



FireSim

Instrumenting and Debugging FireSim-Simulated Designs

<https://firesim.com>



@firesimproject

MICRO 2019 Tutorial

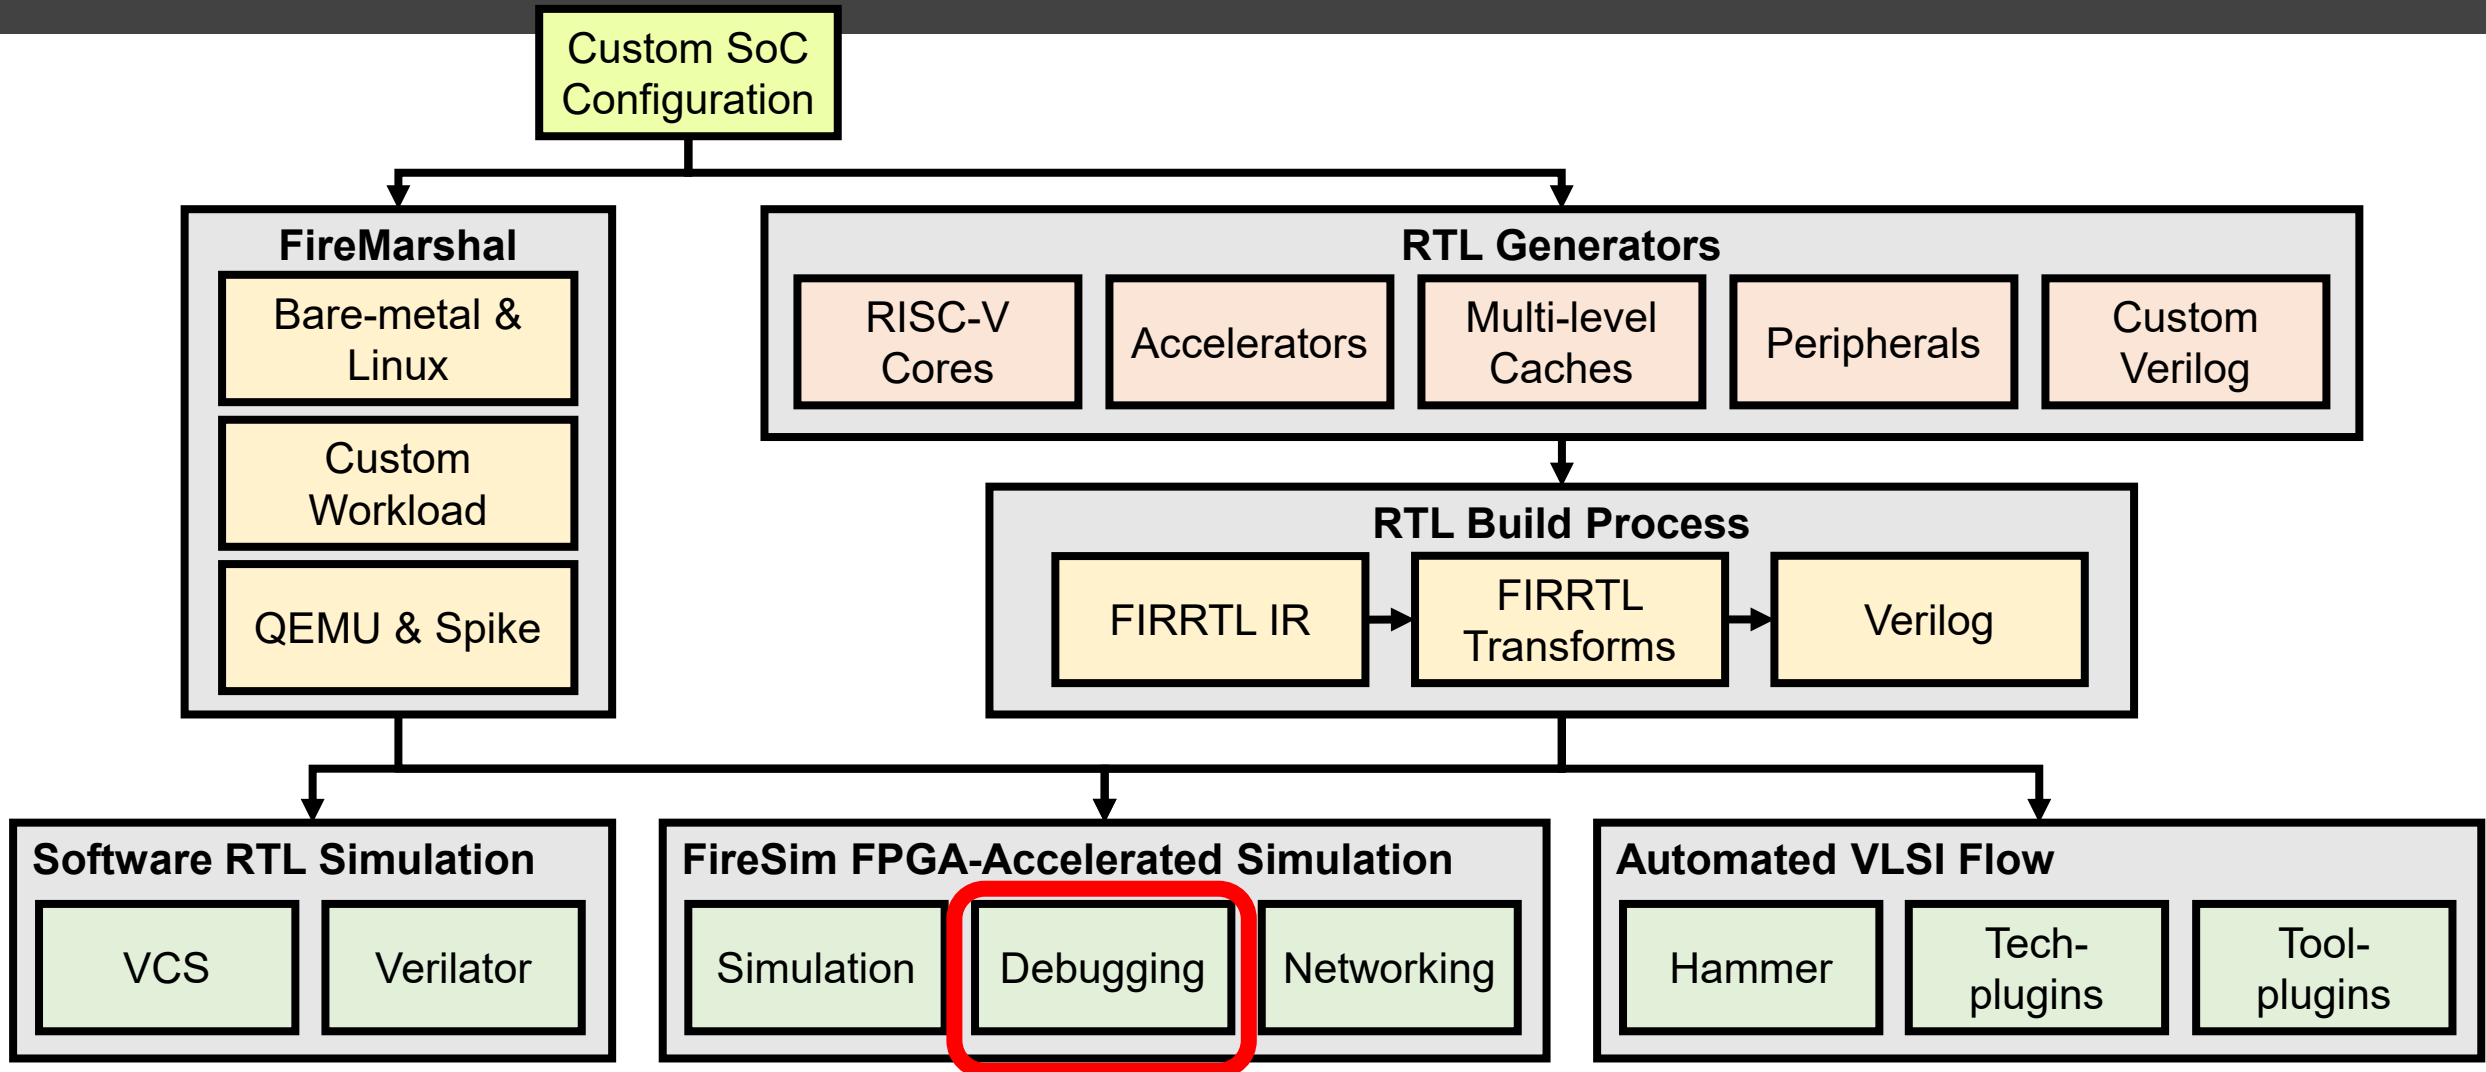
Speaker: Alon Amid



Berkeley Architecture Research



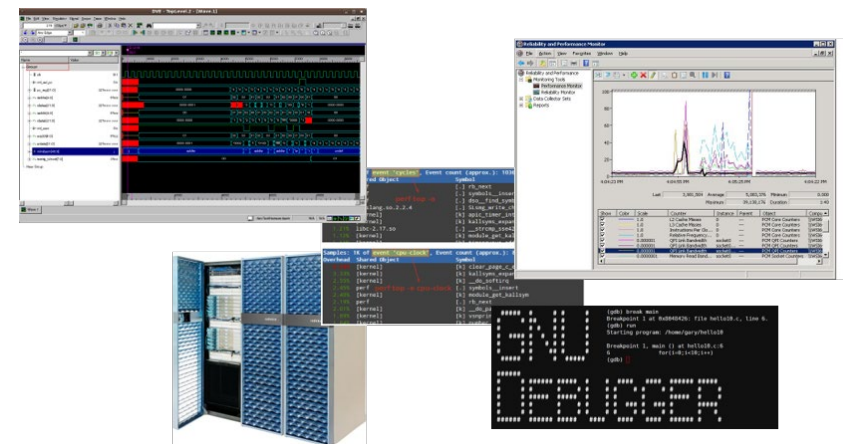
Tutorial Roadmap





Agenda

- FPGA-Accelerated Deep-Simulation Debugging
 - Debugging Using Integrated Logic Analyzers
 - Trace-based Debugging
 - Synthesizable Assertions/Prints
 - Hands-on example
- Debugging Co-Simulation
 - FireSim Debugging Using Software Simulation

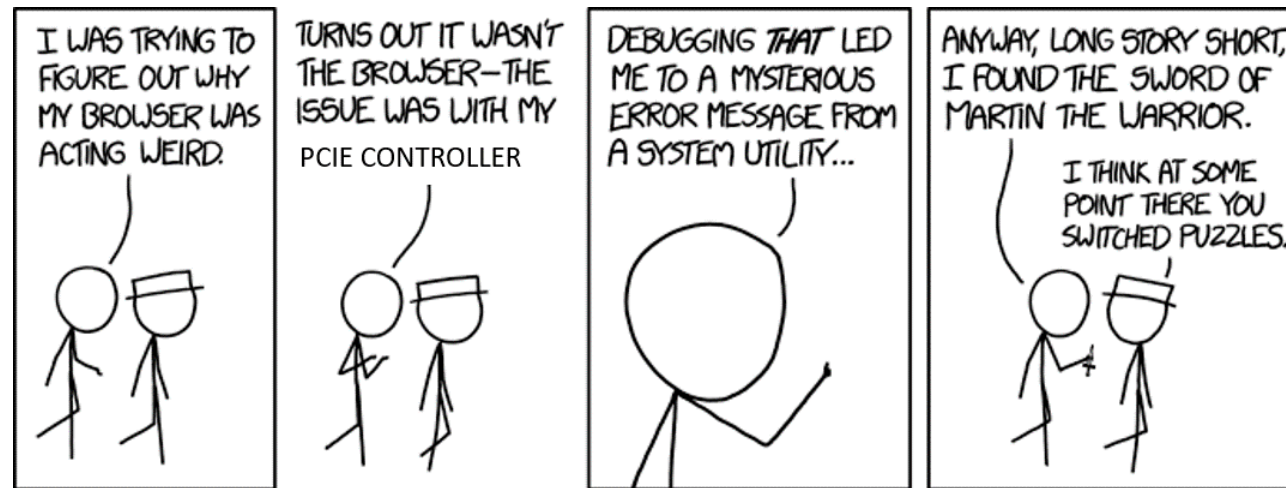




When SW RTL Simulation is Not Enough...

“Everything looks OK in SW simulation, but there is still a bug somewhere”

“My bug only appears after hours of running Linux on my simulated HW”





FPGA-Based Debugging Features

- High simulation speed in FPGA-based simulation enables advanced debugging and profiling tools.
- Reach “deep” in simulation time, and obtain large levels of coverage and data
- Examples:
 - ILAs
 - TracerV
 - Synthesizable assertions, prints



SW
Simulation

FPGA-based
Simulation

Simulated
Time



Debugging Using Integrated Logic Analyzers

Integrated Logic Analyzers (ILAs)

- Common debugging feature provided by FPGA vendors
- Continuous recording of a sampling window
 - Up to 1024 cycles by default.
 - Stores recorded samples in BRAM.
- Realtime trigger-based sampled output of probed signals
 - Multiple probes ports can be combined to a single trigger
 - Trigger can be in any location within the sampling window
- On the AWS F1-Instances, ILA interfaced through a debug-bridge and server

```
// Integrated Logic Analyzers (ILA)
ila_0 Cl_ILA_0 (
    .clk      (clk_main_a0),
    .probe0   (sh_ocl_awvalid_q),
    .probe1   (sh_ocl_awaddr_q ),
    .probe2   (ocl_sh_awready_q),
    .probe3   (sh_ocl_arvalid_q),
    .probe4   (sh_ocl_araddr_q ),
    .probe5   (ocl_sh_arready_q)
);

ila_0 Cl_ILA_1 (
    .clk      (clk_main_a0),
    .probe0   (ocl_sh_bvalid_q),
    .probe1   (sh_cl_glcount0_q),
    .probe2   (sh_ocl_bready_q),
    .probe3   (ocl_sh_rvalid_q),
    .probe4   ({32'b0,ocl_sh_rdata_q[31:0]}),
    .probe5   (sh_ocl_rready_q)
);

// Debug Bridge
cl_debug_bridge Cl_DEBUG_BRIDGE (
    .clk(clk_main_a0),
    .S_BSCAN_drck(drck),
    .S_BSCAN_shift(shift),
    .S_BSCAN_tdi(tdi),
    .S_BSCAN_update(update),
    .S_BSCAN_sel(sel),
    .S_BSCAN_tdo(tdo),
    .S_BSCAN_tms(tms),
    .S_BSCAN_tck(tck),
    .S_BSCAN_runtest(runtest),
    .S_BSCAN_reset(reset),
    .S_BSCAN_capture(capture),
    .S_BSCAN_bscanid_en(bscanid_en)
);
```

From: aws-fpga cl_hello_world example

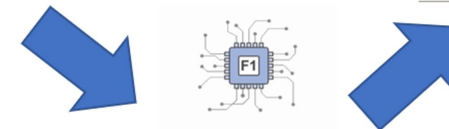
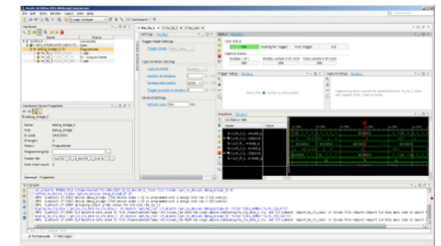


Debugging Using Integrated Logic Analyzers

AutoILA – Automation of ILA integration with FireSim

- Annotate requested signals and bundles in the Chisel source code
- Automatic configuration and generation of the ILA IP in the FPGA toolchain
- Automatic expansion and wiring of annotated signals to the top level of a design using a FIRRTL transform.
- Remote waveform and trigger setup from the manager instance

```
import midas.targetutils.FpgaDebug  
  
class SomeModuleIO(implicit p: Parameters) extends SomeIO()(p){  
  val out1 = Output(Bool())  
  val in1 = Input(Bool())  
  FpgaDebug(out1, in1)  
}
```





BOOM Example

- Debugging an out-of-order processor is hard
 - Throughout this talk, we'll have examples of FPGA debugging used in BOOM.
- Example from `boom/src/main/scala/lisu/dcache.scala`
- Debugging a non-blocking data cache hanging after Linux boots

```
class BoomNonBlockingDCacheModule(outer: BoomNonBlockingDCache) extends LazyModuleImp(outer)
  with HasL1HellaCacheParameters
{
  implicit val edge = outer.node.edges.out(0)
  val (tl_out, _) = outer.node.out(0)
  val io = IO(new BoomDCacheBundle)

  FpgaDebug(tl_out)
  FpgaDebug(io.req)
  FpgaDebug(io.resp)
  FpgaDebug(io.s1_kill)
  FpgaDebug(io.nack)
  ...
}
```





Debugging using Integrated Logic Analyzers



Pros:

- No emulated parts – what you see is what's running on the FPGA
- FPGA simulation speed - $O(\text{MHz})$ compared to $O(\text{KHz})$ in software simulation
- Real-time trigger-based

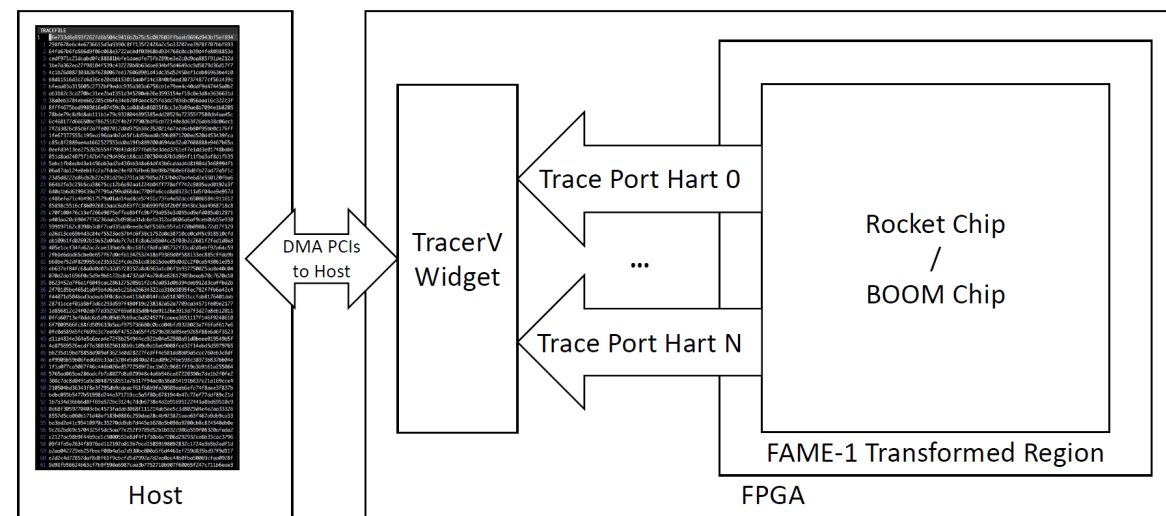
Cons:

- Requires a full build to modify visible signals/triggers (takes several hours)
- Limited sampling window size
- Consumes FPGA resources



TracerV

- **Out-of-band** full instruction execution trace
- Bridge connected to target trace ports
- By default, large amount of info wired out of Rocket/BOOM, per-hart, per-cycle:
 - Instruction Address
 - Instruction
 - Privilege Level
 - Exception/Interrupt Status, Cause
- TracerV can rapidly generate several TB of data.

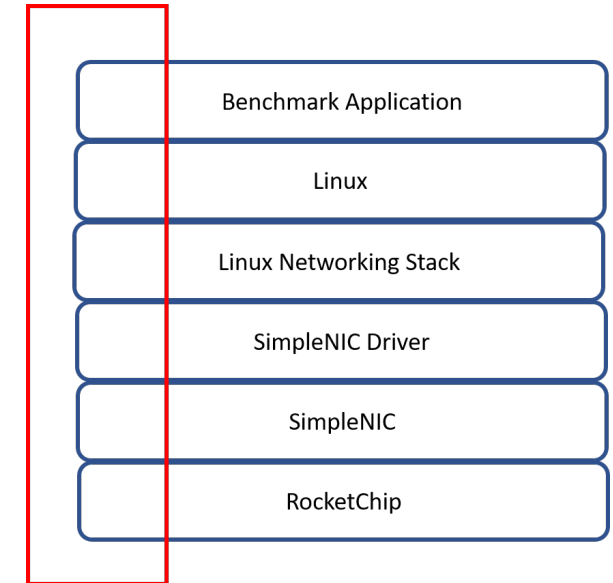




TracerV

- Out-of-Band: profiling does not perturb execution
- Useful for kernel and hypervisor level cycle-sensitive profiling
- Examples:
 - Co-Optimization of NIC and Network Driver
 - Keystone Secure Enclave Project
 - High-performance hardware-specific code (supercomputing?)
- Requires large-scale analytics for insightful profiling and optimization.

Why
Is
This
Slow
?





TracerV



Pros:

- Out-of-Band (no impact on workload execution)
- SW-centric method
- Large amounts of data

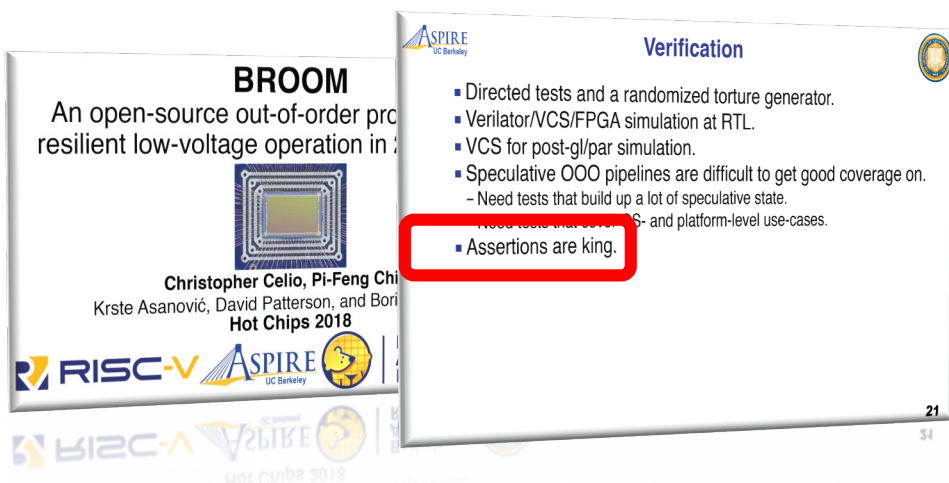
Cons:

- Slower simulation performance (40 MHz)
- No HW visibility
- Large amounts of data



Synthesizable Assertions

- Assertions – rapid error checking embedded in HW source code.
 - Commonly used in SW Simulation
 - Halts the simulation upon a triggered assertion. Represented as a “stop” statement in FIRRTL
 - By default, emitted as non-synthesizable SV functions (\$fatal)



```
class Count extends Module {  
  val io = IO(new Bundle {  
    val en = Input(Bool())  
    val done = Output(Bool())  
    val cntnr = Output(UInt(4.W))  
  })  
  // count until 10 when `io.en` is high  
  val (cntnr, done) = Counter(io.en, 10)  
  io.cntnr := cntnr  
  io.done := done  
  
  // assertion for software simulation  
  // `cntnr` should be less than 10  
  assert(cntnr < 10.U)  
}
```

From: BROOM: An open-source Out-of-Order processor with resilient low-voltage operation in 28nm CMOS, Christopher Celio, Pi-Feng Chiu, Krste Asanovic, David Patterson and Borivoje Nikolic. HotChip 30, 2018

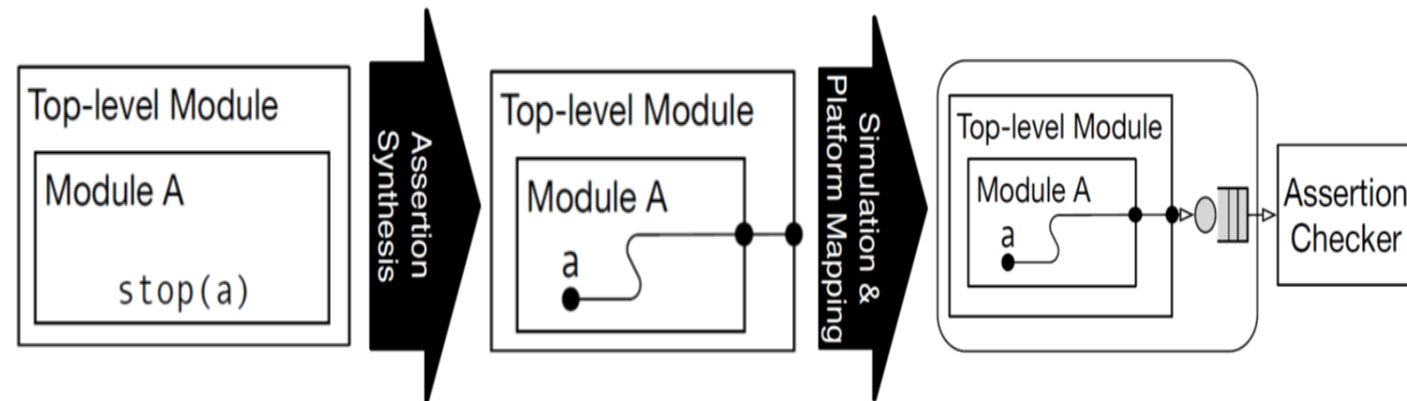
From: Trillion-Cycle Bug Finding Using FPGA-Accelerated Simulation Donggyu Kim, Christopher Celio, Sagar Karandikar, David Biancolin, Jonathan Bachrach, Krste Asanović. ADEPT Winter Retreat 2018





Synthesizable Assertions

- Synthesizable Assertions on FPGA
 - Transform FIRRTL `stop` statements into synthesizable logic
 - Insert combinational logic and signals for the `stop` condition arguments
 - Insert encodings for each assertion (for matching error statements in SW)
 - Wire the assertion logic output to the Top-Level
 - Generate timing tokens for cycle-exact assertions
 - Assertion checker records the cycle and halts simulation when assertion is triggered





BOOM Example

- Example from `boom/src/main/scala/exu/rob.scala`
- Assert is the ROB is behaving un-expectedly
 - Overwriting a valid entry

```
assert (rob_val(rob_tail) === false.B, "[rob] overwriting a valid entry.")
assert ((io.enq_uops(w).rob_idx >> log2Ceil(coreWidth)) === rob_tail)
assert (!(io.wb_resps(i).valid && MatchBank(GetBankIdx(rob_idx)) &&
!rob_val(GetRowIdx(rob_idx))), "[rob] writeback (" + i + ") occurred to an
invalid ROB entry.")
```



BOOM Example

- How it looks in the UART output (while Linux is booting):

```
[ 0.008000] VFS: Mounted root (ext2 filesystem) on device 253:0.
[ 0.008000] devtmpfs: mounted
[ 0.008000] Freeing unused kernel memory: 148K
[ 0.008000] This architecture does not have kernel memory protection.
mount: mounting sysfs on /sys failed: No such device
Starting syslogd: OK
Starting klogd: OK
Starting mdev...
mdev: /sys/dev: No such file or directory
[id: 1840, module: Rob, path: FireBoom.boom_tile_1.core.rob]
Assertion failed: [rob] writeback (0) occurred to an invalid ROB entry.
    at rob.scala:504 assert (!(io.wb_resps(i).valid && MatchBank(GetBankIdx(rob_idx)) &&
    at cycle: 1112250469

*** FAILED *** (code = 1841) after 1112250485 cycles
time elapsed: 307.8 s, simulation speed = 3.61 MHz
FPGA-Cycles-to-Model-Cycles Ratio (FMR): 2.77
Beats available: 2165
Runs 1112250485 cycles
[FAIL] FireBoom Test
SEED: 1569631756
    at cycle 4294967295
```

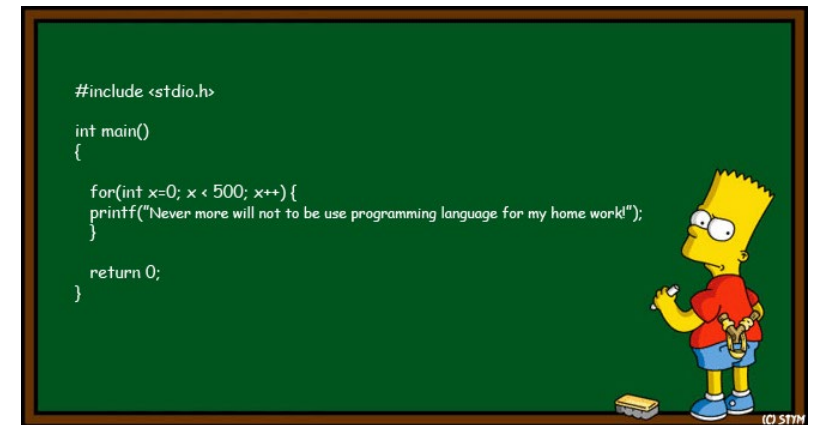
It would take ~62 hours to hit
this assertion is SW RTL
simulation (at 5 KHz sim rate),
vs. just a few minutes in FireSim





Synthesizable `printf`

- Research feature presented in DESSERT [1] (together with assertions)
- Enable “software-style” debugging using `printf` statements
- Convert Chisel `printf` statements to synthesizable blocks
 - Appropriate parsing in simulation bridge
 - Including signal values
- Impact on simulation performance depends on the frequency of `printf`s.
- Output includes the exact cycle of the `printf` event
 - Helps measure cycles counts between events



<https://www.deviantart.com/stym0r/art/Bart-Simpson-Programmer-134362686>



BOOM Example

- Example from `boom/src/main/scala/lsu/lsu.scala`
- Print a trace of all loads and stores, for verifying memory consistency.

```
if (MEMTRACE_PRINTF) {  
  when (commit_store || commit_load) {  
    val uop      = Mux(commit_store, stq(idx).bits.uop, ldq(idx).bits.uop)  
    val addr     = Mux(commit_store, stq(idx).bits.addr.bits, ldq(idx).bits.addr.bits)  
    val stdata  = Mux(commit_store, stq(idx).bits.data.bits, 0.U)  
    val wldata  = Mux(commit_store, stq(idx).bits.debug_wb_data, ldq(idx).bits.debug_wb_data)  
    printf(midas.targetutils.SynthesizePrintf("MT %x %x %x %x %x %x %x\n",  
      io.core.tsc_reg, uop.uopc, uop.mem_cmd, uop.mem_size, addr, stdata, wldata))  
  }  
}
```



Synthesizable `printf`/Assertions

Pros:

- FPGA simulation speed
- Real-time trigger-based
- Consumes small amount of FPGA resources (compared to ILA)
- Key signals have pre-written assertions in re-usable components/libraries

Cons:

- Low visibility: No waveform/state
- Assertions are best added while writing source RTL rather than during “investigative” debugging
- Large numbers of `printf`s can slow down simulation



Hands-on Synthesizable `printf` Example

- We would like to observe when the SHA3 algorithm completes a round, and some details about the round. This is represented by the
- [chipyard-afternoon/generators/sha3/src/main/scala/dpath.scala](https://github.com/berkeley-architecture/chipyard-afternoon/blob/master/generators/sha3/src/main/scala/dpath.scala)
 - Line 103

```
when(io.absorb) {
  state := state
  when(io.aindex < UInt(round_size_words)) {
    state((io.aindex%UInt(5))*UInt(5)+(io.aindex/UInt(5))) :=
      state((io.aindex%UInt(5))*UInt(5)+(io.aindex/UInt(5))) ^ io.message_in
  }
}
```



Hands-on Synthesizable `printf` Example

- We would like to observe when the SHA3 algorithm completes a round, and some details about the round. This is represented by the
- [chipyard-afternoon/generators/sha3/src/main/scala/dpath.scala](https://github.com/ucb-bar/chipyard-afternoon/blob/master/generators/sha3/src/main/scala/dpath.scala)
 - Line 103

```
when(io.absorb) {
  state := state
  printf(midas.targetutils.SynthesizePrintf("SHA3 finished an iteration with
index %d and message %x\n", io.aindex, io.message_in))
  when(io.aindex < UInt(round_size_words)) {
    state((io.aindex%UInt(5))*UInt(5)+(io.aindex/UInt(5))) :=
      state((io.aindex%UInt(5))*UInt(5)+(io.aindex/UInt(5))) ^ io.message_in
  }
}
```



Hands-on Synthesizable `printf` Example

- Since it takes 4 hours to rebuild an FPGA image, and we have only 1 hour left, we have prepared an FPGA image with this example synthesizable `printf` (using a parameterized configuration)

```
when(io.absorb) {
  state := state
  if(p(Sha3PrintfEnable)) {
    printf(midas.targetutils.SynthesizePrintf("SHA3 finished an iteration with
index %d and message %x\n", io.aindex, io.message_in))
  }
  when(io.aindex < UInt(round_size_words)) {
    state((io.aindex%UInt(5))*UInt(5)+(io.aindex/UInt(5))) :=
      state((io.aindex%UInt(5))*UInt(5)+(io.aindex/UInt(5))) ^ io.message_in
  }
}
```





Hands-on Synthesizable `printf` Example

- For reference, the build recipe for this FPGA image (in `deploy/config_build_recipes.ini`) is:

```
[firesim-singlecore-sha3-no-nic-l2-1lc4mb-ddr3-print]  
DESIGN=FireSimNoNIC  
TARGET_CONFIG=DDR3FRFCFSLLC4MB_FireSimRocketChipSha3L2PrintfConfig  
PLATFORM_CONFIG=WithPrintfSynthesis_BaseF1Config_F120MHz  
instancetype=c5.4xlarge  
deploytriplet=None
```



Hands-on Synthesizable `printf` Example

Update our workload to copy the output `printf` file:

- `vim $FDIR/deploy/workloads/sha3-bare-rocc.json`
- **Add the `synthesized-prints.out` to our simulation output**

```
{
  "benchmark_name": "sha3-bare-rocc",
  "common_simulation_outputs": [
    "uartlog", "synthesized-prints.out"
  ],
  "common_bootbinary": "../../../../../sw/firesim-
software/workloads/sha3/benchmarks/bare/sha3-rocc.riscv",
  "common_rootfs": "../../../../../sw/firesim-software/wlutil/dummy.rootfs"
}
```



Hands-on Synthesizable `printf` Example

- Setup the `config_runtime.ini`
`vim $FDIR/deploy/config_runtime.ini`
 - Select the AGFI that was synthesized with the `printf`
 - Select the bare-metal SHA3 test workload
- Boot the simulation by running the following sequence of commands:
 - `$ firesim infrasetup`
 - This should take about 3 minutes
 - `$ firesim runworkload`
 - This should take about <1 minute

```
f1_16xlarges=0  
m4_16xlarges=0  
f1_4xlarges=0  
f1_2xlarges=1
```

```
runinstancemarket=ondemand  
spotinterruptionbehavior=terminate  
spotmaxprice=ondemand
```

```
[targetconfig]  
topology=no_net_config  
no_net_num_nodes=1  
linklatency=6405  
switchinglatency=10  
netbandwidth=200  
profileinterval=-1
```

```
defaulthwconfig=firesim-singlecore-  
sha3-no-nic-12-11c4mb-ddr3-print
```

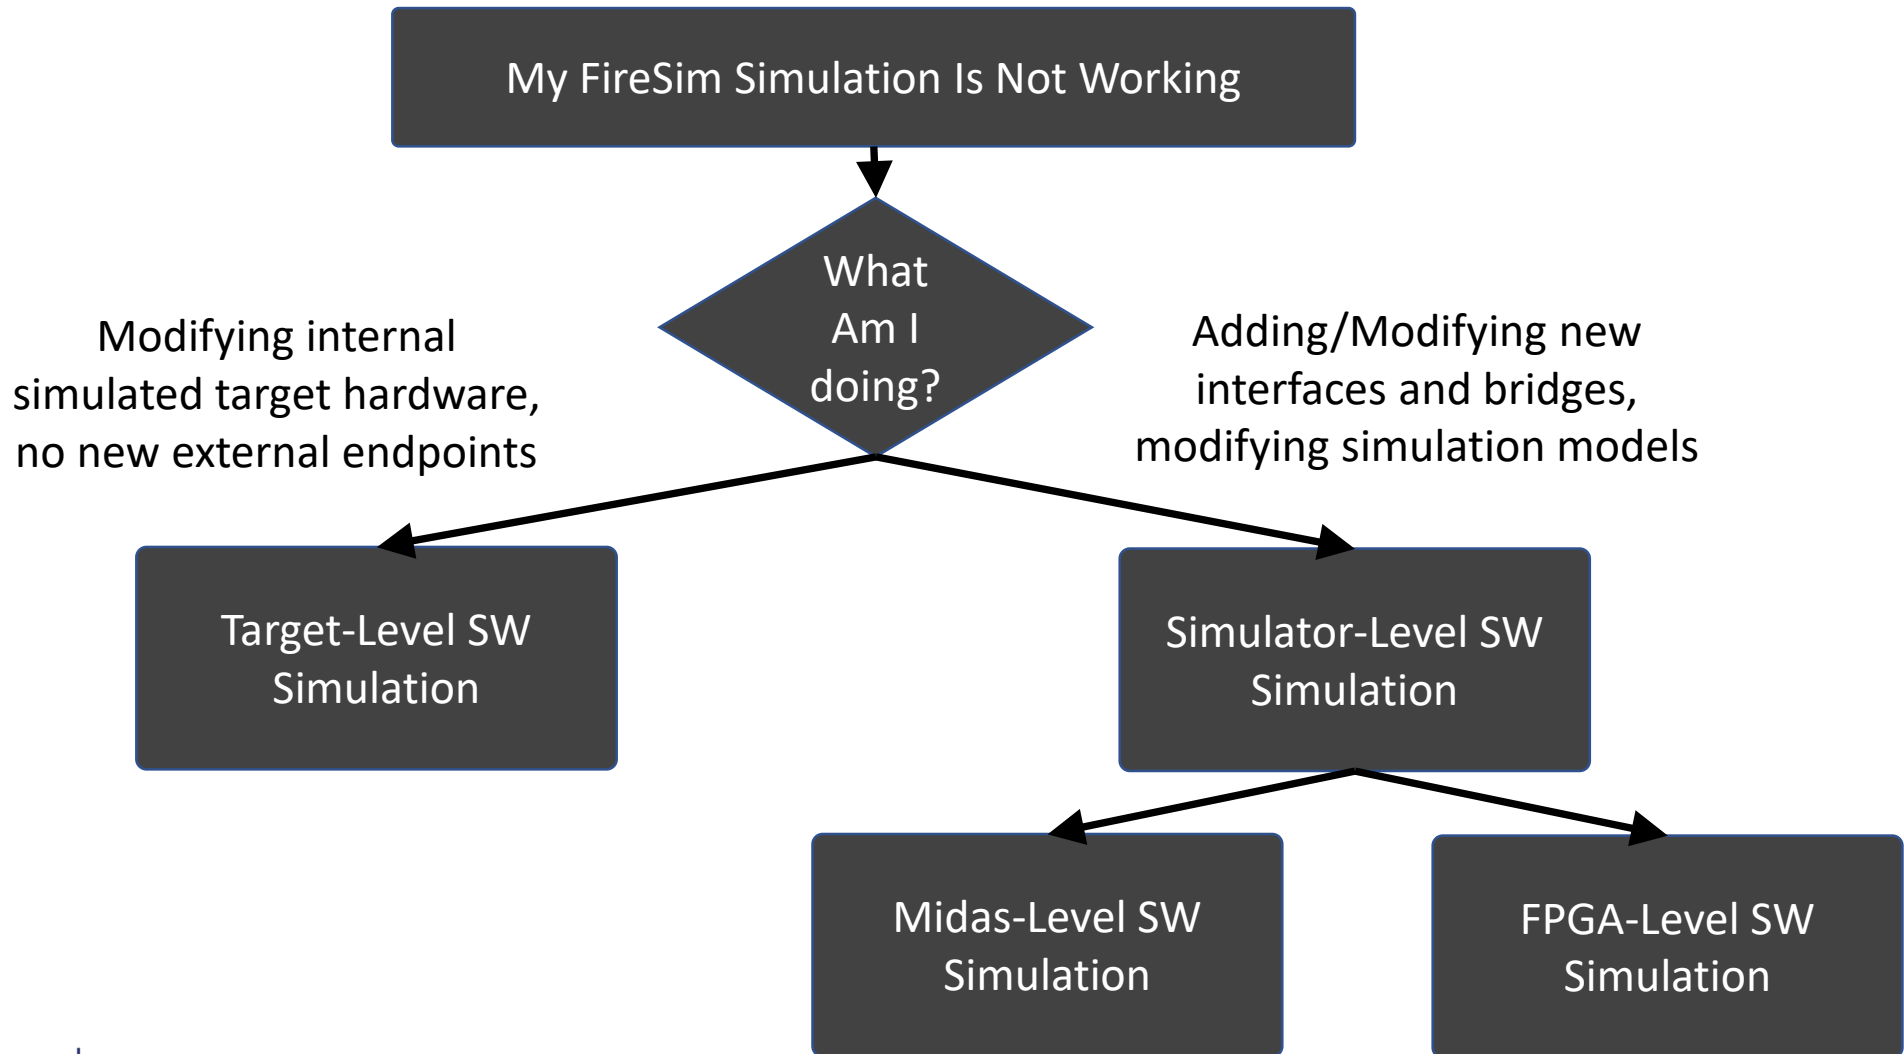
```
[workload]  
workloadname=sha3-bare-rocc.json
```



While this is running...



Debugging Using Software RTL Simulation





Debugging Using Software RTL Simulation

Target-Level Simulation

- Software Simulation
- Target Design Untransformed
- No Host-FPGA interfaces

MIDAS-Level Simulation

- Software Simulation
- Target Design Transformed by Golden Gate
- Host-FPGA interfaces/shell emulated using abstract models

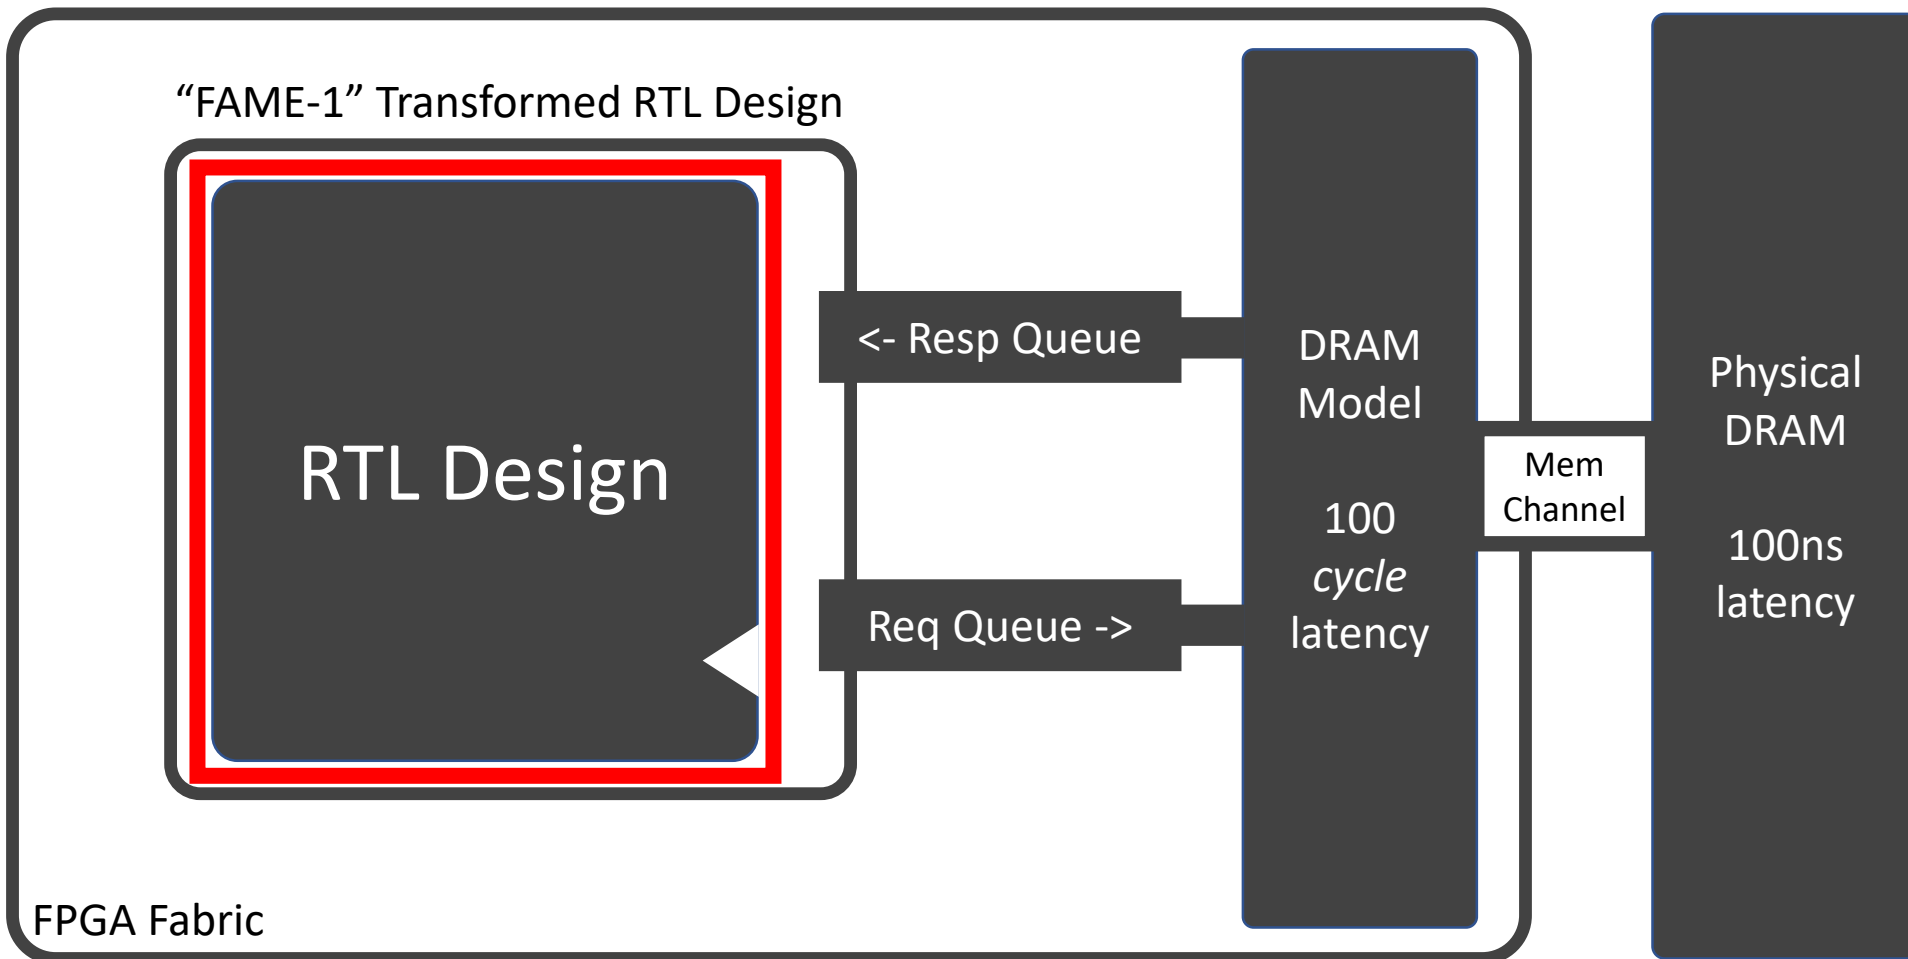
FPGA-Level Simulation

- Software Simulation
- Target Design Transformed by Golden Gate
- Host-FPGA interfaces/shell simulated by the FPGA tools



Debugging Using Software RTL Simulation

Target-Level
SW Simulation

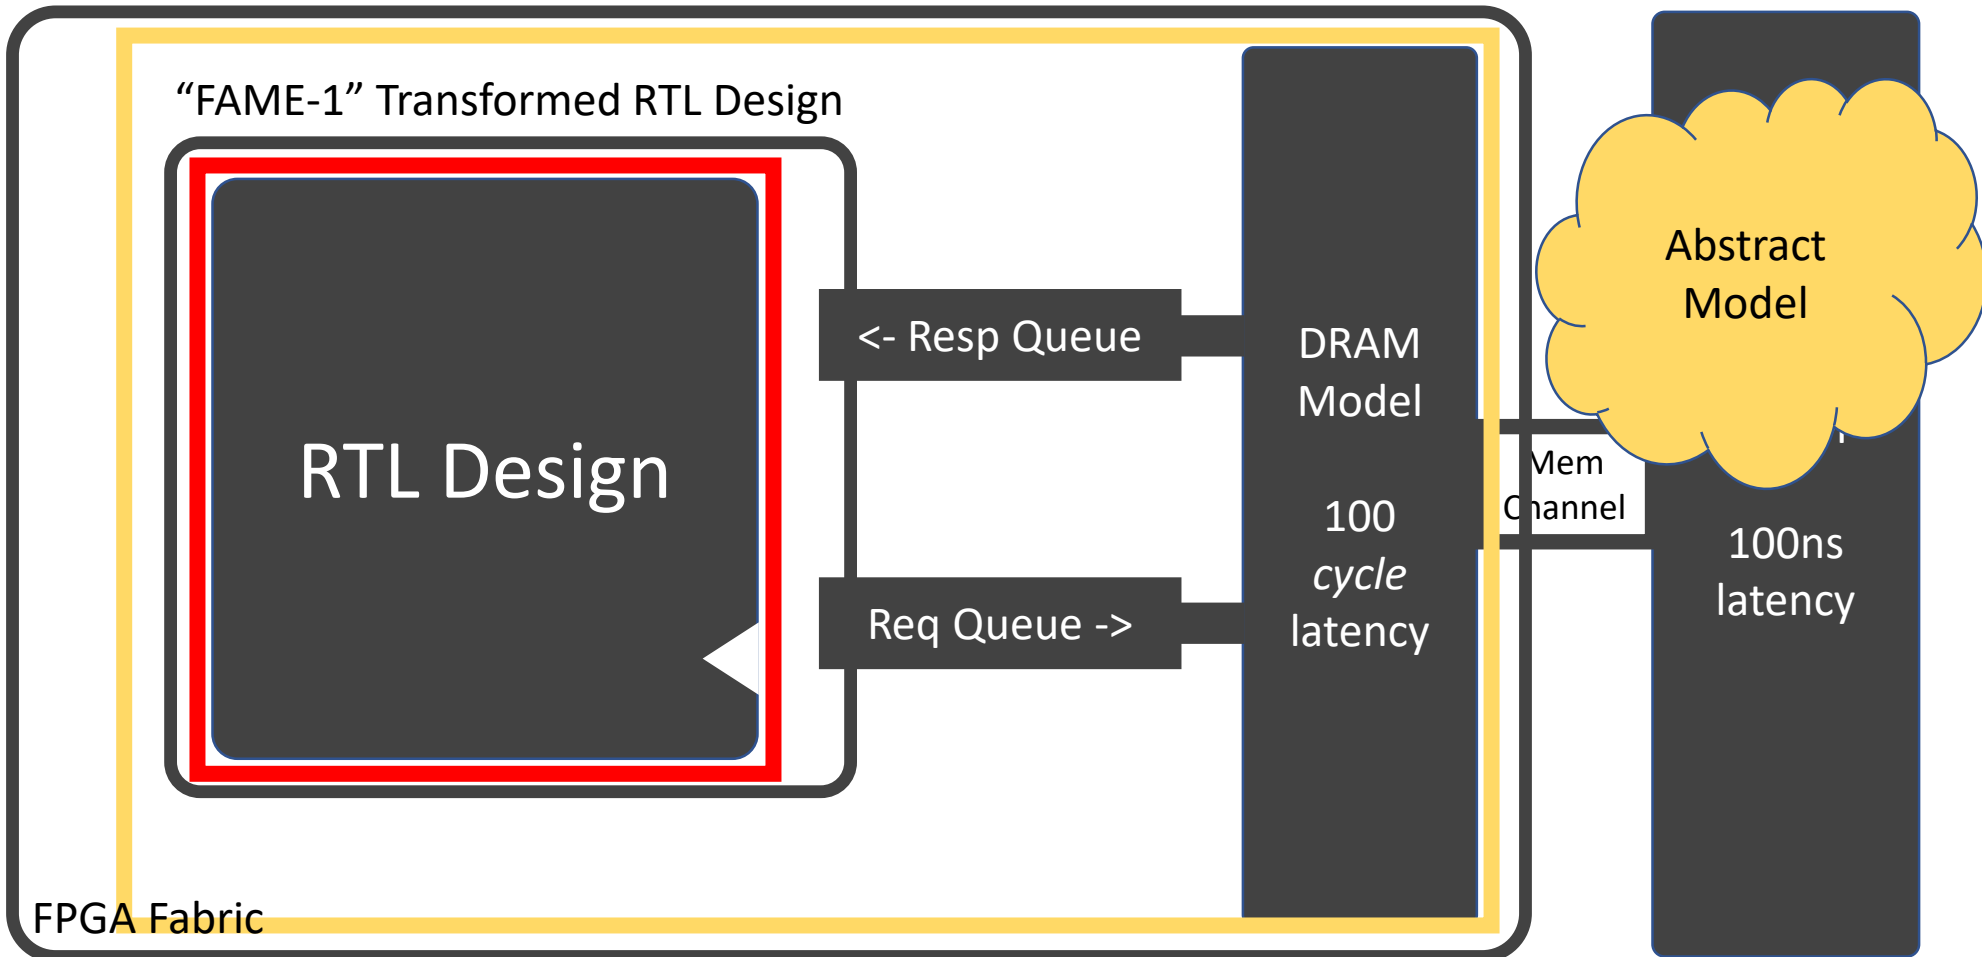




Debugging Using Software RTL Simulation

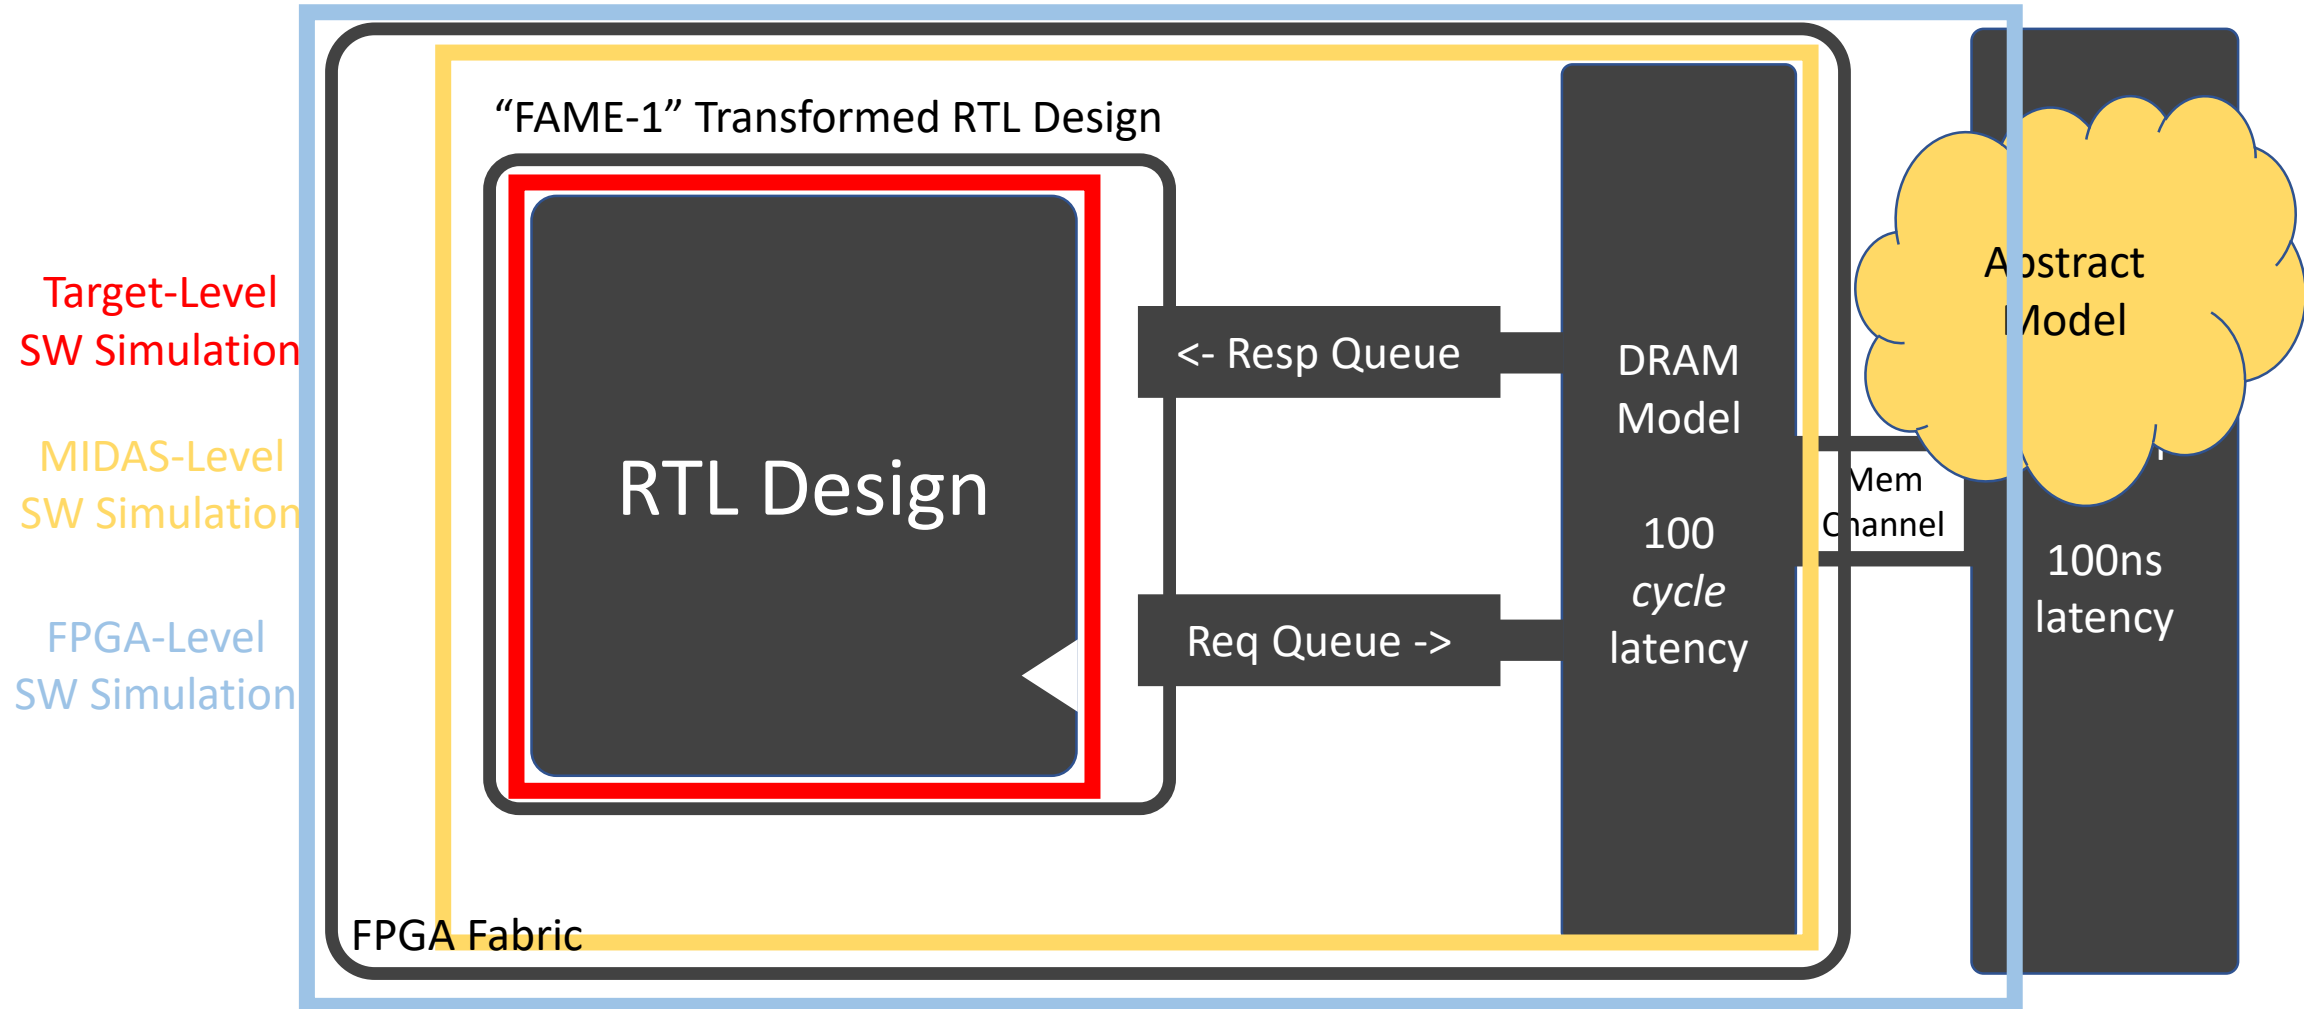
Target-Level
SW Simulation

MIDAS-Level
SW Simulation





Debugging Using Software RTL Simulation





Debugging Using Software RTL Simulation

Level	Waves	VCS	Verilator	XSIM
Target	Off	~5 kHz	~5 kHz	N/A
Target	On	~1 kHz	~5 kHz	N/A
MIDAS	Off	~4 kHz	~2 kHz	N/A
MIDAS	On	~3 kHz	~1 kHz	N/A
FPGA	On	~2 Hz	N/A	~0.5 Hz



Back to our hands-on example



Hands-on Synthesizable Printf Example

Output file in

\$FDIR/deploy/results-workload/<timestamp>-sha3-bare-rocc/sha3-bare-rocc0/synthesized-prints.out

```
CYCLE: 36086158 SHA3 finished an iteration with index 0 and message 0000000000000000
CYCLE: 36086159 SHA3 finished an iteration with index 1 and message 0000000000000000
CYCLE: 36086160 SHA3 finished an iteration with index 2 and message 0000000000000000
CYCLE: 36086161 SHA3 finished an iteration with index 3 and message 0000000000000000
CYCLE: 36086162 SHA3 finished an iteration with index 4 and message 0000000000000000
CYCLE: 36086163 SHA3 finished an iteration with index 5 and message 0000000000000000
CYCLE: 36086164 SHA3 finished an iteration with index 6 and message 0000000000000000
CYCLE: 36086165 SHA3 finished an iteration with index 7 and message 0000000000000000
CYCLE: 36086166 SHA3 finished an iteration with index 8 and message 0000000000000000
CYCLE: 36086167 SHA3 finished an iteration with index 9 and message 0000000000000000
CYCLE: 36086168 SHA3 finished an iteration with index 10 and message 0000000000000000
CYCLE: 36086169 SHA3 finished an iteration with index 11 and message 0000000000000000
CYCLE: 36086170 SHA3 finished an iteration with index 12 and message 0000000000000000
CYCLE: 36086171 SHA3 finished an iteration with index 13 and message 0000000000000000
CYCLE: 36086172 SHA3 finished an iteration with index 14 and message 0000000000000000
CYCLE: 36086173 SHA3 finished an iteration with index 15 and message 0000000000000000
CYCLE: 36086174 SHA3 finished an iteration with index 16 and message 0000000000000000
CYCLE: 36086175 SHA3 finished an iteration with index 17 and message 0000000000000000
CYCLE: 36086203 SHA3 finished an iteration with index 0 and message 0000000000000000
CYCLE: 36086204 SHA3 finished an iteration with index 1 and message 0006000000000000
CYCLE: 36086205 SHA3 finished an iteration with index 2 and message 0000000000000000
CYCLE: 36086206 SHA3 finished an iteration with index 3 and message 0000000000000000
CYCLE: 36086207 SHA3 finished an iteration with index 4 and message 0000000000000000
```



Hands-on Synthesizable Printf Example

Don't forget to terminate your runfarms (otherwise, we are going to pay for a lot of FPGA time)

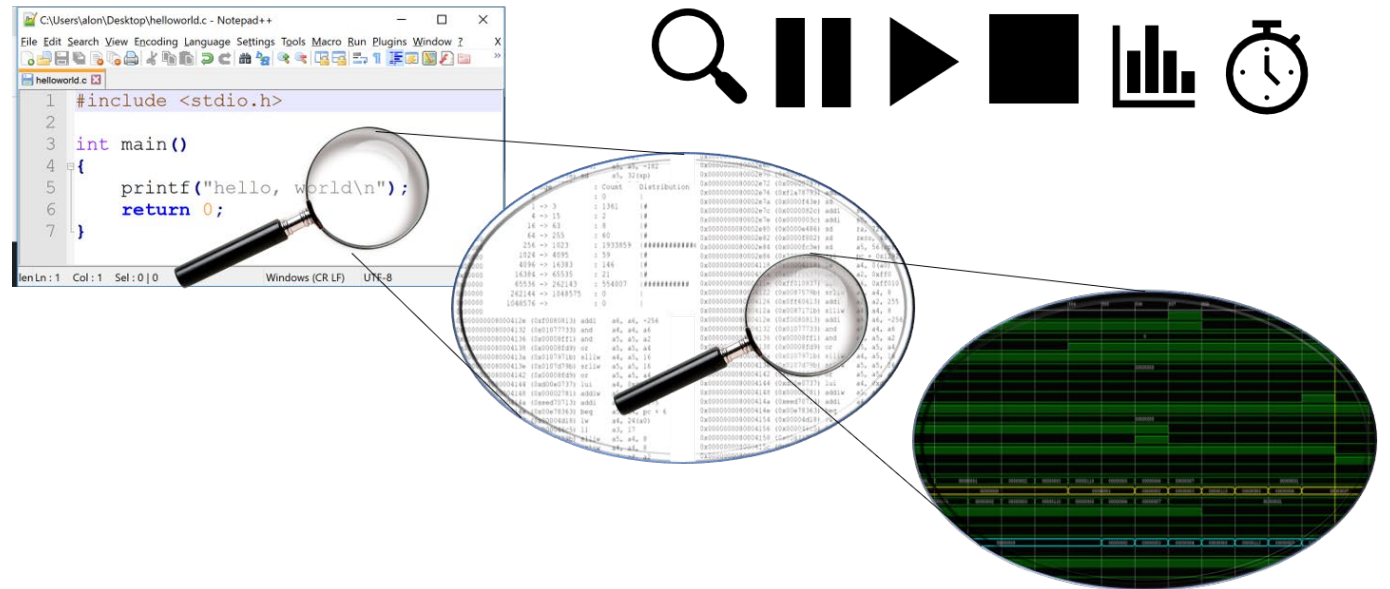
```
$ firesim terminatorunfarm
```

Type **yes** at the prompt to confirm



The FireSim Vision: Speed and Visibility

- High-performance simulation
- Full application workloads
- Tunable visibility & resolution
- Unique data-based insights





Summary

- Debugging Using Integrated Logic Analyzers ([docs](#))
- Advanced Debugging and Profiling Features
 - TracerV ([docs](#))
 - Assertion and Print Synthesis ([docs](#))
- Debugging Using Software Simulation ([docs](#))
 - Target-Level
 - MIDAS-Level
 - FPGA-Level
- FireSim Debugging and Profiling Future Vision

Check out <https://docs.fires.im/>
for more usage details