

CovVise: How We Stopped Throwing Away Interesting Coverage Data

http://www.veripool.org/papers/CovVise_SNUGBos09_pres.pdf

Wilson Snyder
Cavium Networks
wsnyder@wsnyder.org

Robert Woods-Corwin
NVIDIA
covvise@rwoodscorwin.fastmail.fm

September 21, 2009

CovVise was created at and for SiCortex, Inc.


SiCortex formally closed in June 2009.

R.I.P.



SiCortex





-
- Tradition Dictates
 - And Our Deviations
 - CovVise Language Extensions
 - SystemC Extensions
 - SystemVerilog Extensions
 - CovVise Post-Simulation
 - CovVise Database
 - CovVise Web Interface
 - Metrigator
 - Conclusions
 - Q&A

- Traditionally,
 - Only the Verification Team adds coverage
- Let the Designers also add coverage!
 - They already add “line coverage”
 - Just as they now add assertions, **when writing RTL**
 - Fifo full, empty, unlikely cross products, bypasses
- Avoid duplication
 - Some coverage is much easier when in RTL
 - And some best left to the verification team (interfaces)
- Keep RTL simple... Later slides

Tradition: Coverage is Done Late

- Traditionally,
 - Coverage is done near the end of the project
 - Quantifies that little was missed
- Learn from “Test Driven Development”
 - (Test Driven Development == Write tests before code)
 - Write the “test” of verification code, I.E. the coverage, FIRST!
 - Saves writing focused test when random hits unexpected bins
 - Focuses effort on big missing coverage items
 - So find important bugs faster
 - Reduces chance that interesting coverage cases are forgotten
 - Provides good metric for management




Tradition: All Hits Are Equal

- Traditionally,
 - Per bin, all test hits against that bin contribute to a single sum
- Which is better?
 - 100 tests that hit a bin once?
 - One test that hits a bin 100 times?
 - Both count the same after aggregation some tools!
- Require some number of hits per test to count this bin as “covered”
 - Prevents initialization-only from covering bins
 - Insures good random or strong focused coverage

$$1*100 \neq 100*1$$

“The Big Takeaway”

- Traditionally,
 - Coverage is only collected on passing tests
 - Failures shouldn't count towards coverage goals
- But learn from the Challenger disaster
 - Didn't graph failing cases, only successful ones
- If a bin is hit **only** by a failing test
 - This bin is unlikely to be impossible
 - This bin may indicate a bug is hiding behind it
 - Conversely, fixing the failing test would improve coverage
 - Focus effort on testing around this bin



-
- Tradition Dictates
 - And Our Deviations
 - CovVise Language Extensions
 - SystemC Extensions
 - SystemVerilog Extensions
 - CovVise Post-Simulation
 - CovVise Database
 - CovVise Web Interface
 - Metrigator
 - Conclusions
 - Q&A

- We extended SystemC to provide coverage ala SystemVerilog, via the [SystemPerl](#) pre-processor

```
SC_MODULE(myModule) { // A SystemC Module
    SP_COVERGROUP myGroup (
        coverpoint myCoverPoint {
            bins seven = 7; // single value
            bins three_to_five = [3:5]; // range
            bins members_of_enum = enum_type; // enum
        };
    };
    ...
    void process() { // a method of the class
        if (sampling_signal) { // when to sample
            SP_COVER_SAMPLE(myGroup); // increment
        }
    }
}
```

- We also allow crosses, illegals (asserts) and ignores

```
SC_MODULE(myModule) {  
    ...  
    SP_COVERGROUP myGroup (  
        coverpoint first {  
            bins three_to_five = [3:5];  
        };  
        coverpoint second[8] = [0:7];  
        cross myCross {  
            rows = {first};  
            cols = {second};  
        };  
        illegal_bins_func = myCross_illegal()  
        ignore_bins_func = myCross_ignore()  
    );  
};
```

- RTL is procedural
 - Coverage should be too... Alas the language doesn't allow this
- Allowed Designers to use \$ucover_* macros, for example

```
always @* begin
    if (...) begin
        $ucover_clk(clock, label)
    end
end
```


- Vpassert (part of [Verilog-Perl](#)) expands this to:

```
reg _temp;
label: cover property (@(posedge clock) _temp)
always @* begin
    _temp_sig = 0;
    if (...) begin
        _temp_sig = 1;
    end
end
```

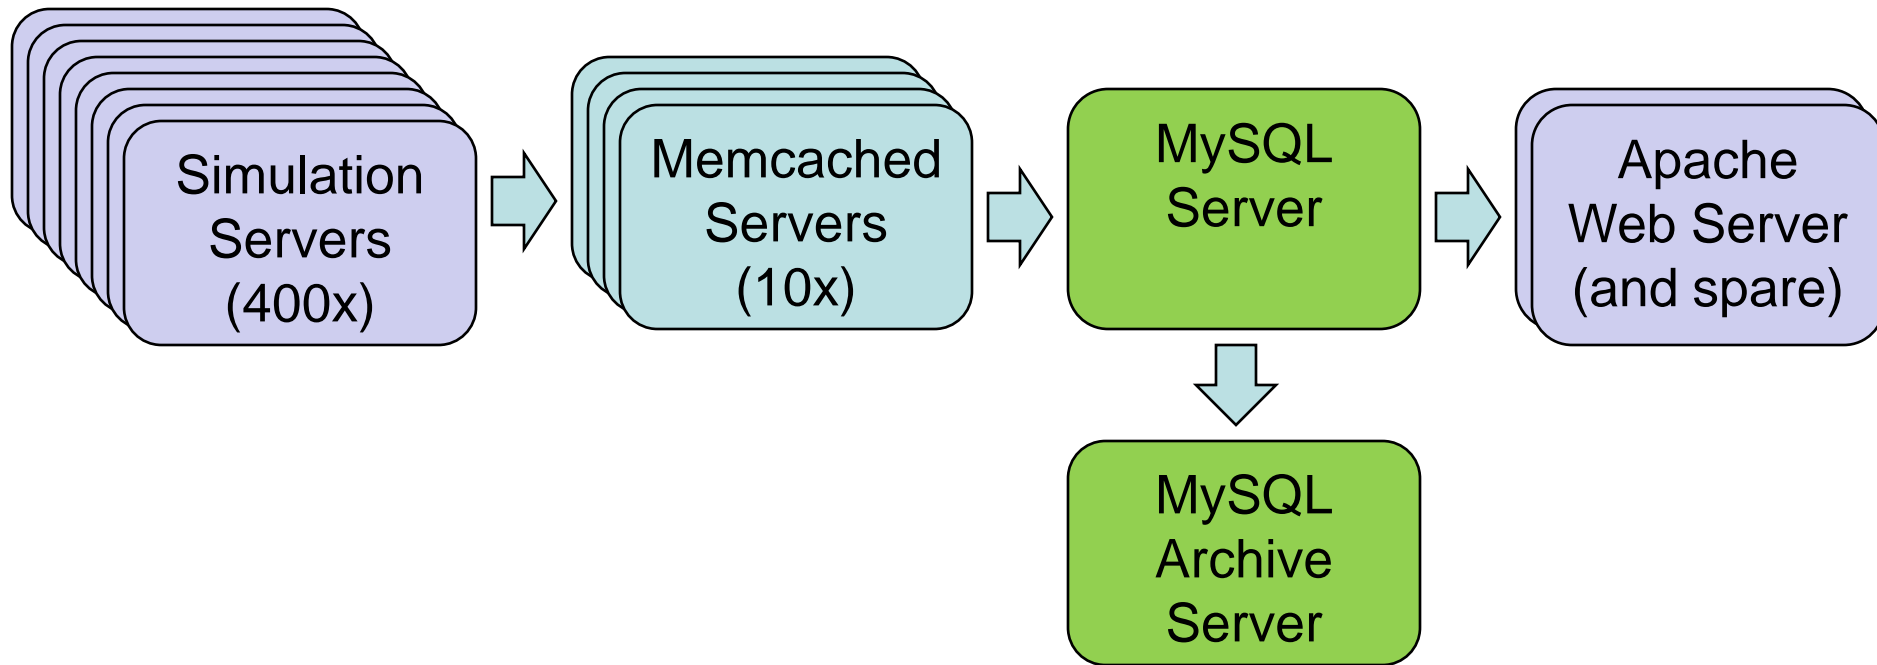
- Works with SV formal tools, too

-
- Tradition Dictates
 - And Our Deviations
 - CovVise Language Extensions
 - SystemC Extensions
 - SystemVerilog Extensions
 - **CovVise Post-Simulation**
 - CovVise Database
 - CovVise Web Interface
 - Metrigator
 - Conclusions
 - Q&A

- High demands: 100k tests * 100k bins
- 10B inserts per day




- High demands: 100k tests * 100k bins
- 10B inserts per day
- Old data auto-pushed to archival database



- Users interface to CovVise data through the web
- “Simple enough even a VP can use it”™ ☺
- Data is presented as hierarchy of coverage “pages”
- The interface begins with CovVise home page, which list “ensembles” of test runs...





CovVise

Home Ensembles [Vtestweb](#) [Manpage](#) [Help](#)

Ensembles

- Ensembles

Only show rows containing:

! – Only-failures bins
 W – Waived bins
 71.1% - Low coverage
 58.0% - High coverage

Ensemble Name	Testseries	Rev	Start Time	Coverage	Bins	Tests	Files
chip_nightly	51689	77547	05-20 00:46:48	! (71.1%) 58.0%	130,729	3,374	178
chip_random	51543	77511	05-19 14:26:26	! (74.0%) 66.6%	129,016	20,286	155
chip_random	51436	77469	05-19 11:46:08	! (73.2%) 65.5%	129,016	6,630	155
brett	51432	77469	05-19 11:39:50	! fails	109,363	1	103
brett	51384	77358	05-19 10:07:15	! fails	116,697	3	104
jason	51377	77468	05-19 10:02:31	(13.7%) 8.9%	116,850	1	104
brett	51363	77468	05-19 09:30:44	(14.3%) 9.6%	116,850	3	104
chip_nightly	51328	77459	05-19 00:23:56	! (70.9%) 57.9%	130,724	3,376	178

Other formats: [Detailed HTML](#)

- Page tree under Ensemble 6Lycug

- Cpu ⚠ (66.1%) 54.7% 107,379 bins
- + Ebox ⚠ (89.6%) 76.4% 1,052 bins
- Fbox ⚠ (40.5%) 19.3% 56,256 bins
- + bypassing ⚠ (16.2%) 1.9% 31,684 bins
- cc_bits ⚠ (93.7%) 5.5% 144 bins

Page: [sp_group/Cpu/Fbox/cc_bits](#)

Source: [CpuFPredictor.sp 433](#)

Cross of consumers of cc-bits are 2/3/4 cycles after updating the relevant cc-bits

consumer	between_instructions							
	relevant_cc_bit							
	dly3							
	cc0	cc1	cc2	cc3	cc4	cc5	cc6	cc7
MOVF_D	1	5	4	3	2	3	2	3
MOVF_PS	3	4	0	0	3	1	4	11
MOVF_S	4	4	30	5	0	6	4	7
MOVT_D	2	fails	2	4	2	3	2	1
MOVT_PS	1	3	4	3	2	1	3	3
MOVT_S	3	1	7	2	2	2	0	1

- + exponents_in_add ⚠ (76.0%) 18.5% 550 bins
- + Vbox ⚠ (61.0%) 46.5% 8,580 bins

Listing of project/ver/chip/cpu/rbox/RboxCov.sp

Line	Count	C++ Text
160		SP_COVERGROUP RboxMulBypass (
161		description = "Rbox Bypass Coverage";
162		page = "Cpu/Rbox";
163		
164		coverpoint cov_mul_mux_operand(MuxOp) {
165		bins MulA = RBoxMux::Ea;
166		bins MulB = RBoxMux::Eb;
167		illegal_bins unlucky = default;
168		};
169		coverpoint cov_mul_bypass_src(Bypass) {
170		auto_enum_bins = RBoxBy;
171		ignore_bins_func = bypass_src_ignore();
172		};
173		
174		cross MulBypass {
	(100.0%) 88.8%	18 Bins
		Show Table
		1,703 col0=M3, hier=top.board.pred.cpuRPred0.coverage, row0=MulA
		546 col0=M3, hier=top.board.pred.cpuRPred0.coverage, row0=MulB
		85 col0=M4, hier=top.board.pred.cpuRPred0.coverage, row0=MulA
		112 col0=M4, hier=top.board.pred.cpuRPred0.coverage, row0=MulB
		16 col0=MW, hier=top.board.pred.cpuRPred0.coverage, row0=MulA
		8 col0=MW, hier=top.board.pred.cpuRPred0.coverage, row0=MulB
		70,093 col0=REGFILE, hier=top.board.pred.cpuRPred0.coverage, row0=MulA
		70,389 col0=REGFILE, hier=top.board.pred.cpuRPred0.coverage, row0=MulB
175		rows = {MuxOp};
176		cols = {Bypass};
177		};
178); /*RboxMulBypass*/

➔ Coverage Details for Binrun k1YyKA

CumCover	Category	Tests	Count
0.0%	Pass High-Count	0	0
100.0%	Pass Low-Count (< 10)	3	9
0.0%	Pass Waived	0	0
0.0%	Pass Zero-Count	1,642	0
	Fail Non-Zero-Count	0	0
	Fail Zero-Count	686	0
100.0%	TOTAL	2,331	9

Similar table is also popped-up when mouse hovers over any coverage data.

+ Other Ensembles With Bin vWot1Q

+ Tests with Binrun k1YyKA

+ Coverage Details for Binrun k1YyKA

- Other Ensembles With Bin vWot1Q

Binrun Id	Ensemble Name	Start Time	Coverage	Count	Tests
1LPAFg	chip_nightly	05-15 00:30:29	100.0%	27	2,319
3SxIGw	chip_random	05-17 00:46:39	(100.0%) 0.0%	20	22,450
k1YyKA	THIS	05-20 00:00:48	(100.0%) 0.0%	9	2,331
LkWDrA	chip_random	05-14 13:29:00	(100.0%) 0.0%	9	13,061
SouRCQ	chip_nightly	05-17 00:32:27	0.0%	8	2,924

Page: 1 [2](#) [3](#) [4](#) [All](#)

→ + Tests with Binrun k1YyKA

Solves:

“I know I’ve seen this bin hit some previous night”

+ Coverage Details for Binrun k1YyKA

- Other Ensembles With Bin vWot1Q

Binrun Id	Ensemble Name	Start Time	Coverage	Count	Tests
1LPAFg	chip_nightly	05-15 00:30:29	100.0%	27	2,319
3SxIGw	chip_random	05-17 11:46:39	(100.0%) 0.0%	20	22,450
k1YyKA	THIS	05-20 00:46:			
LkW0rA	chip_random	05-14 13:29:			
SouRCQ	chip_nightly	05-17 00:32:			

Page: 1 [2](#) [3](#) [4](#) [All](#)

What tests hit this bin, and did they pass or fail?

- Tests with Binrun k1YyKA

Binrun Id	Test Name	Seed	Status	Count
k1YyKA	cpu_idefm_rtl/avp=add.d,mode=EL-M64R2-U	628065	Pass	4
k1YyKA	cpu_idefm_rtl/avp=add.d,mode=EL-M64R2-K	410295	Pass	4
k1YyKA	cpu_idefm_rtl/mg,src=fboxAddSubMul	572271	Pass	1
k1YyKA	cpu_fpipe/cpu_cpp,src=fbox_demo_s	983400	Pass	0
k1YyKA	cpu_fpipe/cpu_cpp,src=fboxa_madd_bug	449504	Pass	0
k1YyKA	cpu_subrtl/mg,src=everything,instrs=5k	317175	Pass	0
k1YyKA	cpu_idefm_rtl/cpu_cpp,src=llsc1	606318	Fail	0
k1YyKA	cpu_idefm_rtl/cpu_cpp,src=cp0_error_dcach	541852	Fail	0
k1YyKA	cpu_subrtl/mg,src=everything,instrs=5k	91087	Fail	0

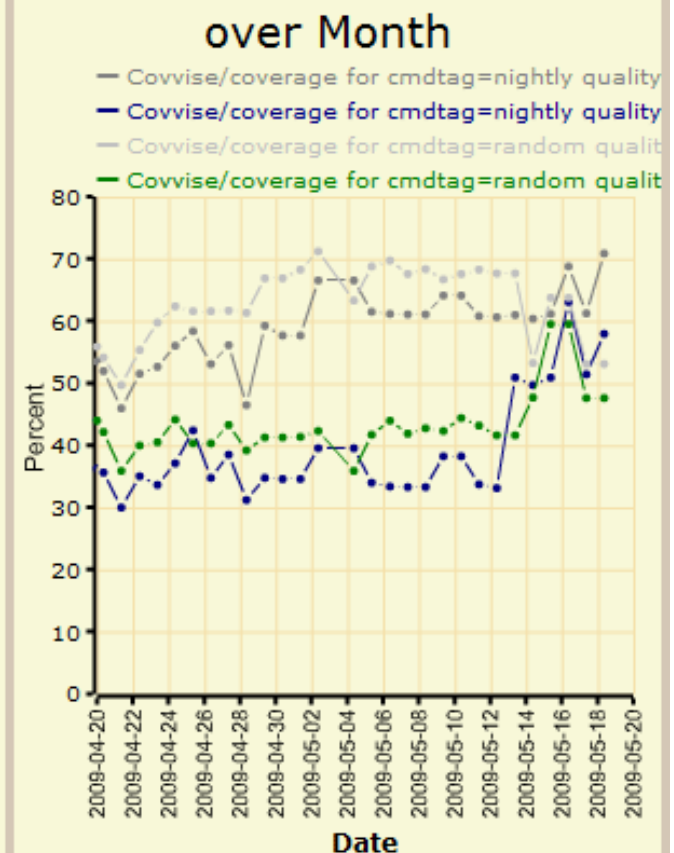
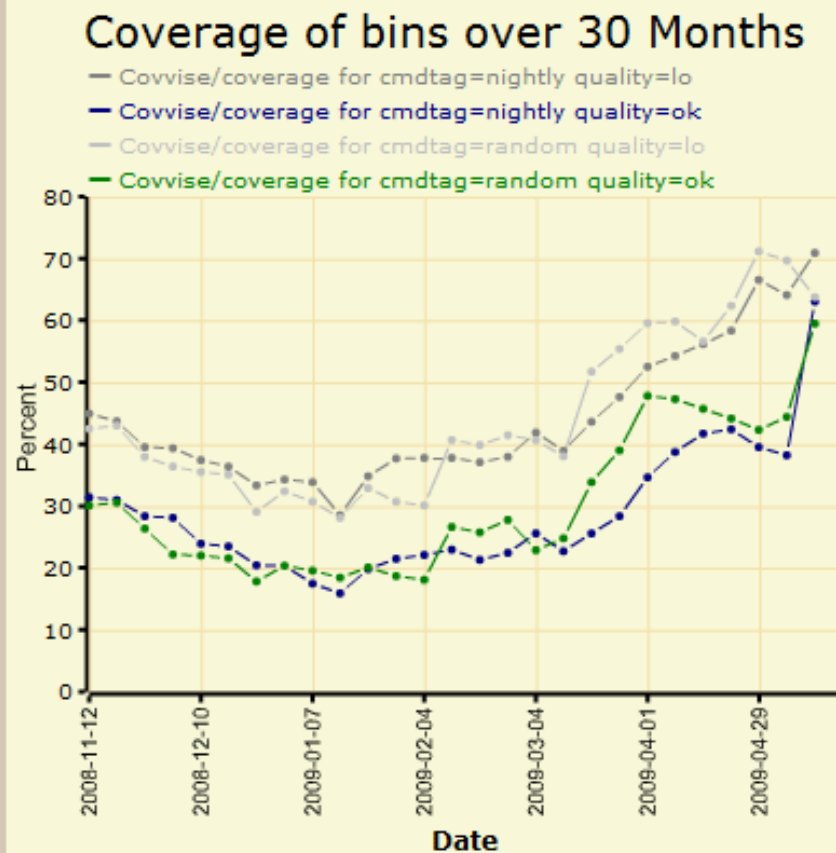
- Metrigator: Our verification metrics database
 - CovVise coverage (percent low coverage, ok coverage, number of bins)
 - Bug count (total, closed, per-component, by priority, etc)
 - Bug closure rate (total, per-component, by priority, etc)
 - Source code commits (size, number of edits)
 - Verification test success (number of tests, failures)
- Spots Correlated Trends
- Appeases Management 😊



Metrigator Coverage Graph

- Report Actions

- Graph 1: Coverage of bins over 30 Months



- Tradition Dictates
 - And Our Deviations
- CovVise Language Extensions
 - SystemC Extensions
 - SystemVerilog Extensions
- CovVise Post-Simulation
 - CovVise Database
 - CovVise Web Interface
 - Metrigator
- Conclusions
- Q&A

- Verification Engineers got
 - SystemC extended with SystemVerilog-ish coverage
 - Data collected on failing tests, to easily detect interesting bins
- Designers got
 - Coverage as part of normal RTL procedural statements
 - Easy browsing of data
- Management got
 - Early coverage for progress tracking and work reduction
 - Pretty graphs
- If we were to do it again?
 - Start coverage insertion even earlier
 - [SiCortex,] Don't run out of money ☺

- The open source design tools are available at <http://www.veripool.org>
 - These slides + paper at <http://www.veripool.org/papers/>
 - CovVise – Have you been paying attention?
 - SystemPerl – /*AUTOs*/ for SystemC
 - Verilog-Perl – Toolkit with Preprocessing, Renaming, etc
 - Verilator – Compile SystemVerilog into SystemC
- Additional Tools
 - Make::Cache - Object caching for faster compiles
 - Schedule::Load – Load Balancing (ala LSF)
 - Verilog-Mode for Emacs – /*AUTO...*/ Expansion
 - Vregs – Extract register and class declarations from documentation

