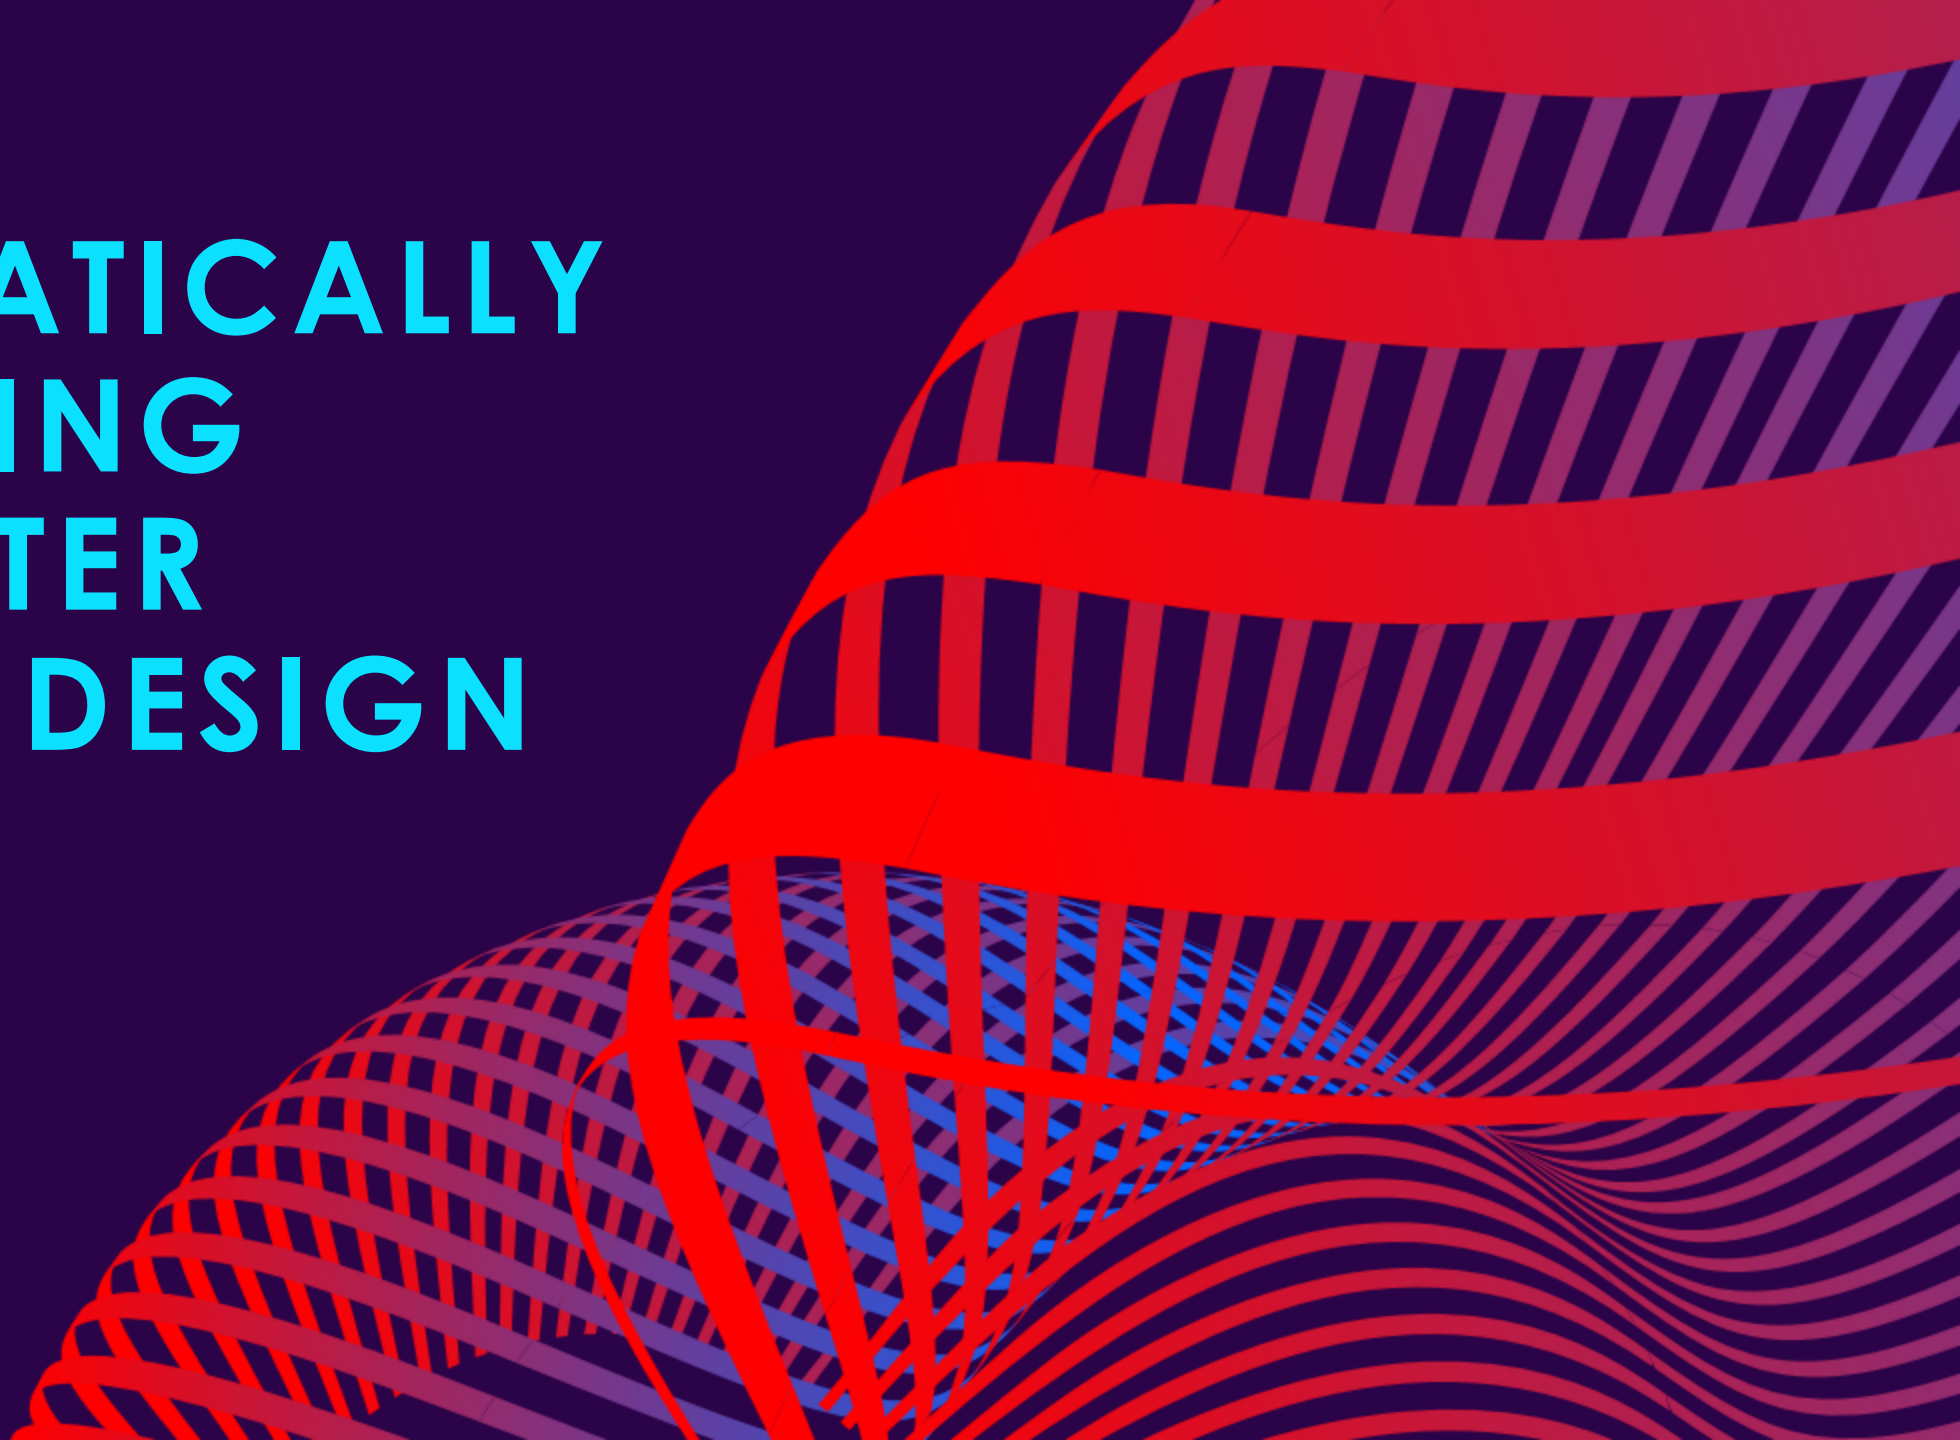
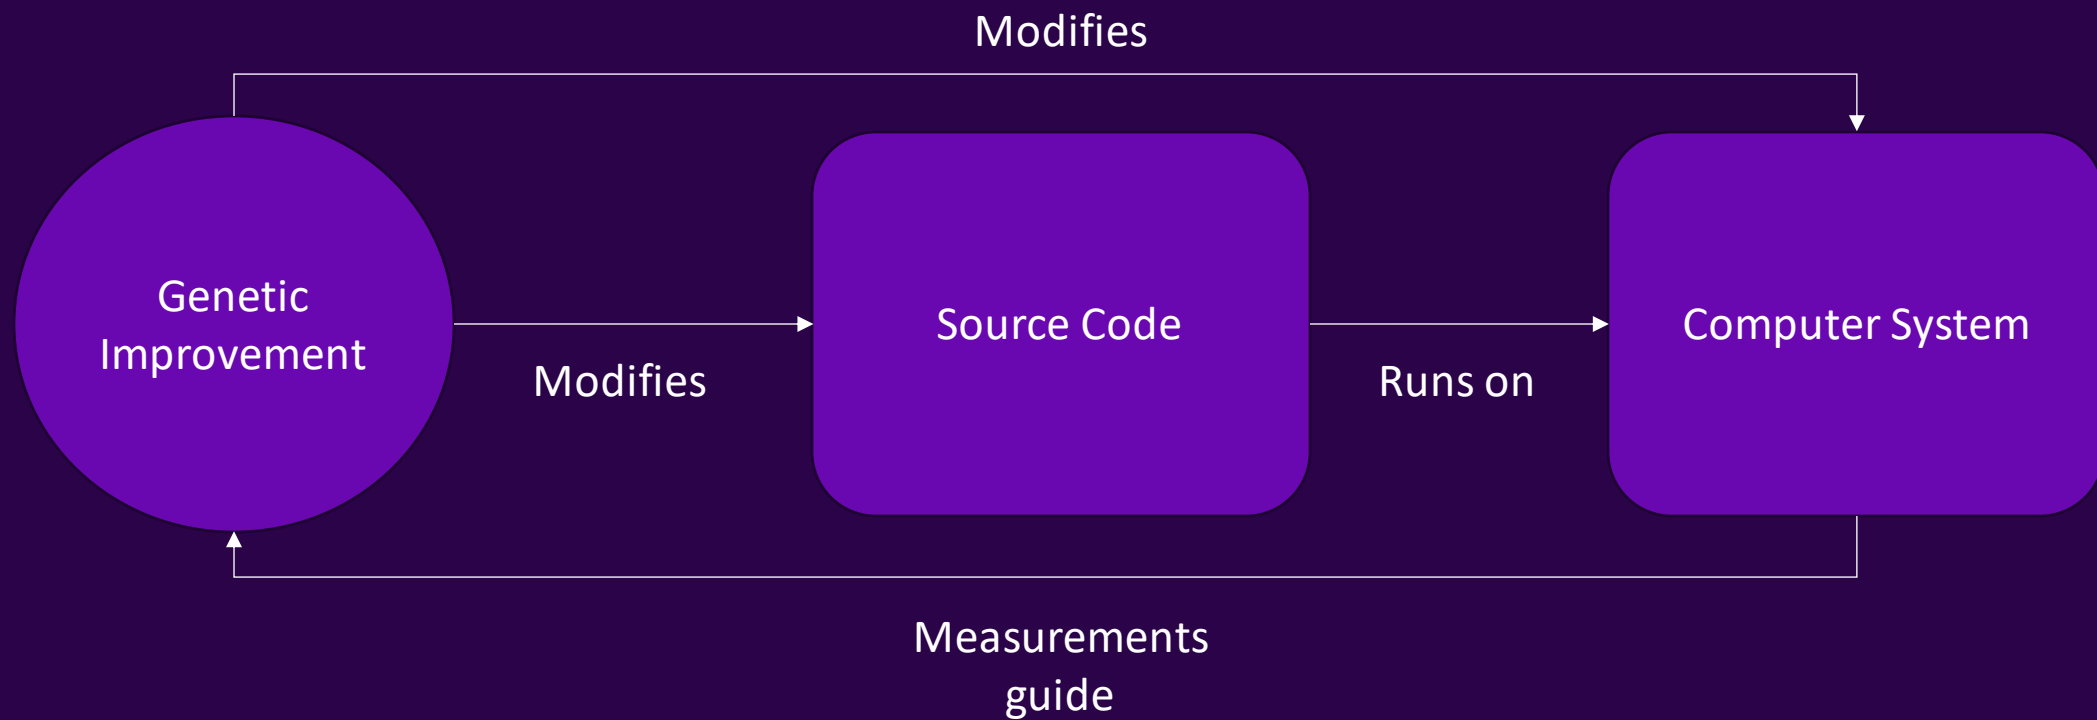


# AUTOMATICALLY EXPLORING COMPUTER SYSTEM DESIGN SPACES

Bobby R. Bruce



# THE GENERAL IDEA





# THE GENERAL IDEA




Genetic  
Improvement

Typically, in GI we optimize software  
with respect to a static target  
computer system.



# THE GENERAL IDEA

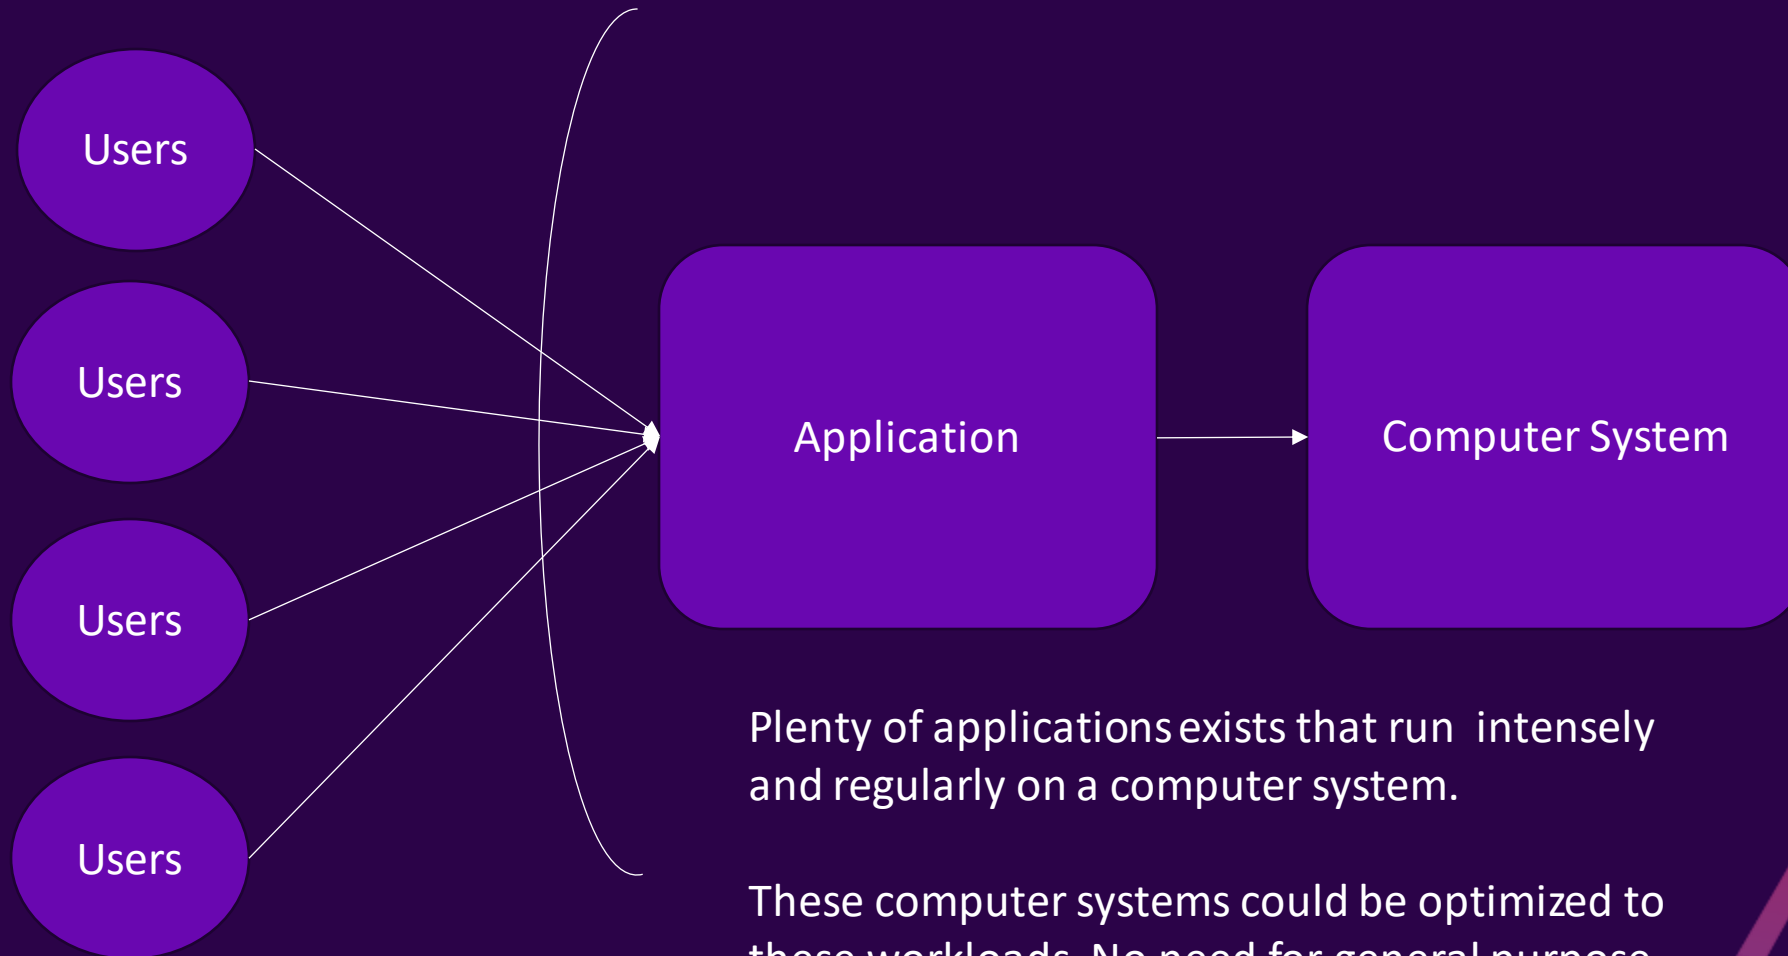


Genetic  
Improvement

We *could* also use GI to optimize a computer system with respect to a static software target

Isn't source-code the definition of computer system behavior?

# A SIMPLE MOTIVATING EXAMPLE



Plenty of applications exists that run intensely and regularly on a computer system.

These computer systems could be optimized to these workloads. No need for general purpose systems. They could be specialized.

# OTHER MOTIVATIONS

The screenshot displays the SiFive Core Design Tool interface. On the left, a sidebar lists configuration categories: Modes & ISA, On-Chip Memory, Ports, Security, Debug & Trace, Interrupts, Design For Test, Clocks and Reset, Branch Prediction, and RTL Options. The main area is divided into two panels. The left panel, titled 'ISA Extensions', shows the 'Multiply (M Extension)' selected with a performance of '1 Cycle'. Below this, 'Floating Point' options are shown with 'Double FP (F & D)' selected. Further down, 'Atomics (A Extension)' and 'Bit Manipulation (B Extension)' are also selected. The 'Extensions' section at the bottom has 'SiFive Custom Instruction Extension (SCIE)' unchecked. The right panel, titled 'Untitled S7 Core Core Complex', shows a configuration for 'S7 SERIES CORE 1 Core RV64IMAFDCB'. It lists features like 'Machine Mode + User Mode', 'Multiply (1 Cycle) + Atomics + FP (F & D)', 'No SCIE', 'Area Optimized Branch Prediction', 'Clock Gating', 'PMP 8 Regions', 'Instruc. Cache 32 KiB + 2-way', 'Data Cache 32 KiB + 4-way', 'Instruc. TIM 32 KiB', 'Data Loc. Store 32K', 'Adr. Remappers None', and 'No Raw Trace Port + 2 Perf Counters'. On the right side of this panel, four ports are listed: 'Front Port 64-bit AXI4', 'System Port 64-bit AXI4', 'Peripheral Port 64-bit AXI4', and 'Memory Port 64-bit AXI4'. Below these, the 'L2 Cache' is set to 'None'. At the bottom, three modules are shown: 'Debug Module' (JTAG + SBA, 4 HW Breakpoints, 0 Ext Triggers), 'PLIC' (7 Priority Levels, 127 Global Int.), and 'CLINT' (0 Local Interrupts). The base is noted as 'S76 Standard Core'.

**Modes & ISA**

- On-Chip Memory
- Ports
- Security
- Debug & Trace
- Interrupts
- Design For Test
- Clocks and Reset
- Branch Prediction
- RTL Options

**ISA Extensions**

☒ Multiply (M Extension) ?

Multiply Performance

1 Cycle

Floating Point ?

No FP Single FP (F) Double FP (F & D)

☐ Half-Precision Hardware Floating Point ?

☒ Atomics (A Extension) ?

☒ Bit Manipulation (B Extension) ?

**Extensions**

☐ SiFive Custom Instruction Extension (SCIE) ?

**Untitled S7 Core Core Complex**

**S7 SERIES CORE 1 Core RV64IMAFDCB**

Machine Mode + User Mode  
Multiply (1 Cycle) + Atomics + FP (F & D)  
No SCIE

Area Optimized Branch Prediction

Clock Gating PMP 8 Regions

Instruc. Cache 32 KiB + 2-way Data Cache 32 KiB + 4-way

Instruc. TIM 32 KiB Data Loc. Store 32K

Adr. Remappers None

No Raw Trace Port + 2 Perf Counters

Front Port 64-bit AXI4

System Port 64-bit AXI4

Peripheral Port 64-bit AXI4

Memory Port 64-bit AXI4

L2 Cache None

Debug Module

JTAG + SBA  
4 HW Breakpoints  
0 Ext Triggers

PLIC

7 Priority Levels  
127 Global Int.

CLINT

0 Local Interrupts

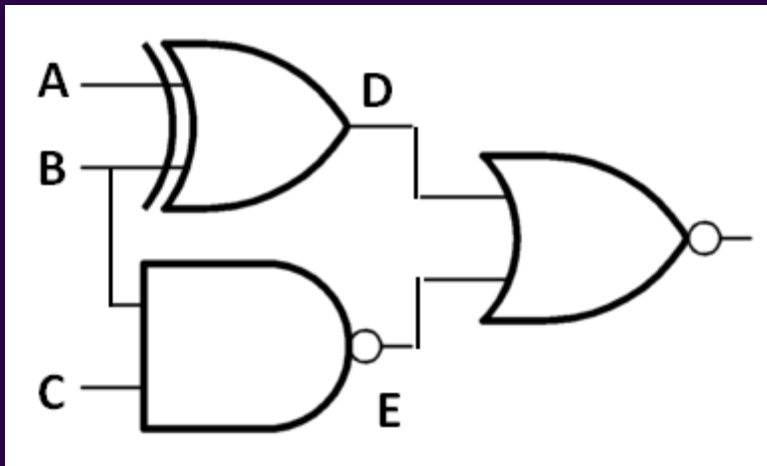
Base: [S76 Standard Core](#)

This is the SiFive Core Design Tool

Open Architecture projects such as RISC-V are making this easier.

Costs of silicon customization, “design to tape-out” are dropping dramatically

# AT WHAT LEVEL SHOULD WE WORK?



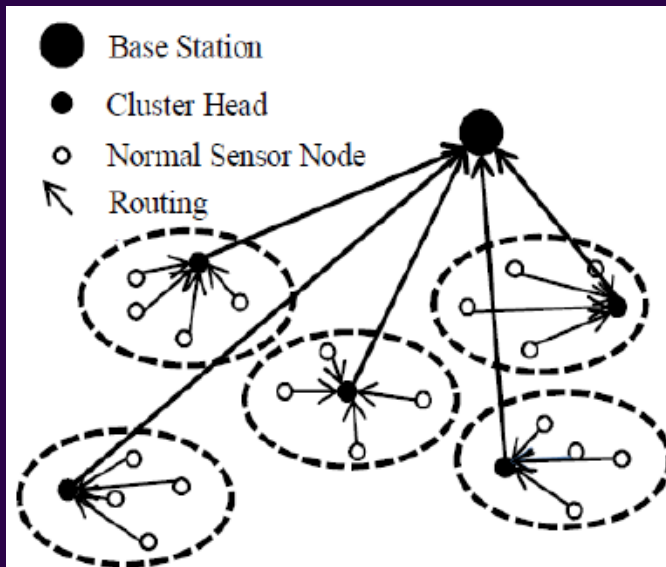
Logic gates, transistors, circuit-level?

A few cons:

- It's been done!
- Run into scalability problems
- Customizing at this level is very expensive

# AT WHAT LEVEL SHOULD WE WORK?

8

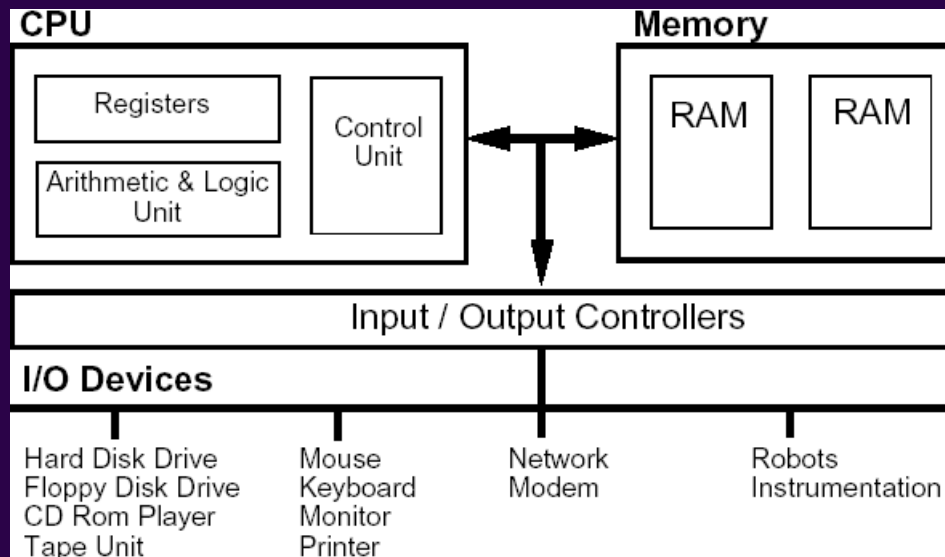


Clusters, computer networks, etc.?

A few cons:

- Difficult to simulate workloads sufficiently for optimization
- (I'm not really all that interested!)

# AT WHAT LEVEL SHOULD WE WORK?



## Computer architecture?

Some Pros:

- Common APIs and standards allow for interchanging components (think genes and alleles)
- We can utilize off the shelf components and designs.
- Plenty of research opportunity --- computer architecture designing and optimization is very manual.



# HOW DO WE EVALUATE ARCHITECTURE DESIGNS?

SIMULATION



# MODERN SIMULATORS

They provide a language for GI to modify



# MODERN SIMULATORS

That language follows a grammar like any other

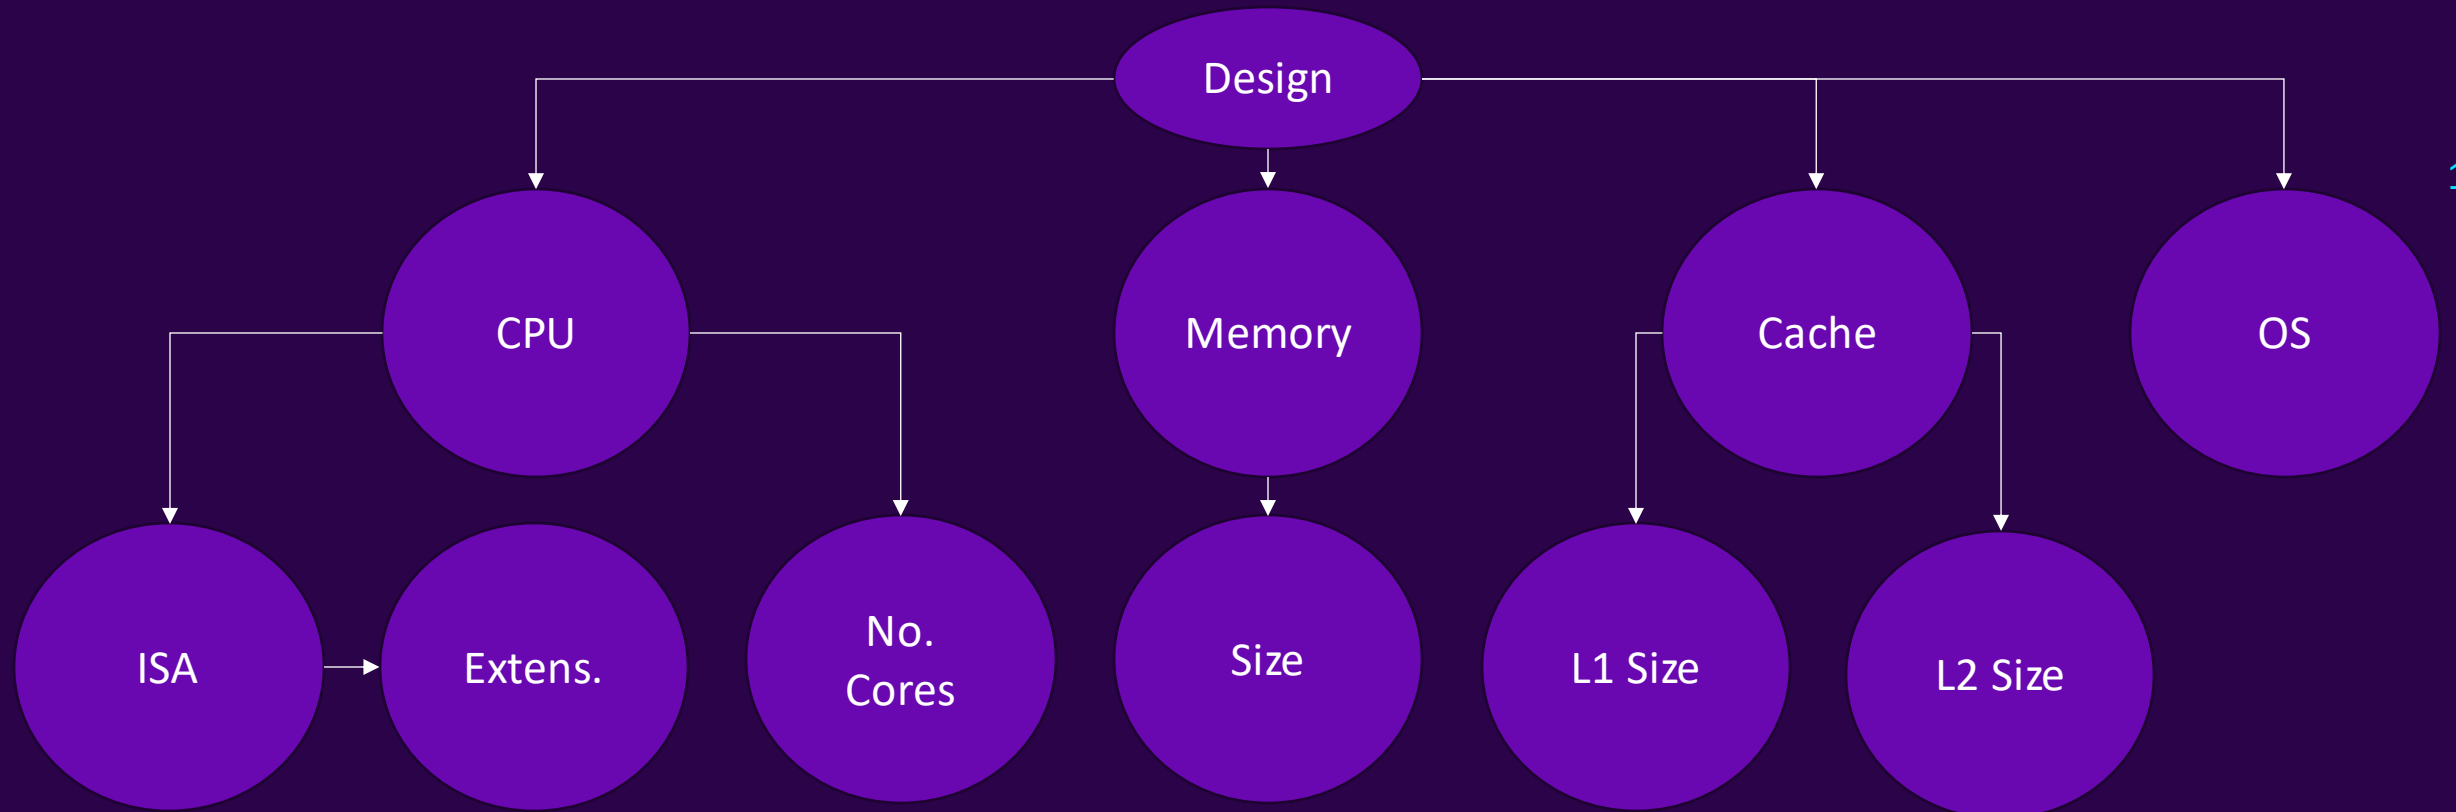
```
1 from gem5.components.boards.simple_board import SimpleBoard
2 from gem5.components.cachehierarchies.classic.no_cache import NoCache
3 from gem5.components.memory.single_channel import SingleChannelDDR3_1600
4 from gem5.components.processors.simple_processor import SimpleProcessor
5 from gem5.components.processors.cpu_types import CPUTypes
6 from gem5.resources.resource import Resource
7 from gem5.simulate.simulator import Simulator
8
9 # Obtain the components.
10 cache_hierarchy = NoCache()
11 memory = SingleChannelDDR3_1600("1GiB")
12 processor = SimpleProcessor(cpu_type=CPUTypes.ATOMIC, num_cores=1)
13
14 #Add them to the board.
15 board = SimpleBoard(
16     clk_freq="3GHz",
17     processor=processor,
18     memory=memory,
19     cache_hierarchy=cache_hierarchy,
20 )
21
22 # Set the workload.
23 binary = Resource("x86-hello64-static")
24 board.set_se_binary_workload(binary)
25
26 # Setup the Simulator and run the simulation.
27 simulator = Simulator(board=board)
28 simulator.run()
```

The “variables” in a design are the components and their properties

For example, we can swap out the memory system with a different type (single channel or dual) or configure its size

# MODERN SIMULATORS

A good target for Grammar-based GP



# A FRAMEWORK



## WHAT'S NEEDING DONE?

**STATS**

Can we create simulations with good enough fidelity?

**ACCURACY**

What stats do we need from our simulator, what are we optimizing for??

**COST MODEL**

How do we estimate the cost of a design so we can determine the trade off?

**SPEED**

How do we do 1000s of evaluations when 1 can take hours?

**BENCHMARKS**

What workloads should we optimize and are they meaningful?

# ANY (NICE) QUESTIONS?

Bobby R. Bruce

[bbruce@ucdavis.edu](mailto:bbruce@ucdavis.edu)

<https://www.bobbybruce.net>