

Top-Down Design

2

1 Mixed-Signal Design Productivity

The mixed-signal design process has changed relatively little over the past two decades, and in comparison to the digital design process, is slow, labor intensive, and error prone. While digital designers have improved their design methodology and adopted design automation, analog and mixed-signal designers by and large have not.

There are two reasons why digital designers are far ahead of analog designers in improving their design processes. First, digital designers confronted the need to design very large and complex systems much earlier than analog designers. Consider that large digital chips today consist of tens of millions of transistors, while complex analog chips contain only tens of thousands of devices. Second, the digital design problem is much more amenable to automation than the analog problem.

Consider a digital design. In most cases digital systems are implemented as finite-state machines (FSM) and constructed from standard cell libraries. Using a FSM formulation acts to unify and homogenize digital design and gives it a well-understood mathematical foundation. This foundation was thoroughly explored in the late '80s resulting in the commercial logic synthesis tools of the early '90s. These tools take a register-transfer level description (RTL), a relatively high-level description of a digital system that is created by designers and can be verified with the help of logic simulators, to produce an optimized gate-level description of the system. This transformation is possible because digital systems are constructed from a limited set of relatively simple and well-behaved building blocks. The building blocks of digital systems are gates and registers. The blocks, generally referred to as cells, all share a common I/O model and so are easily interconnected, are derived from a relatively small number of cell types that have very simple and easily described behavior, are easily parameterized in terms of the number of inputs and outputs, and have a simple and easily adjusted performance trade-off that involves only speed and power. Logic synthesizers operate by creating a complete mathematical description upon which it performs transformations in order to create an optimal design in terms of speed, power, and

area. This is a two step process. First, equivalence transformations are applied to the mathematical descriptions in order to reduce the total number of gates, which minimizes the area, and the depth of the logic, which roughly maximizes the speed. This is possible because each block has a simple logical description and a common interface model. Then, the drive ability of each gate is adjusted to provide the lowest power while still meeting speed requirements. This is possible because this speed-power trade-off is easily made in each gate.

Now consider analog design. Analog design has no equivalent to finite-state machines, and so has no unified formulation and no common mathematical foundation. It also has no universal equivalence transformations that allow the topology of a circuit to be easily modified without risk of breaking the circuit. These problems prevent a topological mapping from a behavioral description to hardware. Even if one were mathematically possible, the lack of a common I/O model for analog blocks would prevent the topological modifications that are needed for either mapping or topology optimization.

It might be possible to try to enforce a common I/O model for analog circuits, but doing so would be very expensive. For example, one might simply specify that the signals at the inputs and outputs of analog blocks center around a particular value, have the same maximum swing, and that outputs have zero output impedance and inputs have zero input admittance. The problem is that doing so would necessitate extra circuitry in each analog block that is there simply to assure compliance to the I/O model. That circuitry reduces the overall performance of the circuit by increasing power dissipation, increasing noise, decreasing bandwidth, etc. This differs from the digital world where the common I/O model was achieved naturally and without significant cost. In addition, it is not possible to achieve these ideals at high frequencies. Instead, some common reference impedance would have to be specified, such as the 50Ω used at the system level, but driving such loads greatly increases power dissipation.

Finally, there is no simple way to trade-off the various performance metrics that are important with analog blocks, which makes it difficult to perform a parametric optimization. Sensitivity-based local optimizations can be used, but the improvement provided by these approaches is usually small. Monte Carlo-based global optimizers offer better improvements, but require substantially more in the way of computer resources.

The end result is that analog designers have no equivalent to RTL, a relatively high-level language in which they can describe their design and from which they can synthesize an implementation that is guaranteed to be functionally correct and have near optimal performance. As such they must transform their designs from concept to

implementation by hand, and so the design process is naturally much slower and more error prone than the design process for digital circuits.

The outlook for providing the equivalent to logic synthesis for analog designers is bleak. However, things cannot continue as they are; the current situation is becoming untenable. While a complex digital chip can be designed correctly on the first try in a few months, designing a complex analog chip can require 3-4 re-spins and up to a year and a half to get right. This is problematic for many reasons:

1. The tremendous mismatch in schedule and risk between the analog and digital portions of a mixed-signal design makes it difficult to justify combining analog and digital on the same chip.
2. The high risk makes planning difficult. It is hard to predict when product will be available, and when valuable analog designers will free up.
3. A long time-to-market makes it tough to react to changes in market trends and competitive pressures.
4. Analog and mixed-signal product development demands large investments of time and money. This makes it difficult to justify developing new analog products, especially in tough economic times.
5. Analog and mixed-signal designers are scarce and hard to recruit. Compounding this problem is the inherently low-level of productivity of the current mixed-signal design process, which makes it difficult for small design houses that are not focused on analog to field an analog design capability.
6. Some mixed-signal designs are becoming so large that, with the low productivity of the analog design process, a team of analog designers that is large enough to take on the project and complete it in a timely manner simply cannot be assembled.
7. To compensate for semiconductor processes that are increasingly unfriendly to analog designs results in an increasing use of auto calibration and adaptive filtering. This substantially increases the complexity of both the design and the verification of the design, which magnifies the problems already mentioned.

Clearly a change is needed. It is interesting to note that when digital designers were trying to design systems of a size comparable to today's mixed-signal designs, their design process was not that different from what analog designers are using today. But it was at that point that they began to transition to a more structured and more automated design methodology. For analog designers, substantial automation may not be in the cards in the near future, but the need to transition to a more structured design methodology that is both more efficient and that allows designers to handle the growing size of analog and mixed-signal circuits is clearly needed.

The availability of logic synthesis tools was not the only enabling factor for digital designers to move to more efficient design methodologies. By moving to FSM and RTL, digital designers also gave up considerable performance in terms of speed and power. They made this sacrifice to be able to design the larger and more complex systems quickly. This sacrifice was a critically important enabling factor. Analog and mixed-signal designers have not demonstrated the willingness to make a similar sacrifice. In those cases where performance is not critical, the tendency is to instead convert the circuit to a digital implementation in order to gain flexibility. In the remaining cases sacrificing performance is not an option; however it is also not clear that such a sacrifice is needed. Analog designers do not have the equivalent of logic synthesis, so they will continue to use custom design methodologies. While moving to IP (intellectual property) reuse may entail some sacrifice in overall system performance, changing to a top-down design methodology does not inherently imply lower system performance. In fact, the opposite is usually the case, using top-down design results in higher performance. Rather, the sacrifice that is demanded of analog and mixed-signal designers is that they must learn new skills, such as behavioral modeling, and they must be more disciplined in the way they design.

It is unlikely that analog and mixed-signal designers will ever be allowed on a large scale to trade any substantial amount of performance and power for a reduction in design time. Rather, in those cases where the performance and power requirements are not demanding, a digital implementation is usually preferred.

2 Traditional Approaches to Mixed-Signal Design

At the Design Automation Conference in 1998, Ron Collett of Collett International presented findings from a 1997 productivity study in which his firm analyzed 21 chip designs from 14 leading semiconductor firms. The study revealed a productivity gap of 14 \times between the most and least productive design teams. The study also revealed that developing analog and mixed-signal circuitry requires three to seven times more effort per transistor than designing digital control logic, though this factor was normalized out of the 14 \times ratio.

The reason for the poor productivity of those at the bottom end of the scale are increasingly complex designs combined with a continued preference for a bottom-up design methodology and the occurrence of verification late in the design cycle, which leads to errors and re-spins. There's a huge disparity in productivity between those mixed-signal designers who have transitioned to an effective "top-down" design methodology, and those who practice "bottom-up" design and rely solely on SPICE.

2.1 Bottom-Up Design

The traditional approach to design is referred to as bottom-up design. In it, the design process starts with the design of the individual blocks, which are then combined to form the system. The design of the blocks starts with a set of specifications and ends with a transistor level implementation. Each block is verified as a stand-alone unit against specifications and not in the context of the overall system. Once verified individually, the blocks are then combined and verified together, but at this point the entire system is represented at the transistor level.

While the bottom-up design style continues to be effective for small designs, large designs expose several important problems in this approach.

1. Once the blocks are combined, simulation takes a long time and verification becomes difficult and perhaps impossible. The amount of verification must be reduced to meet time and compute constraints. Inadequate verification may cause projects to be delayed because of the need for extra silicon prototypes.
2. For complex designs, the greatest impact on the performance, cost and functionality is typically found at the architectural level. With a bottom-up design style, little if any architectural exploration is performed and so these types of improvements are often missed.
3. Any errors or problems found when assembling the system are expensive to fix because they involve redesign of the blocks.
4. Communication between designers is critical, yet an informal and error prone approach to communication is employed. In order to assure the whole design works properly when the blocks are combined, the designers must be in close proximity and must communicate often. With the limited ability to verify the system, any failure in communication could result in the need for additional silicon prototypes.
5. Several important and expensive steps in the bottom-up design process must be performed serially, which stretches the time required to complete the design. Examples include system-level verification and test development.

The number of designers that can be used effectively in a bottom-up design process is limited by the need for intensive communication between the designers and the inherently serial nature of several of the steps. The communication requirements also tend to require that designers be co-located.

2.2 Moving to Top-Down Design

In order to address these challenges, many design teams are either looking to, or else have already implemented, a top-down design methodology [4, 18]. In a primitive top-

down approach [3], the architecture of the chip is defined as a block diagram and simulated and optimized using a system simulator such as Matlab or Simulink. From the high-level simulation, requirements for the individual circuit blocks are derived. Circuits are then designed individually to meet these specifications. Finally, the entire chip is laid out and verified against the original requirements.

This represents the widely held view of what top-down design is. And while this is a step towards top-down design, it only addresses one of the issues with bottom-up design (point 2 in Section 2.1). In essence, these design groups have not fundamentally changed their design process; they have simply added an architectural exploration step to the front. The flaw in this approach is that there is an important discontinuity in the design flow that results because the representation used during the architectural exploration phase is incompatible with the representation used during implementation. This discontinuity creates two serious problems. First, it leaves the block designers without an efficient way of assuring that the blocks all work together as expected. One could assemble transistor-level representations of the blocks and run simulations, but the simulations are too slow to be effective. The first time the blocks can be thoroughly tested together is first silicon, and at that point any errors found trigger a re-spin. Second, the discontinuity makes communication more difficult and ad hoc and so acts to separate the system designers from the circuit designers, and the circuit designers from each other. Without a reliable communication channel, designers resort to using verbal or written specifications, which are often incomplete, poorly communicated, and forgotten half way through the project. It is the poor communication process that creates many of the errors that force re-spins, and the separation that allows the errors to hide until the design is available as silicon.

To overcome these issues, one needs a design methodology that

1. Improves communication between designers (between system and block designers, between block designers, and between current designers and future designers (to support reuse)).
2. Eliminates the discontinuity that acts to hide errors and separate the system designers from the block designers.
3. Improves verification so that it finds the errors that cause re-spins, and finds them earlier so that they are less disruptive and easier to fix.
4. Improves designer effectiveness.
5. Reorganizes the design tasks, making them more parallel and eliminating long serial dependencies.
6. Reduces the need for extensive transistor-level final verification.
7. Eliminates re-spins!

RF designers typically use this type of primitive top-down design approach. They begin with the system design. Typically using a spreadsheet, the gain, noise figure and distortion budget is explored; and with the help of guides like the Friis equation, is distributed amongst the various blocks of the receiver. The design is then iterated until the performance of the system as predicted by the spreadsheet is met and the performance requirements on the blocks are reasonable. At this point, the design proceeds bottom-up relying solely on transistor-level design. Eventually, the spreadsheet is updated with the actual values coming from the transistor-level simulation, and if the system performance is not satisfactory, the process repeats. The problem is that even when using the updated results, the performance predicted by the spreadsheet will not match the results achieved in silicon. This happens because of miscommunications, either in the meaning or the actual values of the block specification, and because the system-level description is crude and does not account for things like loading effects. When designing non-integrated receivers, this is not as problematic because all the stages are generally designed for power matching and the voltage supply is reasonably high ($V_{dd} \geq 5$ V). In CMOS design the voltage supply is low (1.2 V in a 0.13 μ m process) and the blocks do not share matched impedances. The result, of course, is that multiple silicon iterations are needed to achieve the required system performance levels.

3 Principles of Top-Down Design

A well designed top-down design process methodically proceeds from architecture- to transistor-level design. Each level is fully designed before proceeding to the next and each level is fully leveraged in the design of the next. Doing so acts to partition the design into smaller, well defined blocks, and so allows more designers to work together productively. This tends to reduce the total time required to complete the design. A top-down design process also formalizes and improves communications between designers. This reduces the number of flaws that creep into a design because of miscommunication. The formal nature of the communication also allows designers to be located at different sites and still be effective.

Following a top-down design methodology also reduces the impact of changes that come late in the design cycle. If, for whatever reason, the circuit needs to be partially redesigned, the infrastructure put in place as part of the methodology allows the change to be made quickly. The models can be updated and the impact on the rest of the system can be quickly evaluated. The simulation plan and the infrastructure for mixed-level simulations is already available and can be quickly applied to verify any changes.

An effective top-down design process follows a set of basic principles.

1. A shared design representation is used for the entire length of the project that allows the design to be simulated by all members of the design team and in which all types of descriptions (behavioral, circuit, layout) can be co-simulated.
2. During the design process each change to the design is verified in the context of the entire previously verified design as dictated by the verification plan.
3. A design process that includes careful verification planning where risks are identified up-front and simulation and modeling plans are developed that act to mitigate the risks.
4. A design process that involves multiple passes, starting with high level abstractions and refining as the detail becomes available. In effect, running through the entire process very quickly at the beginning with rough estimates and guesses to get a better understanding and better estimates, and then refining the design as the process progresses.
5. To the degree possible, specifications and plans should be manifested as executable models and scripts, things that are used in the design process on a daily basis, rather than as written documents.

3.1 A Shared Design Representation

In the primitive top-down design process commonly used today, the system designers use a different design representation than the circuit designers. For example, the system designers might use a spreadsheet, Matlab, Simulink, SPW, or SystemView while the circuit designers would use Verilog, VHDL, or SPICE. This causes a myriad of problems, perhaps the most important being that they are using different tools to explore the design and that make it difficult for them to share what they learn during the design process. As mentioned before, this leads to communication problems and eventually to design errors that are generally not caught until first silicon.

If instead a common simulatable design representation is used, such as Verilog-AMS, then the system engineers can build an architectural-level description of the design constructed from behavioral models of each of the blocks that can be evaluated by each of the circuit designers. In effect, the circuit designers start by receiving an executable example of what they are expected to design. If they have trouble meeting their assigned specifications, they can go back to the system engineers with simulations that show how the system is affected by the shortfall. Since both types of engineers are working in a familiar environment, communication is enhanced and potential resolutions can be explored together. The ready availability of behavioral models of the blocks that act as executable examples greatly reduces the need for

onerous specifications that describe the desired behavior of each block, specifications that are often poorly written and that frequently go unread.

3.2 Every Change is Verified

In a primitive top-down design methodology, the architectural description of the system is usually thoroughly verified using simulation. However, the design is then re-created at the circuit level during the implementation phase and this version of the design is never checked against the original architectural description. This discontinuity is where many of the errors creep in that are not found until first silicon. In effect, the verification that was done in the architectural phase is not leveraged during the implementation phase. Verification in the implementation phase is generally not as effective because it is slow and so cannot be as comprehensive. In addition, the test benches used during the architectural design phase often cannot be reused during the implementation phase, and are generally difficult to re-create.

It is important instead to use a common simulatable representation for the design that allows both the system-level and circuit-level descriptions of the various blocks to be co-simulated, as is possible with Verilog-AMS. This capability is referred to as mixed-level simulation [14,22]. With it, individual blocks or small sets of individual blocks can be represented at the transistor- or even layout-level and be co-simulated with the rest of the system, which is described with high-level behavioral models. While these simulations are often considerably slower than simulations where every block is described at the high-level, they are also considerably faster than simulations where every block is described at the transistor level. And they allow the individual blocks to be verified in a known-good representation of the entire system. In effect, the system simulations are leveraged to provide an extensively verified test bench for the individual blocks.

Consider a simple example. It is not uncommon for a system to fail at first silicon because of a miscommunication over the polarity of a digital signal, such as a clock, enable, or reset line. Such errors cannot survive in the high-level description of the system because of the extensive testing that occurs at this level. They also cannot survive during mixed-level simulation because the individual block, where the error is presumed to be, is co-simulated with shared models for which the polarity of the signal has already been verified. They can, however, survive in either a bottom-up or primitive top-down design process because the test benches for the individual blocks are created by the corresponding block designers. Any misunderstanding of the required interface for the block will be reflected both in the implementation of the block and in its test bench, and so will not be caught until first silicon.

3.3 Verification Planning

Generally users of bottom-up or primitive top-down design methodologies find that the errors detected at first silicon are a result of rather mundane mistakes that occur at the interfaces between the various blocks. These errors are generally caused by communication breakdowns and would have been easy to find with simulation had anyone thought to look for them. The fact is that the focus of verification efforts in these methodologies is on guaranteeing the performance of the individual blocks, and not on identifying the problems that result when the blocks are interconnected. Some effort is generally spent on trying to verify the system as a whole, but it comes late in the process when the system is represented largely at the transistor level. At this stage, the simulations are quite slow and the amount of functionality that can be verified is very limited.

In a well-conceived top-down design process a verification planning step occurs that focuses on anticipating and preventing the problems that occur when assembling the blocks into a system. In order to be effective, it must move the verification to as early in the design process as possible and occur with as much of the system described at a high level as possible. Moving the chip-level verification up in the design process means that errors are caught sooner, and so are easier and less expensive to fix. Using high-level models means that the simulations run faster, and so can be substantially more comprehensive.

In a zealousness to accelerate the simulation, care must be taken to assure that enough of the system is at the right level to assure that the desired verification is actually occurring. Thus, the verification plans must include both simulation plans, that describe how the verification is to occur, and modeling plans, that indicate which models need to be available to support the verification plan and which effects should be included in the models. The modeling plan is very important. Without it behavioral models may be written that do not include the desired effect while including many effects that are unrelated to what is being verified. If they do not model the desired effect, then the verification will not be effective, if they model too many effects, then the verification runs unnecessarily slow and the models become more difficult and expensive to develop. The goal with the modeling plan is to identify a collection of simple models along with directions as to when they should be used, rather than trying to develop one complex model that is used in all cases.

An important benefit of the verification plan is that it allows the design team to react to late changes in the design requirements with confidence. When a change to the requirements occurs, it is possible to quickly revisit the verification plan, modify the design, update the models, and then apply it to the amended design to assure it satisfies the new requirements. Since it spells out all the simulations that need to occur to

verify the design, there is little chance that a change needed by the new requirements that happens to break some other part of the design will go unnoticed.

Another important benefit of the up-front planning process used when developing the verification plan is that it tends to sensitize the design team to possible problem areas, with the result that those areas are less likely to become problems.

3.4 Multiple Passes

To reduce the risk of design iterations that result from unanticipated problems, it is important to take steps to expose potential problems early by working completely through an abstract representation of the design, using estimates as needed. As the design progresses and more detailed and reliable information becomes available, the abstract representation is successively refined. This process begins by developing a top-level behavioral model of the system, which is refined until it is believed to be an accurate estimate of the desired architecture. At this point, there should be reasonable understanding as to how the blocks will be implemented, allowing size estimates to be made for the blocks, which leads to an initial floorplan. Top-level routing is then possible, which leads to parasitic estimation, with the effect of the parasitics being back annotated to the top-level. Simulations can then expose potential performance problems as a result of the layout, before the blocks are available. This may result in early changes to the architecture, changes in block specifications, or perhaps just an improvement of the verification plan. However, these changes occur early in the design process, which greatly reduces the amount of redesign needed.

As the blocks are implemented and more information becomes available, the process is continually repeated while updating and refining the design.

3.5 Executable Specifications and Plans

When a design fails because of miscommunications between engineers, it is a natural reaction to insist that in future designs, formal specifications and plans be written in advance as a way of avoiding such problems. In practice, this does not work as well as generally hoped. The act of writing things down is beneficial as it gets the engineers thinking more deeply about their designs up-front, and so they develop a better understanding of what is expected and what could go wrong. However, as for the written specifications and plans themselves, they can take a long time to write, are usually not very well written or maintained, and are often not read by the other engineers. The fact is, the documents themselves are rarely effective at improving the communications between the engineers. Rather it is the better understanding that comes from writing them that acts to improve communications.

If instead, specifications and plans took the form of executable models and scripts that would be used and valued on a daily basis, perhaps with a small amount of accompanying documentation, then they would be naturally well written, well used, and well maintained. The models and scripts are also inherently very specific, which eliminates the ambiguity that occurs in written documents and that can result in misunderstandings that lead to re-spins. These models and scripts should be maintained with the design data and shared between all designers. This avoids another of the problem with written specifications; the situation where one engineer is unaware that another has updated a specification.

Use of executable specifications and plans in the form of models and scripts both substantially improves the design process for the initial version of the chip, as well as greatly easing reuse of either the design as a whole, or the blocks used in constructing the chip. IP reuse, or reuse of the blocks, is made considerably easier because validated high-level models of the blocks are available at the end of the design process. These models would then be used to easily evaluate the blocks as to their suitability for use in other designs. Derivatives, or system reuse, are greatly simplified by the existence of all of the models and scripts. It makes it much easier for either a new team, or new team members, to get a quick understanding of an existing design and initiate the process of making changes to retarget the design to a new application. Furthermore, having models and verification scripts that have been refined by the experiences of the first design team make it more likely that the follow-on designs will debut without surprises.

4 A Rigorous Top-Down Design Process

The rigorous top-down design methodology described here is a substantial refinement of the primitive top-down process described in Section 2.2. It follows the principles described in Section 3 in order to address all of the problems associated with bottom-up design, as identified in Section 2.1.

4.1 Simulation and Modeling Plans

An important focus in a good top-down design methodology is the development of a comprehensive verification plan, which in turn leads to the simulation and modeling plans. The process begins by identifying particular areas of concern in the design. Plans are then developed for how each area of concern is to be verified. The plans specify how the tests are performed, and which blocks are at the transistor level during the test. For example, if an area of concern is the loading of one block on another, the plan might specify that one test should include both blocks represented at the tran-

sistor level together. For those blocks for which models are used, the effects required to be included in the model are identified for each test. This is the beginning of the modeling plan. Typically, many different models will be created for each block.

It is important to resist the temptation to specify and write models that are more complicated than necessary. Start with simple models and only model additional effects as needed (and as spelled out in the modeling plan). Also, the emphasis when writing models should be to model the behavior of the block, not its structure. A simple equation that relates the signals on the terminals is preferred to a more complicated model that tries to mimic the internal working of the block. This is counter to the inclination of most designers, whose intimate knowledge of the internal operation of the block usually causes them to write models that are faithful to the architecture of the block, but are more complicated than necessary.

It is also not necessary to model the behavior of a circuit block outside its normal operating range. Instead, you can add code to a model that looks for inappropriate situations and reports them. Consider a block that supports only a limited input bias range. It is not necessary to model the behavior of the block when the input bias is outside the desired range if in a properly designed circuit it will never operate in that region. It is preferable to simply generate a warning that an undesirable situation has occurred.

Following these general rules results in faster simulations and less time spent writing models. However, the question of how much detail is needed in each model is a delicate one that must be answered with great care. It is important to understand the imperfections in the blocks and how those imperfections affect the overall performance of the system before one can know whether the effects should be included in a model. Also, it is not always true that a pure behavioral model is superior to a more structurally accurate model. Often making the model more structurally accurate makes it more predictive, and also may make it easier to include some secondary effects due to parasitics.

The simulation plan is applied initially to the high-level description of the system, where it can be quickly debugged. Once validated, it can then be applied to transistor level simulations.

A formal planning process generally results in more efficient and more comprehensive verification, meaning that more flaws are caught early and so there are fewer design iterations.

4.2 System-Level Verification

System-level design is generally performed by system engineers. Their goal is to find an algorithm and architecture that implements the required functionality while providing adequate performance at minimum cost. They typically use system-level simulators, such as Simulink [21] or SPW [6], that allow them to explore various algorithms and evaluate trade-offs early in the design process. These tools are preferred because they represent the design as a block diagram, they run quickly, and they have large libraries of predefined blocks for common application areas.

This phase of the design provides a greater understanding of the system early in the design process [12,13]. It also allows a rapid optimization of the algorithm and moves trades to the front of the design process where changes are inexpensive and easy to make. Unworkable approaches are discarded early. Simulation is also moved further up in the design process where it is much faster and can also be used to help partition the system into blocks and budget their performance requirements.

Once the algorithm is chosen, it must be mapped to a particular architecture. Thus, it must be refined to the point where the blocks used at the system level accurately reflect the way the circuit is partitioned for implementation. The blocks must represent sections of the circuit that are to be designed and verified as a unit. Furthermore, the interfaces must be chosen carefully to avoid interaction between the blocks that are hard to predict and model, such as loading or coupling. The primary goal at this phase is the accurate modeling of the blocks and their interfaces. This contrasts with the goal during algorithm design, which is to quickly predict the output behavior of the entire circuit with little concern about matching the architectural structure of the chip as implemented. As such, Verilog-AMS becomes preferred during this phase of the design because it allows accurate modeling of the interfaces and supports mixed-level simulation.

The transition between algorithm and architecture design currently represents a discontinuity in the design flow. The tools used during algorithm design are different from the ones used during architecture design, and they generally operate off of different design representations. Thus, the design must be re-entered, which is a source of inefficiencies and errors. It also prevents the test benches and constraints used during the algorithm design phase from being used during the rest of the design process.

On the digital side, tools such as SPW do provide paths to implementation via Verilog and VHDL generation. Similar capabilities do not yet exist for the analog or mixed-signal portions of the design. An alternative is to use Verilog-AMS for both algorithm and architecture design. This has not been done to date because the simulators that support these languages are still relatively new. It will probably take a while for this approach to become established because of the absence of the application specific

libraries needed for rapid system-level exploration. Alternatively, simulators like AMS Designer from Cadence that supports both algorithm and architecture development by combining SPW with Verilog-AMS can be used [2].

4.3 Mixed-Level Simulation

Without analog synthesis, analog design is done the old fashioned way, with designers manually converting specifications to circuits. While this allows for more creativity and gives higher performance, it also results in more errors, particularly those that stem from miscommunication. These miscommunications result in errors that prevent the system from operating properly when the blocks are assembled even though the blocks were thought to be correct when tested individually.

To overcome this problem, mixed-level simulation is employed in a top-down design methodology for analog and mixed-signal circuits. This represents a significant but essential departure from the digital design methodology. Mixed-level simulation is required to establish that the blocks function as designed in the overall system.

To verify a block with mixed-level simulation, the model of the block in the top-level schematic is replaced with the transistor level schematic of the block before running the simulation. For this reason, all of the blocks in the architectural description of the system must be “pin-accurate”, meaning that they must have the right number of pins and characteristics of each pin must be representative of the expected signal levels, polarities, impedances, etc.

The pin-accurate system description, described at a high level, acts as a test bench for the block, which is described at the transistor level. Thus, the block is verified in the context of the system, and it is easy to see the effect of imperfections in the block on the performance of the system. Mixed-level simulation requires that both the system and the block designers use the same simulator and that it be well suited for both system- and transistor-level simulation.

Mixed-level simulation allows a natural sharing of information between the system and block designers. When the system-level model is passed to the block designer, the behavioral model of a block becomes an executable specification and the description of the system becomes an executable test bench for the block. When the transistor level design of the block is complete, it is easily included in the system-level simulation.

Mixed-level simulation is the only feasible approach currently available for verifying large complex mixed-signal systems. Some propose to use either timing simulators (sometimes referred to as fast or reduced accuracy circuit simulators) or circuit simulators running on parallel processors. However, both approaches defer system-level

verification until the whole system is available at transistor level, and neither provides the performance nor the generality needed to thoroughly verify most mixed-signal systems. They do, however, have roles to play both within the mixed-level simulation process and during final verification.

Successful use of mixed-level simulation requires careful planning and forethought (provided during the verification planning process). And even then, there is no guarantee that it will find all the problems with a design. However, it will find many problems, and it will find them much earlier in the design process, before full-chip simulations, when they are much less costly to fix. And with mixed-level simulation, it is possible to run tests that are much too expensive to run with full-chip simulation.

4.3.1 Mixed-Level Simulation Example

Though this example is several years old, it is representative of the type of circuit complexity that is common today. It is a PRML channel chip that is difficult to simulate for two reasons. First, it is a relatively large circuit that involves both analog and digital sections that are closely coupled. Second, the architecture involves complex feedback loops and adaptive circuits that take many cycles to settle. The combination of many transistors and many cycles combines with the result being a simulation that is so expensive as to be impractical. In this case, the expected simulation time was predicted to be greater than a month.

The traditional approach to simulating a complex circuit like this is to simulate the blocks individually. Of course this verifies that the blocks work individually, but not together. In addition, for this circuit it is difficult to verify the blocks when operating outside the system and it is difficult to predict the performance of the system just knowing the performance of the individual blocks.

When the architecture was simulated at a high level with each block represented by a pin-accurate behavioral model, the simulation time was less than 10 minutes. Then, when a single block was run at the transistor level, the simulation ran overnight. Even though the full system was never simulated at the transistor level, when fabricated it worked the first time because this methodology verified the blocks in the context of the system and it verified the interfaces between the blocks.

4.4 Bottom-Up Verification

Once a block is implemented, one could update the models that represent it to more closely mimic its actual behavior. This improves the effectiveness of mixed-level and system-level simulation and is referred to as bottom-up verification. To reduce the chance of errors, it is best done during the mixed-level simulation procedure. In this way, the verification of a block by mixed-level simulation becomes a three step pro-

cess. First the proposed block functionality is verified by including an idealized model of the block in system-level simulations. Then, the functionality of the block as implemented is verified by replacing the idealized model with the netlist of the block. This also allows the effect of the block's imperfections on the system performance to be observed. Finally, the netlist of the block is replaced by an extracted model. By comparing the results achieved from simulations that involved the netlist and extracted model, the functionality and accuracy of the extracted model can be verified. From then on, mixed-level simulations of other blocks are made more representative by using the extracted model of the block just verified rather than the idealized model.

Bottom-up verification should not be delayed until the end of the design process, but should rather be done continuously during the entire design process. Once a block has been implemented to the degree that a more representative model can be extracted, that model should replace the idealized top-level model as long as it does not evaluate substantially slower. Doing so tends to improve the effectiveness of mixed-level simulation and the accuracy of the extracted models. And, as a side benefit, the models that would be needed if the block were to be made into a shared IP block are already available and tested at the end of the project. If the model development for bottom-up verification were postponed to the end of the design process, the natural pressure to meet schedule targets as designs near tape-out often result in some of the verification, and perhaps all of the modeling, being skipped. This increases the chance of error and decreases the opportunity for reuse.

When done properly, bottom-up verification allows the detailed verification of very large systems. The behavioral simulation runs quickly because the details of the implementation are discarded while keeping the details of the behavior. Because the details of the implementation are discarded, the detailed behavioral models generated in a bottom-up verification process are useful for third-party IP evaluation and reuse.

4.5 Final Verification

In a top-down design process, SPICE-level simulation is used judiciously in order to get its benefits without incurring its costs. All blocks are simulated at the transistor level in the context of the system (mixed-level simulation) in order to verify their functionality and interfaces. Areas of special concern, such as critical paths, are identified up front in the verification plan and simulated at the transistor level. The performance of the circuit is verified by simulating just the signal path or key pieces of it at the transistor level. Finally, if start-up behavior is a concern, it is also simulated at the transistor level. The idea is not to eliminate SPICE simulation, but to reduce the time spent in SPICE simulation while increasing the effectiveness of simulation in general by careful planning.

It is in this phase that the dynamic timing simulators (fast reduced-accuracy transistor-level simulators) play an important role. They often have the capacity to simulate large mixed-signal systems at the transistor level for a reasonable period of time. Again, even with timing simulators the simulations are generally only fast enough to provide limited verification. So use of a timing simulator does not offset the need for mixed-level simulation.

4.6 Test

During the design phase, the test engineers should use top-level description of the design as a simulatable prototype upon which to develop the test plan and test programs. The availability of a working model of the system early in the design process allows test engineers to begin the development and testing of test programs early. Moving this activity, which used to occur exclusively after the design was complete, so that it starts at the same time the block design begins significantly reduces the time-to-production [9,10,26]. Bringing test development into the design phase can reduce post-silicon debug time by 50% and can eliminate a turn by finding chips that are untestable early. It can also improve tests, which then improves yield.

5 Further Benefits of Top-Down Design

Besides the benefits described in the previous two sections, a rigorous top-down design methodology addresses all of the various needs and issues described in Sections 1-2, which includes the following.

5.1 Improves Communications Between Engineers

Communications between the designers is improved in two substantial ways. First, the use of a shared high-level model of the system that everyone verifies their designs in eliminates most of the miscommunication that occurs when following either bottom-up or primitive top-down design processes. In addition, the executable specifications and plans (models and scripts) are more specific and less ambiguous, and considerably reduce the time spent writing and reading formal specifications, providing a more efficient and effective replacement.

5.2 Improves Productivity

The improved productivity that results with a rigorous top-down design process is due mainly to the elimination of mistakes and re-spins. A more formal, less error-prone design process with better communication between engineers means that less time is

spent making and recovering from mistakes, and more time is spent on productive design tasks.

5.3 Improves Ability to Handle Complex Designs

The ability of a design team following a rigorous top-down design methodology to handle more complex designs follows from the better system exploration and from the increased understanding of the design that comes from it, and from the improved communications. In addition, the use of mixed-level simulation dramatically improves the team's ability to verify complex circuits.

5.4 Allows Parallel Execution of Design Tasks

Reducing time-to-market is an important way in which design teams can increase their chance of success and the returns of their products. Part of the reduction in time-to-market is a result of the improved productivity and effectiveness of the design team, as described above. However, a rigorous top-down design methodology also has the benefit in that it allows more engineers to be effectively engaged in the development process at the same time, resulting in a further decrease in time-to-market.

As pointed out earlier, the existence of a shared executable high-level model of the system allows the test program development to be done in parallel with the block design and assembly, thereby eliminating a large contributor to the delay between when the design team and when manufacturing think the chip is ready to go. In addition, many of the final verification tasks that are needed with a bottom-up design style are moved forward in the form of mixed-level simulations and performed in parallel by the block and top-level designers. The block developers can also get started developing models or evaluating IP while the system designer is finalizing the overall system architecture.

The improved and more formal communication that results in a rigorous top-down design methodology allows more engineers to be involved in the design process without overstressing the shared members of the team: the team lead and the top-level and system designers. There is also a natural support for hierarchy on large projects. Only two levels have been described in this chapter, but a large chip can be partitioned into major sections (ex. RF, analog, digital, etc.), with overall leaders for the whole chip, as well as leaders for the individual sections.

5.5 Supports IP Reuse

Not only does the top-down design process described in this document improve the communication between the members of the design team, but when the design is

being reused, it also improves communication between design teams. If the design is being sold, then the process also improves the communications between different companies: seller and buyer.

A rigorous top-down design process creates as a natural by-product a thoroughly validated high-level description of the design, which is a critical enabler of IP reuse. This description is used by potential customers when evaluating the IP and by customers when integrating the IP. To see the value of having the needed model fall out of the design process as a by-product, consider the case where it does not. Using a bottom-up design process requires that the model be developed after the design is complete. This creates several barriers to the success of the IP. First, with the model not being used as an integral part of the design process it does not get much in the way of incidental testing. Substantial extra effort is required to field a high quality model, resulting in extra cost and delay. Furthermore, it is unlikely that the same quality model would be developed with an adjunct process. Second, with the model not being leveraged during the design process, the total cost of developing the model offsets any revenue from the IP, requiring higher levels of market success to break even. Finally, the model development process delays the release of the IP. This is especially troublesome as the price of IP drops dramatically as it becomes a commodity. Time-to-market is especially critical in the IP market as the price can drop by a factor of ten within a year of its release. Delay of even a month dramatically affects the total revenue of a product.

6 Final Words on Top-Down Design

Many design groups currently claim to be following a top-down design process, yet experience most of the problems attributed to the use of a bottom-up design style. This is because they are basically employing a bottom-up style with a few relatively cosmetic changes that serve to give the appearance of top-down design. This chapter lays out a series of principles that must be followed to realize all of the benefits associated with a rigorous top-down design methodology, with Verilog-AMS being the foundation upon which that methodology is built.

A rigorous top-down design methodology requires a significant investment in time and training and a serious commitment throughout the design process if it is to be successful. However, it is much easier the second time around and once mastered provides dramatic returns. Fewer design iterations and silicon re-spins are needed, which results in a shorter and more predictable design process. More optimal designs are produced that are better verified. It allows design teams to be larger and more dispersed, giving the option of trading a higher initial investment rate for a shorter time-

to-market. And it is relatively tolerant of changes in the requirements that occur late in the design cycle.

Employing a rigorous top-down design methodology dramatically increases the effectiveness and productivity of a design team. If a design team fails to move to such a design style while its competitors do, it will become increasingly ineffective. It eventually will be unable to get products to market in a time of relevance and so will be forced out of the market.

Given the high pressure world that most designers live in, it is difficult for them to acquire the skills needed to be successful in a rigorous top-down design methodology. In addition, there is little training available from continuing education centers. This suggests that the transition to top-down design will be slow. The best hope for accelerating the move to top-down design is for universities to give designers the necessary background and training in the benefits and practice of rigorous top-down design. There are some signs that this is beginning [11], but it is not as aggressive or as widespread as it needs to be in order for there to be a smooth and timely transition.

What's Next

With Chapters 1-2 as a motivation as to the importance of Verilog-AMS, we are now ready to present the language itself. Verilog-A is introduced in the next chapter. With Verilog-A you can model purely analog components and blocks. Chapter 4 introduces Verilog-AMS. These two chapters should provide you with an understanding of the fundamental concepts of the two languages, but not all of the details. Those are presented in Chapter 5.