(DAC 03 Item 30) -- [01/20/04]

Subject: Sequence Physical Studio, Golden Gate, Synplicity Iota, ChipVision

RIGHT PLACE, RIGHT TIME (Part I): If you were an EDA company that focused on signal integrity and power issues, this was your DAC. It's a going to 130 nm and lower thing. Since Sequence has four tools in that space, lots of people had lots to say about them! Talking 14 pages here. (Be sure to check out the Apache section of this trip report, too!)

> I am very skeptical where the entire signoff is based on a single company's check -- sort of like the fox watching the hen house. Because of this, my group used a flow of Silicon Ensemble/Magma followed by by CadMos. CadMos was painful to integrate as they did not have a built in timing engine. Closing the loop took several days. My group has not tried Nano Encounter where CadMos has been tied with PKS and NanoRoute.
>
> Most recently I used Magma and BlastNoise, and followed that up with Sequence's Physical Studio. We did discover problems that escaped Magma and were latter reported by Studio. The problem was conveying the information back to Magma. Magma was not willing to cooperate with Sequence -- they did not want to give Sequence the MTCL format to do the necessary repair. We kind of came up with a solution of our own. I pushed Sequence very hard to go for a single pass solution, where the repair is done in their database itself. They are working with a startup to develop this ECO capability, and I have not seen this result yet.
>
> Physical Studio was fairly easy to come up to speed with, and Sequence prepared the library for us, as it was a generic Artisan 0.13G library. I also used their Columbus extraction engine for signoff. The static timing engine ShowTime, to which it was tied, had some issues but they had it resolved quickly. (Source synchronous timing.)
>
> My final verdict is this: If Sequence gets the ECO capability to work, I would highly recommend Physical Studio for 0.13 and 0.09. The accuracy is comparable to that of Celtic, but easier to use. They are the only ones who can extract inductance, and can also do GDS-based extraction.
>
> - Ravi Selvaraj of Sinett Corp

> We use Sequence's Physical Studio for glitch fixing and analysis. We use Cadence's Simplex and have tried AstroRail. We have not done instance derated IR drop effect on timing, just IR drop analysis.

- Jim Vanaria of TranSwitch Corp.

We have build a solid flow based on Physical Studio and Apollo. Now that we are doing with the first set of chips, we are looking to enhance the flow for the next generation 0.13 um designs. Physical Studio has significantly enhanced supports for hierarchical design flow. Their PNM model preserves all the parasitic for SI analysis at the chip level. It seems that they are ahead of Synopsys's ILM model which does not preserve all the SI characteristics for chip level analysis.

- [An Anon Engineer]

I do use Sequence's Physical Studio a lot, but not CoolTime. Physical Studio would best be compared with PrimeTime-SI and CeltIC. We did provide test cases to CoolTime but only saw the presentation and analysis results. But for IR drop induced timing violations, CoolTime does share the same optimization platform as Studio (to fix xcoupling-induced timing and glitch violations).

Strengths:

1. To us, Physical Studio is a point tool to help us automate the fixing of SI-related violations. It generates ECO change files that we can feed back into Astro to implement the fixes. Its fixing capability was something not available in any other tools as off the end of last year. But since then, CeltIC and Magma has been able to catch up on this. PTSI-generated constraints directives to Astro are still not a good way for fixing, since the changes may break other things. Also, the timing engine between PrimeTime-SI and Astro are not very well-correlated.

2. Glitch analysis is also a strength, and Studio provides ELMO, an HSPICE wrapper tool for characterizing libraries for glitch analysis. Synopsys, on the other hand, does not provide glitch libraries, but instead waits for cell vendors to provide this. However, CeltIC seems strongest in this area because it uses a pseudo HSPICE engine and also pre-characterized glitch libraries, and is capable of glitch propagation analysis. Studio will only support this in the August release.

3. Analysis speed is another strength, but other tools are catching up. However, we did realize that some of this speed may be due to some shortcuts (which they claim to not to affect correlation to HSPICE) such as the "variation" mode xcoupling timing analysis. But again, Sequence does not claim strength as a crosstalk timing signoff tool, although some of their customers do use Studio (Showtime engine) as signoff.

Weaknesses:

1. I think Studio is weak in general robustness, because bugs and new issues are uncovered regularly. Part of these are due to the fact that Studio interfaces with many pieces of data from other tools. For example, Studio is always trying to catch up on SDC support. We end up having to switch releases many times over.

2. The documentation is also not updated regularly and the fact that the tool commands are still undergoing many changes does not help.

3. Studio is able to fix max_cap and max_trans violations, but their interpretation of max_cap and max_trans are a bit unorthodox and therefore may not track well with PrimeTime-SI and Astro analysis.

4. Timing correlation with backend and other tools. This is an issue because the fixes that Studio help inplement may not cover all violations seen in PrimeTime-SI and Astro, for example. For tools like Magma, for example, where the toolsuite uses the same delay calculator, the flow is more integrated and consistent.

- Paul Pua of TranSwitch Corp.

We had used Sequence's Physical Studio tool on a project that was taped out in September 2002. Our task in this project was limited to the physical implementation of a single digital block which was instantiated four times in the full chip. The block had about 170 K placed instances with 16 memory instances. The memory instances occupied about 40% of the block area. The design was implemented using Artisan TSMC 0.18 um FSG process standard cells and memory cells. The clock frequency was 150 MHz.

Following the flow recommended by Sequence, we did congestion driven placement and clock tree synthesis using Silicon Ensemble. At this stage we used Physical Studio to post-clock optimization and setup optimization. This was followed by detailed routing using WRoute, parasitic extraction using Columbus and then the post route optimization stage in Studio where we fixed setup, hold and cross talk violators. Studio generated an ECO repair file which was imported back into the SE environment and the ECO placement and final routing was redone.

The Studio tool performed quite well in the optimization stage where we fixed setup and hold violators. Cross talk analysis and the impact on the path delays was also quite good. The biggest headache that we faced with Studio was the fact that it did not fix max transition violators. We had to use QPOPT in the SE environment to fix the max trans violators after Studio optimization. It seems that the latest version of Studio has resolved this problem.

The other major issue we faced was moving data back and forth between the Cadence SE environment and the Sequence Studio environment. Due to some problem with the tools, we had to read in the post-route optimized ECO DEF from Studio into Silicon Ensemble and perform a full route. This meant that Wroute would build the database from Groute stage and the "Build Design" phase would run for several hours (on a 4 CPU, 4 GB, E420R SUN).

On the whole we were able to make the design work at the required frequency and we had silicon out that worked first time. But the turnaround time and the need to use QPOPT to fix max trans were the major concerns. We hope that with the max trans fix and perhaps an internal ECO router, the new version of Physical Studio will perform much better.

- Derric Lewis of Qualcore Logic

Sequence's Physical Studio commands are easy to use. Sequence's estimation parameters for noise avoidance step are easy to compute. You need to provide cap_per_len, cpl_per_len adn res_per_len. We update these numbers for each project. We compare point to point numbers with actual PrimeTime and SPICE numbers for longest paths after a route is done. It varies depending on the tool used in your flow. For example, if you introduce PhysOpt then the cap number would significantly differ from without PhysOpt being used.

1. Post-route optimization algorithms are good for setup fixes. They are several commands that one can try and fix all setup violations.
2. Sequence's GUI is a delight to look at. Two years ago, it was not there.
3. Initially there were many hand-shaking issues with Avanti database, but now it is a solid gold. Sequence's engineers were taking problems seriously and resolved quickly for us.
4. Post-route SPEF based noise numbers match very well with PrimeTime. PrimeTime numbers are pessimistic. Our noise sign-off tool was Physical Studio. We bought peace-of-mind using this tool for SI issues after tapeout. All our four chips worked on the first silicon passing several thousands of test on the FIRST day.
5. We had problems with hold fixes done in especially for blocks with many macros. Either it was over-doing it and thus violating setup or not solving all hold problems. At the end, we manually need to add or remove hold buffers.
6. Physical Studio Pro was not working initially to our expectations. Now, in the current release 2003.1.22.5, it seems to be better. MinMax analysis is working fine now.

We would like to evaluate their on-chip variation commands. It seems like the PhysOpt+Astro+AstroXTALK flow seems to achieve the same results we obtained with a Physical Studio/Apollo flow.

- Subramanian Sganesan of Silicon Access

Physical Studio? We used it in an ASIC flow for CAM products design. After floorplanning in Encounter, PhysOpt is used for initial optimization, followed by clock tree synthesis and detailed route. Physical Studio is used for post-route timing closure with SI. Overall, I found Studio works well with post-route optimization, despite some minor issues we had to hack through.

Physical Studio is capable of doing pre-route optimization and noise avoidance, in addition to the post-route feature we use. The pre-route avoidance worked well for ASIC designs I did with previous companies, but underperformed in comparison with PhysOpt for our memory products, in which we deal with large chips with large memory arrays surrounded with long and narrow channels for control logic standard cell placement. Most tools have hard time dealing with such channel-based floorplans, but PhysOpt seems to do a better job in comparison.

Since our timing sign-off is based on Star-XT and PrimeTime, we did extensive correlation study against Studio's extraction and STA features (Columbus and ShowTime). The results are fairly good, which allows us to close timing in Studio and directly proceed to timing sign-off. Sometimes, we do find some discrepancies between Studio and

PT, which requires us to either adjust sdc to remove PT pessimism or add extra margin in Studio to achieve PT closure. We didn't use PrimeTime-SI. Instead, we rely on Studio for noise fixes and SI sign-off.

Things I like about Studio:

a. Tcl interface and high-level commands that make batch job easy. We can do optimization with and without SI at SS and FF corners for min and max SPEF in one shot.
b. A comprehensive set of commands that allow user to interactively query, change, and optimize the design. Sometimes, I use it just for manual ECO of the netlist (add/remove/upsize cells).
c. SI analysis and fixes are concurrent.
d. Post-route optimization converge well.
e. Correlation with XT/PrimeTime is satisfactory (in non-SI mode).
f. It directly reads .lib, STAMP models, and SDC with commonly used constraint syntax.
g. STA is fast with large designs.
h. HTML-based reporting is very useful and friendly.
i. AE support is excellent.

Issues:

a. Compatibility problem with the latest LEF/DEF formats (5.4, 5.5). Besides some syntax issues that needed to be fixed, we had to convert the DEF to the early formats to make things work.
b. The optimization sequence seems to be design dependent and requires some trial-and-error experiments to achieve the best results. It's definitely not a push-button tool, nor one should expect so.
c. As I said above, pre-route optimization doesn't work so well for channel based floor plans.
d. I tend to get a lot SI violations, most of which got fixed automatically, but am not sure if how much pessimism is involved.

    - Xi-Wei Lin of Micron Technology

Sequence had a paid 3 hour hands-on tutorial for CoolTime, which did not meet my expectations. It was closer to an informercial than to the paid tutorial where I expected to get deeper insight to the technology.

I can compare CoolTime against VoltageStorm-TL (transistor level) in power-grid analysis and against CeltIC in signal integrity. I cannot comment about the performance and accuracy of Sequence tools because this was not evaluated.

The major advantage of VoltageStorm-TL is the accuracy of transistor-level and the vector dependency of static analysis. However as a static tool it relies on the ability of the user to estimate the various current signatures. Its dynamic option is not good. The ability of running dynamic analysis at Sequence in gate level is very promising because the simulation is fast, and quantity here becomes quality. The support of capacitances at the supply network is important, too.

In signal integrity, Cadence CeltIC is the perfect tool -- runs fast, produces accurate results, very stable. All that others can do is to try to compare with it.

- [An Anon Engineer]

Sequence Design has a new tool, CoolTime, which does dynamic voltage drop analysis. The inputs are LEF, DEF, and either vectors or the usual vectorless information (toggle rates and dominant logic states of the ports, which are propagated inwards). You can also set toggle rates on internal registers. The output is worst case IR drop and SDF files for STA. It currently has no "what if" analysis capability. They also have the old Sente tools which analyze power more at the RTL level. Their overall low power flow, utilizing PowerTheater, CoolTime and Physical Studio is called NanoCool, and features low power clock tree insertion, dual threshold voltage design and cell resizing for power minimization.

- John Weiland of Intrinsix

I saw demos of Sequence CoolTime and Apache RedHawk.

This year every company was offering some type of power/rail analysis tool. Apache Redhawk and Sequence CoolTime are both very competitive in the field. The ability to do "vectorless" simulations, provide recommendations for fixes, and fast run times are what impress me the most. Traditional flow for power analysis requires VCDs that can be ten, twenty, or thirty GB in size. It consumes tremendous amount of simulation time, and disk space (even though disks are cheap, but a handful of VCD per layout, plus the layout database you run out of space quickly).

The ability to perform vectorless simulation may be powerful, but can also be its drawback. There needs to be correlation done on designs between switching activity annotated power analysis, and using vectorless power analysis. The instant, and dynamic IR drop abilities are going to be very useful, especially in the larger designs. It would be even better if tools can not only recommend fixes to power rail design, but also recommend placement of storage elements like capacitors in unused silicon areas to help alleviate temporary power surges.

- [An Anon Engineer]

I think the strengths of CoolTime and Redhawk are that we don't try to prepare the representative vectors for power estimation. They genetate vectors automatically to try to guarantee the peak power consumption. That's really fascinating me. Actually my engineers are using AstroRail and VoltageStorm and they are usually suffering from vector preparation for inducing peak power consumption. They are not sure those vectors are adequate for peak IR drop analysis.

I feel that CoolTime is superior to Redhawk in speed, while Redhawk is superior to CoolTime in accuracy. So, for a very large SOC design (say, more than 10M gates design?) I would like to prefer CoolTime rather than Redhawk!

- Dongsoo Cho of Samsung

My opinion of CoolTime is based just on DAC impression, not tool testing, so it depends a lot on how well things were presented. Intel uses mostly internal tools for timing and reliability.

Among the other presentations I saw, CoolTime's was very professional and the concurrent electrical analysis is the right concept. Especially the instantaneous current and voltage drop analysis are very valuable, if it works as well as presented. A second important feature is glitch analysis, if indeed all pruning techniques work well. From the detail level of the presentation it was difficult to guess if timing windows are re-calculated based on glitch analysis or how timing and noise converge, this is a concern that can be validated only in real-life usage.

At the abstraction level that DAC offers I think CoolTime and Apache RedHawk are better than others.

- [An Anon Engineer]

On paper, CoolTime looks better than any other rail analysis products out there today. The things we like about CoolTime is:

Strengths:

1. Fast timing and SI analysis engine/algorithm inherit from ShowTime
2. Simple AC current model to enable fast dynamic power/rail analysis
3. Close-loop timing, noise and IR-drop analysis within one tool
4. Dynamic IR-drop back-annotation per instance changes along time!! (vs. one VDD constant per instance in PrimeTime-SI, for example.)
5. Elegant vector-less toggle analysis based on Monte Carlo method
6. Even consider the small decoupling caps on rail such as N-well caps
7. Interconnect caps through open-channel PMOS, etc.

Weaknesses:

1. The No.1 problem of the tool is that their design team build the tool based on a huge assumption: all standard cells have functional descriptions in the Synopsys .lib file. This may be true for 3rd party vendor library. But definitely not true for proprietary libraries with complex cells.
2. Sequence is still a small company, and their technical support team seems under-staffed ... or just too many customers are jumping on CoolTime all a sudden.
3. Below average document quality.
4. Can't directly take PrimeTime constraints/commands (TCL). It's a legal issue. But maybe SOMEONE can write a translator to automate the translation process.

No automated flow to incorporate power meshes within custom blocks (GDS) such as memories to the top level.

- Wilson Chan of QualComm

CoolTime was pretty cool!!

- [An Anon Engineer]

About 2 years ago we needed the ability to extract resistive, capacitive and inductive parasitics. At that time, the only tool the foundry PDK supported for RCL extraction was Sequence's ColumbusRF. So that's the tool we began using, and continue to use.

In general, ColumbusRF/AMS meets our needs. There are some minor problems with it but my experience is that just about any tool is going to have some issues. The one significant limitation of ColumbusRF/AMS is that it has difficult extracting large power planes. However, most extractor of this nature are likely to suffer from this problem; field solvers are required to address this issue. This is a place where AssuraRCX might be superior because it has, or will have, a built-in field solver that can be used on selected nets.

When our subscription for Columbus ran out, I recommended a renewal because I felt it was our best option for RCL extraction with the PDKs we have. I've been watching AssuraRCX for the past two years. The growth of the tool and the support from our foundry have reach the point where an evaluation makes sense, probably in the next couple of months.

- Scott Witherspoon of TelASIC Communications

We use Sequence's Columbus for extraction (gate + transistor) and ShowTime for noise analysis. Our analysis was done mostly at the block level for noise signoff. The extracted format was SPEF for easy integration with PrimeTime and Magma. We had some custom designed blocks that went through transistor extraction and SPICE. Our design was over 100M transistors with several high speed interfaces running at a dual data rate of >800 Mbps

Strengths:

1. Gate level extraction is quite fast, blocks with ~350K+ gates would extract in a few hrs.
2. Sequence was also very supportive in providing the process and library data for the various process flavors we needed.
3. Columbus Gold was also easy to work with and integrate into our SPICE flows.
4. Noise analysis is very detailed and the data output very easy to parse. The tool has several features for limiting pessimism in the noise calculations which allows you to focus on real violations.
5. Close correlation on timing front with PrimeTime.
6. Support is excellent. R&D turnaround for a couple of issues we flagged was very quick. The timing and noise analysis is very fast. We were surprised by the speed (1/3) or less compared to PrimeTime.

Weaknesses:

1. The user interface is through HTML which can be both a blessing and a major pain as parsing timing info for other tools becomes tedious.
2. Documentation is barely adequate.
3. Interactive mode has a limited command set, with very limited help available in the shell itself.
4. This severely restricts customization of flows to provide timing analysis as per design requirements.
5. White box extraction (preserving details at the block interface as

in PrimeTime ILM's) wasn't quite ready when we needed (which was around Nov 02), consequently we went with STAMP models (blackbox extraction) which is supported.

Overall Sequence's extraction and noise sign-off tools worked very well with the Magma tools and PrimeTime.

- [An Anon Engineer]

After extensive evaluation we decided on the Power Compiler, so far we are happy with the results.

- Nicco Bhabu of Chip Express

We have a Power Compiler license. It is getting integrated into our flow but is still being tested out. We have used Powermill for most of our power calculations and seems to work well for us. We will switch to the new Synopsys Nanosim now I guess."

- [An Anon Engineer]

In principle, if you do not know how to calculate up-front the power mesh to fit your design needs, then these tools can be used. I managed to fit the whole set of equations that are required to perform this calculation into a simple Excel file. We are using Simplex to validate our calculations as a validation to prove that the design considerations were correct. So far we did not see any difference larger than 3 mV at 0.13 um.

- Yuval Itkin of Metalink Ltd.

I cannot comment on particular tools, just that there is still the feeling that power tools are inadequate.

- Luke Simonson of UCLA

Power designing tools like Synplicity Iota need to be confronted with chips after fabrication.

- Ahmed Jerraya

I think power issues was one of the big themes at DAC this year. I think Orinocco looks interesting since the biggest impact on power will be at the architectural level. Other interesting power related products I found were Virage's power optimized libraries and Golden Gate's power optimized synthesis. I think they will be useful to get a little bit extra power savings but for applications such as ours at Elliptic we can't rely on the tools to give us the power savings we need. The key is still to select the right architecture.

- Anders Nordstrom of Elliptic Semiconductor

ChipVision sells Orinoco, a tool for RTL and algorithmic power analysis. Like some of the design planning tools, it first creates a library of big RTL building blocks with power characterized from a Synopsys .lib file, then reads the activity data from C or SystemC simulation. They help minimize power via algorithmic optimization and also by placing tightly coupled blocks close together.

ASC was about to release a tool for behavioral level power optimization, based on work done at Princeton. They claim up to 10X power reduction.

Golden Gate Technology now has a tool that optimizes cell placement for power consumption, which they claim provides 20%-25% power reduction without any speed penalty (I'm just repeating what they told me). They are working on a signal router. They also have a tool for power grid design.

Library Technologies sells a power simulator that you hook into your Verilog via PLI.

- John Weiland of Intrinsix

Golden Gate's demo made me curious. Their claim of 30% reduction in power sounds too good to be true, however, we may give it a try.

- Eli Assoolin of Transchip Israel Research Center

Golden Gate PowerPlace - Didn't look. Must have missed them.

- [An Anon Engineer]

We don't use power designing tools. Now we do gated clock. It works great.

- Ji Li of Via Tech

I spent a lot of time at both Synplicity and Synopsys, and neither one mentioned much about power designing tools.

- [An Anon Engineer]

Sequence's PowerTheater seems to be the leader.

- [An Anon Engineer]

We use PowerTheater. It's surprisingly accurate. We couldn't believe it ourselves. Given RTL, VCD, power library, clock info, etc. It comes up with power consumption numbers that are within 25% of actual power consumption. We did not benchmark it against other power tools.

- [An Anon Engineer]

I have been using Sequence's PowerTheater for many large and small gate-level designs. I use it for general power estimation and for instance-based power consumption with ALF library as input to VoltageStorm for IR-drop analysis.

It is an easy tool to use and well documented. Many of its power calculation results have also been close to silicon.

The main problem I see with this tool is its capacity. It runs out of process memory when the netlist has more than 2M placeable instances. There is also issue with using very large simulation data but there are workarounds for that at the expense of longer runtimes.

- Bijan Panahi of NEC Electronics America

We've been using the PowerTheater/WattWatcher tool for over three years now with success. We use it for RTL, Gate, and post-P&R (full-chip) power analysis. At RTL we use it mostly for architectural trade-offs. At gate-level we use it to give us accurate power numbers. We back-annotate parasitic information to get even more accuracy. From our last chip the silicon power was within about 10% of what PowerTheater predicted, not bad! This is pretty good considering our transistor-level simulators will get us to within 5% but at a cost of about 100-1000x of simulation time.

The strengths of PowerTheater/WattWatcher is its easy to read power reporting. It generates html power reports that divide the power into different categories. The speed is similar to a Verilog simulation. This makes gate-level simulation possible all the way up to chip level. The tool is versatile with netlist support. The tool also uses common library formats such as Liberty and ALF. The stimulus is from a Verilog simulation output and can be a standard VCD file or a PLI-linked PowerTheater specific file.

I haven't used the GUI much. We use PowerTheater/WattWatcher mostly in batch mode. The GUI can be used to help navigate through the setup. It's also helpful to cross-probe between power numbers and the RTL code.

These are the weakness of PowerTheater/WattWatcher as I see them. The tool is highly dependent on how well the library is characterized for power. ALF modeling is the best but not well-supported by some characterization tools. Some ASIC libraries are poorly characterized for power, if at all, and, as a result, the accuracy of tool can diminish. We use custom libraries so we have very good correlation between simulation and actual silicon. The accuracy of the RTL number is highly dependent on the RTL coding style. I wouldn't put too much emphasis of the accuracy from RTL. Another weakness is that power modeling of black-box elements such as memories and analog circuits is very simplistic.

Power Compiler touts many of the same power analysis capabilities as PowerTheater. The flow for estimation is similar; the stimulus is an activity file from simulation. We mostly use Power Compiler for power optimized synthesis. The power reports from Power Compiler are not as easy to read. The html reports PowerTheater produces are really nice to report power to the rest of the team, managers, etc. Power Compiler produces text reports similar to timing reports. As far a I know there isn't a GUI for Power Compiler it all runs within the DC shell environment. Although to do only power estimation there is a separate shell (pe_shell) that doesn't tie up a DC license. I've never checked the analysis accuracy of Power Compiler.

The way I see it is that Power Compiler and PowerTheater complement each other in our flow. Similar to DC and VCS. One is for synthesis the other is dedicated to simulation/analysis.

- [An Anon Engineer]

Tensilica provides configurable microprocessor cores as soft-IP to its customers. We perform power analysis on our designs to ensure that they meet our low power design goals and to generate characterization data for different configurations of the processor. We have traditionally used Synopsys Power Compiler on a gate-level netlist for doing this analysis.

For over a year now, we have been using Sequence PowerTheater (and its predecessor, the Sente WattWatcher tool). The primary motivation for this was to do power analysis at the RTL design stage of the project, where there is much more flexibility in changing things, then wait until we had a synthesized netlist. An RTL test bench is up and running on most projects a long time before getting synthesized netlists that successfully run gate simulations. The PowerTheater tool was easy to integrate in our testbench environment, so we could start running power analysis simulations soon after getting our verification testbench up and running.

Pros:

One good feature of the PowerTheater tool is its report generation capability. It can generate hierarchical power dissipation reports, in which you can view the power being dissipated at any level of your RTL hierarchy from "full chip" to individual flops. This feature is very useful to identify the modules and sub-modules that exceed their power budget. Such an analysis is much more difficult on a flat gate level netlist. The HTML formatted reports are easy to navigate and publish on our Intranet for various designers to look at.

We have successfully used the PowerTheater tool to identify power saving opportunities in our design. As a result, we have set up a weekly RTL power regression run that is run every weekend. This helps us continuously monitor power dissipation on our design throughout the design flow, and identify and fix issues quickly.

Cons:

Until a recent release, the tool only ran on Solaris and not on Linux. Since we have increasingly been using Linux for our development, this was a problem. PowerTheater is now supported on the Linux platform, but our experience has been that it is not yet as robust running on Linux as it is on Solaris. For that reason, we have continued to use it on Solaris.

The PowerTheater tool is not very good at identifying logic that will be optimized by synthesis tools. Thus if the RTL design has non-trivial amounts of logic which will be removed by synthesis tools, PowerTheater results will be overly pessimistic.

```
The "activity" files generated by the PowerTheater tool are huge.  It
would be nice if the dump/compression algorithm was improved to create
smaller files.

    - Himanshu Sanghavi of Tensilica
```

Got a better one in mind?

Top Home

"This here ain't no one's opinion 'cept my own."
This Web Site Is Modified Every 2 to 3 Days

| Contact John Cooley | Webmaster | Legal | Feedback Form |