WHITE **PAPER**

cadence®

Liberate Characterization Portfolio

A complete solution for fast and accurate characterization and validation

The Cadence® Liberate™ Characterization Portfolio delivers the industry's most comprehensive and robust solution for the characterization and validation of your foundation IP—from standard cells, I/Os, and complex multi-bit cells, to memories and mixed-signal blocks. Its patented Inside View technology delivers better correlation to silicon by improving library throughput and ensuring timing, power, noise, and statistical coverage of your IP.

Contents

A Complete Portfolio

The Liberate Characterization Portfolio is a collection of high-performance tools for nominal characterization, statistical characterization, validation, and generation of Liberty libraries for standard cells, I/Os, memories, and mixed-signal blocks. The portfolio achieves both accuracy and high speed through the powerful combination of the Inside View approach—patented technology for generating and optimizing characterization stimulus—coupled with a parallel processing capability that takes advantage of enterprise-wide compute resources (Figure 1). The portfolio includes Liberate™ Characterization, Liberate LV Library Validation, Liberate Variety™ Statistical Characterization, Liberate MX Memory Characterization, and Liberate AMS Mixed-Signal Characterization.

The Liberate Characterization Portfolio also integrates with the Spectre® Circuit Simulator, the industry-standard SPICE simulator, delivering even greater throughput than when used with standalone third-party simulators (Figure 2).

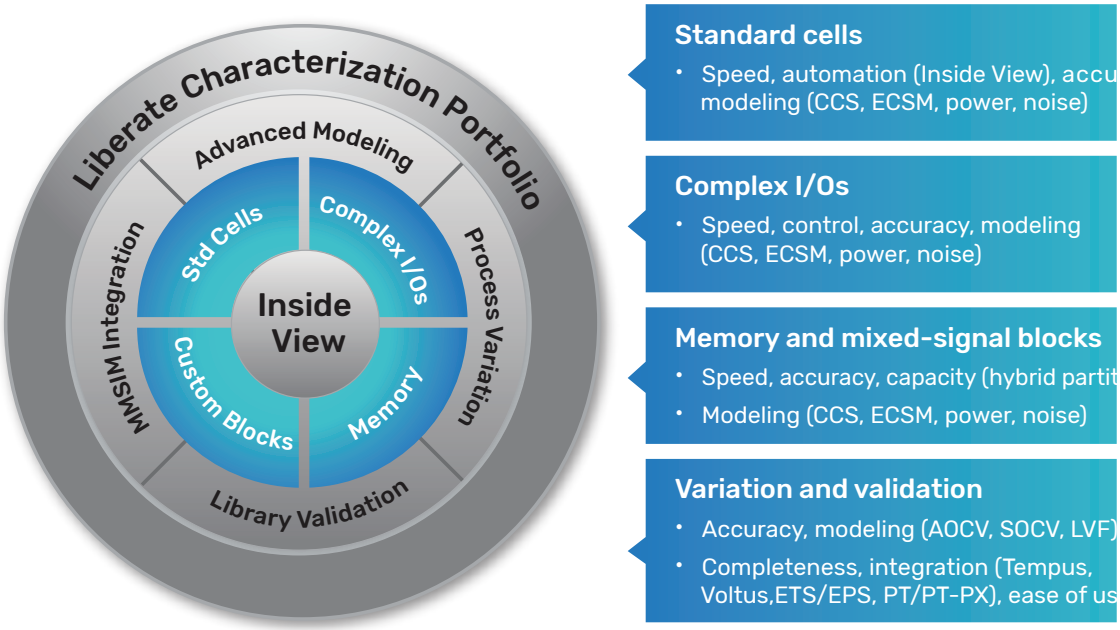

Figure 1: The Liberate Characterization Portfolio with Inside View technology

Benefits

- Ultra-fast cell library characterization for standard cells and complex I/Os
- Automatic pre-characterization of each cell using the Inside View transistor-level circuit analysis technology to learn all internal logic states and enable automatic vector generation
- Complex cell characterization for low-power and/or high-speed designs
- Support for larger macro blocks, such as memories and custom mixed-signal blocks, employing a unique "hybrid partitioning" technology to optimize runtime
- Ultra-fast throughput to complete library validation overnight on a small number of multi-core computers
- Variation-aware timing model creation accounting for process variations (systematic and random) for any set of correlated or uncorrelated process parameters

Liberate Characterization

Liberate Characterization is an ultra-fast standard cell and I/O library creator. As part of the complete Liberate Characterization Portfolio (Figure 2), it generates electrical cell views for timing, power, and signal integrity including advanced current source models (CCS and ECSM).

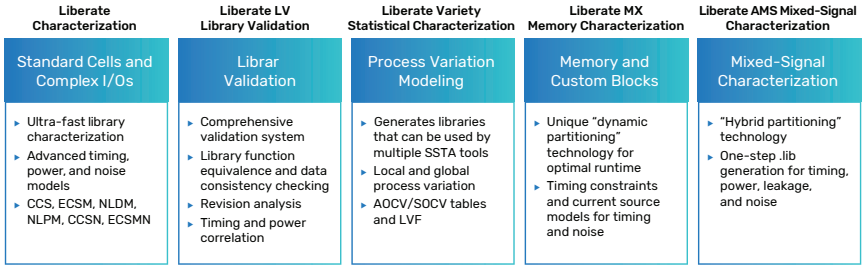

Figure 2: The portfolio comprises five characterization methods

Our Inside View approach automatically pre-characterizes each cell using transistor-level circuit analysis, which yields all the necessary stimulus and internal logic states to ensure a complete, accurate, and highly efficient characterization of that cell (Figure 3).

Figure 3: Pre-characterization circuit analysis

Liberate characterization supports complex cells including those required for high-speed and/or low-power design such as pulse latches, multi-bit flip-flop arrays, custom cells, state retention flip-flops, level shifters, power switches, and cells with sleep modes.

Creation and upkeep of library views

Designing in advanced-node nanometer process technologies requires many additional library views to achieve high-quality silicon and avoid silicon re-spins due to inaccurate signoff analysis. To manage leakage power, it is common to have low-, nominal-, and high-threshold cells, each with different power and performance characteristics.

Furthermore, for accurate modeling of instance-specific voltage variation or temperature gradients, it is necessary to characterize each library process corner over many voltages and temperatures. For the most advanced processes, it is becoming common to offer alternative cell libraries that improve yield at the expense of area and performance.

Consequently, creation and upkeep of all these library views is becoming a major bottleneck in the design flow.

Inside View

Liberate characterization uses Inside View pre-characterization circuit analysis to perform vector generation and pruning (binning), and automatic indices selection, as well as optimization of timing-constraint characterization. This results in an order-of-magnitude speedup over traditional characterization flows, enabling fully automated library creation overnight.

Parallel characterization

Liberate characterization can fully exploit a large network of multi-core CPUs via intelligent job distribution to achieve almost linear speedup per CPU (Figure 4). Liberate characterization is autonomous, greatly reducing network traffic, file I/Os, and queue requests. Furthermore, multiple characterization tasks consisting of arcs from a number of cells can be grouped into "packets" of work for optimal use of CPU resources.

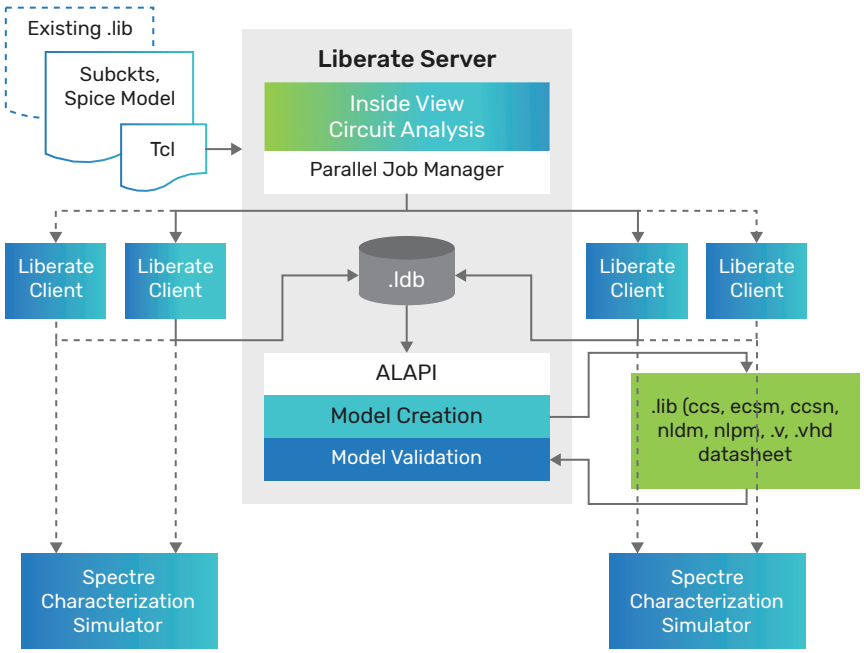

Figure 4: Liberate flow with multiple clients demonstrates parallel characterization

Characterization tasks can be performed using our tightly integrated circuit simulator, Spectre Accelerated Parallel Simulator (APS), or an external simulator, at any level of granularity, from a single arc to a complete cell. Liberate characterization supports commercial job management systems such as LSF, Sungrid, and FlowTracer.

Complex cell and model support

Liberate characterization can characterize highly complex cells (including I/Os such as DDR, USB, and PCIX), clock dividers, pipeline multi-bit flip-flops, "one-hot" muxes, and custom cells with domino logic. It supports a user-specified truth table to drive characterization in addition to automatic vector generation. Complex termination conditions and differential inputs and outputs are also supported, as well as simultaneous input switching for creating best-case corners.

It natively generates current source models for both CCS and ECSM, automatically adjusting the waveform segments to minimize the volume of data while ensuring accuracy and consistency with non-linear delay models (NLDMs). For multi-latch cells, Liberate characterization automatically determines internal probe points for characterizing timing constraints. For noise view generation, it automatically determines the input and output channel-connected logic stages and all the intermediate internal probe points.

Model generation

Liberate characterization generates Liberty, Verilog, Vital, and IBIS models, supporting the latest approved format updates. The models are generated from a central library characterization database (.ldb). Multiple versions of library models can be generated from the database to support tools that use older versions of the formats without re-characterization.

All the data in the database is accessed via a Tcl API that can generate proprietary model formats and user-specific datasheets by modifying provided examples. The database supports incremental updates and can be used to recover from characterization failures caused by network problems.

Advanced characterization algorithms

Liberate characterization supports advanced characterization algorithms providing models with better correlation and minimized pessimism within static timing analysis (STA) tools (Figure 5). These algorithms include set-up and hold pessimism reduction, minimized delay to output, dependent set-up and hold, optimized internal power controls, and many others. These algorithms are activated easily through user-selectable characterization controls.

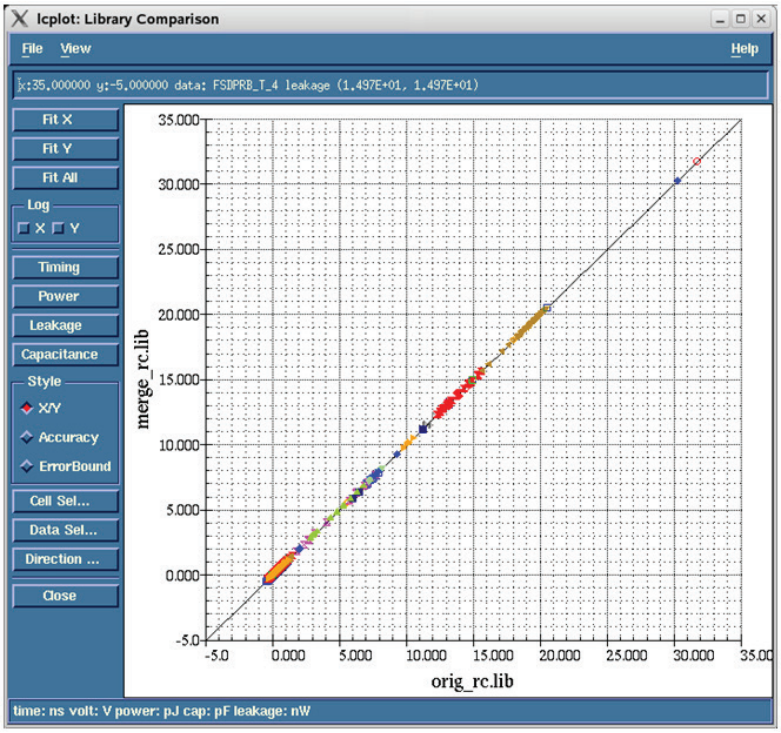

Figure 5: Advanced characterization algorithms comparison

Liberate LV Library Validation

Liberate LV Library Validation provides a collection of utilities for validating libraries including functional equivalence checking, data consistency checking, revision analysis, and correlation with various electrical analysis tools for timing, noise, and power (Figure 6).

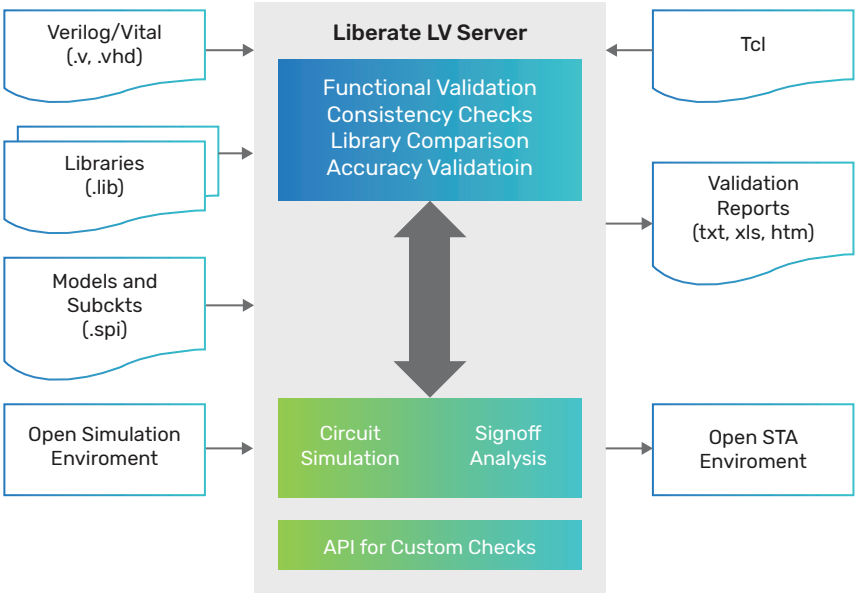

Figure 6: Liberate LV validation, with library verification inputs and outputs

Using Liberate LV validation, a complete validation of a library can be completed overnight on a small number of multi-core computers. For library providers, this ensures library quality before the library is shipped. For library users, it allows cross-checking of the incoming library and provides a clear understanding of the impact of any changes due to revisions of extracted cell netlists or process models.

Library characterization requires a complex combination of circuit simulations, data measurements, data collection, and formatting—often distributed across a large computer network. Since each library view is used for multiple chip designs, it is paramount that the library data is correct and not undermined by measurement inaccuracies or incorrect user input. Liberate LV validation provides the means to validate and verify the final library, ensuring consistency, completeness, and accuracy.

Function and state coverage

When characterizing a cell library, the input directives, vectors, and stimuli often come from a user, provided either as a previous library or hand-coded in the characterization tool's input language. However, these vectors and assumptions may not be consistent with the underlying transistor-level circuits that comprise the current library to be characterized.

Liberate LV validation checks all function descriptions in the input library directly against the transistor-level circuit and reports any differences, thus preventing potential functional errors that may occur later when the design is being formally verified or tested after manufacturing.

It provides the means to ensure that all the functional information stored in a library (.lib) is consistent with the transistor-level SPICE sub-circuits and the library Verilog and/or Vital descriptions.

In addition, Liberate LV validation ensures that all the necessary timing, noise, power (both switching and hidden power), and leakage arcs and states are represented in the library, and it will report any that are missing. It will warn where not enough distinct states exist so that potential inaccuracies in downstream tools can be avoided.

Consistency checks

Liberate LV validation provides a number of data consistency checks such as comparing table-based NLDMs against current (CCS) and voltage (ECSM) data, as well as checking for non-monotonic delay values.

For Verilog and Vital models, consistency with the Liberty models (.lib) can be checked by automatically testing SDF back-annotation onto a high-level design, which instantiates each cell in the library. Multiple SDF generation tools are supported, along with multiple gate-level logic simulators, such as Cadence Incisive® Simulator and other commercial simulators.

Library revisions

Liberate LV validation provides the means to compare a new library against an existing golden library, generating graphical, HTML, and text reports. It includes comparing libraries with different indices, function syntax, states, and cell names, and allows verification of libraries created with different characterization systems. It also highlights the impact of each new library revision, detailing changes in delay, capacitance, constraints, switching power, noise, and leakage.

Library validation and correlation

To verify that the library data is accurate, Liberate LV validation performs a correlation using the library data in the appropriate analysis tool against results obtained from circuit simulation. To ensure delay accuracy, it invokes a static timing analyzer and compares the resulting values against simulation of a test circuit using a SPICE simulator (Figure 7). The test circuits are automatically created; for example, as a variable-length chain of cells with interconnect parasitic elements. Every input-to-output arc will be verified for each logic state, input slew, and load condition. Statistical static timing analysis (SSTA) tools are also supported by comparing the path mean delay and standard deviation to Monte Carlo SPICE simulations.

Liberate Variety characterization also includes the means to measure the accuracy of timing constraints, switching power, leakage, and noise. It supports multiple analysis tools for timing and noise analysis (e.g., Cadence Tempus™ Timing Signoff Solution, Cadence Encounter® Timing System, SPICE simulation with Spectre Circuit Simulator, and third-party simulators.

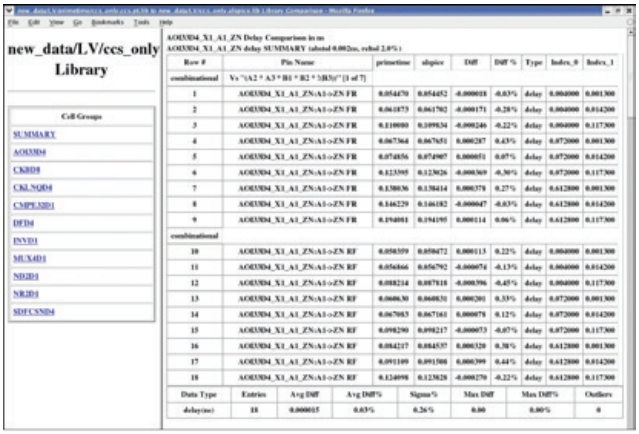

Liberate LV Timing Validation Comparison HTML Report

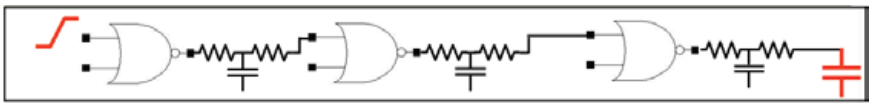

Liberate LV Sample Test Circuit

Figure 7: Library validation (Liberate LV timing validation comparison HTML report, top and correlation (Liberate LV sample test circuit, bottom)

Liberate Variety Statistical Characterization

Liberate Variety Statistical Characterization provides an ultra-fast standard cell characterizer of process variation-aware timing models (Figure 8). It generates libraries that can be used with multiple SSTAs without requiring re-characterization for each unique format. Liberate Variety characterization also generates advanced on-chip variation (AOCV) tables, statistical on-chip variation (SOCV) tables, and Liberty Variation Format (LVF).

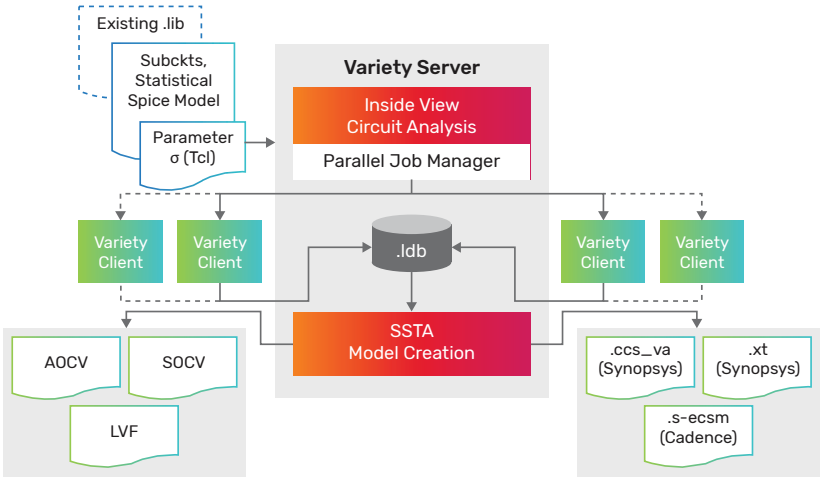

Figure 8: Liberate Variety characterization

Liberate Variety characterization calculates non-linear sensitivity, accounting for systematic and random variation for any set of correlated or uncorrelated process parameters. The resulting libraries can be used to model both local (within-cell and within-die) variations and global die-to-die variations.

SSTA provides a more realistic estimation of timing relative to actual silicon performance, often reducing worst-case timing margins by 10-15%, resulting in a higher performing, higher yielding silicon.

To accurately predict variation, SSTA needs variation-aware timing models that account for both systematic process variations (due to lithography) and random process variations (due to doping fluctuations between transistors).

Statistical timing models

Liberate Variety characterization creates models for SSTA consumption by characterizing each cell for a given set of process parameter variations where the amount of variation is based on statistical SPICE models or actual process measurements. The non-linear sensitivity to process variation for all relevant timing constructs is captured, including delay tables, slew tables, pin capacitance, and timing constraints. Advanced current source models (CCS and ECSM) are also supported.

Liberate Variety characterization can generate multiple SSTA formats from a single characterization database (ldb). Custom SSTA formats can be easily supported using a Tcl API to the characterization database.

Process parameter variation

Parameter variations can be characterized as uncorrelated, correlated, or partially correlated. Uncorrelated parameter sets are simulated independently while correlated parameter sets are simulated together. Partial correlation is supported through the use of a correlation matrix provided by the foundry.

Any process parameter present in the input SPICE model can be characterized including physical parameters such as XL or Vth, or intermediate parameters that have been derived from principle component analysis (PCA).

Systematic and random variation

For systematic inter-cell variation, the process varies in the same direction by the same amount for each transistor inside a cell. Systematic variation can be used to model both on-chip (local) and off-chip (global) variation.

Random intra-cell variation models the process variations that apply to each transistor independently (also known as mismatch). To characterize random variation efficiently, Liberate Variety characterization deploys Inside View pre-analysis technology to avoid characterizing every transistor uniquely for every table entry. This proprietary method has been validated to be highly accurate against traditional Monte Carlo simulations (Figure 9).

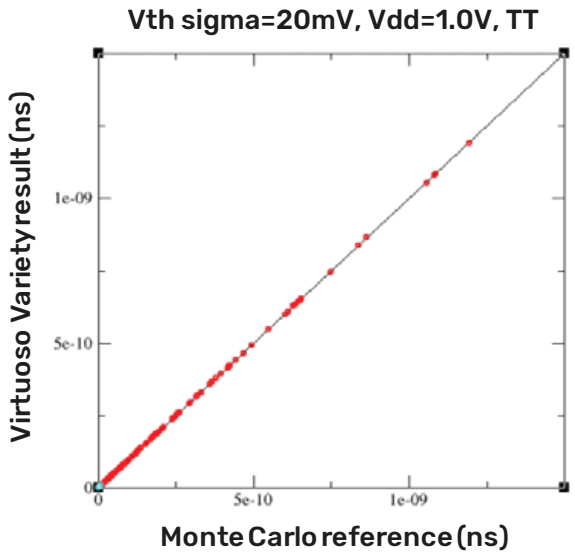

Figure 9: Process parameter variation

Liberate Variety characterization's overhead for random variation characterization is typically less than 3X the nominal characterization. Without the Inside View technology, it would take 25-50X the nominal characterization for each random parameter for a typical standard cell library.

Liberate Variety characterization can fully exploit a large network of multi-core CPUs via intelligent job distribution to achieve almost linear speedup per CPU. Characterization tasks can be performed using the natively integrated Spectre circuit simulator or a third-party simulator, at any level of granularity from a single arc to a complete cell.

Liberate MX Memory Characterization

Liberate MX Memory Characterization extends the ultra-fast standard cell and I/O library characterization capabilities of Liberate characterization to cover large memory cores. Macro blocks require additional pre-analysis steps to make fast and accurate characterization feasible.

Leveraging a network of distributed CPUs and utilizing Inside View technology for optimizing characterization runtime, memory cores can be characterized quickly and easily with the same accuracy and methods as standard cells, including the generation of timing constraints and modeling of current source models for timing, power, and noise.

Memories and large custom macros comprise a large percentage of silicon area on most chips and, consequently, can often be major contributors to chip performance and power consumption. To validate a design's electrical performance, it is essential to have a highly accurate electrical model for each macro equivalent in accuracy of the electrical models used for standard cells and I/Os.

Using pre-packaged models provided by an IP provider or memory compiler may not provide sufficient accuracy, especially as the exact context of the macro is not known until it is placed on the chip. The model may be too pessimistic, causing overruns in schedule and increased usage of larger or leakier cells to close timing. It is also common to operate a macro block at a lower voltage to save power. To get an accurate electrical model that reflects the exact usage of the macro, a design-specific and/or instance-specific macro block characterization is required.

Bi-modal view characterization

Liberate MX characterization uses both a full-block view and partitioned sub-block views to characterize large macro blocks efficiently and accurately. The full-block view is used to characterize power and to drive the creation of sub-block partitions. Typically a FastSPICE simulator, such as Spectre XPS, is used for full-block circuit analysis, followed by an accurate SPICE simulator such as the Spectre circuit simulator, Spectre APS, or external simulator for accurate characterization of sub-blocks.

Liberate MX characterization optionally supports a single-block view characterization using only a FastSPICE simulator.

Circuit partitioning

To partition each macro, Liberate MX characterization first determines internal measurement points via propagation of clocks signals and recognition of internal storage elements (Figure 10). It uses two distinct techniques for partitioning: static and dynamic.

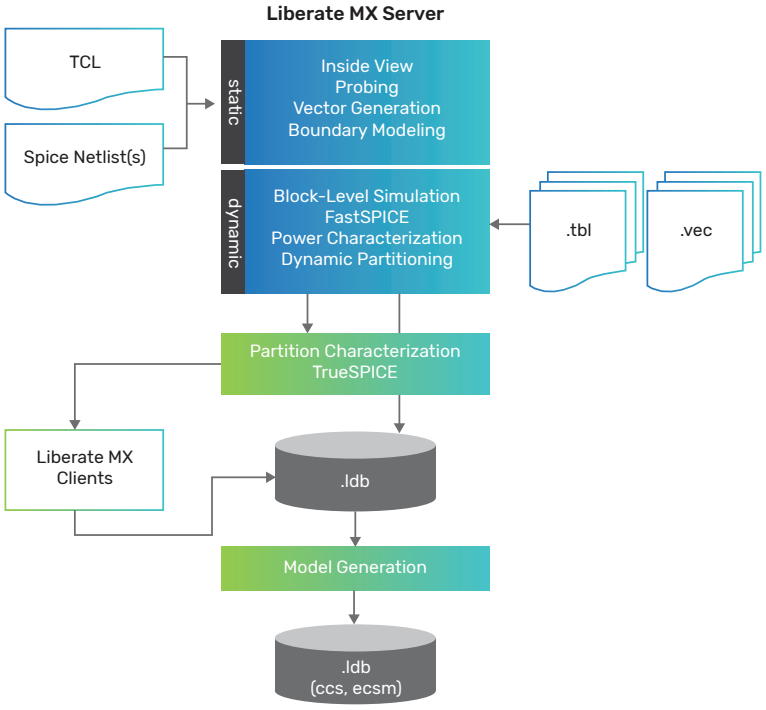

Figure 10: Memory characterization flow

Static partitions are based solely on circuit topology and are used for fully digital sections of the macro. A typical static partition includes all of the channel-connected logic components on the path between a primary input and a first-level flip-flop, including all the required clock-generation circuitry.

Dynamic partitions are derived from a full-block transistor-level simulation using a pattern sequence or a high-level truth-table description. From the simulation, an active circuit snapshot corresponding to the worst-case timing can be extracted. Dynamic partitioning is useful for characterizing timing arcs that contain analog circuitry, such as the clock-to-output-data arcs in embedded memories.

Once a macro has been partitioned into sub-blocks comprising typically a few hundred transistors, and armed with the worst-case vector set required to characterize each arc of each partition, Liberate MX characterization submits each partition for characterization using the integrated Spectre circuit simulator, Spectre APS, or an external simulator.

Library generation

Leveraging the same characterization techniques and the same command options as Liberate characterization arcs from each Liberate MX partition are characterized across a distributed network of computers utilizing all available CPUs. The distribution can be controlled using a job management system. Current source models (CCS/ECSM) for timing, power, and noise can also be generated.

Timing constraints are calculated using the same bi-sectional method as used for standard cells or, even more efficiently, as a difference in clock and data-path delay.

After characterization, all the characterized library data for each partition is assembled and compressed into a single output library representing the macro in Liberty format.

Liberate AMS Mixed-Signal Characterization

Liberate AMS Mixed-Signal Characterization extends the ultra-fast standard cell and I/O library characterization capabilities of Liberate AMS characterization to cover large mixed-signal macros such as phase-locked loops (PLLs), data converters (ADCs, DACs), SerDes, high-speed transceivers, and I/Os. Macro blocks require additional pre-analysis steps to make fast and accurate characterization feasible.

Leveraging a network of distributed CPUs and utilizing the "hybrid partitioning" technology for optimizing characterization runtime, mixed-signal macros can be characterized quickly and easily with the same accuracy and methods as standard cells, including the generation of timing constraints and modeling of current source models for timing, power, and noise.

To validate a design's electrical performance, it is essential to have a highly accurate electrical model for each mixed-signal macro equivalent in accuracy to the electrical models used for standard cells and I/Os. It is also common to operate a macro block at a lower voltage to save power. To get an accurate electrical model that reflects the exact usage of the macro, a design-specific and/or instance-specific macro block characterization is required.

Hybrid partitioning

Liberate AMS characterization uses a "hybrid partitioning" technology comprised of a full-block view and partitioned sub-block views to characterize large mixed-signal macro blocks efficiently and accurately. The full-block view is used to characterize power and to drive the creation of sub-block partitions. Typically, a designer will use a FastSPICE simulator such as Spectre XPS, for full-block circuit analysis followed by an accurate SPICE simulator such as the Spectre circuit simulator for accurate characterization of sub-blocks.

To partition each macro, Liberate AMS characterization first determines internal measurement points via propagation of clocks signals and recognition of internal storage elements (Figure 11). It uses two distinct techniques for partitioning: static and dynamic.

Static partitions are based solely on circuit topology and are used for fully digital sections of the macro. A typical static partition includes all of the channel-connected logic components on the path between a primary input and a first-level flip-flop, including all the required clock-generation circuitry.

Dynamic partitions are derived from a full-block transistor-level FastSPICE simulation using a pattern sequence, or a high-level truth-table description, or a testbench. From the simulation, an active circuit snapshot corresponding to the worst-case timing can be extracted. Dynamic partitioning is useful for characterizing timing arcs that contain analog circuitry, such as input-data-to-internal-clock and internal-clock-to-output-data arcs in high-speed mixed-signal macros.

Once a macro has been partitioned into sub-blocks comprising typically a few hundred transistors, and armed with the worst-case vector set required to characterize each arc of each partition, Liberate AMS characterization submits each partition for characterization using the integrated Spectre circuit simulator, Spectre APS, or an external simulator.

Liberate AMS characterization optionally supports a single-block view characterization using only a FastSPICE simulator.

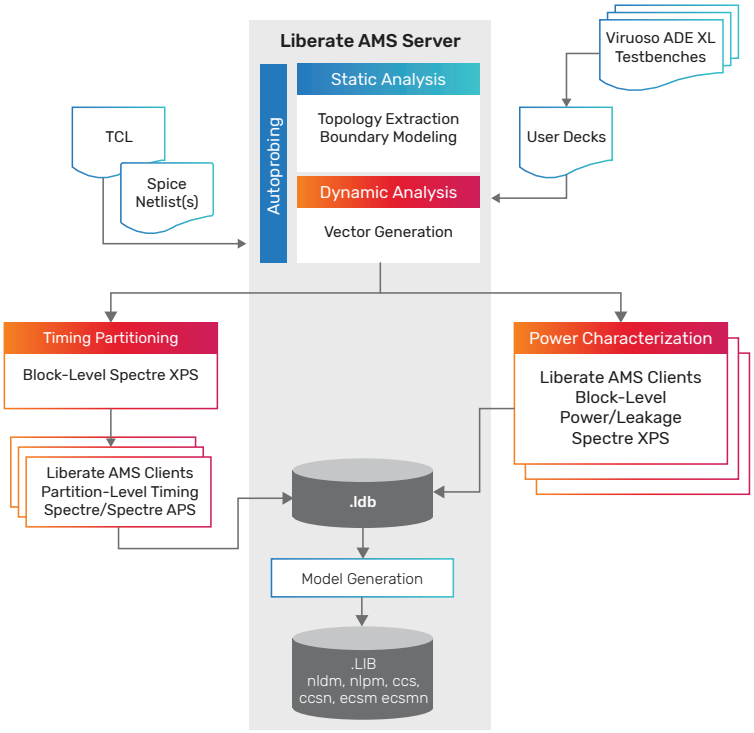

Figure 11: Mixed-signal macro characterization flow

Virtuoso ADE/ADE-XL integration

Liberate AMS characterization is tightly integrated to the Virtuoso environment, enabling deck-driven characterization using/reusing Virtuoso Analog Design Environment (ADE/ADE-XL) testbenches and setup to drive the characterization flow. Thus, no major changes to the implementation/verification flows are required and analog circuit designers are able to quickly move from circuit design validation into library generation.

Library generation

Leveraging the same characterization techniques and the same command options as Liberate characterization arcs from each Liberate AMS partition are characterized across a distributed network of computers utilizing all available CPUs. The distribution can be controlled using a job management system. Current source models (CCS/ECSM) for timing, power, and noise can also be generated.

Timing constraints are calculated using the same bi-sectional method as used for standard cells or, even more efficiently, as a difference in clock and data path delay.

After characterization, all the characterized library data for each partition is assembled and compressed into a single output library representing the macro in Liberty format.

cadence®

Cadence is a pivotal leader in electronic design and computational expertise, using its Intelligent System Design strategy to turn design concepts into reality. Cadence customers are the world's most creative and innovative companies, delivering extraordinary electronic products from chips to boards to systems for the most dynamic market applications. **www.cadence.com**