**Bruce Jacob** 

#### **United States Naval Academy**

SLIDE 1



# **Ridiculously Simple** Computers (& simulation thereof) **Bruce Jacob Cyber Science Department United States Naval Academy**

Disclaimer: The views you are about to hear are the ravings of a madman and do not necessarily represent the views or opinions of the U.S. Naval Academy, Department of the Navy, or Department of Defense (DoD) or any of its components.



- Fundamental understanding
- Distillation of components to essence
- Everything is there to serve a purpose (education => explain that purpose)

**Observations:** 

**RIDICULOUSIY** SIMPLE **COMPUTERS** 

**Bruce Jacob** 

**United States Naval Academy** 

SLIDE 2

## **My Goal = Teaching Platform\*\***

Same-ish goals as HW/SW Co-Design \*\*Results echo yesterday's panel session



Bit:	15	14	13	12	11	10	9	8	7
ADD		0000			reg A				
SUB	0001								
AND	0010								
OR		00	11			reg	Α		
JALR		01	11			reg	Α		
Bit:	15	14	13	12	11	10	9	8	7
Bit: ADDI	15	14 10	13 00	12	11	10 reg	9 A	8	7
Bit: ADDI LW	15	14 10 10	13 00 01	12	11	10 reg	9 A A	8	7
Bit: ADDI LW SW		14 10 10	13 00 01 10	12	11	10 reg reg	9 A A	8	7
Bit: ADDI LW SW BEQ		14 10 10 10 10	13 00 01 10 11	12	11	10 reg reg	9 A A A	8	7
Bit: ADDI LW SW BEQ HALT		14 10 10 10 10 11	13 00 01 10 11 11	12		10 reg reg reg	9 A A A 10	8	7

**Bruce Jacob** 

**United States** Naval Academy

SLIDE 3

#### 9 Instructions + HALT

## **RiSC-16 — Student Model**



reg B	0000	imm16
reg B	0000	imm16
reg B	0000	imm16
reg B	0000	imm16
1111	1111	imm16
	· · ·	



## **RiSC-16 + (**

Bit:	15	14	13	12	11	10	9	8	7	6	5	4	3
ADD		00	00			reg	Α			reg	B		
SUB		00	01			reg	Α			reg	B		
AND		00	10			reg	Α			reg	B		
OR		00	11			reg	Α			reg	В		
-		01	00			reg	A			reg	В		
_		01	01			reg	A			reg	В		
_		01	10			reg	A			reg	В		
JALR		01	11			reg	Α			reg	B		
BOOL		01	11			reg	Α			reg	B		
INV		01	11			reg	Α			reg	B		
S.LT		01	11			reg	Α			reg	B		
S.LTE		01	11			reg	Α			reg	B		
S.EQ		01	11			reg	Α			reg	B		
TRAP		01	11			000	00			11	11		

RIDICULOUSLY SIMPLE COMPUTERS

**Bruce Jacob** 

United States Naval Academy

SLIDE 4

#### 20 Instructions + HALT

Co	mpi	ler	+ K	ern		
	Bit:	15 14 13 12	11 10 9 8	7 6 5 4	3 2 1 0	15 14 13
3 2 1 0	ADDI	1000	reg A	reg B	0000	
reg C	SHRI	1000	reg A	reg B	0001	
reg C			_			
reg C	LW	1001	reg A	reg B	0000	
reg C	LW.PA	1001	reg A	reg B	0001	
	- SW	1010	reg A	reg B	0000	
reg C	SW.PA	1010	reg A	reg B	0001	
reg C				1		
reg C	BEQ	1011	reg A	reg B	0000	
	– BNE	1011	reg A	reg B	0001	
0000	JALI	1011	reg A	0000	0010	
0001			-		-	
0010	-	1100	reg A	reg B	reg C	
1100	_	1101	reg A	reg B	reg C	
1101	_	1110	reg A	reg B	reg C	
1110	RFI	1111	0001	1111	1111	
1111	HALT	1111	1110	1111	1111	



## C Compilers for Dummies int is the only data type bahahahahaha ...

int ip = &data\_structure;
\*ip = value;

int variable; <- `variable' = its address int array[5]; <- `array' = its address int ip = &array; <- what does ip[3] mean?</pre> => ip[3] refers to (&ip)+3, not (\*ip)+3

int ip = &array; ip->[3] = value;

etc ...

**RIDICULOUSLY** SIMPLE **COMPUTERS** 

**Bruce Jacob** 

**United States** Naval Academy

SLIDE 5

## ridiculously simple but ~1000 lines of code



- Future of main memory = nonvolatile (price per bit + stay-alive power)
- ... it will therefore be **BIG** (TBs)
- ... it is also likely to be on-die (monolithic integration)
- ... with many, many access points and 100s of cores using them (multicore is here to stay)

**Bruce Jacob** 

**United States** Naval Academy

SLIDE 6

## Kernels ... w/ Sustainability Hook!!!







RIDICULOUSLY<br/>SIMPLE<br/>COMPUTERSImage: Computer of the second s

SLIDE 8

\*\* adjusted by
the memory
technology's
access latency





#### **DRAM Architecture**

## KNL vs. HBNV

	DRAM	ReRAN
# Mem Controllers	6	1000
L1 Cache size	32KB	32KB
L2 Cache size	1MB	1MB
Maii	n Memory Paramet	ters
Mem Latency	DRAMsim3	SST Mess
	simulated	<b>Read: 20</b>
	DDR4 & HBM	Write: 10
request_width	64 Bytes	8 Byte
(access	bus-width $= 8B$	bus-width
granularity)	burst-length = 8	burst-lengt
Topology	Mesh	Mesh







**Bruce Jacob** 

**United States Naval Academy** 

SLIDE 10

#### **STREAM: DRAM vs ReRAM comparison**

Number of Cores





**Bruce Jacob** 

**United States Naval Academy** 

SLIDE 10

#### STREAM: DRAM vs ReRAM comparison

Number of Cores

**Bruce Jacob** 

**United States** Naval Academy

SLIDE 11





**Bruce Jacob** 

**United States** Naval Academy

SLIDE 12



**Bruce Jacob** 

United States Naval Academy

SLIDE 12

// input: inode of the file to execute new process(inode) { choose ASID initialize a PCB (e.g., PC=0, SP=0xFFFF) put ASID onto runQ return from interrupt

## You jump directly into the raw executable

**RIDICULOUSLY** SIMPLE COMPUTERS

**Bruce Jacob** 

**United States** Naval Academy

SLIDE 13

## **Observations: Simplicity Pseudocode for process invocation:**



**Bruce Jacob** 

**United States Naval Academy** 

SLIDE 14

## **Observations / Conclusions** Verilator + Verilog CPU + kernel + shell is doable in real time (most of you guys already knew this)

#### [demo]



**Bruce Jacob** 

**United States** Naval Academy

SLIDE 15

Jumping into binaries is kinda cool :D

## **Observations / Conclusions**

## Simplicity — Kernel in ~1000 lines of C

- VM + preemptive multitasking + named files + int-driven I/O
- Merged file system + virtual memory system (way less code)

## —> New Way to Use Modeling/Simulation

- Running in first-person as opposed to batch-mode enables exploratory design, vs. design-space exploration
- Study ANY microarchitecture, in ANY amount of detail, running ANY operating system ... in real-time on your laptop



**Bruce Jacob** 

United States Naval Academy

SLIDE 16

## Shameless Plug

#### The Design and Security of Modern Computing Systems

How a Computer Is Built, From Ground Up How It Is Designed To Provide Security How its Defenses Can Be Overrun

Bruce **Jacob** William **Casey** Anthony **Melaragno** 

**United States Naval Academy Cyber Science Department** 



**Bruce Jacob** 

**United States** Naval Academy

SLIDE 17

### Another Shameless Plug www.memsys.io Washington DC Sep/Oct, 2024

# MEMSYS 2023 Attendees

## MEMSYS 2024

#### The International Symposium on Memory Systems **\*** Fall 2024, Washington DC

#### **Important Dates**

Abstract Deadline: 2 June, 2024 Submission: 16 June, 2024 Notification: 15 July, 2024 Camera-Ready: 1 September, 2024

ats gth of an /thing from 1– for example: t. ACM 'sigconf proceedings template, blind submission (no authors listed), up to 16 pages long

#### Organizers

Bruce Jacob (Naval Academy) Kenneth Wright (AMD) Chen Ding (U. Rochester) Hameed Badawy (NMSU) uce Childers (U. of Pittsburgh)



etmar Fey (U. Erlangen) chael Jantz (University of Tennessee) on Dreslinski (University of Michigan) arc Reichenbach (U. Rostock) Zhang (ICT) nathan Beard (Google) shan Chishti (Intel) uce Christenson (Intel) wid Donofrio (TTCL) omas Vogelsang (Rambus) hn Leidel (Tactical Computing Labs) In F Rodrigues (Samsung) bert Trout (Sadram) lliam Wang (ARM) tar Radojkovic (BSC) rbert Wehn (RPTU) atthias Jung (U. Würzburg) endy Elsasser (Rambus) aya Gokhale (LLNL) Hammond (NNSA) zhak Birk (Technion) ke Ignatowski (AMD)

Memory-device manufacturing, memory-architecture design, and the use of memory technologies by application software all profoundly impact today's and tomorrow's computing systems, in terms of their performance, function, reliability, predictability, power dissipation, and cost. Existing memory technologies are seen as limiting in terms of power, capacity, and bandwidth. Emerging memory technologies offer the potential to overcome both technologyand design-related limitations to answer the requirements of many different applications. Our goal is to bring together researchers, practitioners, and others interested in this exciting and rapidly evolving field, to update each other on the latest state of the art, to exchange ideas, and to discuss future challenges. Please visit memsys.io for more information.

#### Areas of Interest

Previously unpublished papers containing significant novel ideas and technical results are solicited. Papers that focus on system, software, and architecture level concepts specifically memory-related, i.e. topics outside of traditional conference scopes, will be preferred over others (e.g., the desired focus is away from pipeline design, processor cache design, prefetching, data prediction, etc.). Symposium topics include, but are not limited to, the following:

- Memory-system design from both hardware and software perspectives
- Memory failure modes and mitigation strategies
- Memory and system safety and security issues
- Disaggregated Memory (e.g. CXL)
- Memory for embedded and autonomous systems (e.g., automotive)
- Operating system design for hybrid/nonvolatile memories
- Technologies including, DRAM, FLASH, NVM etc.
- Memory-centric programming models, languages, optimization
- Compute-in-memory and compute-near-memory technologies
- Data-movement issues and mitigation techniques
- Algorithmic & software memory-management techniques
- Emerging memory technologies, their controllers, and novel uses
- Interference at the memory level across datacenter applications
- In-memory databases and NoSQL stores •
- Post-CMOS scaling efforts and memory technologies to support them, including cryogenic, neural, and heterogeneous memories
- Negative results, validation of results, invalidation of results

To reiterate, papers that focus on topics *outside* the scope of traditional architecture conferences will be preferred over others.

#### Submissions and Presentations

Our primary goal is to showcase interesting ideas that will spark conversation between disparate groups-to get applications people, operating systems people, system architecture people, interconnect people and circuits people to talk to each other. We accept extended abstracts, position papers, and/or full research papers, and each accepted submission is given a 20-minute presentation time slot.

#### We intend to publish all papers in the ACM Digital Library.

#### Venue

The conference will be in Washington DC area.



**Bruce Jacob** 

United States Naval Academy

SLIDE 18

# Thank You! Questions?

## Bruce Jacob bjacob@usna.edu





choose ASID put ASID onto runQ return from interrupt

**RIDICULOUSLY** SIMPLE COMPUTERS

**Bruce Jacob** 

**United States** Naval Academy

SLIDE 19

## But Wait, How Does argv Work? **Pseudocode for process invocation:**

- // input: inode of the file to execute
- new process(inode, phys arg page) {
  - initialize a PCB (e.g., PC=0, SP=0xFFFF)
  - map( VIRTUAL ARG PAGE, phys arg page );

#### The shell reads input line into ARG\_PAGE, calls detach(ARG\_PAGE) & then new\_proc()

