CCSP Characterization Data and the Exponential growth of Silicon IP Data Volumes

A Fractal whitepaper

Introduction

The amount of data that is needed to support the timing/power/signal-integrity validation of System-on-Chip designs has been increasing with every new technology node introduction. This causes significant headaches for CAD groups putting together state-of-the-art design flows that use the latest, best-in-class, performance verification tools to support tape-outs in new technology nodes. In particular qualification of the characterization data of silicon IP (libraries and hard macros) – where the design-data explosion is most pronounced – requires dedicated tools, compute resources and expertise. Where a few home-brewed scripts could do the job a few technology nodes ago, characterization data qualification now requires dedicated tools, such as Fractal Crossfire. This whitepaper illustrates this trend by using the Liberty CCSP format as a running example.

The exponential data-volume growth

As IC manufacturing technology has progressed by following Moore's law, the amount of data necessary to characterize a piece of silicon IP for the different performance domains (timing, power, noise and reliability) has grown with it. This trend is illustrated in the diagram below:

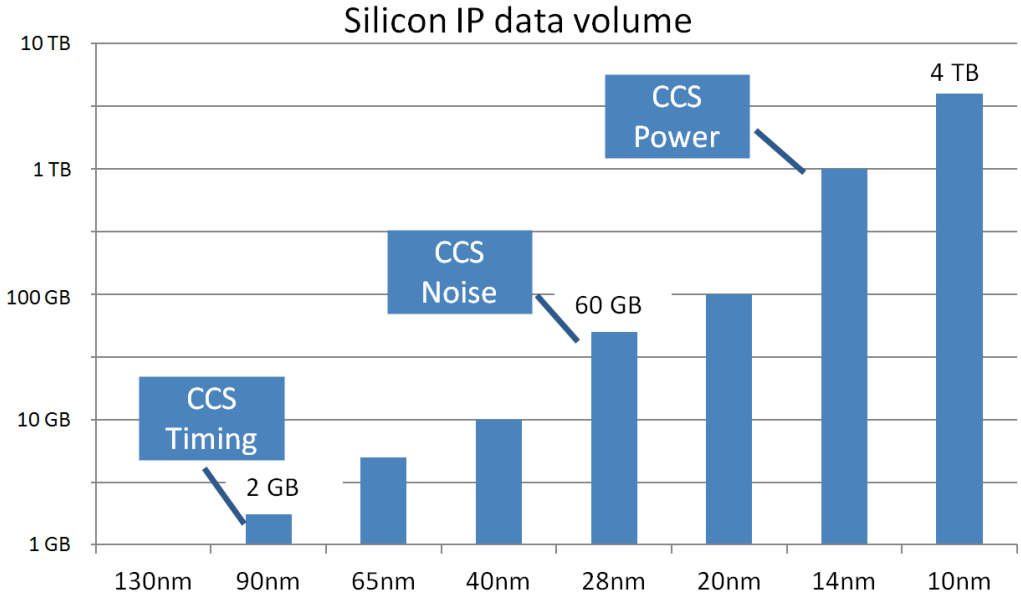

Figure 1, Silicon IP characterization data volume per process node

Figure 1 shows the typical amount of characterization data necessary to describe the silicon-IP (standard-cell library and hard macros) needed to design and verify an SoC at that process node [1]. The logarithmic scale clearly demonstrates the exponential nature of the growth of the data volume. The trend in figure 1 shows that the doubling of the number of transistors

per with successive process nodes is surpassed by the 4X characterization data volume growth factor.

Smaller dimensions on silicon imply that physical phenomena which could safely be ignored in previous nodes now have to be carefully characterized and taken into account when running final design validation. For example, using the traditional NLDM/NLPM models for a 0.25 micron node for timing and power, the current drawn by a cell could be defined by a simple linear dependency on the amount of capacitance attached to the output. Once these current-values were determined from SPICE simulations on a handful of process corners (nominal, slow, fast), this would be sufficient to characterize the power consumption of an entire SoC design in order to reliably dimension the power supply rails. Needless to say the amount of data to needed to describe a state of the art cell-library or IP block was only in the order of 10s of Megabytes.

Feedforward to 2015 where FinFets are manufactured in a 14nm technology, power characterization requires a format like CCSP to accurately model the physics involved. In CCSP, power is no longer modeled based on capacitance alone; instead the current with which an output is able to drive the connected RC network (the Composite Current Source) needs to be modeled in the characterization file. In addition, leakage current and voltage supply dependency of this drive current also need to be taken into account for modeling accuracy.

To make matters even more challenging, the physical effects modeled by CCSP are highly dependent on the process corner, thus requiring different CCSP model files for every different state. In total, a full characterization may now consist of 100's of CCSP files, each for a different process corner. As a consequence, the total amount of data volume explodes, leading to 4TB of characterization data for an advanced 10nm manufacturing process.

The increase in manufacturing variability

As dimensions on silicon are decreasing, small variations that occur during manufacturing need to be taken into account during validation. This phenomenon is known as variability. A well-known effect in physics is the notion that quantum effects will emerge as soon as the dimensions are sufficiently small. A visual appreciation of this quantum variability is the printing of contact-holes, which requires a certain number of photons to create a response from the photo-reactive compounds in the resist on wafer, to create a contact-hole of the desired size. This amount of photons of-course depends on the size of the contact-hole to-be created. As the dimensions of the contact hole shrink, the amount of photons needed becomes less, allowing fabs to produce more dies with the same amount of light: higher densities at the same throughput. When contact-hole dimensions shrink, the required amount of photons decreases, but the variability does not. Thus the

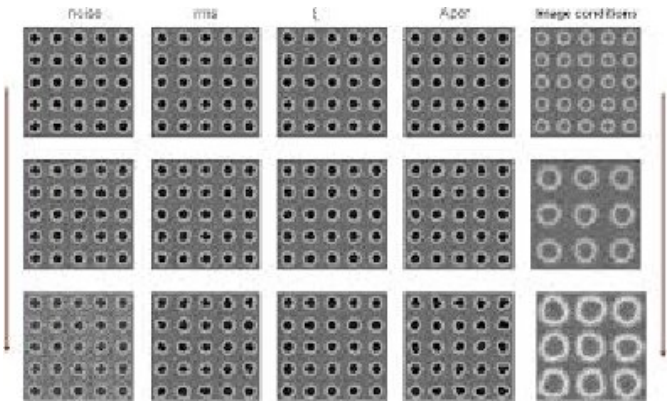

Figure 2, Illustration of photon shot-noise, see [2]

difference in number of photons seen by every contact hole, known as photon shot-noise, starts to have a measurable and visible impact. This leads to contact holes that are all of slightly different size and shape, and with bad luck a few closed contacts per billion holes printed per individual die. A SEM image showing this effect is shown in figure 2.

Such quantized variations happen everywhere during IC manufacturing and depend heavily on the exact processing conditions (layer thicknesses, position of die on the wafer, position of wafer in the lot, etc.)

Also during SoC operation, conditions change all the time: voltage levels and operating temperature can vary depending on the system state. All these variations, and combinations of variations, call for a statistical approach to final timing, power and signal integrity validation. In this approach a validation tool will determine the actual delay or current from sampling from different manufacturing or operating states. Each such state requires its own extensive modeling, hence leading to the total of 4TB of data for the characterization of 10nm silicon IP.

The end of the data volume growth is not in sight yet. Extensions to CCSP, which have been already defined but are not yet heavily used, include characterization of electro-migration (EM) effects and on-chip variability (OCV). Future process nodes (7nm and below) will require extensive usage of these new options and thus fuel more growth in characterization data volume.

New physics, new startups

As we have seen, the progression of Moore's law is driving enhanced modeling of new physical effects. With these new physics come new startups that justify their existence on new tools that can assess the impact of the new physics on the performance of the final SoC design. Such tools then typically employ proprietary data-formats, rather than extending or re-using existing accepted modeling formats with the aim of locking potential customers into their respective universes. Even after these startups have been acquired by the larger EDA fish, these proprietary formats are surprisingly resilient: they stay around long or are never replaced by a uniform modeling language. This is no surprise as the development effort to align tools around a single format is substantial and offers only little added value to customers in terms of enhanced performance or more a reliable sign-off for tape-out.

The need for validation

To summarize the above sections, the data-explosion can be broken down to the following three root-causes:

1. New physics coming from the advancement of Moore's law
2. Covering more process corners due to increased manufacturing variability
3. Divergence of data-formats as the EDA industry comes with new competing solutions

CAD groups that need to put together a design-flow, or IP design groups that build or integrate IP blocks, have a vested interest in assessing that the characterization data they are

receiving is making sense. If it doesn´t, this indicates the characterization, which also relies on automated flows, was not successful. Hence the SoC validation that uses this data also cannot be trusted for tape-out.
In older technology nodes, qualification of characterization data was quite doable in-house. The 50MB of NLDM/NLPM could be read by home-made scripting, and sanity checks on delay/power trends was straightforward (i.e. does a cell with a larger drive capability consume more power?). Establishing that all power/timing arcs in Liberty corresponded to the Verilog description of that same IP was never easy, but could be asserted manually on a few samples.

The arrival of CCSP and its related CCS formats for timing and signal integrity have completely changed this picture. Sanity and trend checks are now much more complex and apply to many more constructs in the CCSP language. For example, Crossfire checks the sum of all currents entering/leaving a cell to be a total of 0, for all different conditions for which the cell was characterized. Ideally, consumers of IP need to see qualification reports that show which trends have been checked for a given dataset, with detailed graphical feedback on pins/components that have failed. IP suppliers must be able to waive violations they perceive as irrelevant, and supply justifications for each such waives for later inspection by their customers. The Fractal Crossfire tool provides IP producers and consumers with the right tool-set to implement CCS hand-off in this level of detail.

Implementing such an automated qualification flow requires dedicated tooling also from a data volume perspective. Even the most simple tests become a nightmare in throughput when they need to be applied on very large datasets. Crossfire has extensive checks on the completeness of CCSP data. The following shortlist is intended to give an indication:

- presence and contents of pg_current_template
- presence, polarity and pin-matching of dynamic_current, incl. pin-matching
- presence and polarity of leakage_current
- presence, polarity and pin-matching of intrinsic_parasitics
- identical characterization tables across all conditions

Today's qualification flows require tooling that is designed from the ground up for managing efficient access to large datasets, typically run on the available high performance computing infrastructure. The Crossfire approach relies on separate processes on dedicated machines that service the access to large data volumes like those found in CCSP. Dedicated checking processes then only need to focus on implementing the individual checks, and are able to process a large check-set in parallel. This separation of concerns between running checks and large data management makes Crossfire the tool of choice for IP design and integration groups.

Conclusions

We have reviewed the underlying trends that lead to the current need for having dedicated IP qualification tools as part of the mixture of tools in an SoC design flow. In the course of following Moore's law, new physical effects have sparked a need for more extensive modeling and more process corners and operating conditions to be characterized separately. The EDA industry has often implemented these new effects into different, competing tools

and data-formats that must be made to cooperate in a consistent SoC design flow. As a result, the design and integration of silicon IP needs dedicated qualification tools in order to assess whether the incoming data for the entire SoC validation is trustworthy. Fractal Crossfire matches this need by supplying support for all standardized and vendor specific models and formats, by its extensive built-in check and waiving mechanism and by its native capability of efficiently dealing with very large volumes of characterization data.

References

[1] Data volumes of customer-data for various process nodes, Fractal Technologies

[2] "Effects of image noise on contact edge roughness and critical dimension uniformity measurement in synthesized scanning electron microscope images", V. Constantoudis; V. Murugesan Kuppuswamy ; E. Gogolides, Micro/Nanolith. MEMS MOEMS. (Jan 25, 2013)