Runtime Verification David Cock – david.cock@inf.ethz.ch





ad

We're Building the Large Program Collider



Collide *instructions* at 0.99*c*, and observe the decay products.



Why is This Useful?

- Formal verification relies on accurate models
- For systems-level HW, these mostly don't exist!
 - Testing lets us build confidence at these low levels
- The hardware is *trying* to tell us what it's doing.
- Further applications:
 - Debugging rack-scale systems.
 - Monitoring control flow (security).



Properties No Double Frees in LTL

```
void *a = malloc();
...
{a is still allocated}
free(a);
```

We can now check this:



Valgrind, with zero overhead!



Our Streaming Verification Engine





Sources HSSTP Testbench: XGene 1, Zynq7000, Custom HW





Capture

Trace Capture on the Zynq7000

- 32b trace port to the FPGA fabric, 250MHz, 8Gb/s.
- Custom TPIU \rightarrow AXI core, with Linux drivers:
 - Integrates with ARM OpenCSAL, interchangeable with ETB.
 - Full-speed capture and FIFO buffering (512kB).
 - Easy to use: trace_launch <bin>; cat /dev/axi_tpiu
 - Coming soon: PCIe & HSSTP output.





Properties No Double Frees in LTL

```
void *a = malloc();
...
{a is still allocated}
free(a);
```

We can now check this:





Processing: Checking LTL with Automata

This is a well-studied problem, and standard algorithms exist:

Gp $free(x) \rightarrow P$!free(x) S x = fmalloc;



Processing

Bound Variables and Multiple Automata



• This is *resolution*.



malloc

Processing Streaming Pipeline

Chained processing stages:

- Capture
 - Using the TPIU \rightarrow AXI core.
 - Using HSSTP: On the Zynq soon, other hardware is buggy.
- Decode
 - We have a custom PTM decoder.
 - Also use Linaro OpenCSD works, but still buggy.
- Checking
 - Run the automaton.
 - Live coverage analysis *demo*.



Directions

- BRISC will support tracing.
- We're moving PTM parsing into the FPGA.
- We need to integrate and test the pipeline. The pieces are all there, they need to work together.
- Analysing live systems.

We have an ongoing survey of cache operations in Linux, seL4 & Barrelfish. We should be able to test these.

More interesting CPU platforms (see. BRISC).



Questions?



Checking LTL with Automata

This is a well-studied problem, and standard algorithms exist:

Gp $free(x) \rightarrow P$!free(x) S x = fmalloc;

Gp P, at t-1 "P was true until t-1"	P, at t "P is still true at t"	Gp P, at t "P has always been true"
0	0	0
0	1	0
1	0	0
1	1	1

