



Custom Sensor-Based Embedded Computing Systems

Frank Vahid

Professor

Dept. of Computer Science and Engineering
University of California, Riverside

Assoc. Director, Center for Embedded Computer Systems, UC Irvine

eBlocks project 2002-present, support provided by the National
Science Foundation and Intel

Ph.D. student: Susan Lysecky (2006, now Asst. Prof. at U. Arizona); several
MS and undergrad students also

The Problem


- What do these problems all have in common?



A working adult with an ageing parent at home – did she get out of bed today, is she moving around?



A small store owner with many employees – are they in the storeroom, breakroom, or out back?



A homeowner who sometimes forgets to close the garage at night



Marines wishing to outfit a building to detect whether someone is inside or when someone was inside

The Problem


- What do these problems all have in common?



Put motion sensors around the house, monitor from the web or cell phone – or even be notified if no motion by certain time in the morning



Put motion and sound sensors throughout, small LEDs (lights) near cash register



Install contact sensor and light sensor, and indicator next to the bed

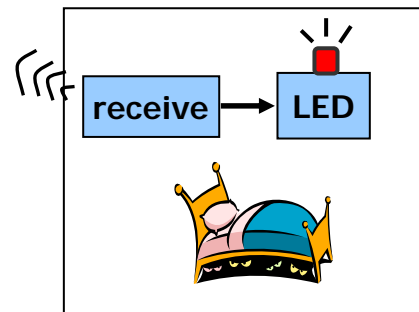



Place motion, heat, and sound sensors in rooms, halls, doorways




Why Can't We Just Do This?



- Widely usable “Lego”-like sensors don't exist today
 - Costly, hard to use, plugged into wall...
- But new technology makes Lego-like sensor blocks possible...



Shrinking Processor Size/Cost Enables New Solution






Courtesy of Joe Kahn



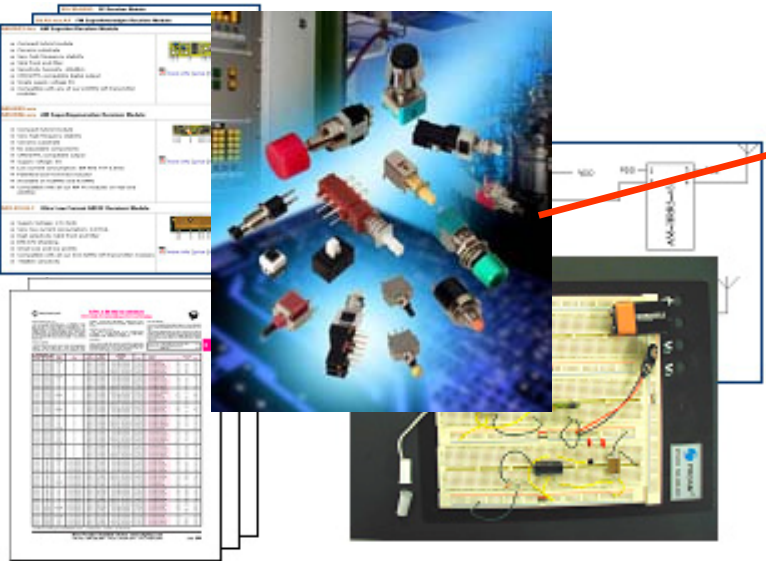
<http://www.templehealth.org>

- Make sensors smarter
 - By adding processor+battery
 - Today, tiny and cheap
 - Becomes a "block" easily connected to other blocks


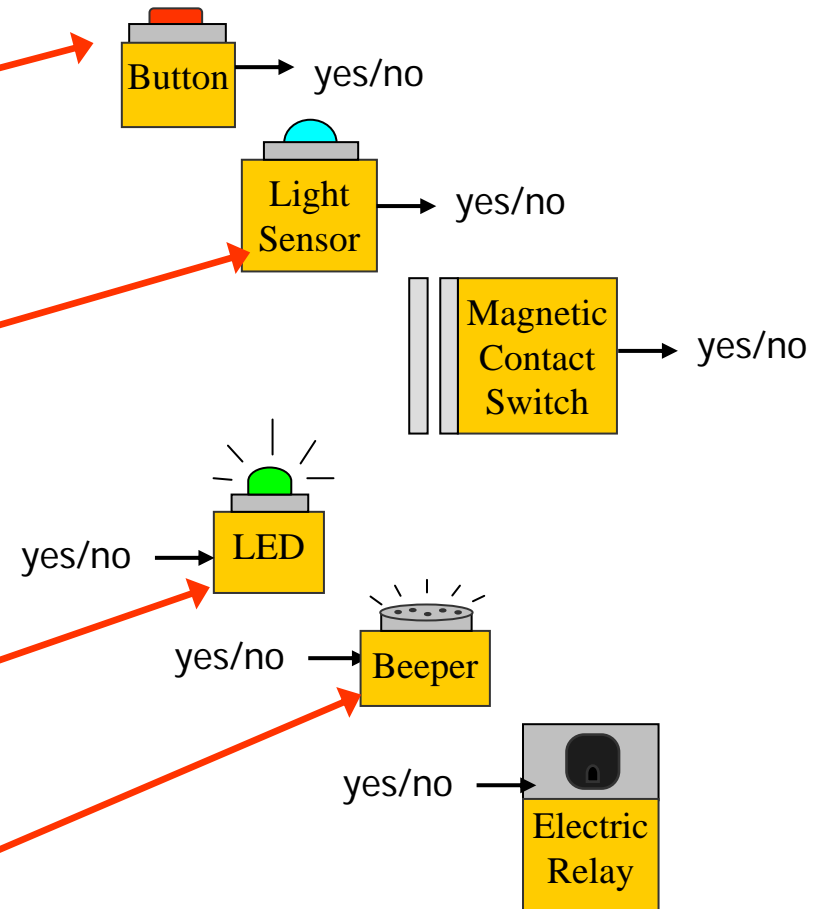


Shrinking Processor Size/Cost Enables New Solution – eBlocks

Existing component view



New "eBlock" view


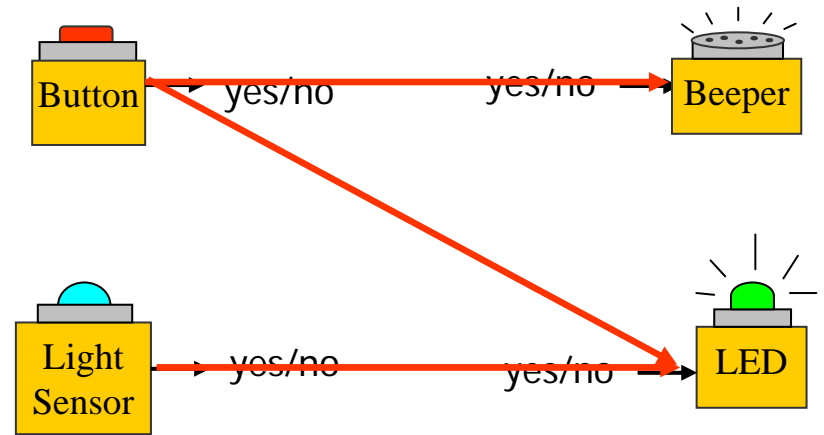


Each LED receives its necessary voltage and the circuit treats each LED equally.



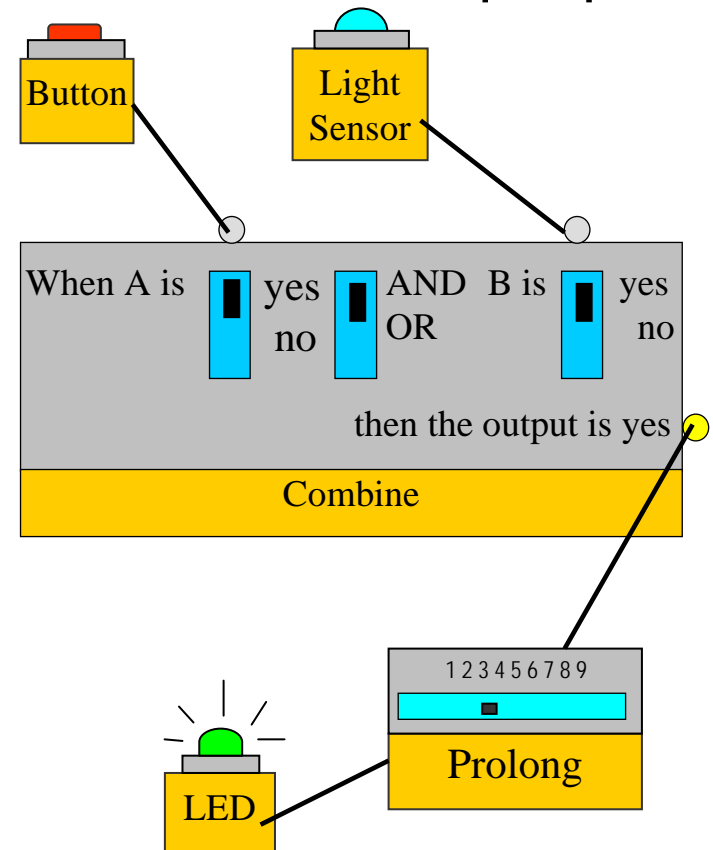
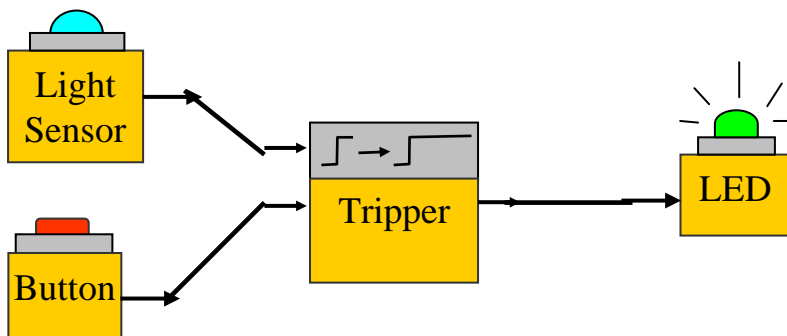

eBlocks

- Just connect blocks, and they work
 - No programming knowledge, no electronics knowledge



eBlocks

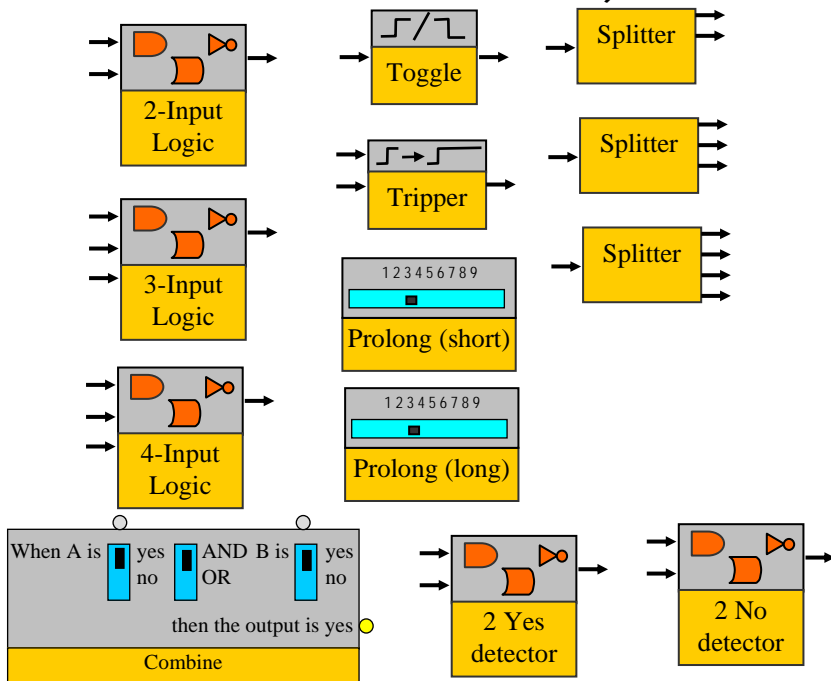
- Add intermediate blocks that compute and maintain state
 - Spatial programming – more intuitive to non-CS people than temporal programming



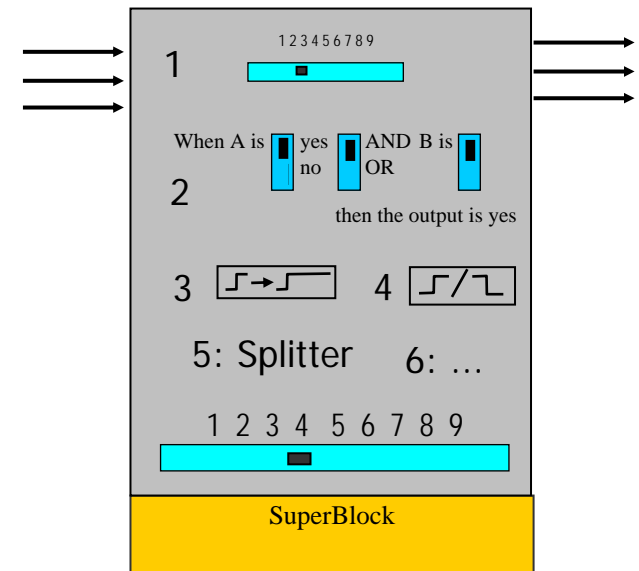
What's Hard

- (1) Finding right set of building blocks

Too many – Overwhelming (too much choice)



Too few – Overwhelming (too much configuration)



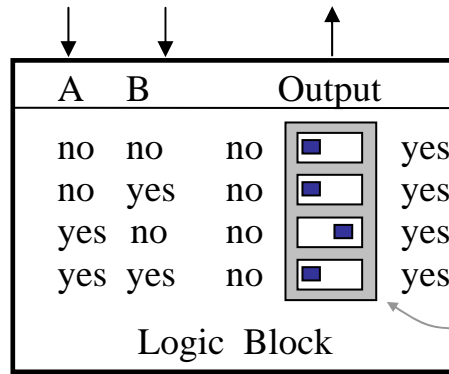
What's Hard

- (2) Making the blocks understandable



- People NOT likely to read directions
 - Those that do are unlikely to understand

Performed extensive user testing (over 500 students, kids, and adults) over two years


Example: *Combine* block




configurable DIP switch



Phrased truth table




Phrased truth table embedded in sentence



Colored truth table embedded in sentence

Most success



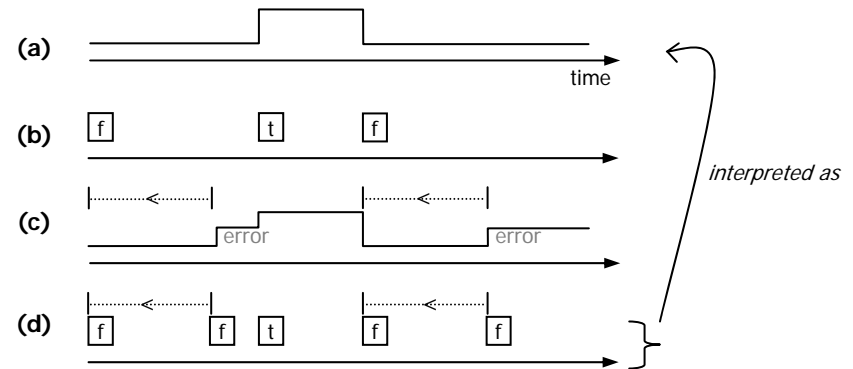
Logic Sentence

What's Hard

- (3) Batteries must last years, yet performance should appear continuous
 - Blocks are off 99.9% of the time

Developed theory to map eBlock events to continuous time

Developed custom CAD tool to automatically find the best block parameter settings out of the billions of possibilities




Reliability

Latency


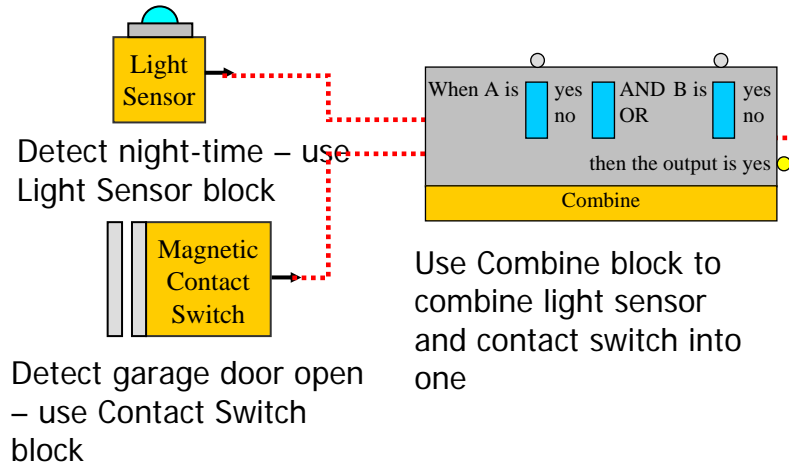
Connect Responsiveness

Disconnect Responsiveness



eBlocks Example

- "Garage Open at Night" detector
- <10 minutes to build



Need to indicate garage open at night – use LED block

Plug pieces together and the system is done!



Graphical Simulator

- ◆ User specifies and tests block design
- ◆ Java-based simulator
 - User chooses between pallets
 - Blocks added by dragging
 - User is able to configure various blocks by clicking on switches
 - Connections created by drawing lines between blocks

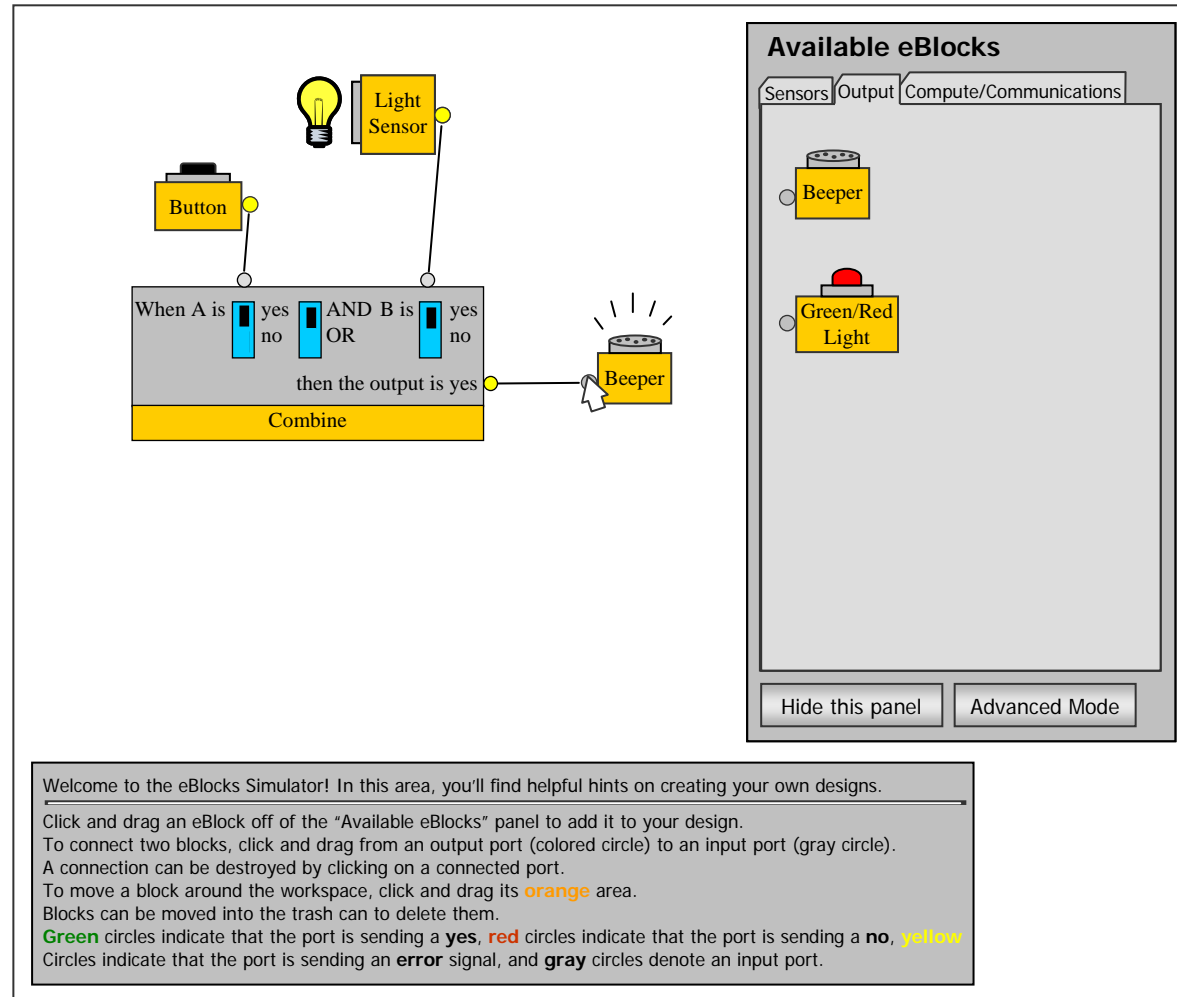
The screenshot displays the eBlocks Simulator interface. The workspace contains several blocks: a yellow 'Button' block, a 'Light Sensor' block (represented by a lightbulb icon), and a 'Combine' block. The 'Combine' block has two input ports, each with a 'yes' and 'no' label, and an output port labeled 'then the output is yes'. A mouse cursor is hovering over the 'no' label of the second input port. To the right is the 'Available eBlocks' panel, which has tabs for 'Sensors', 'Output', and 'Compute/Communications'. Under the 'Output' tab, there are three blocks: a 'Beeper' block, a 'Green/Red Light' block, and another 'Beeper' block. At the bottom of the panel are 'Hide this panel' and 'Advanced Mode' buttons. Below the workspace is a text box with the following text:

Welcome to the eBlocks Simulator! In this area, you'll find helpful hints on creating your own designs.

Click and drag an eBlock off of the "Available eBlocks" panel to add it to your design.
To connect two blocks, click and drag from an output port (colored circle) to an input port (gray circle).
A connection can be destroyed by clicking on a connected port.
To move a block around the workspace, click and drag its **orange** area.
Blocks can be moved into the trash can to delete them.
Green circles indicate that the port is sending a **yes**, **red** circles indicate that the port is sending a **no**, **yellow** circles indicate that the port is sending an **error** signal, and **gray** circles denote an input port.

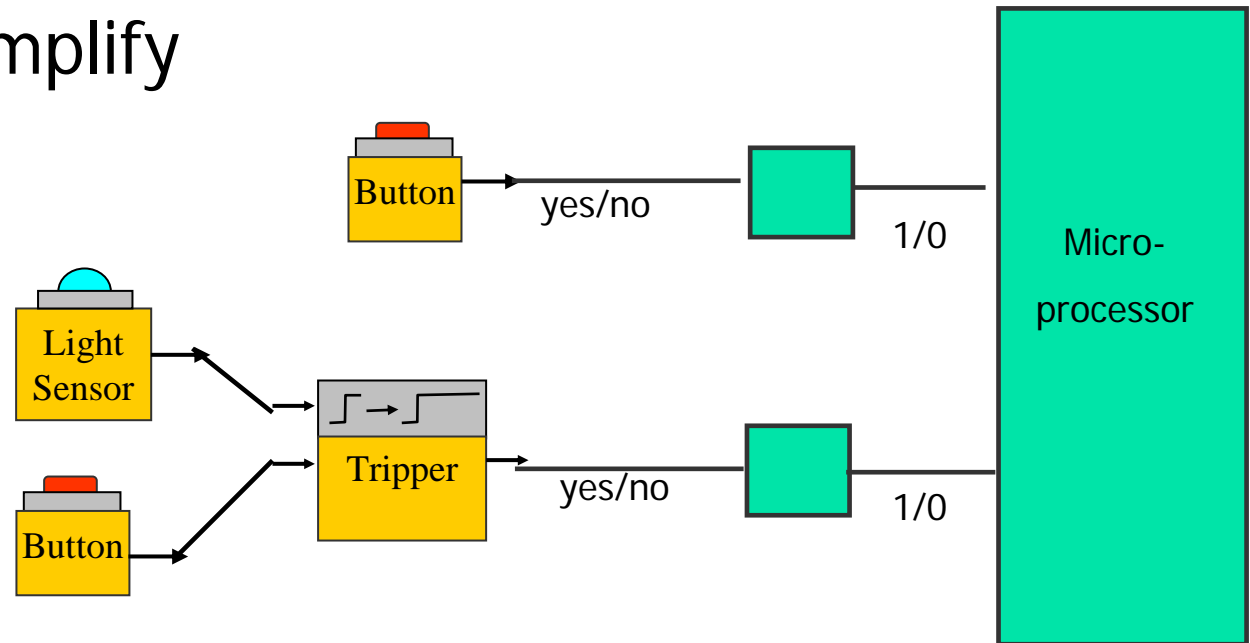
Graphical Simulator

- ◆ User specifies and tests block design
- ◆ Java-based simulator
 - User chooses between pallets
 - Blocks added by dragging
 - User is able to configure various blocks by clicking on switches
 - Connections created by drawing lines between blocks
 - User can create, experiment, test and configure design



eBlocks and Embedded Microprocessors

- Can greatly simplify coding




And now for something completely different...

- Warp processing
 - 2001-present, supported by SRC, Intel, IBM, Freescale, Xilinx
 - Ph.D. students
 - Roman Lysecky – Ph.D., June 2005, now Asst. Prof. at Univ. of Arizona
 - Greg Stitt – Ph.D. June 2007, now Asst. Prof. at Univ. of Florida, Gainseville
 - Ann Gordon-Ross – Ph.D. June 2007, now Asst. Prof. at Univ. of Florida, Gainseville

Circuits on FPGAs Can Sometimes Give Big Speedups



C Code for FIR Filter

```
for (i=0; i < 128; i++)  
  y += c[i] * x[i]  
..  
..  
..
```



- 1000's of instructions
 - Several thousand cycles

Circuit for FIR Filter




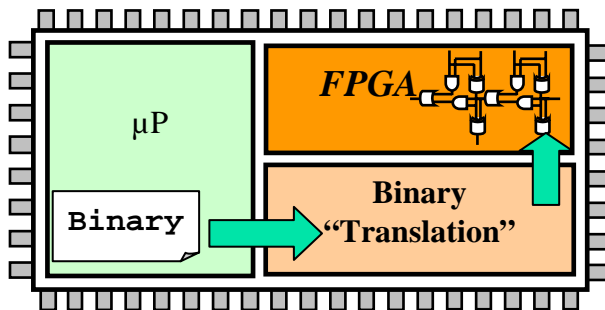


- ~ 7 cycles
 - Speedup > 100x
 - Pipelined -- >500x

Circuit parallelism/pipelining can yield big speedups

Dynamic Translation

- Motivated by commercial dynamic binary translation of early 2000s


e.g.,
Transmeta
Crusoe
“code
morphing”



- **Warp processing** (Lysecky/Stitt/Vahid 2003-2007): dynamically translate binary to circuits on FPGAs

Warp Processing Background

- 1 *Initially, software binary loaded into instruction memory*

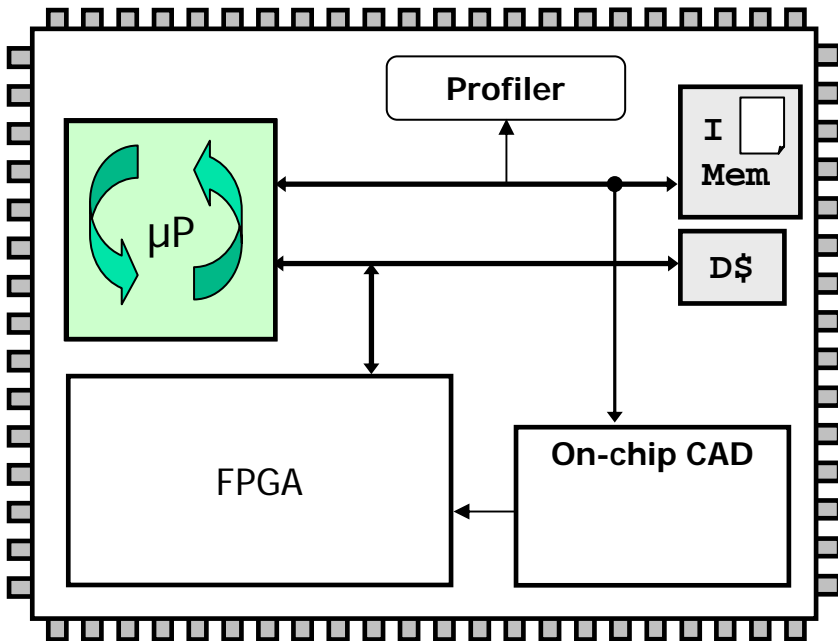


Software Binary

```
Mov reg3, 0
Mov reg4, 0
loop:
Shl reg1, reg3, 1
Add reg5, reg2, reg1
Ld reg6, 0(reg5)
Add reg4, reg4, reg6
Add reg3, reg3, 1
Beq reg3, 10, -5
Ret reg4
```


Warp Processing Background

2 *Microprocessor executes instructions in software binary*



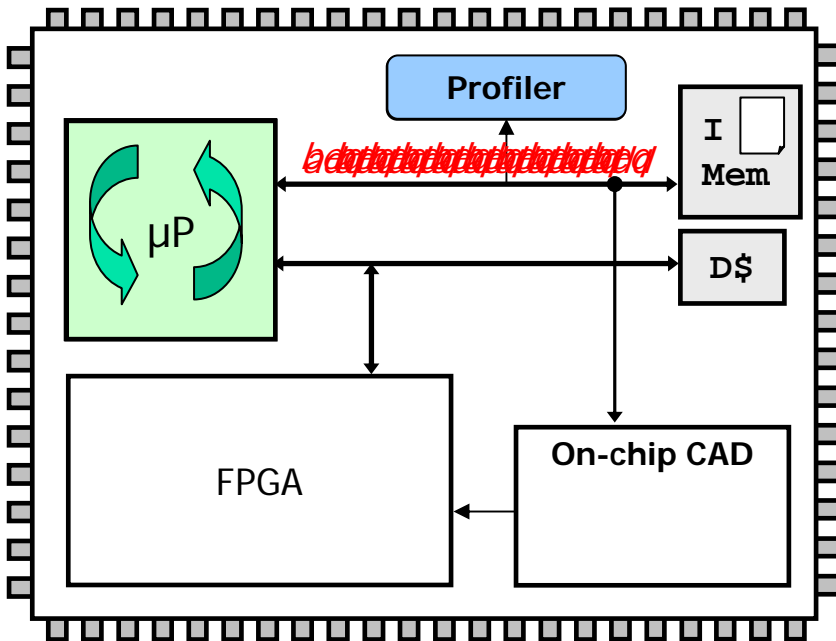
Software Binary

```
Mov reg3, 0
Mov reg4, 0
loop:
Shl reg1, reg3, 1
Add reg5, reg2, reg1
Ld reg6, 0(reg5)
Add reg4, reg4, reg6
Add reg3, reg3, 1
Beq reg3, 10, -5
Ret reg4
```



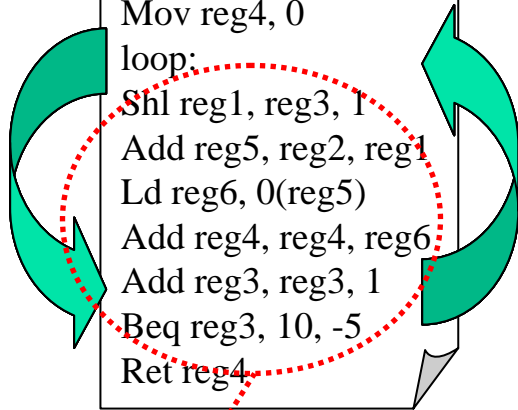
Warp Processing Background

3 Profiler monitors instructions and detects critical regions in binary




Software Binary

```
Mov reg3, 0
Mov reg4, 0
loop:
Shl reg1, reg3, 1
Add reg5, reg2, reg1
Ld reg6, 0(reg5)
Add reg4, reg4, reg6
Add reg3, reg3, 1
Beq reg3, 10, -5
Ret reg4
```




Critical Loop Detected




Warp Processing Background

4 *On-chip CAD reads in critical region*




Software Binary

```
Mov reg3, 0
Mov reg4, 0
loop:
  Shl reg1, reg3, 1
  Add reg5, reg2, reg1
  Ld reg6, 0(reg5)
  Add reg4, reg4, reg6
  Add reg3, reg3, 1
  Beq reg3, 10, -5
Ret reg4
```



Warp Processing Background


5 *On-chip CAD decompiles critical region into control data flow graph (CDFG)*



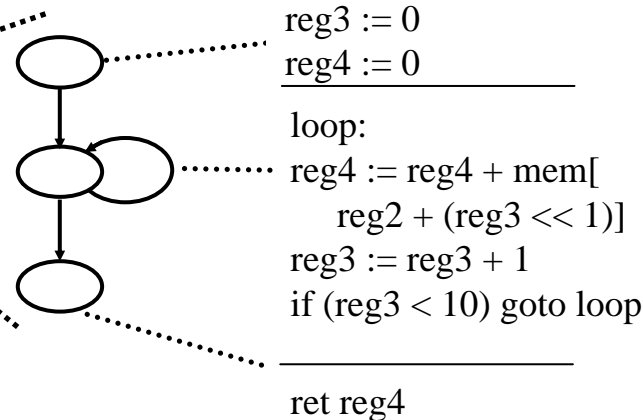
Software Binary

```

Mov reg3, 0
Mov reg4, 0
loop:
Shl reg1, reg3, 1
Add reg5, reg2, reg1
Ld reg6, 0(reg5)
Add reg4, reg4, reg6
Add reg3, reg3, 1
Beq reg3, 10, -5
Ret reg4
    
```




Recover loops, arrays, subroutines, etc. – needed to synthesize good circuits



Decompilation surprisingly effective at recovering high-level program structures
 Stitt et al ICCAD'02, DAC'03, CODES/ISSS'05, ICCAD'05, FPGA'05, TODAES'06, TODAES'07

Warp Processing Background



6 *On-chip CAD synthesizes decompiled CDFG to a custom (parallel) circuit*



Software Binary


```

Mov reg3, 0
Mov reg4, 0
loop:
Shl reg1, reg3, 1
Add reg5, reg2, reg1
Ld reg6, 0(reg5)
Add reg4, reg4, reg6
Add reg3, reg3, 1
Beq reg3, 10, -5
Ret reg4
    
```



Warp Processing Background

7 On-chip CAD maps circuit onto FPGA



Lean place&route/FPGA → 10x faster CAD

(Lysecky et al DAC'03, ISSS/CODES'03, DATE'04, DAC'04, DATE'05, FCCM'05, TODAES'06)


Multi-core chips – use 1 powerful core for CAD

Frank Vahid, UC Riverside


Software Binary

```

Mov reg3, 0
Mov reg4, 0
loop:
Shl reg1, reg3, 1
Add reg5, reg2, reg1
Ld reg6, 0(reg5)
Add reg4, reg4, reg6
Add reg3, reg3, 1
Beq reg3, 10, -5
Ret reg4
    
```



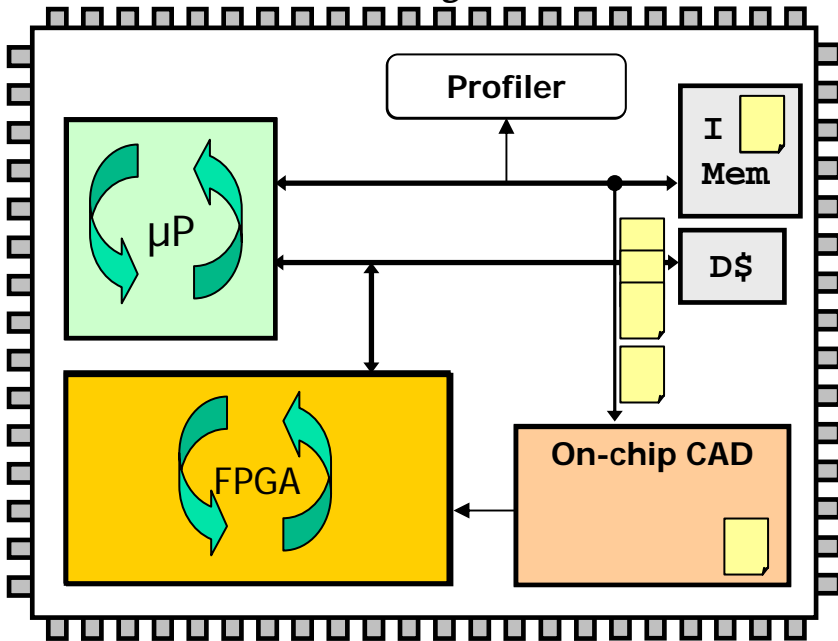
Time Energy



Warp Processing Background

8


On-chip CAD replaces instructions in binary to use hardware, causing performance and energy to “warp” by an order of magnitude or more



Software Binary


```
Mov reg3, 0
Mov reg4, 0
loop:
// instructions that
// interact with FPGA
Ret reg4
```

>10x speedups
for some apps



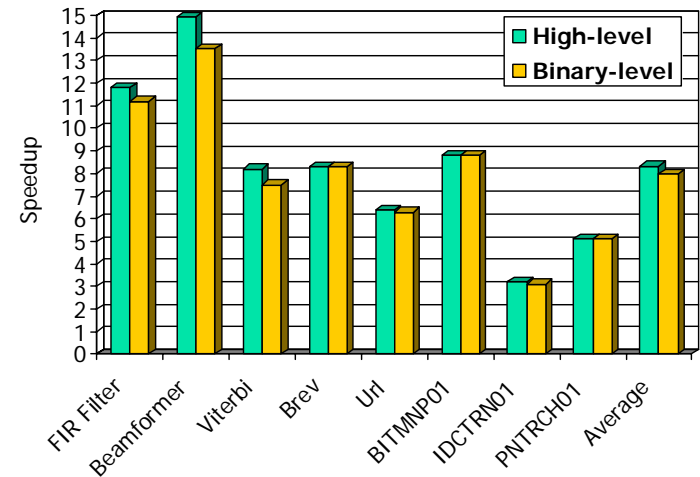
Time Energy

Software-only
“Warped”



Challenge: Decompilation

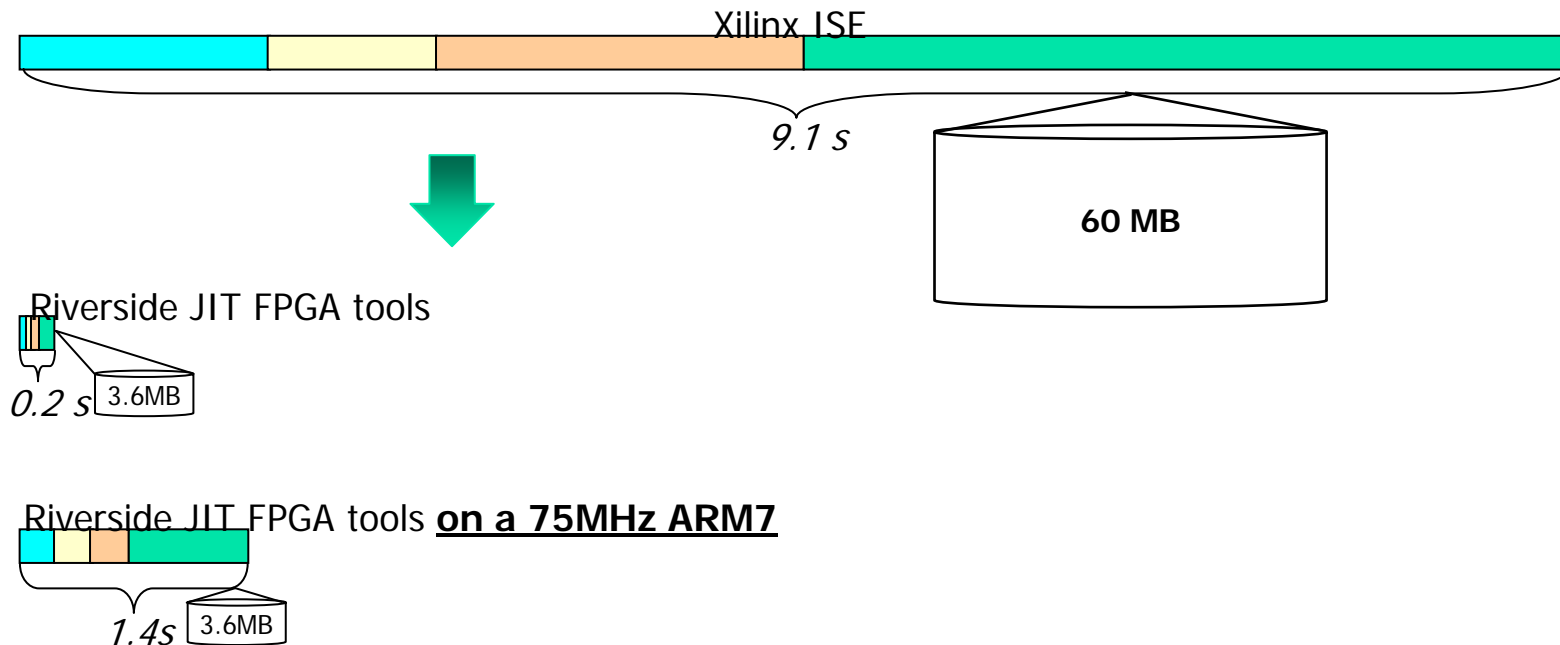
- Requires aggressive decompilation to recover loops, arrays, ..., from binaries
- Results: Competitive with synthesis from C



Example	Synthesis from C Code				Synthesis after Decompiling Binary				%Time _{Overhead}	%Area _{Overhead}
	Cycles	ClkFrq	Time	Area	Cycles	ClkFrq	Time	Area		
bit_correlator	258	118	2.19	15	258	118	2.186	15	0%	0%
fir	129	125	1.03	359	129	125	1.032	371	0%	3%
udiv8	281	190	1.48	398	281	190	1.479	398	0%	0%
prewitt	64516	123	525	2690	64516	123	524.5	4250	0%	58%
mf9	258	57	4.5	1048	258	57	4.503	1048	0%	0%
moravec	195072	66	2951	680	195072	70	2791	676	-6%	-1%
								Avg:	-1%	10%

Challenge: JIT Compile to FPGA

- Developed ultra-lean CAD heuristics for synthesis, placement, routing, and technology mapping; simultaneously developed CAD-oriented FPGA
 - e.g., Our router (ROCR) 10x faster and 20x less memory, at cost of 30% longer critical path. Similar results for synth & placement
 - (EDAA Outstanding Dissertation Award 2006)




Experiments

- Benchmarks: Image processing, DSP, scientific computing
 - Highly parallel examples to illustrate thread warping potential
 - We created multithreaded versions
- Base architecture – 4 ARM cores
 - Focus on recurring applications (embedded)
- TW: FPGA running at whatever frequency determined by synthesis

Multi-core

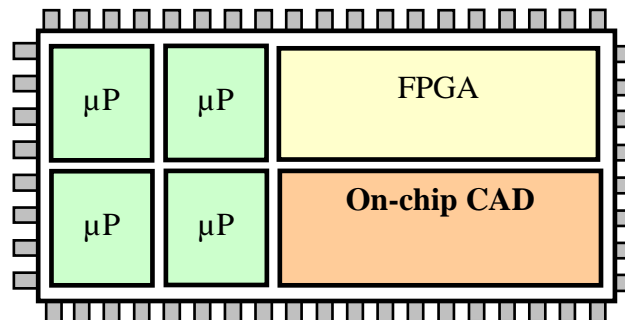
4 ARM11 400 MHz




Compared to

Thread Warping

4 ARM11 400 MHz + FPGA (synth freq)



Speedup from Thread Warping




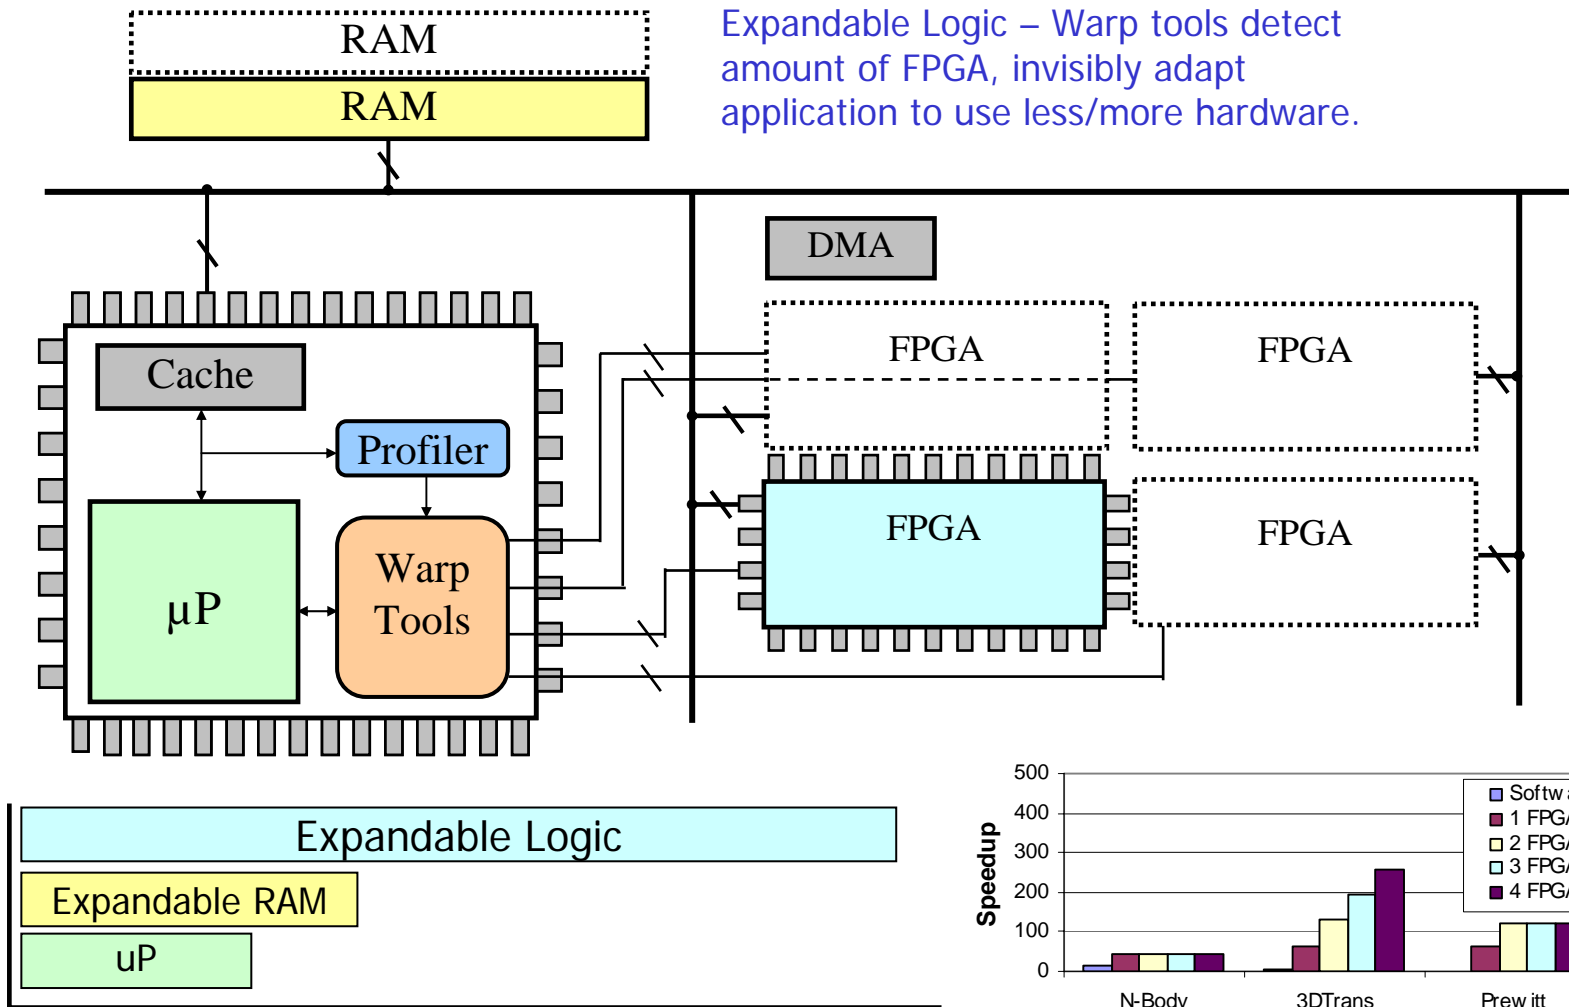
- Average 130x speedup

But, FPGA uses additional area

So we also compare to systems with 8 to 64 ARM11 uPs – FPGA size = ~36 ARM11s

- 11x faster than 64-core system
- Simulation pessimistic, actual results likely better

Dynamic Enables *Expandable Logic* Concept



Virtual Immersion

- eBlocks: Enables customized sensor-based system design by *non-experts*
 - May lead to ...? “Wood and nails of the sensor world”
 - Currently working with hearing-impaired, aging
- Warp processing (featured in this month’s IEEE Computer)
 - Enables large speedups on certain applications (e.g., image processing), user can expand hardware without changing software

