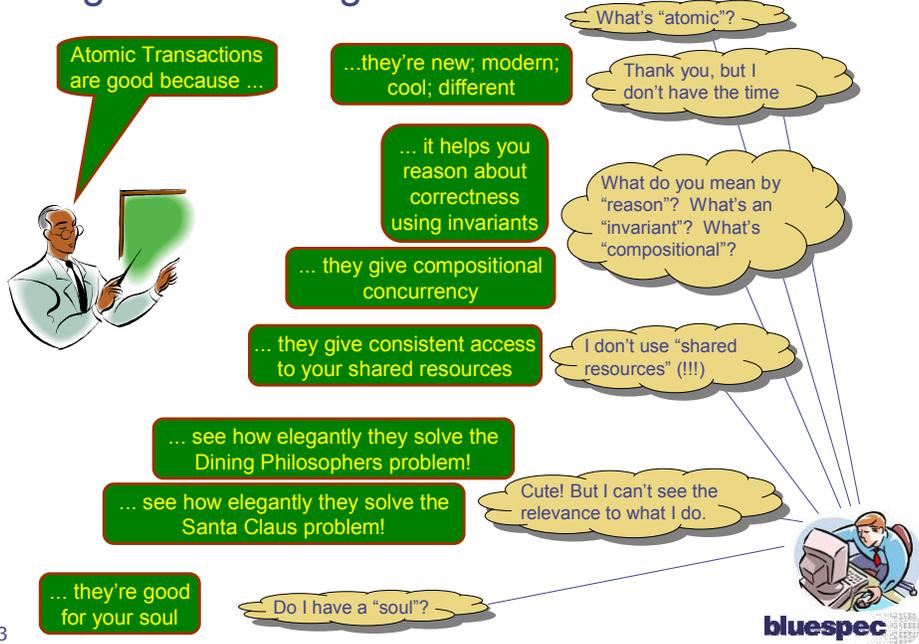
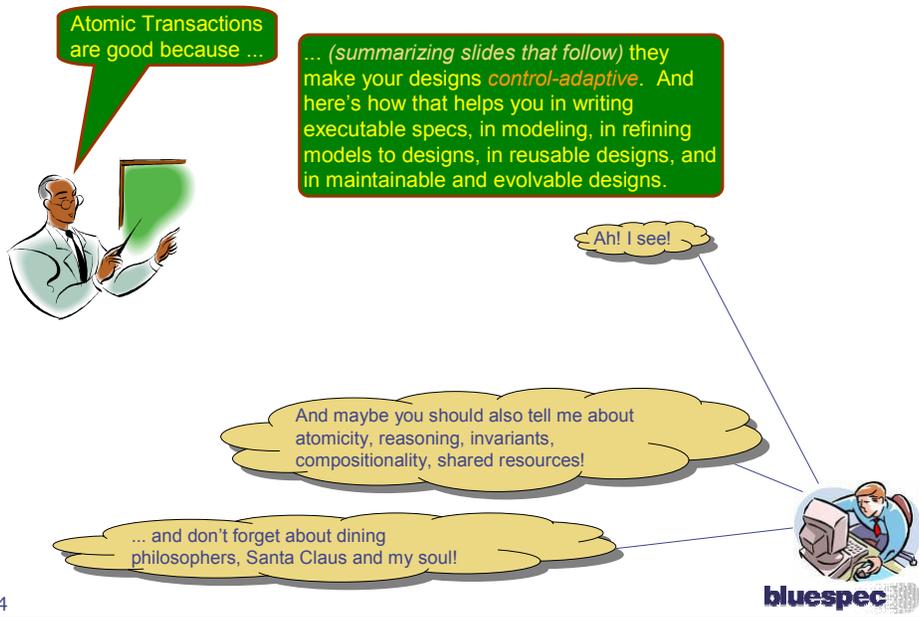




## Arguments that go over like a ton of bricks



## In summary ...



# GCD in BSV

```

module mkGCD (I_GCD);
  Reg#(int) x <- mkRegU;
  Reg#(int) y <- mkReg(0);

  rule swap ((x > y) && (y != 0));
    x <= y; y <= x;
  endrule
  rule subtract ((x <= y) && (y != 0));
    y <= y - x;
  endrule

  method Action start(int a, int b) if (y==0);
    x <= a; y <= b;
  endmethod
  method int result() if (y==0);
    return x;
  endmethod
endmodule

```

State

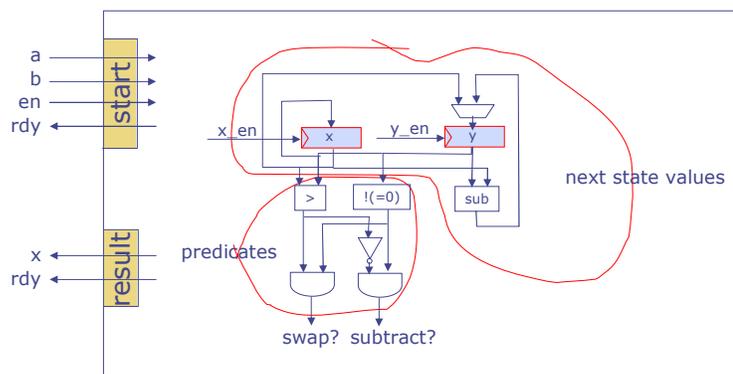
Internal behavior

External Sinterface

5



# Generated Hardware

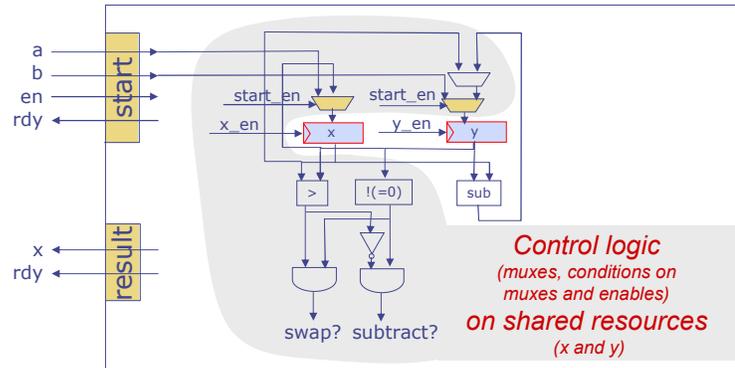


x\_en = swap?  
 y\_en = swap? OR subtract?

6



## Generated Hardware Module

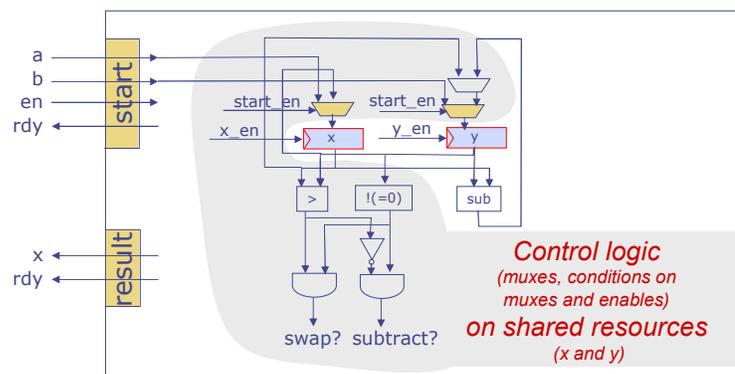


$x\_en = \text{swap? OR start\_en}$   
 $y\_en = \text{swap? OR subtract? OR start\_en}$   
 $rdy = (y==0)$

7

bluespec

## Generated Hardware Module



When coding this in Verilog, this control logic is explicit:

```

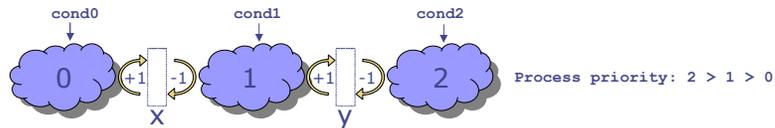
always
  @(posedge CLK)
  if (start_en) x <= a;
  else if (swap_en) x <= y;
  
```

8

bluespec

Summarizing, we have Point 1:

*BSV's Atomic Transactions (rules) are just another way of specifying control logic that you would have written anyway*



```

always @(posedge CLK) begin
  if (cond2)
    y <= y - 1;
  else if (cond1) begin
    y <= y + 1; x <= x - 1;
  end
  if (cond0 && !cond1)
    x <= x + 1;
end

```

```

always @(posedge CLK) begin
  if (cond2)
    y <= y - 1;
  else if (cond1) begin
    y <= y + 1; x <= x - 1;
  end
  if (cond0 && (!cond1 || cond2) )
    x <= x + 1;
end

```

Resource-access scheduling logic i.e., control logic

Better scheduling

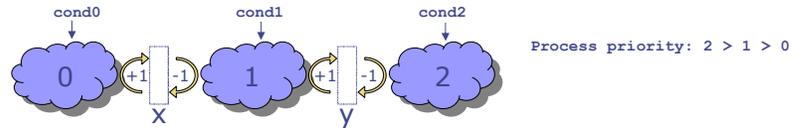
\* There are other ways to write this RTL, but all suffer from same analysis

Fundamentally, we are scheduling three potentially concurrent atomic transactions that share resources.

Note that control of x is affected by a "non-adjacent" condition (cond2), because of atomicity of the intervening process 1.

This is fundamentally what makes RTL so fragile!

## In BSV

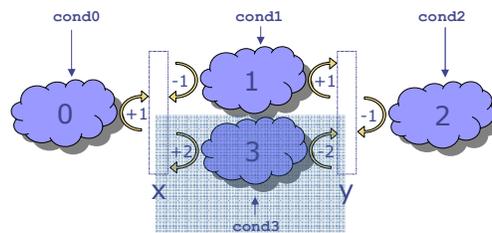


```
rule proc0 (cond0);  
  x <= x + 1;  
endrule  
  
rule proc1 (cond1);  
  y <= y + 1;  
  x <= x - 1;  
endrule  
  
rule proc2 (cond2);  
  y <= y - 1;  
endrule  
  
(* descending_urgency = "proc2, proc1, proc0" *)
```

11

bluespec

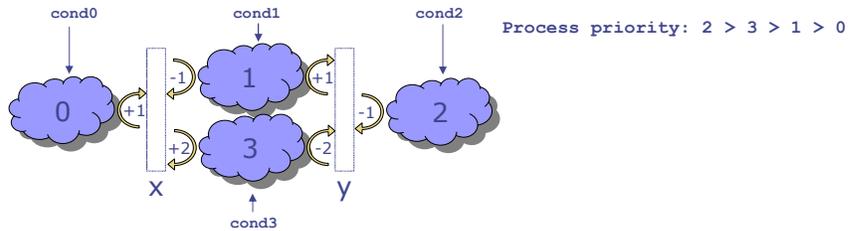
Now, let's make a small change: add a new process and insert its priority



12

bluespec

# Changing the BSV design

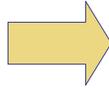


## Pre-Change

```
(* descending_urgency = "proc2, proc1, proc0" *)
rule proc0 (cond0);
  x <= x + 1;
endrule

rule proc1 (cond1);
  y <= y + 1;
  x <= x - 1;
endrule

rule proc2 (cond2);
  y <= y - 1;
endrule
```



```
(* descending_urgency = "proc2, proc3, proc1, proc0" *)
rule proc0 (cond0);
  x <= x + 1;
endrule

rule proc1 (cond1);
  y <= y + 1;
  x <= x - 1;
endrule

rule proc2 (cond2);
  y <= y - 1;
  x <= x + 1;
endrule

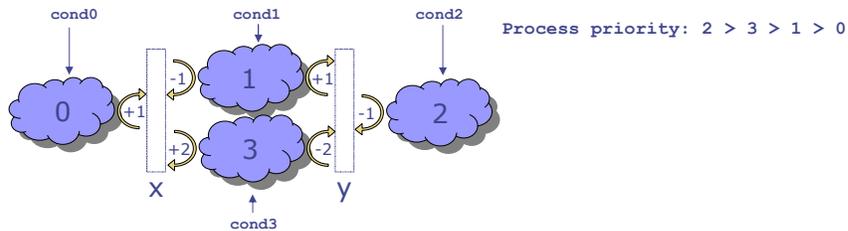
rule proc3 (cond3);
  y <= y - 2;
  x <= x + 2;
endrule
```



13

bluespec

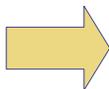
# Changing the Verilog design



## Pre-Change

```
always @(posedge CLK) begin
  if (!(cond2 && cond1))
    x <= x - 1;
  else if (cond0)
    x <= x + 1;

  if (cond2)
    y <= y - 1;
  else if (cond1)
    y <= y + 1;
end
```



```
always @(posedge CLK) begin
  if ((cond2 && cond0) || (cond0 && !cond1 && !cond3))
    x <= x + 1;
  else if (cond3 && !cond2)
    x <= x + 2;
  else if (cond1 && !cond0)
    x <= x - 1;

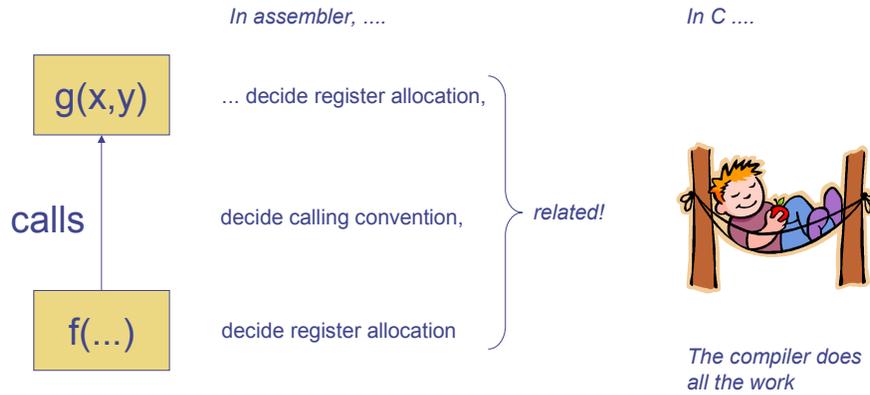
  if (cond2)
    y <= y - 1;
  else if (cond3)
    y <= y - 2;
  else if (cond1)
    y <= y + 1;
end
```



14

bluespec

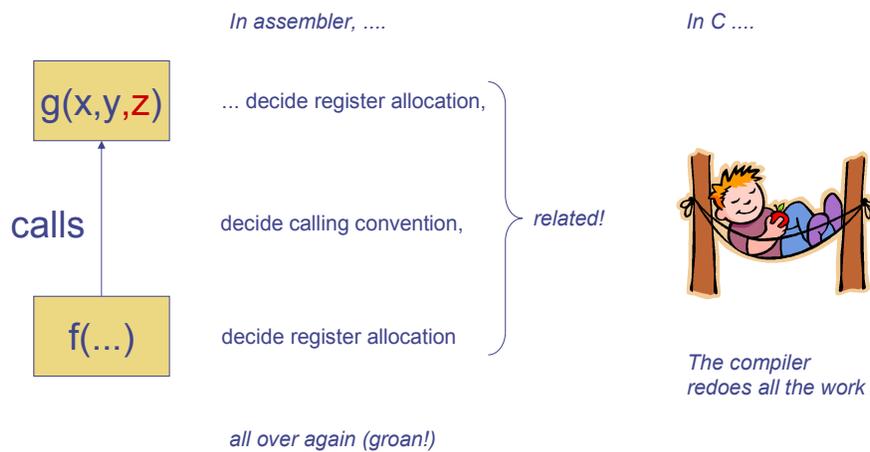
## A software analogy



15

bluespec

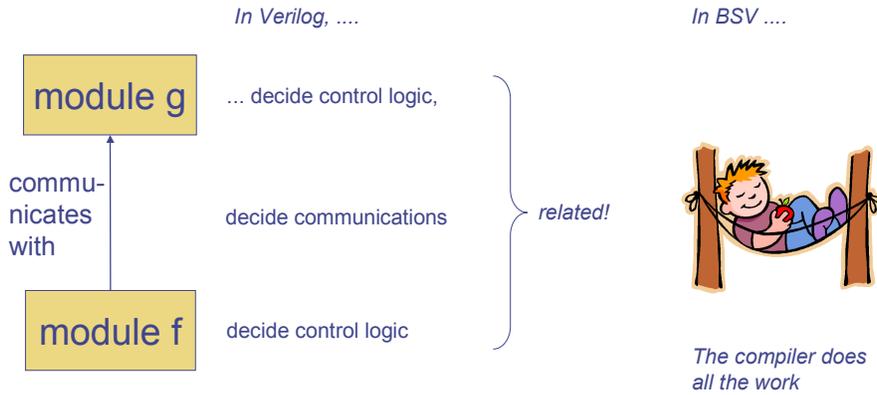
## A software analogy



16

bluespec

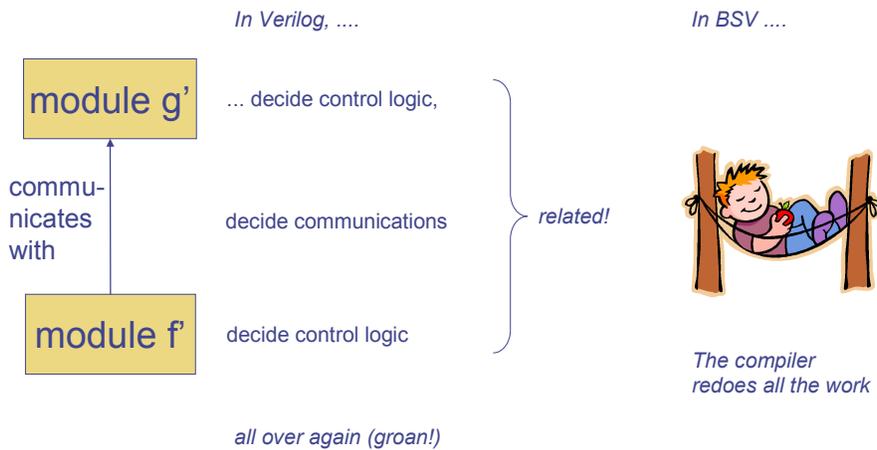
# In HW design



17

bluespec

# In HW design



18

bluespec

Summarizing, we have Point 2:

*In RTL, control logic can be complex to write first time, and complex to change/maintain, because one needs to consider and reconsider non-local influences (even across module boundaries).*

*BSV, on the other hand, is “control adaptive” because it automatically computes and recomputes the control logic as you write/change the design.*

19

bluespec

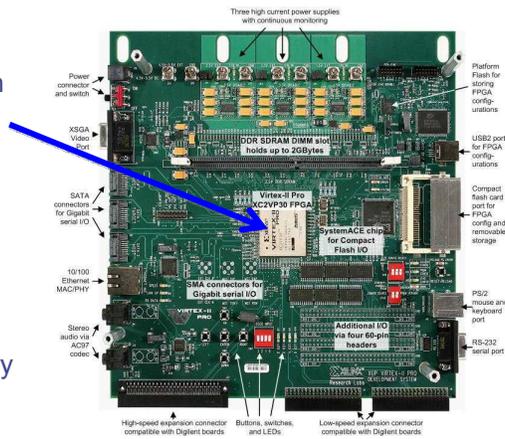
## MEMOCODE 2008 codesign contest

<http://rijndael.ece.vt.edu/memocontest08/>

Goal: Speed up a software reference application running on the PowerPC on Xilinx XUP reference board using SW/HW codesign

Application:  
decrypt,  
sort,  
re-encrypt  
db of records in external memory

Time allotted: 4 weeks



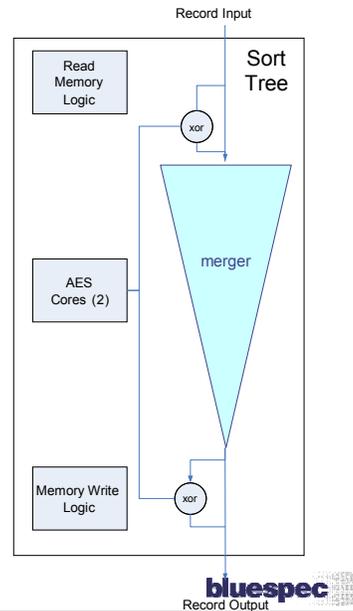
Xilinx XUP

20

bluespec

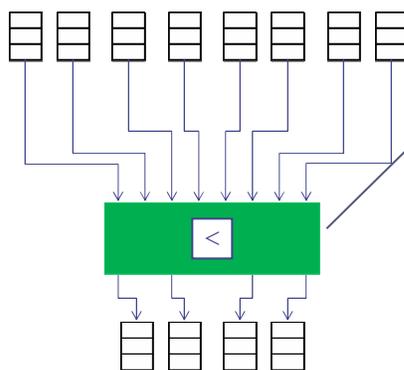
## The winning architecture

- The “merger” has to be used repeatedly in order to mergesort the db
- A control issue: address generation and sequencing of memory reads and writes
- Observation: throughput in merger is limited by apex, so at each lower level, a single comparator can be shared  
 → can fit a deeper merger into available space



21

## Merger: each level



Loop:

1. Choose non-empty input pair corresponding to output fifo with room (scheduling)
2. Compare the fifo heads
3. Dequeue the smaller one and put it on output fifo

22

bluespec

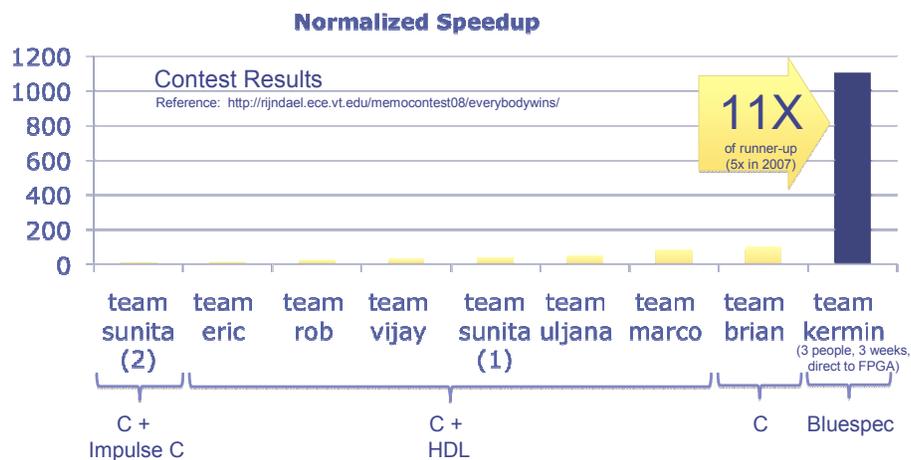
## Scheduler as a Module Parameter

- ♦ Level 1 merger: must check 2 input FIFOs
    - Scheduler can be combinational
  - ♦ Level 6 merger: must check 64 input FIFOs
    - Scheduler must be pipelined
- scheduler is a parameter of level merger

23

bluespec

## MEMOCODE 2008 codesign contest



27 teams started the four week contest.  
8 teams delivered 9 solutions.

24

bluespec

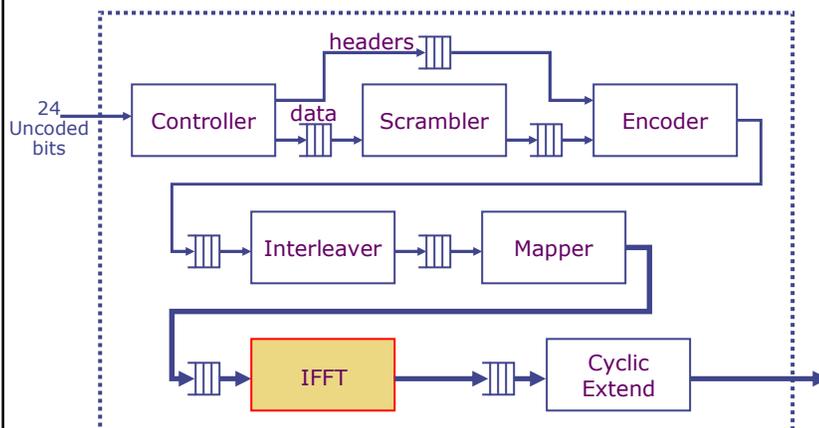
Summarizing, we have Point 3:

*Some control logic is too difficult even to contemplate without the support of atomic transactions*

25

bluespec

### Another application of control adaptivity: parameterizing an 802.11a Transmitter

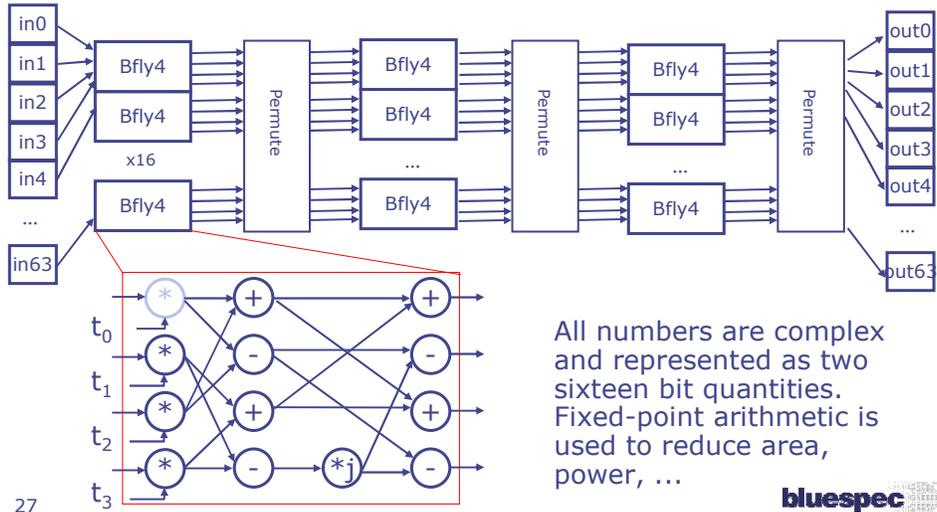


26

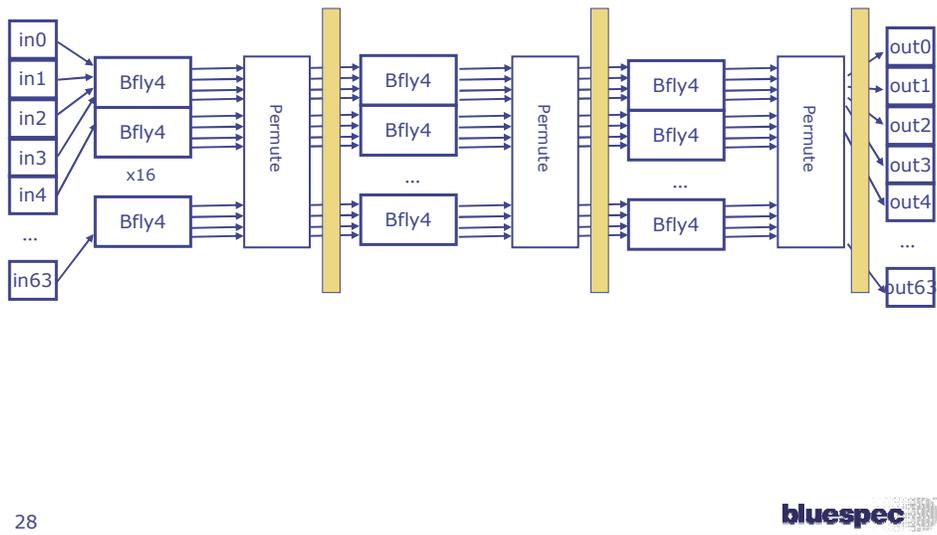
accounts for 85% area

bluespec

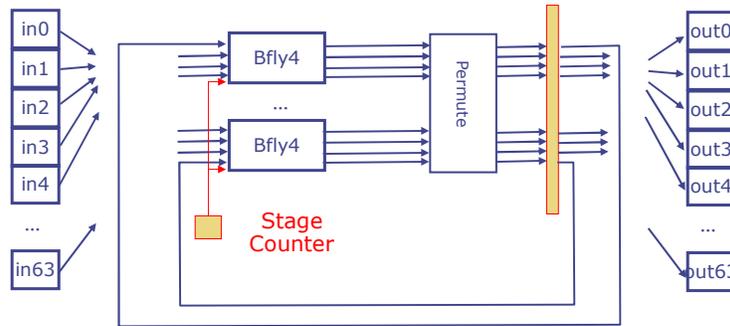
## Combinational IFFT



## Pipelined IFFT



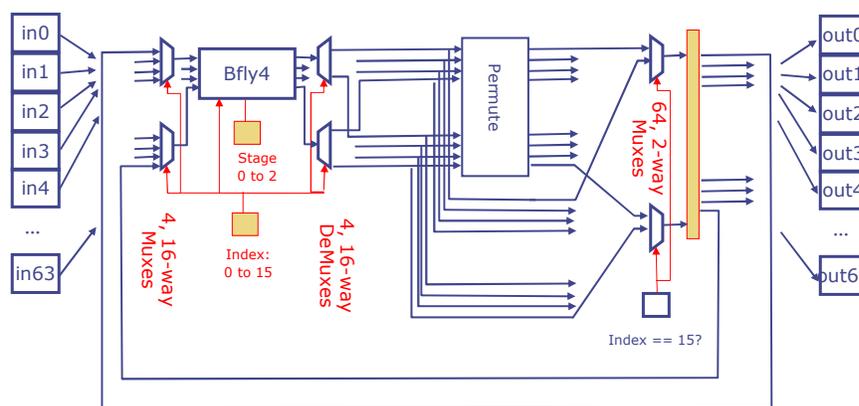
## Circular pipeline: Reusing the Pipeline Stage



29

bluespec

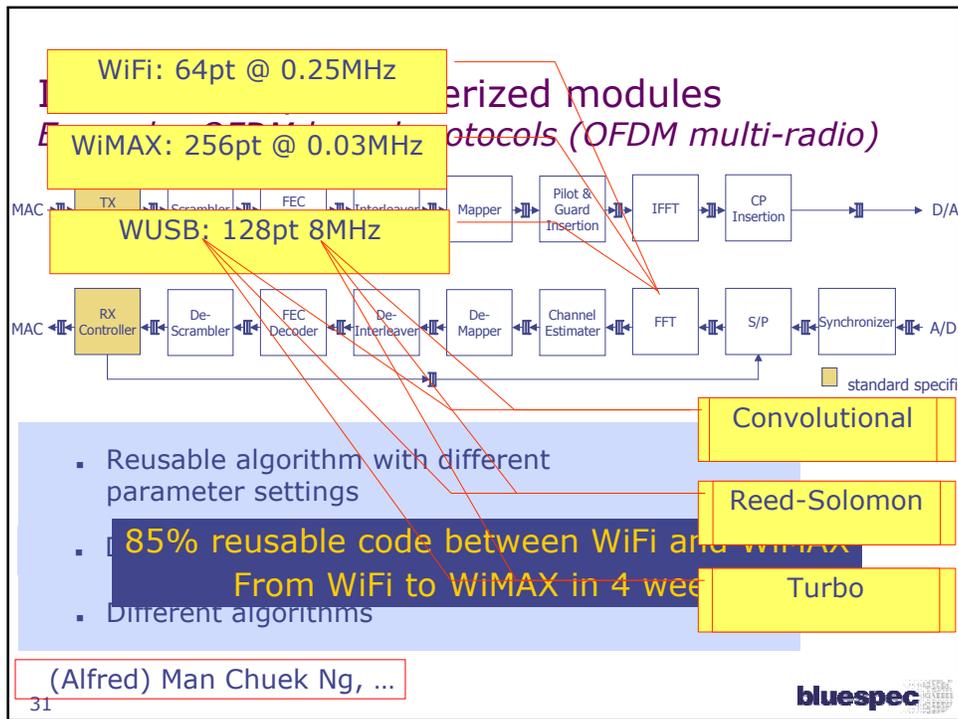
## Superfolded circular pipeline: Just one Bfly-4 node!



*These different micro-architectures will have different (area, speed, power). Each may be the "best" for different target chips.*

30

bluespec



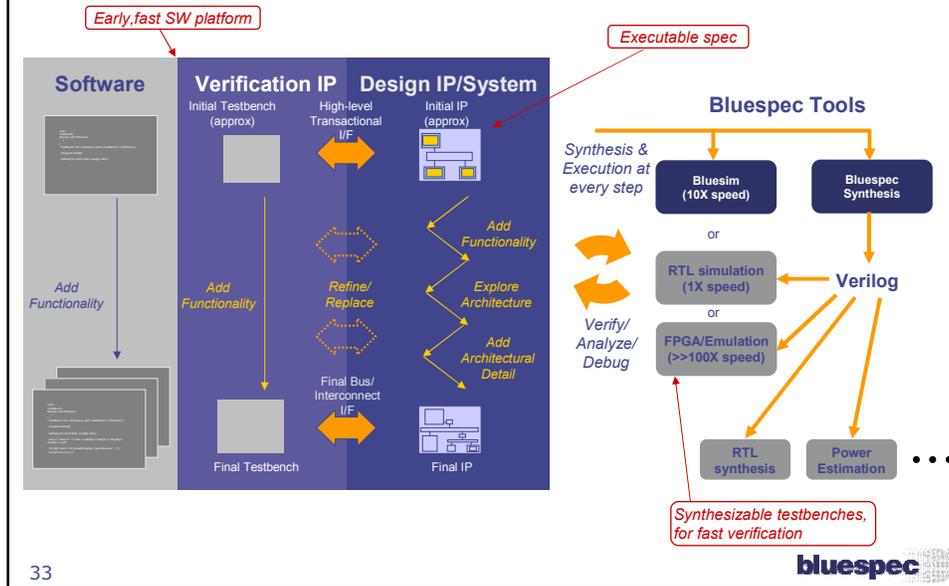
Summarizing, we have Point 4:

Parameterized architectures can deliver tremendous succinctness and reuse ...

... but each instance of a parameterized architecture has different resource contentions → needs different control logic.

Consider the effort to redesign the control logic for each instance (RTL), vs. using regenerating it automatically due to control-adaptivity.

## Control-adaptivity has a profound effect on design methodology



## Gradually expanding ways in which BSV is used

- ◆ Modeling for early SW development
  - Denali
    - Cycle-accurate LPDDR controller
- ◆ Modeling for architecture exploration
  - Intel, IBM, UT Austin, CMU
    - x86 and Power multi-thread/multi-core design
- ◆ Verification
  - Qualcomm
    - Testbenches, transactors
- ◆ and, of course, IP creation (original use)
  - TI, ST, Nokia, Mercury, MIT
    - DMAs, LCD controllers, Wireless modems, Video, ...

## In summary ...

Atomic Transactions  
are good because ...



... (summarizing) they make your designs *control-adaptive*. And this helps you in writing executable specs, in modeling, in refining models to designs, in creating complex, reusable designs, maintainable and evolvable designs.

Ah! I see!

What about Santa Claus\* and my soul?



35

\* ask me if you want to see the Bluespec solution

**bluespec**

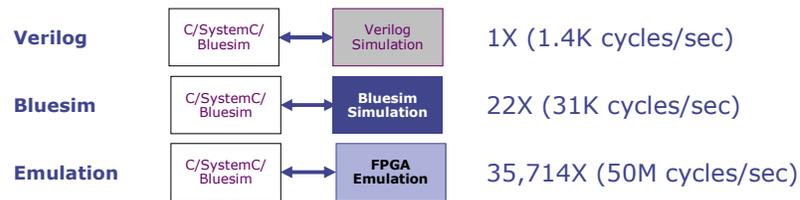
## Extra slides

36

**bluespec**

## Example: **IP verification** (AXI demo on FPGA @ Bluespec)

Bluespec lines of code:	2,000 (including comments)
ASIC gates:	125K
Virtex-4 FX100 slices:	4,723 (10% utilization)

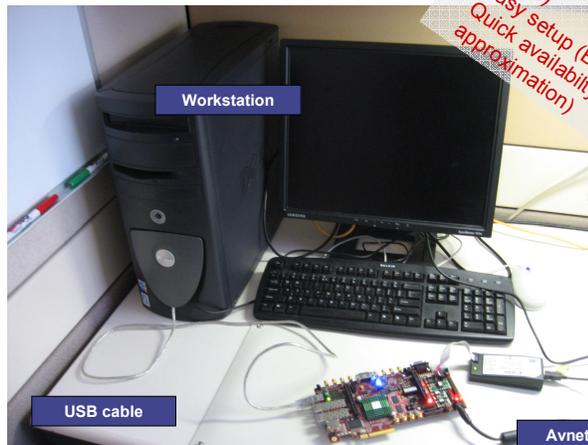


22 day Verilog sim → 1 day Bluesim → 53 sec on FPGA

37

bluespec

## Example: **IP verification** (AXI demo on FPGA @ Bluespec)



- Why BSV?**
- Inexpensive emulation platform (low cost FPGA board)
  - Easy setup (BSV transactors)
  - Quick availability (synthesizability of early approximation)

38

bluespec