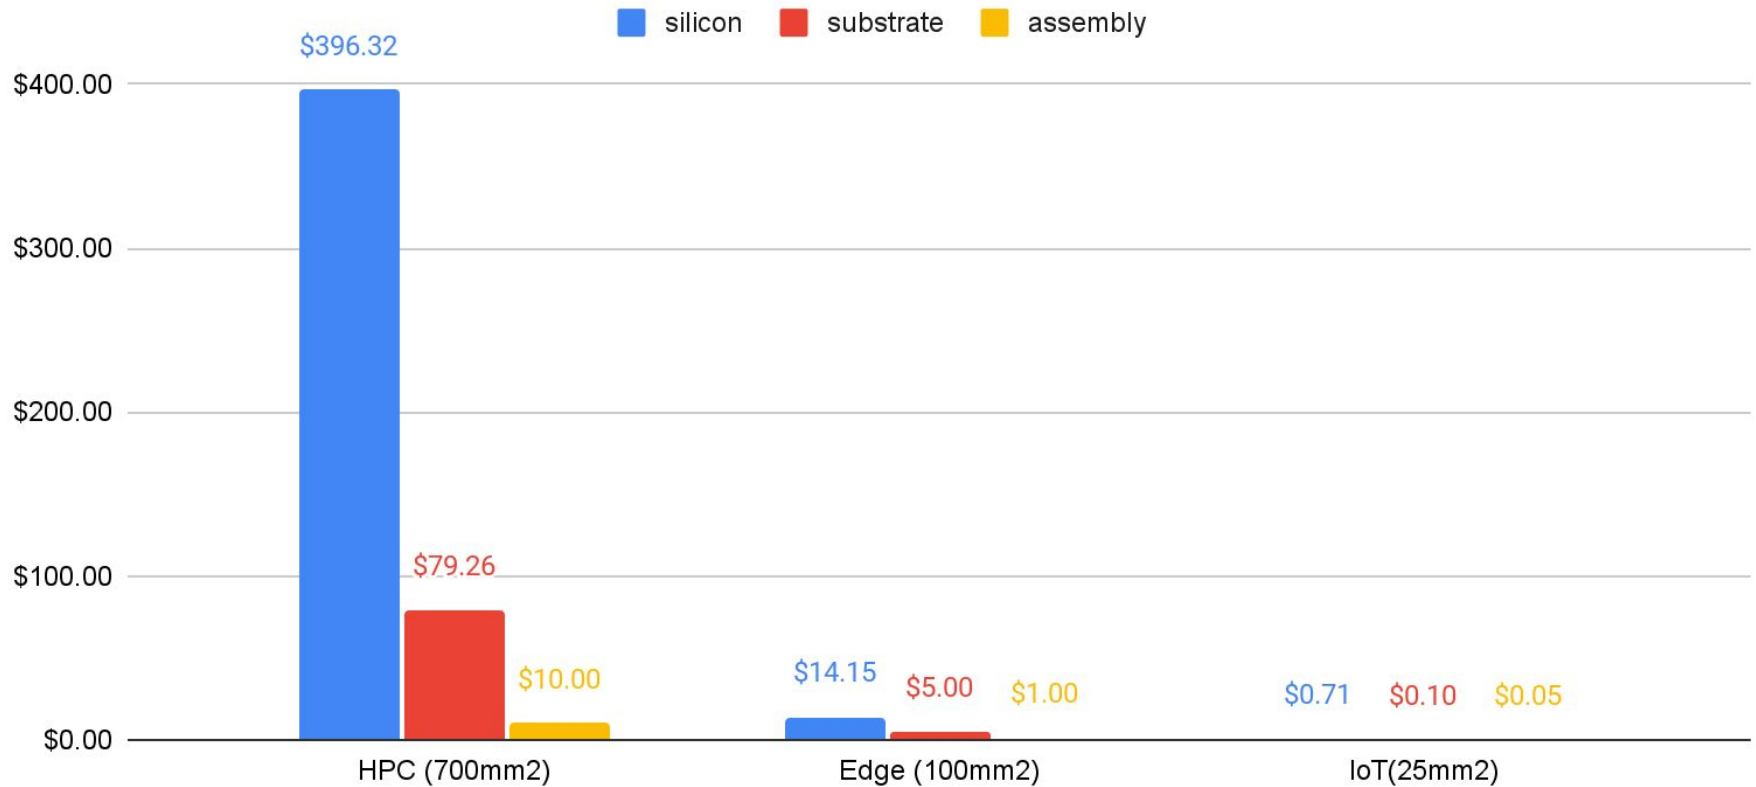


\* MPW drives up unit cost by 10-100X


\*\*Chiplets assume mix and match catalog



# Breaking Down Chip Unit Costs



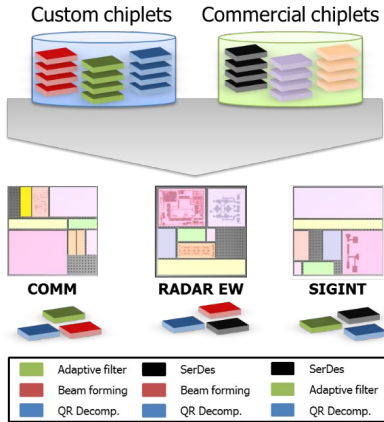

# Breaking Down Chip Design Time



\*\*Chipelets assumes a mix-and-match catalog



# Chipelets to the Rescue

# DARPA CHIPS (2016)



- ✓ A universal efficient interface standard
- ✓ SOTA manufacturing assembly
- ✓ A large and critical set of IP chiplets

REF: DARPA @ Semicom Design West 2019

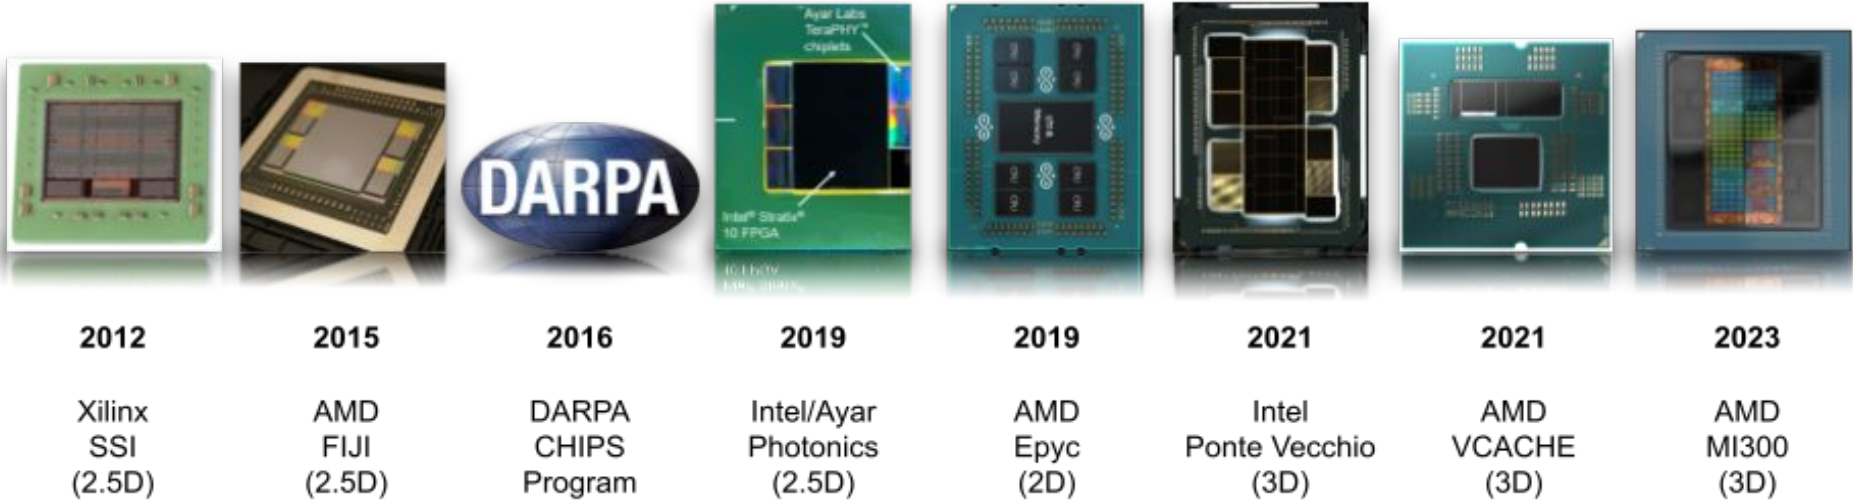


**Extend Moore's law**  
Scale out and scale down while managing yield

**Enable heterogeneous integration**  
Materials/processes, companies, geography, security


**Empower system integrators**  
Democratize access to leading edge silicon for system integrators

# Chiplet Status Report




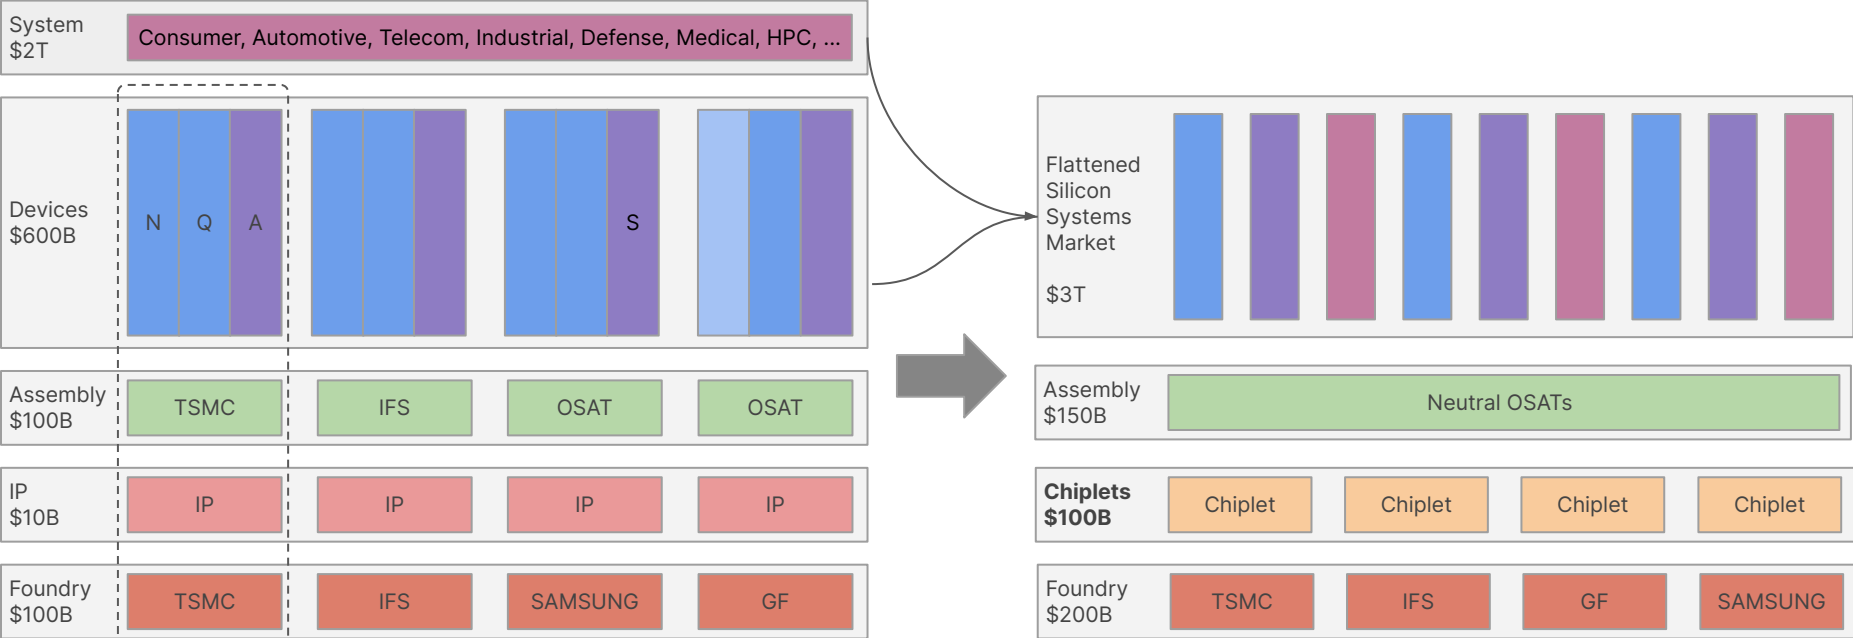
- Focused on unit cost (yield) and performance
- Closed ecosystems and limited reuse/composability

# Lockheed Martin Chiplet Study



*“A multi-year study showed that approximately **90%** of defense modules can be developed using just **7 types of chiplets**” -Lockheed Martin (distro A)*

# Opportunity: Flattening the \$2T Electronics Market



~100% of AI silicon produced at TSMC, a siloed ecosystem

- Chiplets flattens the ecosystem enabling a more efficient semiconductor economy
- Chiplets open up competition in System, Fabless, Assembly, and Foundry markets
- Chiplets offers a solution to the global geopolitical semiconductor supply chain problem
- Chiplets expands the semiconductor market by lowering the barrier to innovation


# Critical Chiplet Question?

- What is the chiplet tax (cost, energy, performance)?
- What functions are needed to cover most applications?
- What is the optimal chiplet size?
- What is the optimal packaging technology?
- When will the tipping point come for chiplets?




# Composable Chiplets

# Chiplets: The New Amino Acids of Silicon Systems




Low Cost FPGA




Heterogeneous FPGA




High Performance FPGA




Mixed Signal FPGA




Old School DSP



Heterogeneous DSP




Microcontroller



Microprocessor




AI FPGA



AI Microcontroller



AI DSP




AI ASIC




Low Cost SoC




High Performance SoC



Von Neumann CPU (CPU)




General Purpose GPU (GGPU)




Manycore CPU




Application Specific Integrated Circuit (ASIC)




Field Programmable Gate Array (FPGA)




Coarse Grained Reconfigurable Array (CGRA)



Processing-In-Memory (PIM)



Application Specific Processor (ASIP)



Heterogeneous System-On-Chip (SoC)

*“Finite investment, infinite possibilities”*

# Mix-and-Match Composability Math

	Repeats Allowed	No Repeats
Permutations (order matters)	$n^r$	$n!/(n-r)!$
Combinations (order doesn't matter)	$(r+n-1)!/r!(n-1)!$	$n!/r!(n-r)!$




n=20, r=16	Repeats Allowed	No Repeats
Permutations	6.55E+20	1.01E+17
Combinations	4.06E+09	4.85E+03

# Our Approach: Composable system of “LEGO” chiplets

## eFabric Substrate

- “Breadboard for chiplets”
- Active silicon (ie 3D)
- Fixed grids
- Fixed sizes (S, M, L, ...)




## Standard Chiplets


- UCIe standard
- Small, discretized
- 2D IO chiplets
- 2D/3D processing

ARM	RV64	DSP	RV32
SRAM	NVM	L2	DRAM
DDR	SRDS	PCIE	MAC
ML	IPU	GPU	H264
FPGA	ADC	JESD	IOP

- **Billions of unique customer configurations**
- First “no-code” chip design platform
- First full stack chiplet interface
- First digital twin chiplet platform
- Four mix-and-match 3D test chips in GF12LP




# eFabric Cross Section




Active silicon interposer with TSVs, 45um 3D microbumps, 110um IO bumps, organic interposer, matched height silicon stacks

# Chiplet RFC #1: Mechanical standards for Chips



**Key Decisions**


- Min Size:** 1 mm<sup>2</sup> is 100MTr (good size)
- Space:** 100um chiplet space (manufacturing)
- Logistics:** Typical MPW sizes ~2x2



**Lego Brick Standard Unchanged Since 1958!**

$P = 8.0 \text{ mm} = \frac{5}{6} \times H = 2.5 \times h$   
 $h = 3.2 \text{ mm} = \frac{1}{3} \times H = 0.4 \times P$   
 $H = 9.6 \text{ mm} = 3 \times h = 1.2 \times P$   
 $2 \times P - 0.2 \text{ mm} = 15.8 \text{ mm}$   
 $P - 0.2 \text{ mm} = 7.8 \text{ mm}$

Name	Size
X1	0.95mm
X2	2mm
X3	3.05mm
X4	4.1mm
X5	5.15mm

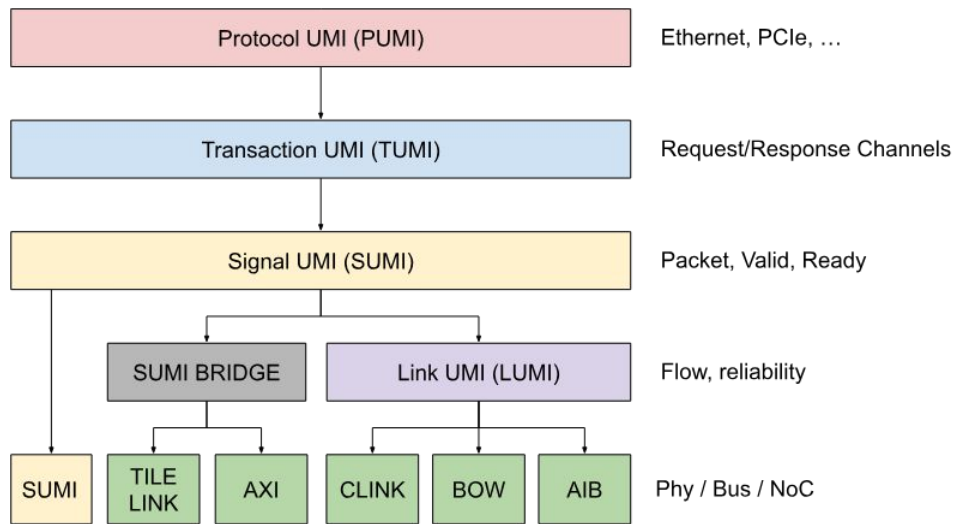
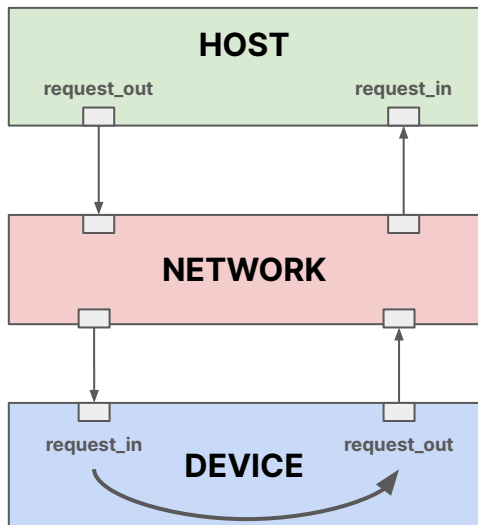


# Chiplet RFC #2: Fixed footprints for chips

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
45	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
A		NIO15	NIO14	NIO13	NIO12	NIO11	NIO10	NIO9	NIO8	VSS	NIO7	NIO6	NIO5	NIO4	NIO3	NIO2	NIO1	NIO0		-405
B	EIO0	VSS	NAN1	VSS	NVCC	NPTN0	NPTP0	VSS	NNC0	NNC1	NNC2	VSS	NPTP1	NPTN1	NVCC	VSS	WAN0	VSS	WIO0	-360
C	EIO1	NAN0	NVDDA	TX0	TX1	TX3	TX6	TX9	TX13	EN3D	RX13	RX9	RX6	RX3	RX1	RX0	WVDDA	WAN1	WIO1	-315
D	EIO2	VSS	RX48	VDD	TX2	TX4	TX7	TX10	TX14	VSS	RX14	RX10	RX7	RX4	RX2	VDD	TX16	VSS	WIO2	-270
E	EIO3	EVCC	RX49	RX50	VSS	TX5	TX8	TX11	TX15	VDDX	RX15	RX11	RX8	RX5	VSS	TX18	TX17	WVCC	WIO3	-225
F	EIO4	EPTN1	RX51	RX52	RX53	VSS	TXS0	TX12	TXC0	VSS	RXC0	RX12	RXS0	VSS	TX21	TX20	TX19	WPTN0	WIO4	-180
G	EIO5	EPTP1	RX54	RX55	RX56	RXS3	VDD	VSS	VDD	STAT0	VDD	VSS	VDD	TXS1	TX24	TX23	TX22	WPTP0	WIO5	-135
H	EIO6	VSS	RX57	RX58	RX59	RX60	VSS	VDD	VSS	CTRL0	VSS	VDD	VSS	TX28	TX27	TX26	TX25	VSS	WIO6	-90
J	EIO7	ENC2	RX61	RX62	RX63	RXC3	VDD	VSS	VDD	CLK0	VDD	VSS	VDD	TXC1	TX31	TX30	TX29	WNC0	WIO7	-45
K	VSS	ENC1	ENC3	VSS	VDDX	VSS	STAT3	CTRL3	CLK3	NRST	CLK1	CTRL1	STAT1	VSS	VDDX	VSS	WNC3	WNC1	VSS	0
L	EIO8	ENC0	TX61	TX62	TX63	TXC3	VDD	VSS	VDD	CLK2	VDD	VSS	VDD	RXC1	RX31	RX30	RX29	WNC2	WIO8	45
M	EIO9	VSS	TX57	TX58	TX59	TX60	VSS	VDD	VSS	CTRL2	VSS	VDD	VSS	RX28	RX27	RX26	RX25	VSS	WIO9	90
N	EIO10	EPTP0	TX54	TX55	TX56	TXS3	VDD	VSS	VDD	STAT2	VDD	VSS	VDD	RXS1	RX24	RX23	RX22	WPTP1	WIO10	135
P	EIO11	EPTN0	TX51	TX52	TX53	VSS	RXS2	RX44	RXC2	VSS	TXC2	TX44	TXS2	VSS	RX21	RX20	RX19	WPTN1	WIO11	180
Q	EIO12	EVCC	TX49	TX50	VSS	RX37	RX40	RX43	RX47	VDDX	TX47	TX43	TX40	TX37	VSS	RX18	RX17	WVCC	WIO12	225
R	EIO13	VSS	TX48	VDD	RX34	RX36	RX39	RX42	RX46	VSS	TX46	TX42	TX39	TX36	TX34	VDD	RX16	VSS	WIO13	270
S	EIO14	EAN1	EVDDA	RX32	RX33	RX35	RX38	RX41	RX45	SNC3	TX45	TX41	TX38	TX35	TX33	TX32	SVDDA	SAN0	WIO14	315
T	EIO15	VSS	EAN0	VSS	SVCC	SPTN1	SPTP1	VSS	SNC2	SNC1	SNC0	VSS	SPTN0	SPTP0	SVCC	VSS	SAN1	VSS	WIO15	360
U		SIO15	SIO14	SIO13	SIO12	SIO11	SIO10	SIO9	SIO8	VSS	SIO7	SIO6	SIO5	SIO4	SIO3	SIO2	SIO1	SIO0		405
V	-405	-360	-315	-270	-225	-180	-135	-90	-45	0	45	90	135	180	225	270	315	360	405	

45um bump pitch	Symmetrical footprint	Passthrough Signals	Multiple clocks
64 bit parallel I/O	Multiple power rails	4 Analog domains	Sideband channel

# Chiplet RFC #3: Universal Memory Interface





<b>Everything is memory</b>	Open source	Packet Based	Low Latency/High Perf
<b>Optimized for chiplets</b>	Latency Insensitive	Built-In Atomics	Bridges to AXI, TileLink

<https://github.com/zeroasiccorp/umi>




# Chiplet RFC #4: A lightweight electrical interface

- **UCle** specification (v1.1)
  - GF12LP
  - Up to 8GT/s maximum, 16b wide
  - **Tightly coupled mode only**
  - Designed to UCle FDI/RDI interface
- **Universal Memory Interface** protocol
- Simulated using open source tools (**Verilator**, **Xyce**, **Icarus**, and **Switchboard**)



# Standards

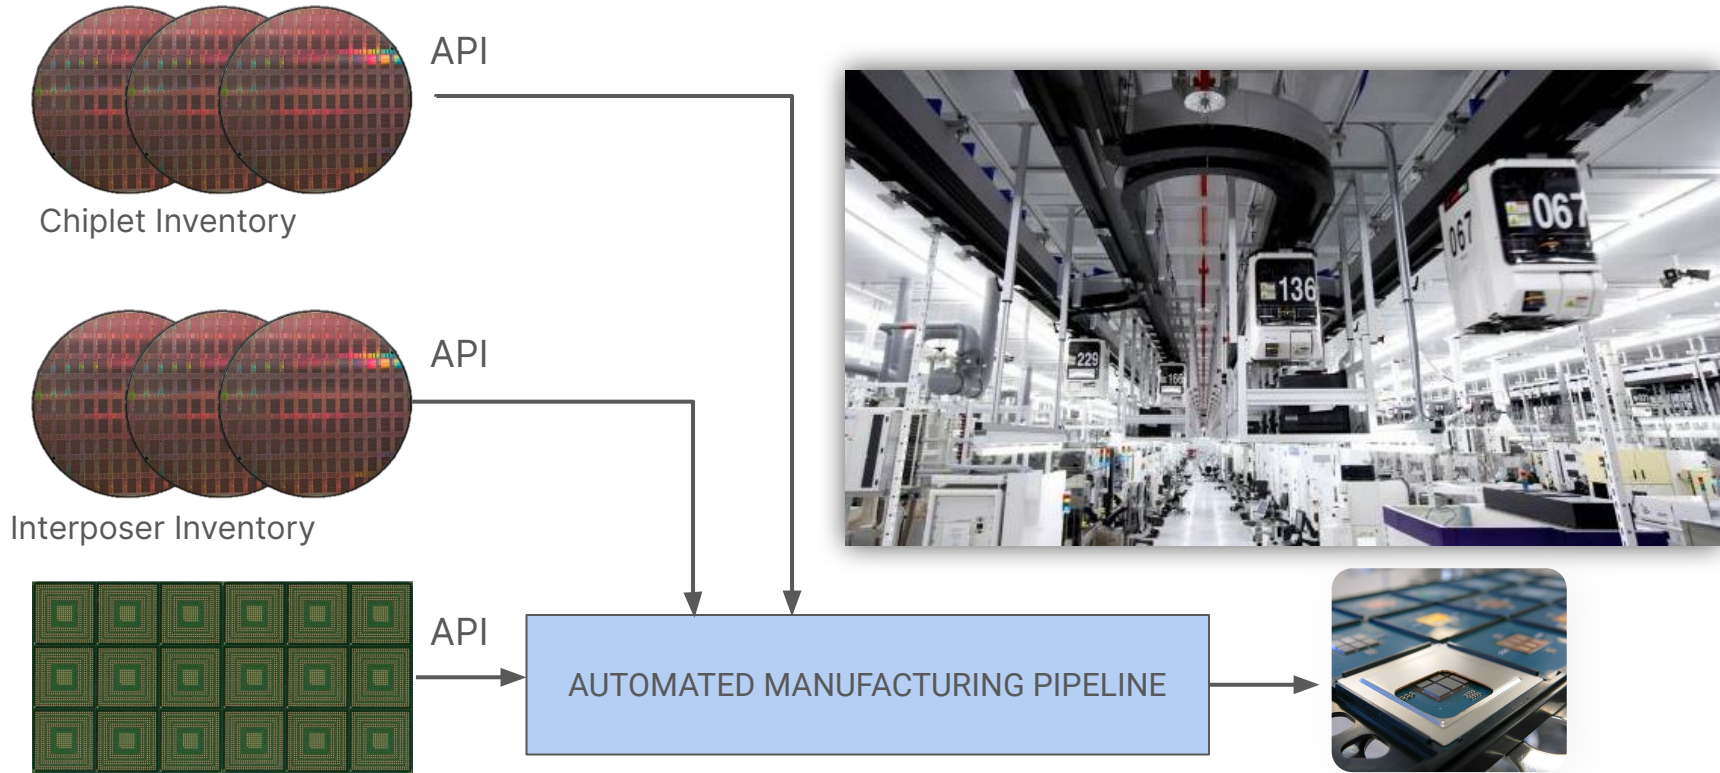


Perception



Reality

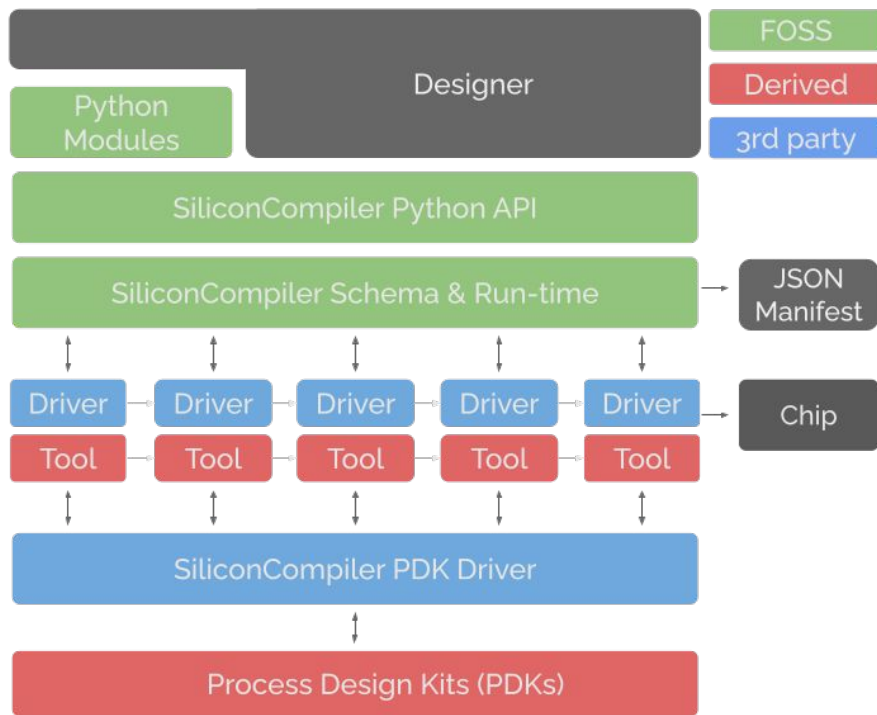
# 10 Year Plan: Built-To-Order Custom Silicon in 24hrs



Standardized automation is the only way to fix the broken economics of on-shore low-volume-high mix manufacturing

# Chiplet Infrastructure

# SiliconCompiler: Automating chiplet layout



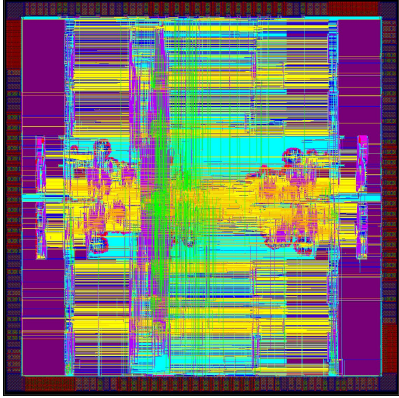


<https://github.com/siliconcompiler>

```
$ pip install --upgrade siliconcompiler
```

```
$ sc -target asic_demo -remote
```

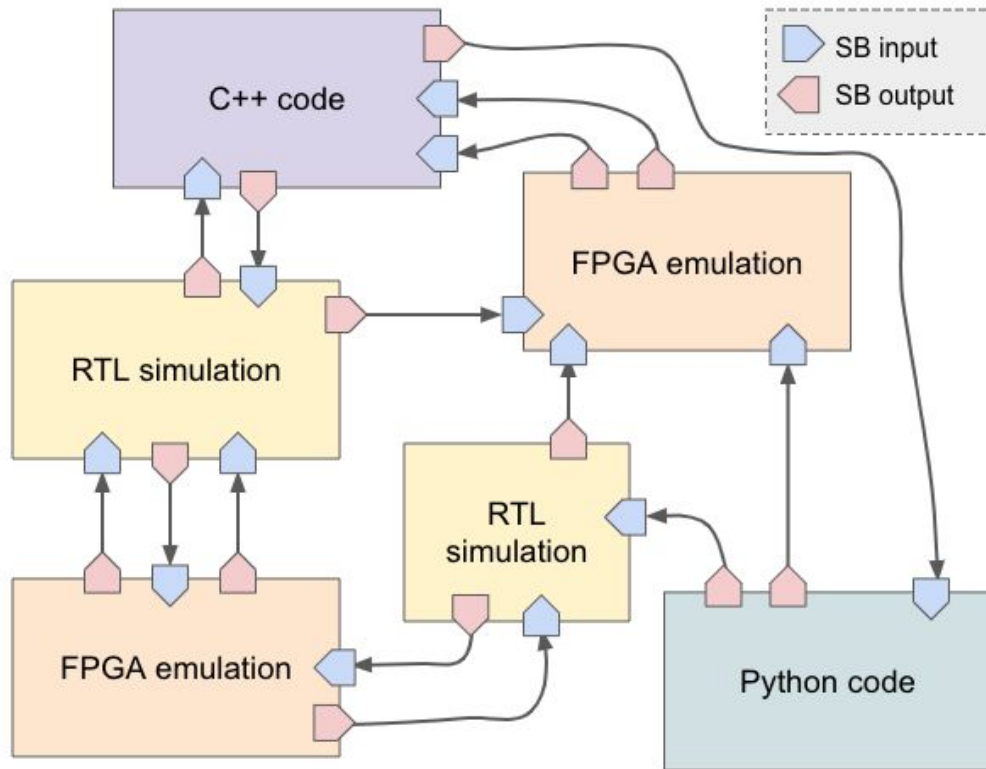
# SiliconCompiler Results



	GOTLAND	MAUI	KODIAK
TYPE	CPU	FPGA	MEMORY
TECHNOLOGY	GF12LP	GF12LP	GF12LP
SIZE	2 x 2 mm	2 x 2 mm	2 x 2 mm
ORIGIN	UCB - ROCKET	ZA	ZA
METRIC	Quad Core RV64GC CPU	-	3MB
DESIGNERS	2	2	2
WALL TIME	< 3 weeks	< 8 weeks	< 8 weeks
RUN TIME	< 24hrs	< 24hrs	< 24hrs

# Switchboard: Addressing the Chiplet Validation Problem

- Heterogeneous simulation framework
- Latency insensitive protocol (ready/valid)
- Fast shared memory queues
- Supports RTL, FPGAs (HIL), SW models)
- UMI implementation
- Python bindings
- 10x faster than commercial emulators
- **Deployed in AWS**
  - 0.2us host latency
  - 4us host-fpga latency
- [github.com/zeroasiccorp/switchboard](https://github.com/zeroasiccorp/switchboard)
- **Demo:** [zeroasic.com/emulation](https://zeroasic.com/emulation)


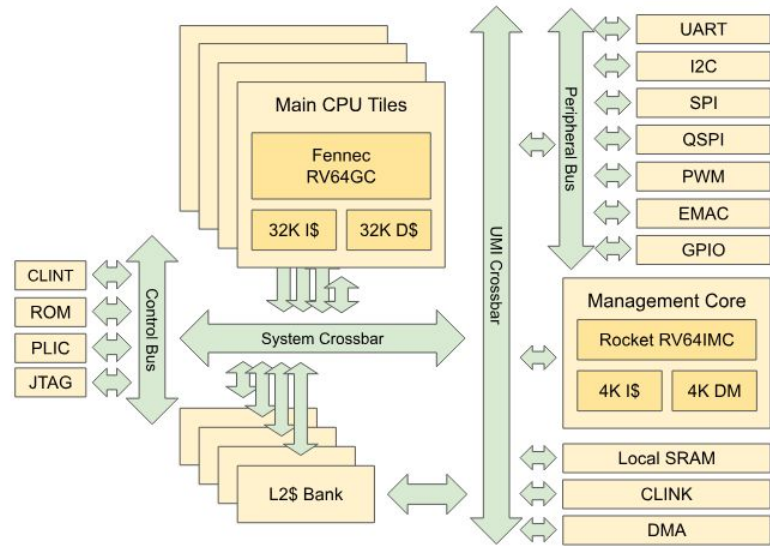


# Chiplet Examples




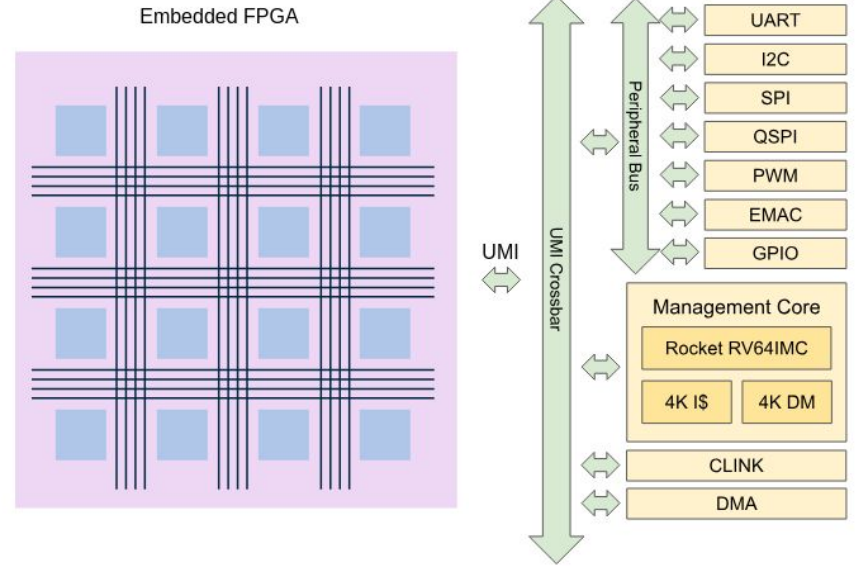
# Quad-Core RV64GC CPU

- Quad-core RISC-V RV64GC cpu cluster
- In-order dual-issue pipeline
- 32KB L1 instruction cache
- 32KB L1 data cache
- 1MB L2 cache
- L2 cache can be reconfigured as scratchpad
- 3.6 CoreMarks/MHz
- Up to 1.5GHz core operating frequency
- On-chip general purpose DMA
- UMI/UCIe I/O ports
- Standard I/O Peripherals
- Open source toolchain/demo:
  - (see RV eco-system)
  - [zeroasic.com/emulation](http://zeroasic.com/emulation)
- GF12LP process
- **2mm x 2mm**

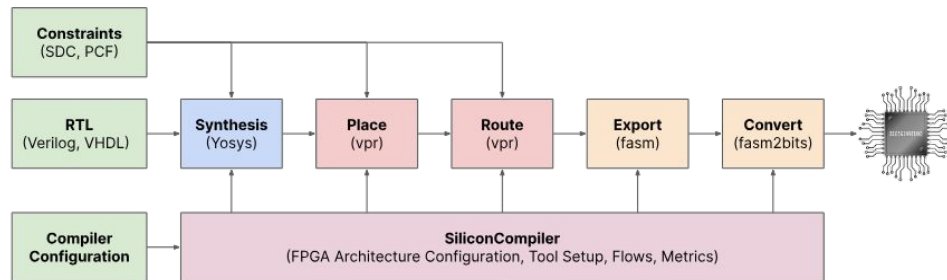


# 26K LUT Heterogeneous FPGA

- 26K LUTs
- 512KB BRAM
- 128 MACs
- Open source toolchain/demo:
  - [github.com/zeroasiccorp/logik](https://github.com/zeroasiccorp/logik)
  - [zeroasic.com/emulation](https://zeroasic.com/emulation)
- UMI/UCle I/O ports
- GF12LP process
- **4mm x 4mm**

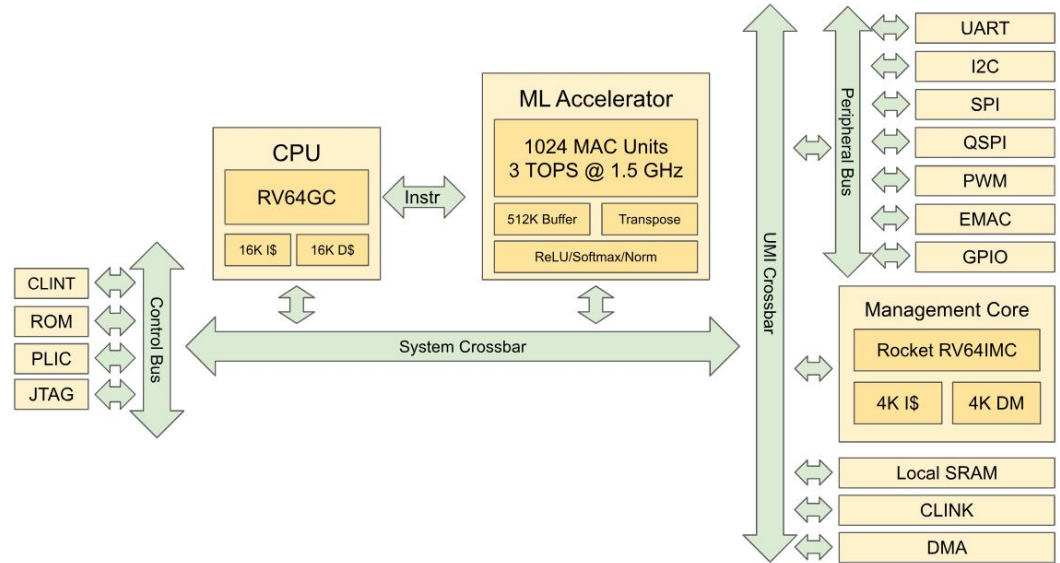


Design	4-LUTs	Flip Flops
picoRV32	3974	1001
UART	217	79




# 3 TOPS ML Accelerator

- **3 TOPS** INT8 AI accelerator
- Hardware accelerated ReLU/Softmax nonlinearities
- Hardware accelerated transpose
- 512KB on-chip buffers
- RISC-V RV64GC CPU
  - 16KB L1 instruction cache
  - 16KB L1 data cache
- UMI/UCIe I/O ports
- Standard I/O Peripherals
- Demo:
  - [zeroasic.com/emulation](https://zeroasic.com/emulation)
- GF12LP process
- **2mm x 2mm**




# Chiplets for HPC

# Advice: Start small....




- Quad-Core 64b RV64GC CPU
- 20K LUTs
- 4 \* 32Gbps SERDES

# Example: Composable FPGAs



	Current	Composable
Min Area	952 mm <sup>2</sup> (est)	4 mm <sup>2</sup>
Min Power	10 W (est)	< 1 W
Project Cost	N/A	\$100K
Project Wall Time	N/A	3 months

# Moore's Law → Olofsson's Law\*



*“The number of chiplets in a package will double every two years.”*

\* See Stiglers' Law