IOSUD - Politehnica University Timişoara
Doctoral School of Engineering Studies

Synthesizable Modeling Methodologies for FPGA-Based Prototyping of Analog and Mixed-Signal Systems

Ph.D. Thesis – Abstract

for obtaining the scientific title of doctor at
Politehnica University Timişoara
in the doctoral field Engineering in Electronics, Telecommunications and Information Technologies

Author Eng. Elisei Ștefan ILIEȘ
scientific supervisor Prof. Dr. Eng. Aurel GONTEAN
month 07 year 2026

This PhD thesis addresses the development and validation of synthesizable modeling methodologies for FPGA-based prototyping of analog and mixed-signal systems. As modern System-on-Chip designs continue to increase in complexity, traditional simulation-based verification becomes increasingly limited by long execution times, especially when analog and mixed-signal components must be included in the validation flow. FPGA-based prototyping provides a practical alternative by enabling high-speed hardware execution and early software/hardware co-validation. However, analog and mixed-signal blocks cannot be directly mapped to FPGA fabric and therefore require synthesizable behavioral models that preserve the relevant system-level behavior while remaining suitable for hardware implementation.

The thesis first reviews existing approaches for hardware prototyping and real-time emulation of analog and mixed-signal systems. The analyzed methodologies include hardware/software co-simulation, custom HDL modeling, specialized automated synthesis frameworks, Simulink/HDL Coder-based flows, High-Level Synthesis, and Field Programmable Analog Arrays. This review highlights that no universal modeling approach exists; instead, the optimal methodology depends on the structure of the target system, the available reference model, the required accuracy, the acceptable development effort, and the constraints imposed by FPGA implementation.

As an initial case study, CIC filters are analyzed as examples of well-structured digital signal-processing blocks for which automatic or semi-automatic HDL generation can be effective. Vendor IP, MATLAB-generated HDL, and an open-source implementation are compared in terms of resource usage, operating frequency, flexibility, and functional equivalence. The results show that such regular architectures can be efficiently generated using existing toolchains, while still requiring careful evaluation of the trade-offs between configurability and hardware efficiency.

The first main mixed-signal case study is the AD9236 analog-to-digital converter. A SystemVerilog real-number model is developed to reproduce the relevant datasheet-level behavior of the converter, including differential input scaling, selectable output coding, reference handling, out-of-range detection, power-down operation, wake-up behavior, and the seven-cycle pipeline latency. This real-number model is validated against the official MATLAB behavioral model provided by Analog Devices using identical stimuli and coherent sampling conditions. The comparison confirms that the proposed model correctly

reproduces the core ADC transfer behavior while also extending the modeled functionality beyond the MATLAB reference by including system-level control features such as OTR and power-down behavior.

The AD9236 model is then translated into a fully synthesizable SystemVerilog floating-point implementation using Synopsys DesignWare arithmetic blocks. The transition from real-valued signals to hardware-synthesizable logic vectors required the selection of a compact custom floating-point format, the replacement of high-level arithmetic expressions with explicit floating-point operators, and the insertion of pipeline registers to close timing on the target FPGA. The initial E5M11 format was found to be insufficient, producing missing-code behavior and reducing the effective ADC resolution. By increasing the mantissa width to E5M12, the final implementation eliminated the missing-code issue, achieved timing closure at 80 MHz, preserved the seven-cycle latency, and reached an SNR of 74.03 dB, close to the theoretical limit of a 12-bit ADC. The optimized E5M12 implementation required 1217 LUTs, 161 registers, and 5 DSP48 blocks on the HAPS-100-1F platform.

The second main case study focuses on dye-sensitized solar cells. A temperature-dependent DSSC SPICE model was first developed as a reference. The DSSC is represented using a single-diode equivalent circuit with series and shunt resistances, and its five model parameters are extracted from experimental measurements using particle swarm optimization. The extracted parameters are then fitted as analytical functions of temperature, including breakpoint behavior represented through logistic functions. The resulting SPICE model reproduces the measured DSSC behavior over the 20–60 °C range, with the relative error at the maximum power point remaining below 2%. This validated SPICE model serves as the reference for the subsequent synthesizable DSSC implementations.

Two distinct FPGA-synthesizable DSSC modeling approaches are proposed and evaluated. The first approach, denoted as DSSC_FP, is an equation-based SystemVerilog floating-point implementation. Since the DSSC equation is nonlinear and contains the unknown current inside an exponential term, the current is computed using an iterative bisection solver. The model separates the temperature-dependent coefficient computation from the iterative current-solving process by using a thermal pipeline and a reusable function-evaluation pipeline controlled by a finite-state machine. Several hardware-oriented optimizations are introduced, including an expanded interpolation form to avoid catastrophic cancellation and an optimized division by two implemented through direct exponent decrement.

Multiple floating-point formats are evaluated for the DSSC_FP implementation. Reduced-precision formats such as FP_E5M10, FP_E7M16, and FP_E8M7 do not provide sufficient accuracy, while higher-precision formats such as FP_E8M23 and FP_E11M52 significantly increase hardware cost without meaningful accuracy improvement. The FP_E8M14 configuration with 13 bisection iterations is selected as the best accuracy–resource–latency trade-off. It achieves a mean absolute error of 21.50 µA, an RMSE of 28.81 µA, a maximum absolute error of 102.81 µA, and a P99-filtered mean relative error of 0.684%, with a latency of 1809.484 ns and an operating frequency of approximately 42 MHz.

The second DSSC approach, denoted as DSSC_NN, replaces the explicit solution of the nonlinear single-diode equation with a neural-network approximation trained on SPICE-generated data. A multilayer perceptron with two inputs, temperature and voltage, and one current output is trained using 450,000 samples generated from the reference SPICE model. The trained model is optimized for hardware implementation by folding the input normalization into the first dense layer and absorbing the output de-scaling into the output neuron. The neural network is then converted to a fixed-point FPGA

implementation using hls4ml and Vitis HLS, with layer-specific precision settings and reuse factors selected through profiling of weights and activations.

The DSSC_NN model is validated using the same 8200-point dataset as the DSSC_FP model. It achieves a mean absolute error of 7.72 µA, an RMSE of 11.38 µA, a maximum absolute error of 81.21 µA, and a P99-filtered mean relative error of 0.156%. The hardware implementation reaches approximately 180 MHz, with a latency of 269 clock cycles, corresponding to 1494.564 ns. The implementation requires 37896 LUTs, 32585 registers, and 141 DSP48 blocks. These results show that the neural-network approach provides higher numerical accuracy and lower absolute latency than the selected equation-based floating-point implementation, but at the cost of significantly higher FPGA resource utilization.

The comparative analysis between DSSC_FP and DSSC_NN highlights two complementary modeling strategies. The DSSC_FP model provides explicit physical interpretability, direct parameter control, lower resource usage, and smoother error surfaces, making it suitable when transparency and parametrizability are essential. In contrast, the DSSC_NN model offers higher accuracy, higher operating frequency, lower absolute inference latency, and a more automated modeling flow, but it requires more FPGA resources and must be retrained if the reference SPICE model changes. Therefore, the choice between the two approaches depends on the target application and on the relative importance of accuracy, resource usage, latency, interpretability, and development effort.

Overall, the thesis demonstrates that synthesizable HDL models can effectively support FPGA-based prototyping of analog and mixed-signal systems. The AD9236 case study confirms that datasheet-level ADC behavior can be translated into a compact synthesizable floating-point implementation while preserving timing and functional characteristics. The DSSC case study demonstrates that nonlinear temperature-dependent photovoltaic behavior can be mapped to FPGA hardware either through an equation-based floating-point solver or through a fixed-point neural-network approximation. The proposed methodologies provide practical alternatives for hardware-in-the-loop validation, design exploration, and future integration of analog and mixed-signal models into complete FPGA-based prototyping systems.