The Demise and Resurrection of Analog Computing

Yipeng Huang

PhD Candidacy Exam

April 8, 2015

My Research

- **Hybrid continuous-discrete computing**

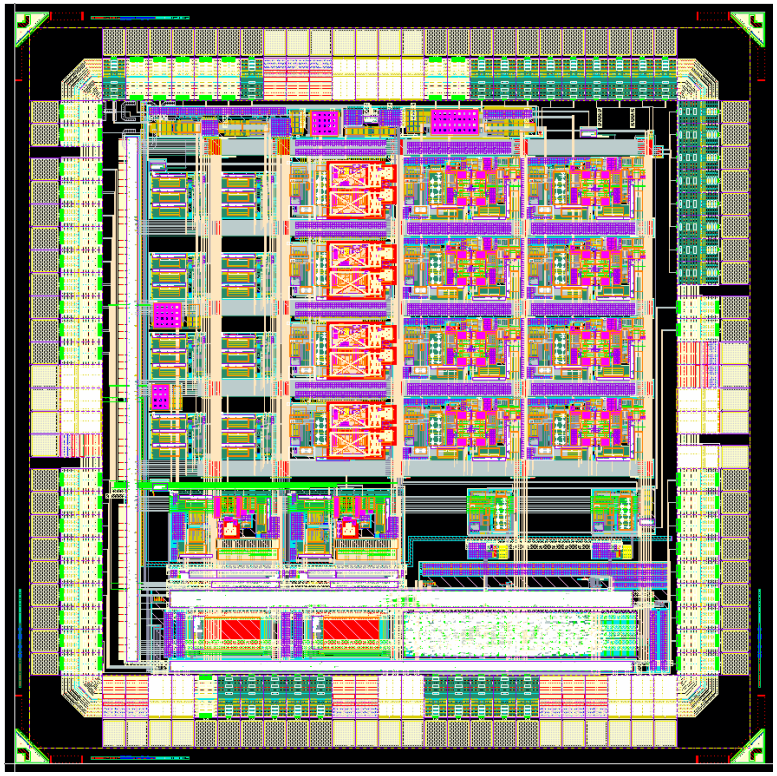

- **Robotics workload characterization**

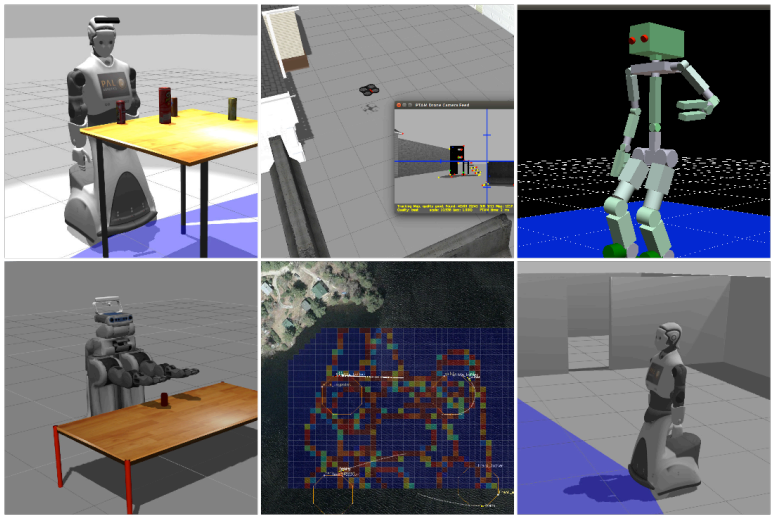

Outline & Purpose

	Digital	Analog
Workloads	[2], [3], [9], [15], [22], [28]	[26]
Programming model	[14], [17], [21], [25]	[7], [8]
Datapath & network	[1], [5], [13], [16], [19]	[4], [8]
Memory hierarchy	[9], [16], [18], [19], [22], [28]	[4], [26]
Core organization	[6], [16], [17], [24], [28]	[4]
Logic & pipelines	[10], [11], [12], [16], [20], [22], [23], [24]	[8], [23], [26], [27]
Physics & devices	[6]	[26]

Outline & Purpose

	Digital	Analog
Workloads	[2], [3],], [22], [28]	[26]
Programming model	[14], [1	[25]	[7], [8]
Datapath & network	[1], [5],], [16], [19]	[4], [8]
Memory hierarchy	[9], [16	19], [22], [28]	[4], [26]
Core organization	[6], [16	24], [28]	[4]
Logic & pipelines	[10], [1 [22], [23	[16], [20],	[8], [23], [
Physics & devices	[6]	

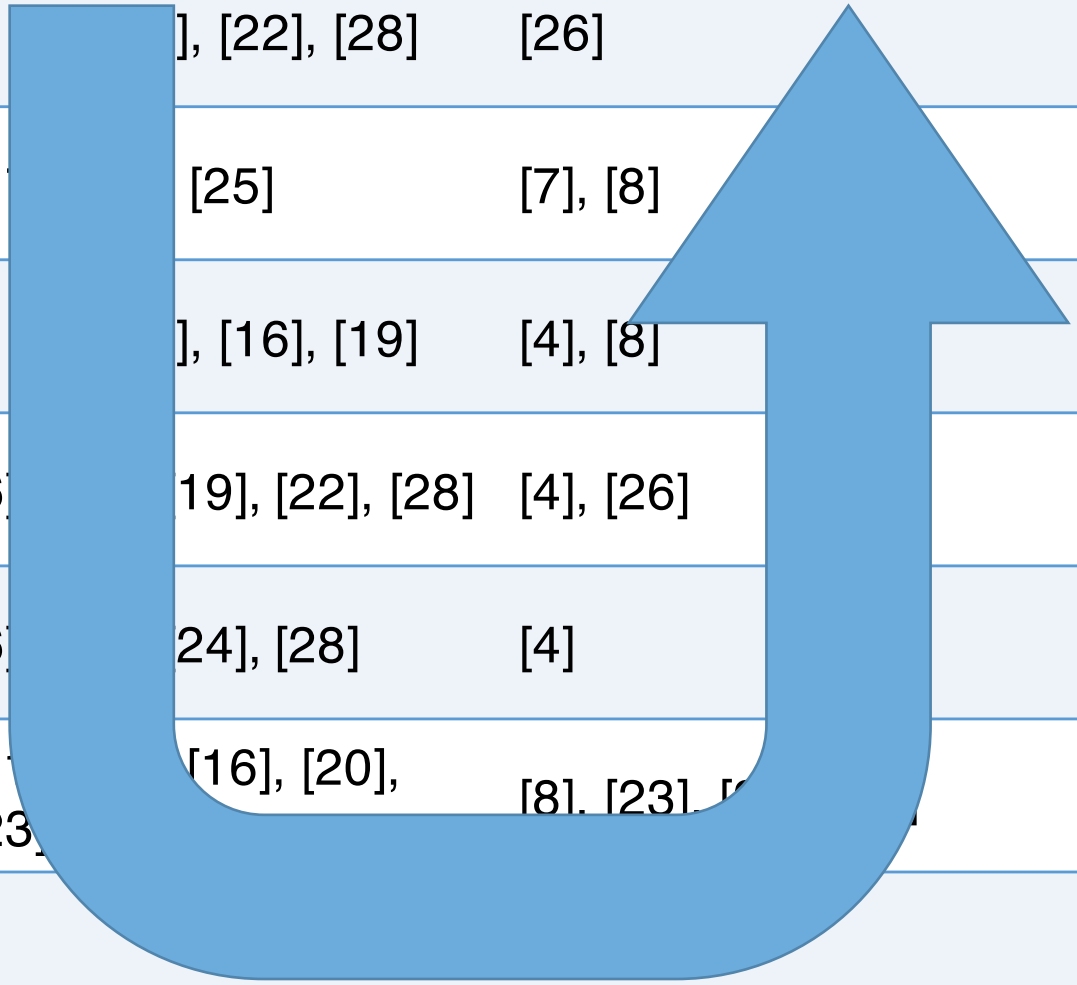

Expose bottlenecks in digital & provide ideas for analog

Digital: Workloads & Programs

- **In cloud workloads, less locality & predictability**
  - Instructions, working data set exceeds capacity of fast cache [9], [19]

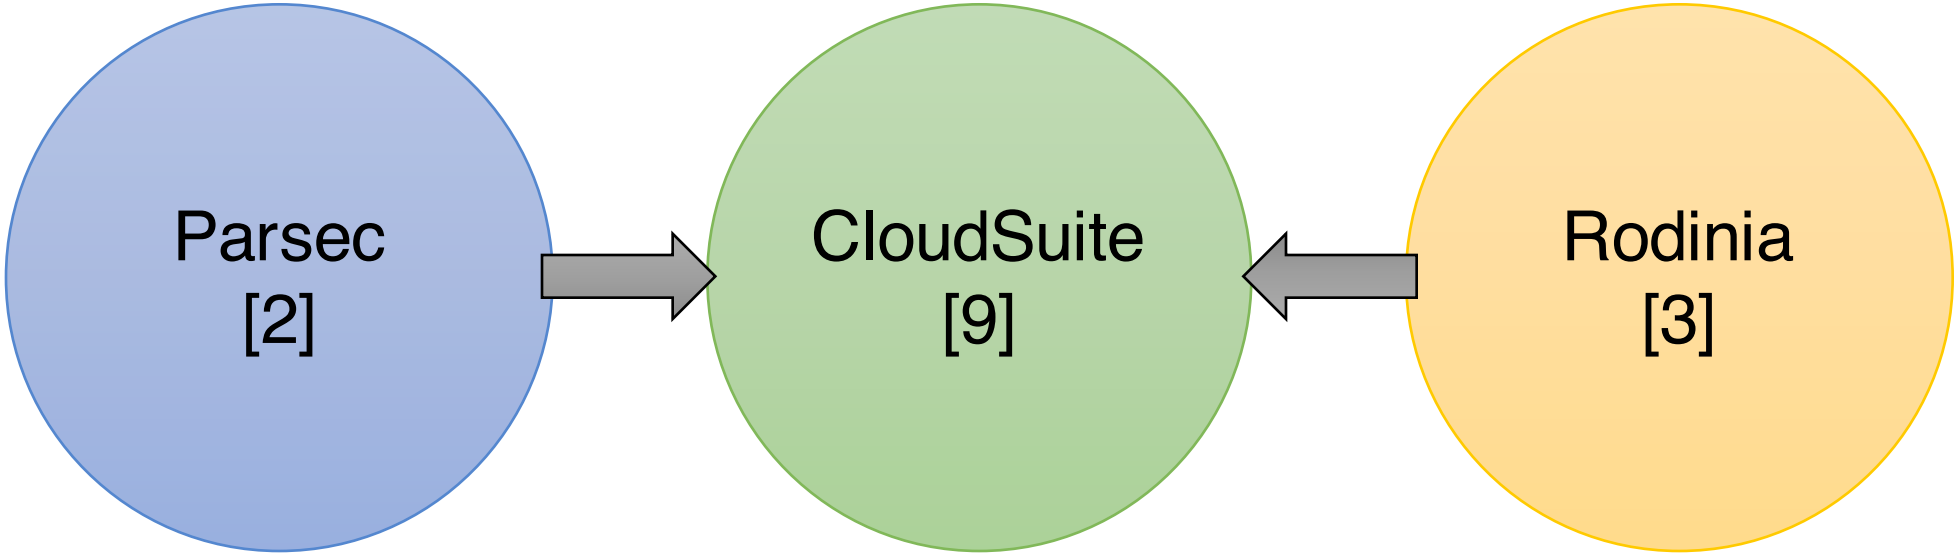

- **Structured Parallelism is Decreasing**
  - Less structured parallelism: scientific computation, linear algebra
  - More amorphous parallelism: graph traversal, dynamic programming
  - Requires new metrics for parallelism [21]

Workloads present decreasing traits that digital computers exploit

Digital: Datapath & Network

- **Optimal network design depends on workload**
  - Complexity may help: Balfour et.al. find meshed, repeated, heterogeneous, torus network optimal [1].
  - Or, it may not: Lotfi-Kamran et al. find a basic crossbar suitable for their design [19].
  - In domain specific or application specific accelerators, datapath can be further refined [22].

While cores are abundant, interconnect is scarce, and no design is best

Digital: Memory & Organization of Cores

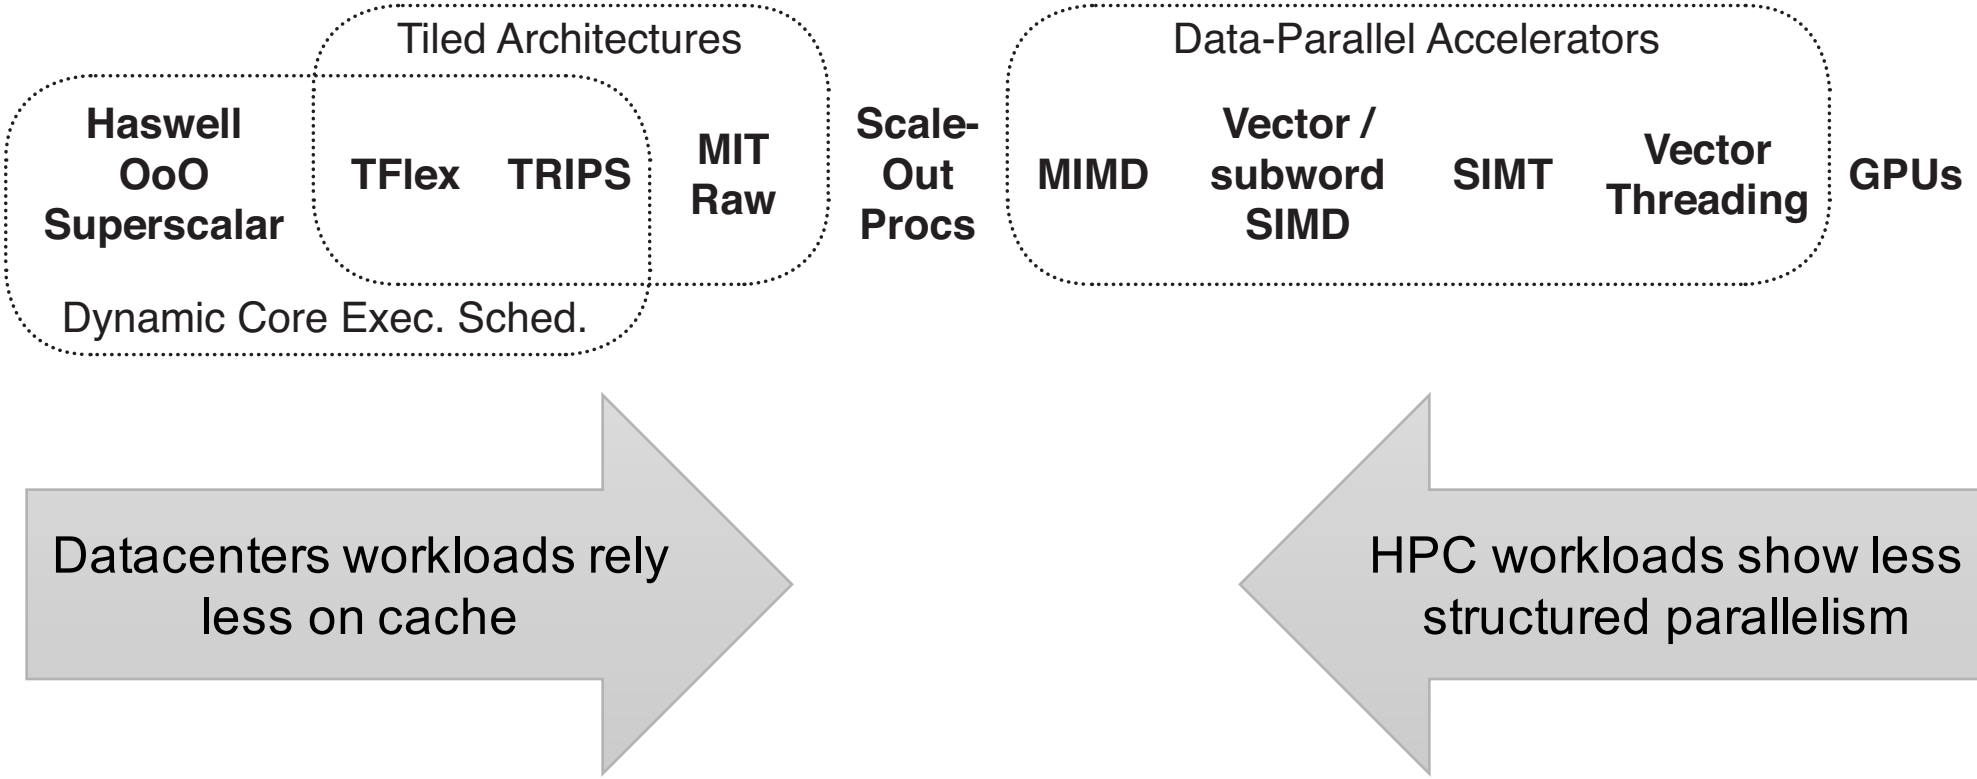

OoO Superscalar vs. GPU dichotomy has broken down

Digital: Logic & Pipeline Static Power

- **To reduce static power, must shut circuits off or use DVFS**
  - Challenge is conveying to microarchitecture the desired throughput, latency, accuracy
- **Throughput**
  - Tao parallelism analysis assists disabling pipelines & whole cores [21]
- **Latency**
  - Slack analysis may guide dynamic voltage frequency scaling [10]
- **Accuracy**
  - Benign data approximation permits using lower power pipelines [23], [26], [27]

Current ISAs communicate little info to guide energy conservation

Digital: Logic & Pipeline Dynamic Power

- **Must minimize pipeline power spent outside of function units**
  - Custom instructions [12] [24], Dyser, ASICs [11] [22], all are attempts to minimize fetch, decode, pipeline, control, register power

Stored program computers incur high overhead outside of ALUs

Outline & Purpose

	Digital	Analog
Workloads	[2], [3], [9], [15], [22], [28]	[26]
Programming model	[14], [17], [21], [25]	[7], [8]
Datapath & network	[1], [5], [6], [13], [16], [19]	[4], [8]
Memory hierarchy	[9], [16], [18], [19], [22], [28]	[4], [26]
Core organization	[6], [16], [17], [24], [28]	[4]
Logic & pipelines	[10], [11], [12], [16], [20], [22], [23], [24]	[8], [23], [26], [27]
Physics & devices	[6]	[26]

Expose bottlenecks in digital & provide ideas for analog

Digital Bottlenecks Summary

Workloads present decreasing traits that digital computers exploit

While cores are abundant, interconnect is scarce, and no design is best

OoO Superscalar vs. GPU dichotomy has broken down

Current ISAs communicate little info to guide energy conservation

Stored program computers incur high overhead outside of ALUs

Physics & Devices

- **Dennard's scaling: engineering with no tradeoffs**
  - Shrinking transistor increased speed and decreased supply voltage
  - Decreased voltage decreased power consumption
- **Digital limitations**
  - Frequency, voltage scaling now only moderate [6]
  - No blockbuster devices in near future
- **Transistors no longer work as perfect switches**
  - But computation isn't limited to using binary switches—
  - Resistor: multipliers for neural networks
  - Capacitor: integrators for equation solvers
  - Memristor: storage devices

Digital vs. Analog Data

- **Benefit of digital data: high signal-to-noise**
  - Allows abstractions expressed as logic
  - Error correction

- **Motivation for analog representation:**
  - More interaction between cyber & physical
  - Not just as peripherals!

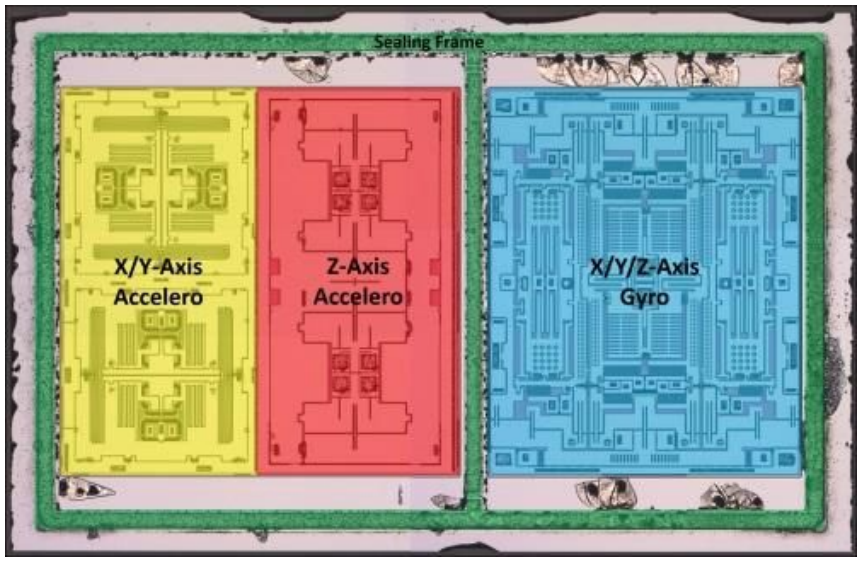

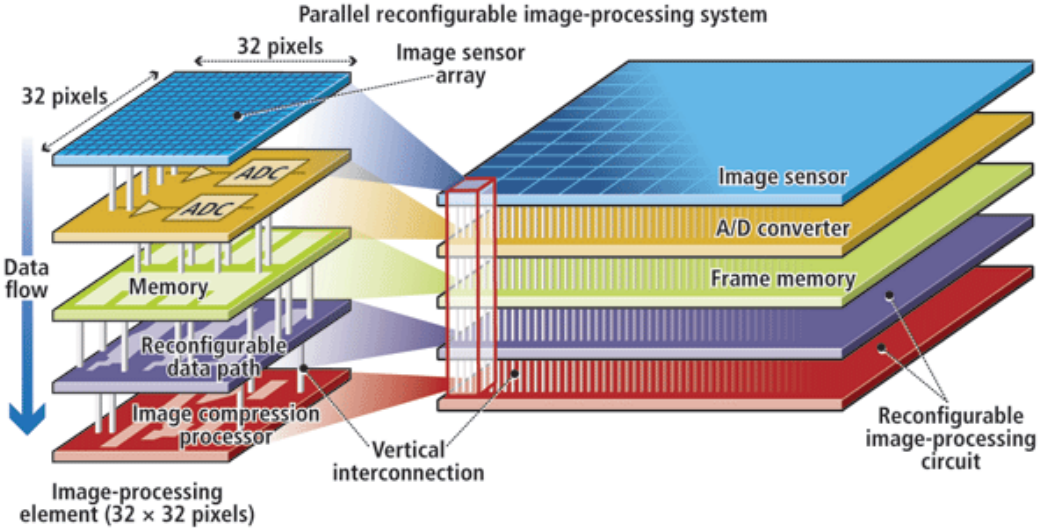

image credits: www.vision-systems.com m.eetindia.co.in

Analog: Logic & Pipelines

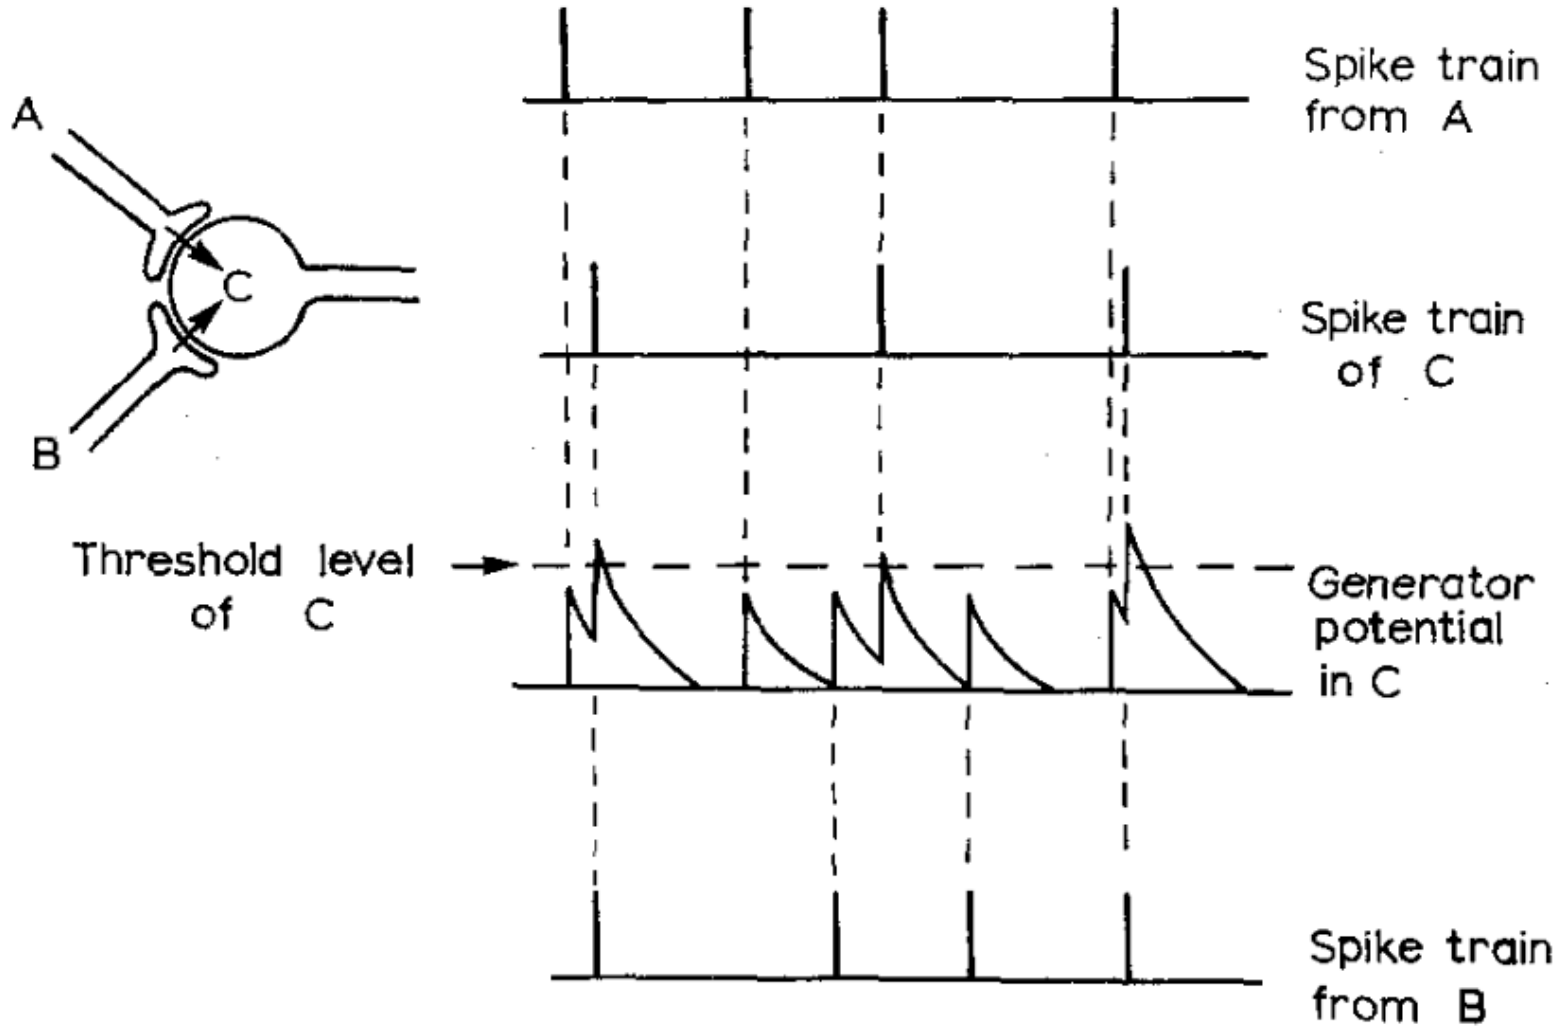

Analog: Logic & Pipelines

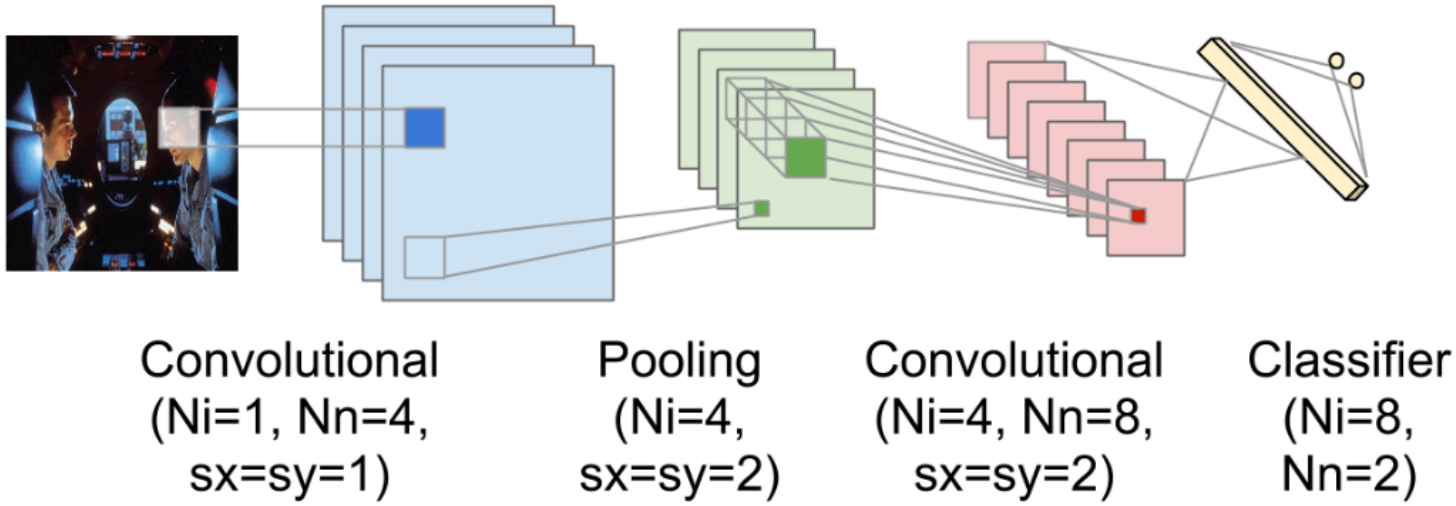

Figure 1. *Neural network hierarchy containing convolutional, pooling and classifier layers.*

Analog: Organization of Cores

- **Analog neural networks support digital computers at various granularities**
  - Pipeline level: Neuflow [8], NPUs [26]
  - Standalone accelerator: DianNao
  - Standalone system: IBM TrueNorth [7]
  - Standalone system, integrating main memory: DaDianNao [4]

Analog: Datapath & Networks

- **Analog datapath & networks are often circuit switched...**
  - Program instructions are embodied in datapath topology at runtime [4], [8]
- **…permitting spike train encoding of data on fewer wires...**
  - Floating point value is encoded in spike train density [7]
  - Makes operations such as multiplication and integration efficient
- **...and prefers asynchronous continuous time data transmission.**
  - When using spike, voltage, or current encodings, there is no notion of clock cycles

Analog: Programs & Workloads

- **Analog requires new programming models**
  - Analog computers lack instruction memory, time mux
  - Program is embodied in datapath setup [8]
  - Solution: compose subnets of neural nets: divide & conquer declarative programming [7]

- **Demonstrated Analog Workloads**
  - Function approximation for floating point programs [26]
  - Neural network pattern recognition for vision / ML [4], [8]
  - …and hardware may drive workloads evolution, too

Digital vs. Analog Summary

	Digital	Analog
Programming model	Imperative, temporal	Declarative, spatial
Network on chip	Packet switched	Circuit switched
Datapath	Synchronous discrete time	Asynchronous continuous time
Registers	Digital registers	Analog capacitors
Data representation	Digital integers, IEEE floating point	Analog voltage, current, spike density
Function units	ALUs	Neurons, etc.
Physical devices	Relies on ideal transistor	Many physical processes

"Analog computing" changes whole stack, each previously explored; which ones are source of benefit?

HCDC Demonstration

$$\Theta = -0.5 * \Theta - 5*\sin\Theta,\ \Theta(0) = 1$$

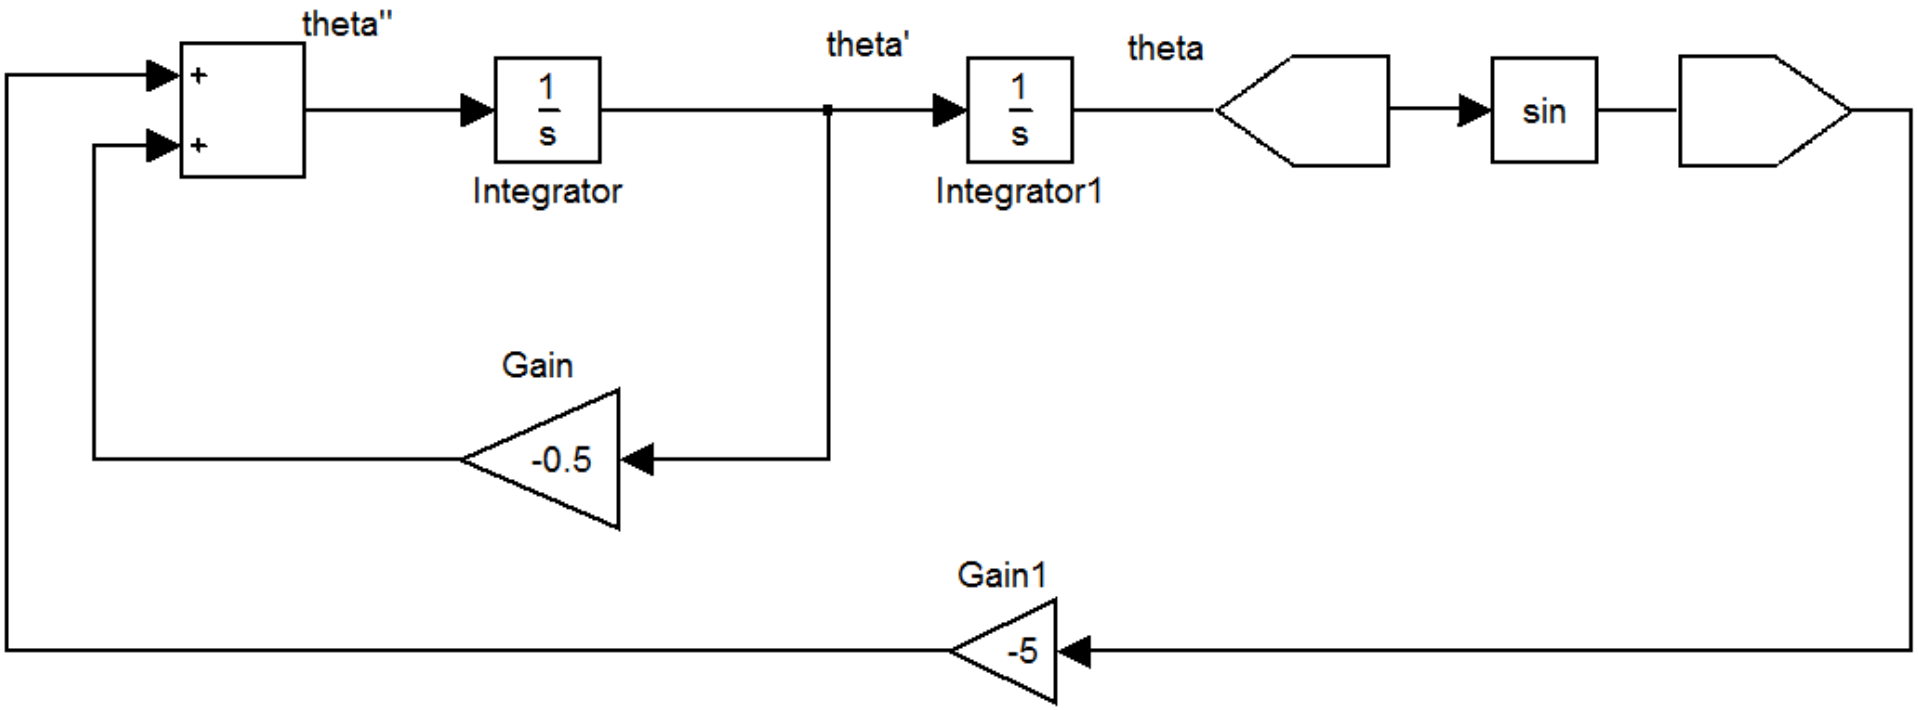

HCDC Demonstration

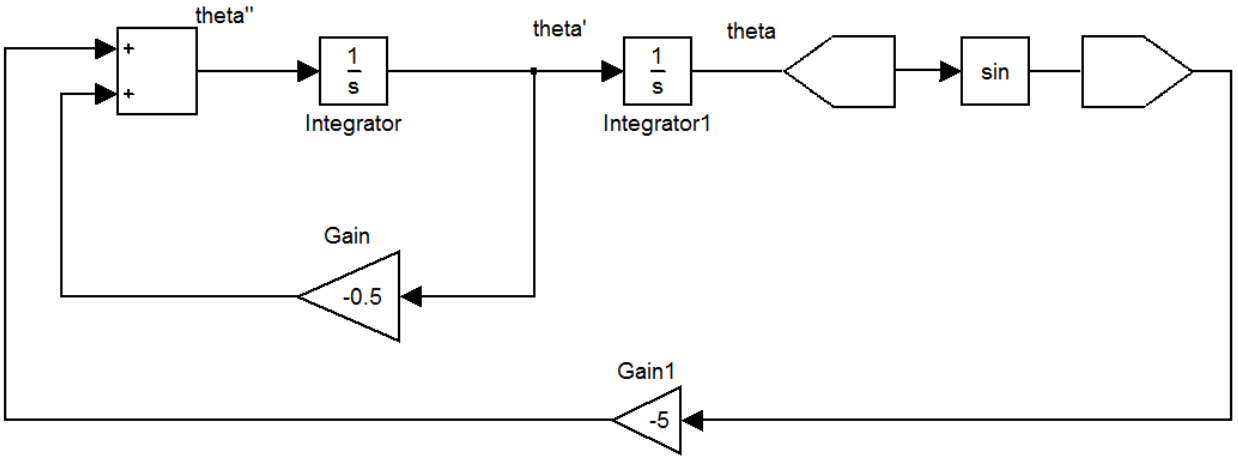

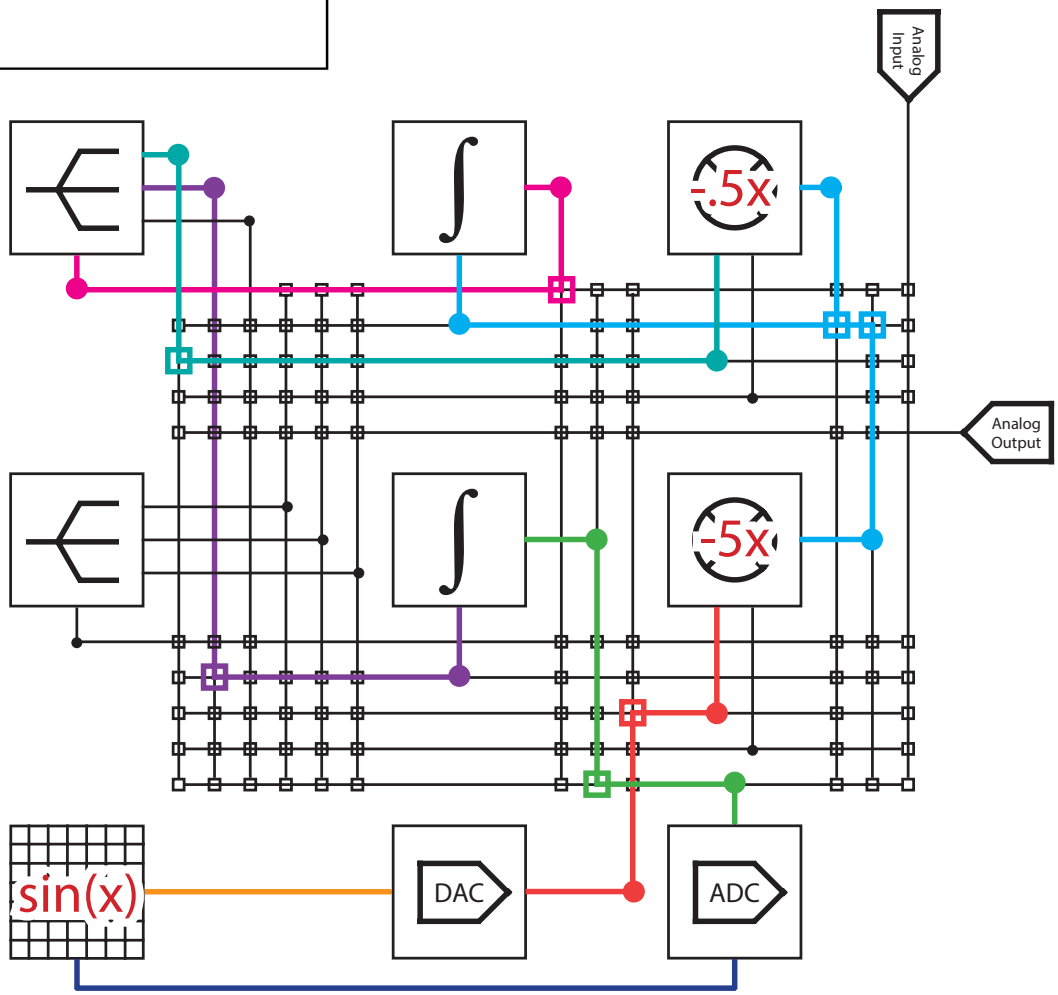

HCDC API Library Calls

```
34 hcdcInit();
35
36 // Set initial integrator values
37 float initial_y0 = 5.000000;
38
39 // Call HCDC wiring instructions
40 setSimpleConn ( {fans[0], out0, muls[0], in0} );
41 setSimpleConn ( {fans[0], out1, muls[0], in1} );
42 setSimpleConn ( {fans[0], out2, ints[0], in0} );
43
44 setSimpleConn ( {muls[0], out0, ints[0], in1} );
45
46 setIntInitial ( ints[0], initial_y0 );
47 setSimpleConn ( {ints[0], out0, fans[0], in0} );
```