

Ravel-XL: A Hardware Accelerator for Assigned-Delay Compiled-Code Logic Gate Simulation

Michael A. Riepe, João P. Marques Silva, Kareem A. Sakallah, *Senior Member, IEEE*,
and Richard B. Brown, *Senior Member, IEEE*

Abstract—Ravel-XL is a single-board hardware accelerator for gate-level digital logic simulation. It uses a standard leveled-code approach to statically schedule gate evaluations. However, unlike previous approaches based on leveled-code scheduling, it is not limited to zero- or unit-delay gate models and can provide timing accuracy comparable to that obtained from event-driven methods. We review the synchronous waveform algebra that forms the basis of the Ravel-XL simulation algorithm, present an architecture for its hardware realization, and describe an implementation of this architecture as a single VLSI chip. The chip has about 900 000 transistors on a die that is approximately 1.4 cm², requires a 256 pin package and is designed to run at 33 MHz. A Ravel-XL board consisting of the processor chip and local instruction and data memory can simulate up to one billion gates at a rate of approximately 6.6 million gate evaluations per second. To better appreciate the tradeoffs made in designing Ravel-XL, we compare its capabilities to those of other commercial and research software simulators and hardware accelerators.

Index Terms—Hardware accelerators, simulation engines, leveled compiled code, digital logic simulation, timing analysis, design verification, special purpose architectures.

I. INTRODUCTION

DESPITE PROMISING advances over the last few years in correct-by-construction logic synthesis [5] and formal (functional) verification [8], logic simulation has yet to be dislodged from its role as an indispensable method for design verification of large digital systems. Logic simulation is utilized by digital integrated-circuit designers at many stages of the design process, from early architectural studies to final foundry sign-off simulations using back-annotated delays and complex switch-level or mixed-signal simulation algorithms. While some simulators, notably those for hardware description languages (HDL's) such as Verilog and VHDL, are flexible enough to be used at all stages of a design, the verification requirements—in terms of abstraction level and accuracy—change at each stage. In general, lowering the abstraction level increases the model's accuracy and reduces simulation speed. It is, therefore, common to use different

simulation point-tools at each stage of the design to address the specific requirements of the designer.

Digital circuit simulators can be classified into two main categories based on the scheduling algorithm they employ for gate evaluation: *statically scheduled leveled-code* (LC) [3], [6], [27], [40] versus *dynamically scheduled event-driven* (ED) [22], [28], [29], [39]. LC algorithms arrange the logic gates so that they are evaluated according to a partial ordering that ensures causality. During simulation, all gates are evaluated in each clock cycle, regardless of whether their inputs have changed since the last cycle. ED algorithms attempt to reduce the number of gate evaluations by dynamically scheduling, at run-time, only those gates whose inputs have changed. Often only a small fraction of the signals in a circuit change state each cycle so the savings is potentially large. Such savings, however, must be offset by the cost associated with the handling and scheduling of these state-change *events*. To maintain efficiency, ED methods require careful design of their data structures and event schedulers; their performance is best at low levels of circuit activity.

Orthogonal to the issue of the gate scheduling algorithm is the question of whether the simulator is *interpreted* or *compiled*. An interpreted simulator steps through the circuit by traversing a data structure representing the circuit graph, generally using time-consuming indirect addressing modes, and alternating between graph traversal and gate evaluation using subroutine calls and returns. As described by Lewis [25], circuit compilation is essentially a preprocessing step that symbolically executes the simulation to “uncover” data structures that can be statically allocated. This eliminates the code required for circuit-graph traversal, which becomes hard-coded into the simulator kernel, and replaces most indirect memory references with direct references to static addresses. Compilation also tends to unroll most loops and “in-line” many function calls, thereby reducing context switch overhead and increasing the amount of instruction-level parallelism available for use by parallel and superscalar processors. Circuit compilation, thus, tends to increase the efficiency and speed of the simulation at the cost of greater pre-processing time and larger code size. Historically, most ED simulators were interpreted, and most LC simulators were compiled. Recent research on threaded-code techniques [22], [28], [29], however, has led to the development of compilers for ED algorithms as well.

The simplest logic simulators incorporate only two-valued logic models and make no attempt to simulate circuit timing (so-called zero-delay models) [3], [40], [41]. This level of

Manuscript received June 24, 1994; revised January 11, 1995. This work was supported in part by the Advanced Research Projects Agency under Grant DAAL03-90-C-0028 and by the National Science Foundation under Grant MIP-9014058.

The authors are with the Advanced Computer Architecture Laboratory, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA.

Publisher Item Identifier S 1063-8210(96)01866-5.

abstraction was traditionally the domain of LC simulators, as the zero-delay model most closely matches the single-pass leveled gate scheduling algorithm (the presence of circuit delays introduces the possibility of hazards on the gate output which cannot be simulated in a single pass through the circuit). Zero-delay simulation is extremely fast but is useful only in the early phases of the design process when the only goal is functional verification. The dominance of LC techniques in this domain is hard to dispute.

ED algorithms are more naturally suited to the task of simulation with more complex timing models. Their ability to follow simulation activity through the circuit allows those gates with hazards to be simulated as often as necessary to obtain complete output waveforms, and arbitrarily complex timing models may be used to calculate the time at which fanout gates must be scheduled. Even so, LC simulation with circuit delays is possible. Maurer [27] has developed an LC algorithm which traces all possible paths through the circuit to obtain, for each gate, the set of *all* times at which the gate *could possibly* change, and schedules the gate for evaluation at each of those times. This allows more complex timing models, such as unit or assigned (multiple) delay, to be used but at the cost of many, often unnecessary, evaluations per gate. Thus, such approaches have little chance of obtaining competitive simulation speed [21].

Because circuits with asynchronous feedback cannot be "leveled," ED algorithms handle circuits with asynchronous feedback much more naturally than LC methods. However, iterative LC evaluation techniques can be used to simulate an asynchronous circuit until it stabilizes [41]. Often, as in the case of the feedback paths in the cross-coupled gates of an RS-latch, only one or two iterations are necessary.

Because of their ability to handle more complex timing models, as well as asynchronous feedback, ED algorithms are dominant late in the design process when circuit timing must be verified. However, this perceived dominance is worth questioning. The ED algorithm produces a complete waveform at each signal, showing the time and value of every transition before the signal stabilizes. Usually this is more information than is needed for design validation. Except on signals that are used to gate primary clocks, the presence of hazards in well-designed synchronous circuits is of little concern. Generally, all a designer is concerned with when verifying correct timing behavior is whether interface signals and latch/flip-flop inputs meet their setup and hold constraints. This implies that there are only two signal events which are of interest during each clock cycle, the *first* and *last*, and any time spent evaluating the transitions in-between is wasted. The application of delay-accurate simulation to verify setup and hold constraints in real circuits also leaves no place for arbitrarily chosen timing models, such as unit-delay, that have no relation to real circuit delays—the simulator must support gate delay values with enough resolution to accurately represent the range of lumped gate/interconnect delays provided by circuit back-annotation tools.

We recently described an LC simulation model and algorithm called Ravel that addresses these observations [31], [32], [37]. The Ravel model is an extension of a timing model that

was developed specifically to analyze and optimize the setup and hold constraints in multiphase synchronous circuits that employ level-sensitive latches [34], [35]. Ravel is based on a synchronous model for logic signals which records two events per cycle, the first and last. Using a "waveform" algebra based on this two event/cycle assumption, it calculates the stable signal values at the beginning and end of each cycle as well as the width of the changing interval in between. The event times at a gate output are calculated by a combination of *min* and *max* functions that depend not only on input event times but also on their logic values. These times are exact (identical to what an ED algorithm computes) as long as all signals in the circuit undergo at *most* two events in each clock period. The calculated event times may still be exact even when some signals experience three or more events in a clock cycle. Generally, though, the computed event times are only *bounds* on the actual event times if the 2 event/cycle assumption is violated.

Historically, the highest performing logic simulation methods rely on custom hardware *accelerators* to boost performance several orders of magnitude beyond what is achievable with software simulators [1], [2], [4], [9], [15], [17], [23], [33], [43]. More recently, hardware *emulators* based on field-programmable gate arrays (FPGA's) [30] have become popular high-end alternatives because of their faster speeds and their reconfigurability. In both cases, however, this performance premium comes at a steep cost, and such options are usually reserved to the verification of high-volume products such as microprocessors.

The Ravel-XL system described in this paper is a single-board hardware realization of the Ravel algorithm designed to maximize simulation speed while remaining simple and inexpensive. The board consists of a custom CPU chip, an asynchronous bus interface to a host processor, and external memory. In contrast to ED-based accelerators which require sophisticated hardware support for event handling [1], [2], the Ravel algorithm leads to a remarkably simple implementation. Similar to modern general-purpose CPU's, the Ravel-XL chip features a pipelined datapath that is supported by a two-level memory hierarchy optimized for the memory requirements of the datapath. In addition, the architecture uses a compact representation for data (one 32 b word per signal) and provides custom hardware instructions to perform the *min* and *max* operations necessary to compute signal waveforms. In its current implementation, Ravel-XL can simulate circuits with up to four distinct clock phases sharing a common cycle time. It has instructions to simulate the basic set of logic gates (AND/NAND, OR/NOR, XOR/XNOR, INV/BUF) with a fan-in limit of 16 inputs. It also models level-sensitive latches as well as edge-triggered flip-flops, and can be enabled to perform setup and hold violation checks. As discussed in Section VII-A, Ravel-XL is currently limited in its ability to model tri-state gates and gated-clocks.

The Ravel-XL board is designed to operate as a dedicated co-processor to a general-purpose host computer using an interrupt-driven asynchronous interface. In this configuration, the host processor is expected to maintain the user interface to the simulation process, to download the "compiled" circuit




Fig. 1. Ravel-XL system board.

and test vectors to Ravel-XL and to read back the resulting output waveforms. Ravel-XL maintains the simulation data and instructions in its own local memory space, enabling it to run at a speed that is independent of the host speed or that of the interface channel. The architecture allows for addressing up to 1 G-word each of physical data and instruction memory allowing designs of up to 1 *billion* gates to be modeled. For example, a million gate circuit such as a modern microprocessor can be accommodated with 16 4 MB DRAM chips on the board.

The custom Ravel-XL chip, designed in a 0.8- μ three-metal CMOS process, consists of about 900 000 transistors—including a 2 K word data cache—and occupies roughly 1.4 cm² of die area in a 256-pin package. Running at 33 MHz, it dissipates about 1.1 W and runs about 30 times faster than the software implementation on a workstation with the same clock rate. A prototype system board, shown in Fig. 1, will consist of the Ravel-XL chip, external code and data memories, an interface to the Digital Equipment Corporation (DEC) TURBOchannelTM bus backplane [16] realized with the DEC TcIATM (TURBOchannel Interface ASIC) chip [14], and a small number of glue-logic chips, initialization ROMs, and bus-driver chips. It is designed to operate as a peripheral device on a DEC workstation.

The remainder of this paper is organized as follows. Section II reviews the Ravel simulation model and algorithm. Section III summarizes the Ravel-XL design goals. Section IV describes the architecture of the Ravel-XL chip, including the instruction set, pipeline and memory-system design and host interface. The implementation of this architecture is discussed in Section V. Section VI analyzes the performance of Ravel-XL and provides comparisons to representative software simulators and hardware accelerators. Section VII discusses our future plans for the Ravel-XL project, and Section VIII closes the paper with some remarks summarizing our contribution.

II. RAVEL MODEL OVERVIEW

A mathematical model of the timing behavior of synchronous sequential circuits was introduced in [34], [35] and used as the basis for efficient timing verification and clock schedule optimization algorithms. This general model views the circuit as a graph whose vertices are clocked state devices—referred to as *synchronizers* to emphasize their role in insuring synchronous operation—which are either edge-triggered *D* flip-flops or level-sensitive *D* latches. Edges in the graph model the combinational logic between synchronizers and are labeled with the minimum and maximum path delays through the logic. The flow of *data* signals through the synchronizers is regulated by a set of periodic signals, collectively referred to as the *clock*, that share a common clock period and

that provide a *time reference* for specifying the event times of the data signals. Each data signal is described in terms of the times of its earliest and latest transition events in one complete period of an appropriate clock signal. Data signals are assumed to have unspecified *stable* logic values at the beginning and end of each clock period; they are assumed to be *changing* and unknown between their earliest and latest event times.

The Ravel LC logic simulator [31], [32], [37] extended the above model for use in logic simulation by requiring the stable values of data signals at the beginning and end of each clock cycle to be completely specified. Ravel models the circuit as a graph whose vertices represent the logic gates as well as the synchronizers. It views each data signal as a “waveform” and provides a set of equations for logically combining such waveforms. The resulting *waveform algebra* is unique in that it explicitly shows the relationship between the logic values and event times of the data signals in a circuit and allows the event times to be calculated accurately by a simple leveled traversal of the combinational logic. The remainder of this section summarizes those features of the Ravel model that must be considered in a hardware implementation of its simulation algorithm.

A. Signal Model

The models for clock and data signals are summarized in Fig. 2. The circuit is assumed to have k clock signals, or *phases*, labeled ϕ_1, \dots, ϕ_k that share a common cycle time T_c . Each clock phase defines a *local* frame of reference—whose origin coincides with its latching edge—for specifying event times of corresponding data signals. Phase ϕ_p is characterized by two parameters: T_p , the width of its active interval and e_p , the occurrence time of its latching edge in a suitably chosen *global* frame of reference.¹ The phases can overlap and are not required to have the same duty cycle, but must be numbered so that their latching edges are totally ordered: $e_1 \leq e_2 \leq \dots \leq e_k$. Furthermore, the global frame of reference is chosen so that $e_k = T_c$. The duration of the time interval between consecutive latching edges of phases p and r is referred to as the phase shift E_{pr} [10]

$$E_{pr} = \begin{cases} (e_r - e_p) & \text{if } (e_r > e_p) \\ (T_c + e_r - e_p) & \text{if } (e_r \leq e_p) \end{cases} = T_c - (e_p - e_r) \bmod T_c \quad (1)$$

and allows for the translation of event times between these two phases. Denoting the occurrence time of a certain event i in the *current* local frame of reference of phase p by $t_i(\phi_p)$, the same event is seen to occur at

$$t_i(\phi_r) = t_i(\phi_p) - E_{pr} \quad (2)$$

in the *next* local frame of reference of phase r . It is important to note that the use of phase-relative frames of reference and modulo arithmetic restricts data event times to a dynamic range with a spread of at most $2T_c$.

¹Without loss of generality, level-sensitive latches are assumed to be active high and flip-flops are assumed to be negative edge-triggered. Under these assumptions, the active interval of a clock phase occurs when the phase is high, and its latching edge is the falling transition.




Fig. 2. Models for clock (a) and data (b) signals.

As shown in Fig. 2(b), the waveform of a data signal x_i is an alternating sequence of stable and changing intervals. In any given cycle of operation this waveform is specified by a four-tuple (v_i, a_i, A_i, V_i) where v_i and V_i are the stable values at the start and end of the cycle, and where a_i and A_i are the event times of the first and last transitions during the cycle in the local frame of reference of some clock signal ϕ_p . The domain of v_i and V_i is the three-valued set $\{0, 1, \text{STABLE}\}$ representing the binary logic constants and a stable but unspecified logic value. Event times, in general, must be modeled as real numbers, but are usually restricted to the integers by choosing a suitable resolution. The two event times must obey the ordering $a_i \leq A_i$ and, for correct synchronous operation, $0 \leq A_i - a_i < T_c$ (the situation $A_i < a_i$ can be used to indicate that a signal is stable throughout the clock cycle, since in this case the event times are ambiguous).

B. Logic Gate Model

Ravel uses a back-end pure propagation delay model for logic gates. Other delay models, such as inertial, rise/fall, and front-end delay, are also possible but will not be elaborated further. Gate delay is specified by two parameters $0 \leq \delta \leq \Delta$ representing the minimum and maximum signal propagation delays through the gate. This delay range can be viewed as a statistical spread over an entire family of gates, or as the deterministic difference between the shortest and longest signals paths within a single gate. A "nominal" delay model is achieved by setting $\delta = \Delta$.

The basic operation performed by Ravel concerns the evaluation of the signal waveform (v_y, a_y, A_y, V_y) at the output y of a logic gate in terms of the n signal waveforms $(v_1, a_1, A_1, V_1), \dots, (v_n, a_n, A_n, V_n)$ at its inputs. It is assumed that the gate's input waveforms have been translated in time to a common frame of reference using (2). Denoting the logic function of the gate by f , gate evaluation can be summarized by the following set of four equations:

$$\begin{aligned} v_y &= f(v_1, v_2, \dots, v_n) \\ V_y &= f(V_1, V_2, \dots, V_n) \\ a_y &= \delta + \max_{1 \leq i \leq n} (c_i a_i \vee \bar{c}_i a_m) \\ A_y &= \Delta + \min_{1 \leq i \leq n} (C_i A_i \vee \bar{C}_i A_M) \end{aligned} \quad (3)$$

where c_i and C_i are Boolean flags indicating the presence or absence of early and late controlling values² on input x_i , and

²A controlling value on a gate input is one which always determines the output value of the logic gate, regardless of its other inputs. A logic-one is the controlling value for AND/NAND gates, and a logic-zero is the controlling value for OR/NOR gates. The XOR gate has no controlling value.

a_m and A_M represent the times of the first and last events over all inputs to the gate:

$$\begin{aligned} a_m &= \min_{1 \leq i \leq n} (a_i) \\ A_M &= \max_{1 \leq i \leq n} (A_i). \end{aligned} \quad (4)$$

To avoid confusion, the "+" and "v" symbols in (3) denote, respectively, arithmetic addition and logical inclusive OR. Juxtaposition in these equations denotes logical AND.

C. Synchronizer Model

The Ravel model of a D -type latch or flip-flop expresses the next-cycle waveform $(v_Q^+, a_Q^+, A_Q^+, V_Q^+)$ at the Q output in terms of the current-cycle waveform (v_D, a_D, A_D, V_D) at the D input. Both waveforms are specified in a frame of reference defined by the controlling clock phase Φ_p . The early and late next-cycle Q values for both latches and flip-flops are obtained using the familiar next-state equation $Q^+ = D$ for D -type memory elements:

$$\begin{aligned} v_Q^+ &= v_D \\ V_Q^+ &= V_D. \end{aligned} \quad (5)$$

On the other hand, the early and late output event times depend on the triggering mechanism. For edge-triggered flip-flops, these times are calculated according to

$$\begin{aligned} a_Q^+ &= \delta + T_c \\ A_Q^+ &= \Delta + T_c \end{aligned} \quad (6)$$

where δ and Δ denote the (back-end) minimum and maximum signal propagation delays through the flip-flop. The output event times for level-sensitive latches require a slightly more complex calculation:

$$\begin{aligned} a_Q^+ &= \delta + \max(a_D, T_c - T_p) \\ A_Q^+ &= \Delta + \max(A_D, T_c - T_p) \end{aligned} \quad (7)$$

where T_p is the width of the active interval of phase ϕ_p .

For either triggering mechanism, the following hold and setup constraints must be satisfied for correct latching of input data:

$$\begin{aligned} a_D &\geq H \\ A_D &\leq T_c - S \end{aligned} \quad (8)$$

where H and S are specified hold and setup parameters.

D. Ravel Code Generation

Equations (1)–(8) form the basis of the Ravel LC simulator. Ravel accepts as input a gate-level synchronous sequential circuit along with a completely-specified multiphase clock schedule, and produces as output a customized "compiled"

simulator for this circuit based on the above equations. The compilation process involves a leveled traversal of the circuit graph from the primary inputs and synchronizer outputs to the primary outputs and synchronizer inputs, and the generation of a “program” that simulates one clock cycle of operation. The code sequence in this program for a single-output combinational circuit fragment sandwiched between a set of source synchronizers and a single destination synchronizer is roughly as follows.³

- Using the phase shift equations (1) and (2), shift each source synchronizer output waveform from its respective frame of reference to the frame of reference defined by the clock phase of the destination synchronizer. This change-of-origin is necessary in order to insure that the waveforms are properly processed by the combinational logic.
- In level order, apply the gate evaluations (3)–(4) to all gates in this circuit fragment.
- Check the hold and setup constraints (8) at the input of the destination synchronizer.
- Evaluate the waveforms at the outputs of the destination synchronizer using (5)–(7).

As described in [34], clock phases are totally ordered based on the occurrence times of their latching edges in a global frame of reference. Within the generated simulation program, the code sequences corresponding to different destination synchronizers are arranged in a partial ordering that is consistent with this total order on the clock phases.

III. RAVEL-XL DESIGN GOALS

The Ravel-XL system implements the Ravel simulation algorithm in hardware. Its design was guided by three objectives. Listed according to their priority, they are the following:

- 1) to maximize performance,
- 2) to maximize capacity, and
- 3) to minimize cost.

The bulk of this paper describes the design choices we made to address the performance objective. Capacity was maximized through the use of bit-efficient data and instruction formats, and the design of a memory system which does not degrade significantly in performance when simulating large circuits, making feasible the simulation of circuits with up to a billion gates. Cost was minimized indirectly by rejecting expensive design options and by requiring the whole system to fit on a single printed-circuit board.

The performance goal is measured in terms of the *effective* number of gates processed per second, EGPS, and is given by

$$\begin{aligned} \text{EGPS} &= \frac{1}{\text{IPG} \times \text{CPI} \times T_c \times A} = \frac{f_c}{\text{CPG} \times A} \\ &= \frac{\text{GEPS}}{A} \end{aligned} \quad (9)$$

where

- IPG is the average number of instructions required to process one gate;

³Primary inputs and outputs can be easily accommodated by inserting fictitious synchronizers.

- CPI is the average number of processor cycles required to complete one instruction;
- T_c and f_c are, respectively, the processor cycle time in seconds and corresponding clock frequency in Hz;
- $\text{CPG} = \text{IPG} \times \text{CPI}$ is the average number of processor cycles required to process one gate;
- $\text{GEPS} = f_c \div \text{CPG}$ is the number of gate evaluations performed each second, and is the most prevalent metric in the literature;
- A is the activity level of the circuit expressed as the percentage of gates that must be processed in each simulated cycle of operation.

Accounting for circuit activity makes (9) a consistent metric for comparing the performance of ED as well as LC simulators and accelerators. For LC techniques, A should be set to 1 to reflect the fact that all gates are processed regardless of the actual circuit activity. In reporting performance figures we will frequently use M-EGPS to denote a million effective gate evaluations per second. We should note that IPG usually depends on the number of gate inputs. Multiplying EGPS by the average number of inputs/gate yields the average number of evaluated inputs per second (EIPS) which is often more meaningful when discussing individual circuits. Unless explicitly stated otherwise, when deriving EGPS figures we will assume that IPG is based on two-input gate.

IV. RAVEL-XL ARCHITECTURE

In this section we develop a hardware architecture for the Ravel algorithm that meets the above goals. Specifically, this architecture reduces CPG: 1) by minimizing the data storage requirements through the use of compact data and instruction formats, 2) by exploiting the inherent concurrency in the algorithm through the use of pipelined parallel functional units in a custom datapath, and 3) by reducing the impact of high memory traffic through careful matching of the design of the memory system to the data and code access patterns. The other factor in the performance equation, namely, the frequency of operation, depends on the implementation of this architecture; implementation issues are discussed in Section V.

A. Signal Representation

The software implementation of Ravel requires four 32 b words to represent the waveform (v_y, a_y, A_y, V_y) of each gate output y : two words to hold the arrival times, and two words to hold the logic values. This liberal use of memory space, particularly for storing logic values, is dictated primarily by the desire to avoid the insertion of performance-degrading bit packing and unpacking operations in the instruction stream. In contrast, a custom-designed accelerator can have compact data formats with no penalty, and possibly some gain, in performance.

Signal waveforms in Ravel-XL are stored as 32 b words with 2 b fields for the logic values and 14 b fields for the arrival times. The 2 b value fields permit the encoding of the binary logic values 0 and 1 as well as the stable unspecified

value according to the following table:

$v_y[1]$	$v_y[0]$	Logic value
0	0	0
0	1	1
1	0, 1	STABLE

The use of 14 b time fields is justified by recalling, from Section II-A, that the dynamic range of signal times is at most $2T_c$. Thus, for $T_c = 10$ ns the minimum resolvable time in a 14 b representation is about 1.2 ps. The time fields are considered to be unsigned integers ranging from 0 to 16384. To represent the negative time values that may arise during the phase shift calculation at the start of each evaluation cycle (see Section II-D), all signal times are *biased* so that the most negative time that must be represented is mapped to 0. It is easy to show that the most negative time value that must be considered is $-(\max_p T_p)$ and that it occurs at the output of level-sensitive latches controlled by the clock phase with the widest active interval. The bias value is calculated from the clock parameters by the host computer which adds it to (subtracts it from) the signal times that are downloaded to (uploaded from) Ravel-XL.

B. Custom Hardware Datapath

The core of the Ravel-XL chip is a gate/synchronizer evaluation unit that implements (1)–(8). The gate evaluations (3)–(4) are “unrolled” and calculated iteratively using the template:

$$\begin{aligned} y &= G(x_1, x_2); \\ \text{for } i &= 3 \text{ to } n \\ y &= G(y, x_i) \end{aligned} \quad (10)$$

where y represents a logic value or event time at the gate output, x_1, \dots, x_n represent the corresponding variables at the gate inputs, and G denotes the appropriate input/output transformation (logical, min, or max). Using this algorithm, the output waveform of an n -input gate can be computed in $2(n-1) + 1$ steps: $(n-1)$ steps to calculate a_m and A_M from (4), and $(n-1) + 1 = n$ steps to calculate the zero-delay output waveform using (10) and to add the appropriate gate delay using (3). A simple manipulation of the arrival time equations in (3) allows a_m and A_M to be factored out of the max and min functions yielding

$$\begin{aligned} a_y &= \delta + [\bar{c}_y a_m \vee c_y \max_{1 \leq i \leq n} (c_i a_i)] \\ A_y &= \Delta + [\bar{C}_y A_M \vee C_y \min_{1 \leq i \leq n} (\bar{C}_i \vee A_i)] \end{aligned} \quad (11)$$

where c_y and C_y are boolean flags indicating, respectively, the presence of one or more inputs with early and late controlling values:

$$\begin{aligned} c_y &= c_1 \vee c_2 \vee \dots \vee c_n \\ C_y &= C_1 \vee C_2 \vee \dots \vee C_n. \end{aligned} \quad (12)$$




Fig. 3. Block diagram of the custom Ravel-XL gate evaluation datapath.

Use of (11) and (12) instead of (3) reduces the number of required computation steps to just⁴ n .

Fig. 3 is a schematic diagram of the gate/synchronizer evaluation unit highlighting its main components. The datapath has several register banks that are used to hold the computation operands and a set of functional units for performing the required operations. The registers can be conveniently divided into two groups based on how they are accessed by the functional units.

- 1) Read-only registers that are loaded with “constant” parameters by the host computer before Ravel-XL starts the simulation. This group includes a single 14 b register T_c that holds the cycle time, four 14 b registers that hold the occurrence time $(T_c - T_p)$ of the enabling edge of each clock phase, and a bank of 16 14 b registers, PSH, that hold the phase shifts between each pair of phases as computed by (1).
- 2) Read/Write registers (shown with a shadow in Fig. 3) that are loaded from the code and data memories and read by the functional units during the simulation. This group includes:
 - a) two 14 b registers δ and Δ that hold, respectively, the minimum and maximum signal delay of the gate or synchronizer being evaluated;
 - b) two 14 b registers that contain, respectively, the hold time H and the difference between the clock period and the setup time $(T_c - S)$ for the synchronizer being evaluated;
 - c) a bank of 16 32 b registers that hold the input waveforms for the gate under evaluation.

The datapath consists of nine independent functional units that implement the gate and synchronizer evaluation equations. Synchronizer evaluation is handled by three units:

- 1) the synchronizer unit which computes the signal waveforms at the outputs of flip-flops and latches using (5)–(7),
- 2) the phase shift unit which implements (2), and
- 3) the violation detection unit which checks for setup and hold violations using (8).

The remaining six units handle the evaluation of logic gates:

- 1) Unit v_y calculates the early logic value at the gate output.
- 2) Unit V_y calculates the late logic value at the gate output.
- 3) Unit MIN computes $\min(\bar{C}_i \vee A_i)$ in (11) and C_y from (12).

⁴Strictly speaking, this is true only when $n \geq 2$. For single-input gates, the minimum number of computation steps is 2.

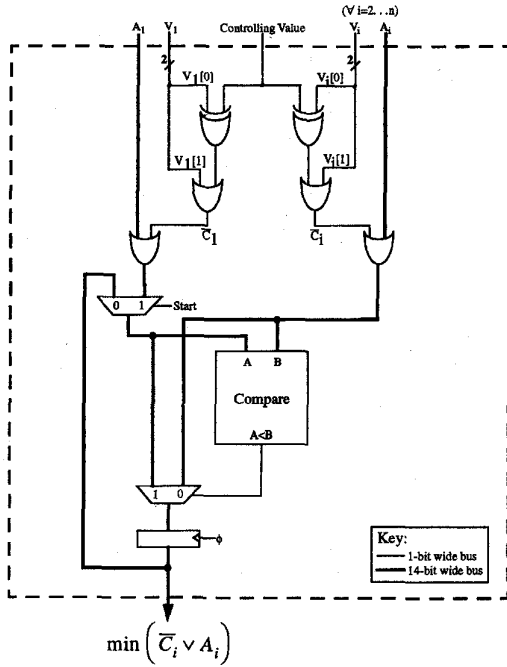


Fig. 4. A schematic of the datapath unit that computes $\min(\bar{C}_i \vee A_i)$ in (11). Here, “Controlling Value” is the binary controlling logic value of the gate type being evaluated. During the first cycle “Start” is enabled and two operands, (V_1, A_1) and (V_2, A_2) , are brought in. During all other cycles, $i = 3 \dots n$, “Start” is disabled and the input (V_i, A_i) is combined with the current cumulative result stored in the output register.

- 4) Unit MAX computes $\max(c_i a_i)$ in (11) and c_y from (12).
- 5) Unit a_m calculates the time of the earliest input event using (4).
- 6) Unit A_M calculates the time of the latest input event using (4).

The gate evaluation units operate in parallel, each using the iterative template (10). As an illustration, Fig. 4 shows the portion of functional unit MIN responsible for computing $\min(\bar{C}_i \vee A_i)$.

C. Instruction Set

Ravel-XL has seven instructions: four to perform the various simulation computations, two to handle communication with the host computer, and a NOP (No Operation) for debugging. Three of the simulation instructions are CISC-style instructions that are in one-to-one correspondence with the equations for gate evaluation, synchronizer evaluation and phase shifting. To reduce code length and still allow full access to a 32 b word-addressable address space these instructions use a base-displacement addressing mode [19]: the address of a word-aligned operand is obtained by concatenating a 16 b value from a base register with the 16 b positive displacement field in the instruction. The chip has 17 16 b base registers that are implicitly paired with the input and output operands of gates and synchronizers. The fourth simulation instruction is used to reload these base registers when it becomes necessary to address operands beyond 64 K-words from the current base.

The remainder of this section provides a detailed description of the instructions; the instruction formats are summarized in Fig. 5.

The four simulation instructions are as follows: GEV for Gate Evaluation, SEV for Synchronizer Evaluation, PSH for Phase Shift calculation, and LDB for Load B are registers.

GEV is a variable-length instruction that computes the output signal waveform for gates with up to 16 inputs. For an n -input gate the instruction is $2 + \lceil n/2 \rceil$ 32 b words long and must be padded with zeros so that it is word-aligned when the number of gate inputs is odd. The instruction can simulate any of the eight basic gate types which are identified by the TYPE field.

SEV computes the signal waveform at the output of a synchronizer in terms of the input waveform and the clock parameters. The synchronizer type (flip-flop or latch) is indicated by a 1 b flag FF, and the controlling clock phase is specified in a 2 b field PH. The instruction can be enabled to perform a setup/hold check by setting the 1 b SHC flag. To avoid propagating false signal departure times from the outputs of synchronizers with setup violations, synchronizer output departure times are clipped by the hardware to a maximum value of $T_c + \Delta$.

PSH implements (2). It subtracts the phase shift value stored in the indicated PSH register from the event times of the indicated signal waveform.

LDB loads a new base address into the indicated base register. When the ALL flag is set the base address is written to all seventeen base registers, which is useful during initialization.

The two instructions used for host communication are ENDS and WAIT. Both cause Ravel-XL to send an interrupt to the host and to pause until the host responds with a suitable command. ENDS is used to indicate the completion of a simulated clock cycle, and that Ravel-XL is ready for the next set of input patterns. WAIT instructions can be inserted in the simulation code to force breakpoints during execution; they are useful for debugging by allowing single-stepping, and can also be used for synchronization in a multiprocessor implementation of Ravel-XL (see Section VII-C).

D. Pipeline Design

For a typical circuit, with many more gates than synchronizers, simulation code based on the above instruction set is clearly dominated by the GEV instruction. This, in turn, implies that the overall performance of Ravel-XL is strongly dependent on an efficient implementation of GEV. In this section we analyze the communication and computational requirements of the GEV instruction and describe the design of a pipeline that minimizes its execution time.

The execution of a GEV instruction for an n -input gate is naturally decomposed into four steps. These steps, and the number of processor clock cycles needed to complete each, are readily shown to be the following:

- instruction fetch, requiring $(2 + \lceil n/2 \rceil)\alpha$ cycles;
- input waveforms fetch, requiring $n\alpha$ cycles;
- output waveform evaluation, requiring n cycles;
- output waveform writeback, requiring α cycles.

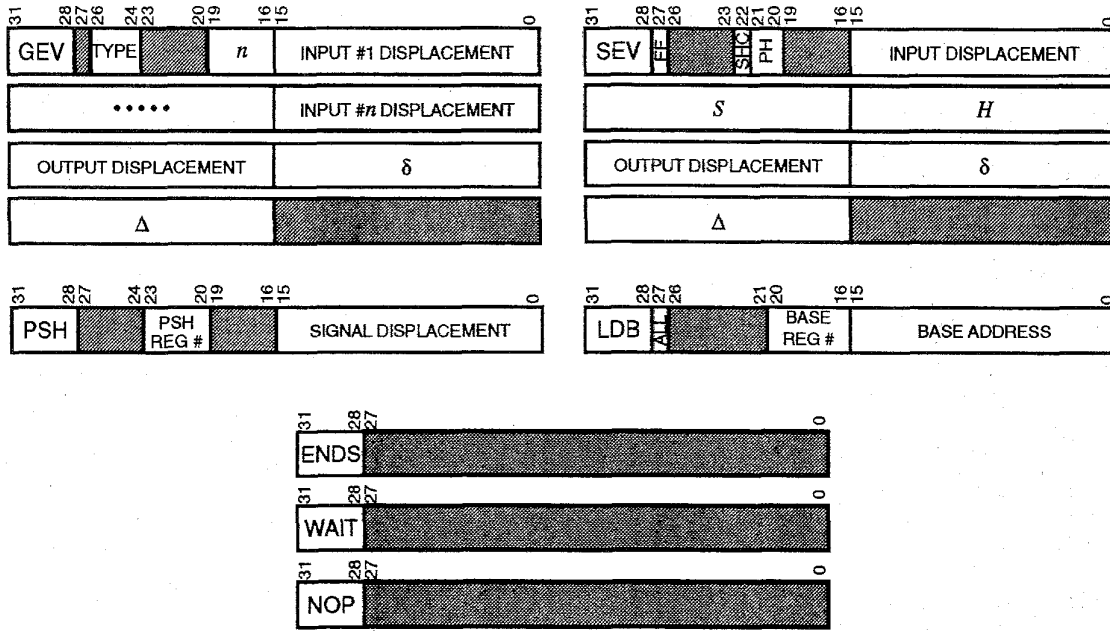


Fig. 5. Instruction formats for the Ravel-XL instruction set. Shaded fields must be set to zero and are reserved for future use.

where α is the normalized memory system cycle time—defined as the ratio between the memory and processor cycle times—and is typically greater than or equal to one. A baseline “serial” execution of the instruction, therefore, leads to a total execution time of $n + (n + 3 + \lceil n/2 \rceil)\alpha$ cycles.

The options available for reducing this execution time are basically as follows:

- 1) overlapping, or pipelining, the execution of the instruction phases;
- 2) minimizing α through proper choice of memory system organization and parameters.

These options are usually considered when designing any type of processor and are not particular to the Ravel-XL design. However, for general-purpose processors the two options are typically intertwined and must be considered simultaneously. Fortunately, the particular “structure” of the GEV instruction in Ravel-XL allows these two options to be considered somewhat independently. This fact becomes evident upon examination of the execution time of a simple four-stage pipeline whose stages are in one-to-one correspondence with the four instruction steps. In such a pipeline, each GEV instruction can be completed in an average of

$$\max[(2 + \lceil n/2 \rceil)\alpha, n\alpha, n, \alpha] = \alpha \max[2 + \lceil n/2 \rceil, n] \quad (13)$$

cycles. Execution time is clearly dominated by the instruction and data fetch steps regardless of the value of α . The rest of this section, thus, is devoted to further exploration of option 1. The tradeoffs involved in option 2 are examined separately in Section IV-E.

This four-stage pipeline implies a three-ported memory system with separate ports for 1) code fetch, 2) data fetch, and 3) data writeback. Recognizing that code and data can be separated into different memory spaces leads to an alternative

design with a single-ported code memory and a double-ported data memory. This split-memory design is simpler, cheaper, and potentially faster than the initial design. Further simplification is possible by noting that, on average, there are n read operations for every write operation to data memory. A dedicated write channel to data memory would, thus, be underutilized. Reducing the data memory to a single read/write port amounts to opting for a three-stage pipeline in which the waveform fetch and instruction writeback phases are *conceptually* combined. The total instruction execution time in this case becomes

$$\begin{aligned} & \max[(2 + \lceil n/2 \rceil)\alpha, (n + 1)\alpha, n] \\ &= \alpha \max[2 + \lceil n/2 \rceil, n + 1] \\ &= \alpha \begin{cases} 3 & \text{for } n = 1 \\ n + 1 & \text{for } n \geq 2. \end{cases} \end{aligned} \quad (14)$$

The operation of such a three-stage pipeline is illustrated in Fig. 6 for a three-input GEV instruction. In this figure, CF, DF, and EW refer, respectively, to the *code fetch*, input waveform *data fetch*, and output waveform *evaluation and writeback* stages. In order to prevent conflicting read and write requests to the data memory, the EW stage is deliberately skewed with respect to the CF and DF stages. Thus, after reading the n input waveforms of gate G_i , the channel to data memory becomes available for writing the output waveform of gate G_{i-1} . This arrangement delays the evaluation of gate G_i by $n - 1$ cycles and increases the latency of the pipeline to $2(n + 1)$. Fortunately, unlike the case of general-purpose instruction processors, such high latency is not detrimental to the performance of Ravel-XL due to the absence of branches in the instruction stream. The only data dependency that may exist in the pipeline occurs when the waveform to be fetched is still being computed in the EW stage (a read-after-write, or

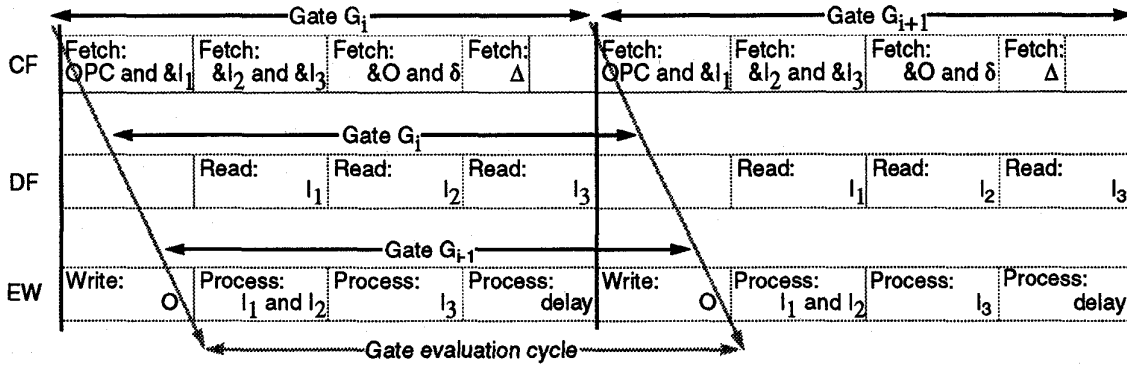


Fig. 6. Pipeline operation for a three-input GEV instruction.

RAW, hazard), and is handled by stalling the pipeline. More sophisticated solutions, such as adding *data forwarding* paths to the pipeline, are unwarranted since careful compilation can eliminate most data dependencies.

E. Memory System Design

Equation (14) shows that, with our three-stage pipeline design, simulation time is directly proportional to α , and minimized when $\alpha = 1$. As can be seen in Fig. 6, for a three-input gate the pipeline makes one reference to the code memory, and one reference to the data memory, each cycle. Our basic goal in the design of the memory system is therefore to match its *effective* cycle time to that of the processor in order to achieve a transfer rate of one instruction word and one data word per processor cycle. Additionally, this transfer rate must be sustained even when simulating large circuits. For processor frequencies below 100 MHz a simple but expensive solution is to use high-speed SRAM's with $\alpha = 1$. However, a more practical, and much cheaper, solution for obtaining single-cycle access is to design appropriate memory structures that allow the use of slower DRAM chips. This goal amounts to reducing a given normalized memory cycle time α , which may be >1 , to an effective normalized memory cycle time $\alpha_{\text{eff}} = 1$.

To obtain $\alpha_{\text{eff}} = 1$ when $\alpha > 1$ the memory system must be organized so that it matches the patterns of *locality* in the code and data streams [19]. Locality is expressed in two ways: temporal and spatial. The split memory system implied by our pipeline design gives us the opportunity to optimize the code and data memory architectures differently. This has proven useful, since the access patterns to the two memory spaces turns out to be markedly different.

In general-purpose processors, the traditional method for capturing locality is with caches. However, Lewis has observed that the straight-line code produced by compiled simulators causes poor hit rates [24]. Instead of instruction and data caches Lewis advocates the use of off-chip memories and a very deep pipeline—which would have no adverse side effects on branchless code—to absorb the long latencies. This design would address the latency issues, but would have difficulty meeting our bandwidth requirements. Ravel-XL requires an average of one memory access to each bank each

cycle—Lewis' solution would require a very large multi-ported off-chip memory to support this requirement.

The poor instruction cache hit rate is caused by a complete absence of *temporal* locality. However, we can take advantage of the high degree of *spatial* locality provided by the branchless nature of the code to obtain $\alpha_{\text{eff}} \cong 1$. Our solution uses an interleaved external code memory with prefetching. As long as the number of interleaved memory banks is greater than or equal to $\alpha + 1$, such a memory structure will be able to deliver consecutive instruction words from the straight-line code-stream at the rate of one per cycle in steady-state. Based on this analysis we chose to set α to 3, and to use a four-way interleaved memory to hold the simulation program instructions. At a target processor cycle time in the 20–40 ns range, this choice requires the use of DRAM memories with cycle times in the 60–120 ns range. Such parts are readily available and are fairly inexpensive.

Lewis also observed that the *data* stream has an irregular access pattern and lacks temporal locality as well. We have carried out a number of architectural studies, however, that indicate otherwise. We will demonstrate that, with proper compiler techniques, the temporal locality in the data stream can be controlled, allowing a cached memory organization to achieve high hit rates. We also examine the spatial locality in the data stream, and its effects on the data cache miss rate. In our discussion of the data cache we will address all four of the main cache parameters: cache size, associativity, line-size, and write policy. Our analysis will decompose the miss rate into its three components: compulsory misses, capacity misses, and conflict misses [19], and discuss the effects of our design decisions on each.

Temporal locality in the data stream results from the re-use of output signal waveforms in the evaluation of fanout gates, and is strongly dependent on the order in which the instructions are scheduled. Our compiler (discussed in more detail in Section VII-B) attempts to schedule the code stream in an order that favors the evaluation of logic gates followed immediately by their fanout gates, thus maximizing the temporal locality of the data waveforms. Temporal locality affects the rate of *capacity misses*, which are, in turn, controlled by adjusting cache *size*. As shown in Fig. 7, architectural studies have demonstrated that a cache of 2 K-words is sufficient to keep miss rates under 20% in a circuit having 35 000 gates.




Fig. 7. Cache miss rates for three different cache sizes as a function of circuit size. Here circuit size is expressed as the total number of gate inputs, since one cache access is required for each input. Cache size is the number of 32 b words.

We expect the miss rate to decrease further as we instrument the compiler with additional optimizations.

Compulsory misses turn out not to be an issue in this design. Since the host processor must download the primary input waveforms at the beginning of each simulation cycle, and since the host interface writes waveforms into the data memory through the cache, no cold misses will occur on the primary inputs. In addition, since all of a gate's inputs must be evaluated before it can be processed, waveforms will never be read before they are written. Thus, all compulsory misses are eliminated.

The final category of cache misses, *conflict* misses, is addressed by the degree of *associativity* in the cache. As shown in Fig. 8, the architectural studies did not seem to indicate that the expense of implementing a set-associative cache was warranted; instead, we chose the simpler option of a direct-mapped cache. This result is due to the absence of looping behavior, and the fact that the order in which addresses are accessed can be controlled by the compiler when it assigns addresses to the operands.

Spatial locality in the data stream, which depends on the order in which the instructions are scheduled, as well as the order in which the compiler assigns addresses to the waveforms, is more difficult to characterize than in the code stream. In a cached memory organization, the use of a *line-size* greater than one can be used to take advantage of spatial locality in the reference stream. Our compiler currently assigns addresses to the data variables in a linear fashion as they are first used. If it were modified to assign them in an order that would maximize spatial locality we might see some benefit from larger line sizes. However, such a cache adds complexity to the design, and would require an interleaved external data memory to support fast line fills. For reasons of simplicity we chose not to explore this option.

Finally, we opted for the simpler write-through, as opposed to a write-back, write policy. This is justified by the availability




Fig. 8. Effects of the degree of associativity on the cache miss rate. The total cache size is constant at 2 K words.

of adequate bandwidth on the memory channel to complete the write requests without conflict: writes occur only once for every n reads and read requests caused by cache misses are expected to be infrequent. According to (14), consecutive write requests are separated by at least 3 clock cycles. Thus, to avoid write conflicts, $\alpha \leq 3$.

The fact that we have been able to obtain reasonable cache hit rates for circuits much larger than the cache size suggests that our choice of using a data cache is justified. We believe that our data supports a claim that miss rates will not get much worse, even for very large circuits. We base this claim on several properties of combinational logic as used in large designs. First, the number of logic levels between synchronizers does not increase, as this directly impacts clock frequency. Second, the "width" of the logic, defined as the number of gate fanouts that must be maintained in the cache at any one time, is bounded by the structured design style used in their construction. Even in large chips, most combinational logic is grouped into relatively small blocks with few external connections. As long as these logic groups fit within the cache, the miss rate will not degrade.

F. Setup/Hold Violation Detection

When setup or hold violations are detected by an SEV instruction, the address of the offending synchronizer input signal is written to a violation table in data memory that can be read by the host at the end of the simulation. Since violation information is diagnostic, and not intended to be re-read during the simulation process, violation reports are written *directly* to data memory without going through the cache. Furthermore, to avoid unnecessary pipeline stalls, violation writeback requests are assigned a lower priority than operand writeback requests. This is accomplished with the use of a four-entry FIFO buffer to queue violation reports waiting to be written back. The violation report at the head of the FIFO is written to data memory during idle cycles on the data bus; the pipeline is stalled only when the FIFO is full. A larger

buffer could be used to reduce the incidence of stalls; this was deemed unnecessary, however, since violations are expected to be infrequent and to be relatively small in number.

G. Ravel-XL Host Interface

The host computer sees Ravel-XL as a memory-mapped peripheral device. The host has read/write access to both the code and data memories as well as to several internal Ravel-XL registers. A 32 b address sent by the host over the address bus is mapped by Ravel-XL to one of four address spaces according to the value of the two most significant bits: code memory, data memory, the setup/hold violation tables, and the internal system registers.

In addition to the datapath registers that are used for storing the clock parameters, the host can access the program counter, a status register, and registers that contain the address of the setup/hold violation table and the total count of violations in the table. The status register has three defined flag bits that are set by Ravel-XL: bit 7 is set when an ENDS instruction is executed; bit 6 is set upon execution of a WAIT instruction; and bit 5 is set by the SEV instruction upon detection of one or more setup/hold violations.

Two pseudo registers, START and CONTINUE, are used by the host to control the simulation process. A write to START resets the program counter and commands Ravel-XL to begin simulating; it is issued at the start of the simulation session in response to ENDS instructions. A write to CONTINUE is used to command Ravel-XL to resume simulation from a breakpoint; it is issued in response to WAIT instructions.

V. RAVEL-XL IMPLEMENTATION

A single-chip VLSI implementation of the Ravel-XL architecture is currently being prepared for fabrication. The implementation was guided by two major objectives: 1) to minimize the likelihood of pipeline stalls and 2) to minimize the clock cycle time. As noted earlier, the lack of significant data dependencies in the Ravel-XL instruction stream makes the incidence of pipeline stalls quite rare. To further reduce the possibility of stalls, deep buffers are sandwiched between the pipe stages to absorb any transient delays in the memory system response. Cycle time minimization was addressed by decomposing the chip into several largely autonomous functional units each consisting of a datapath and an associated controller. Such a "distributed control" approach—as opposed to a single global controller—reduces the possibility of a performance-limiting critical path in the control logic. Additionally, it leads to smaller controllers that are much simpler to design and test.

The design process started with architectural simulations of Ravel-XL using a behavioral model written in the Verilog Hardware Description Language (HDL) [11]. This model was manually partitioned into distinct datapath and control sections to aid the subsequent design synthesis phase. Physical design was performed using the EPOCH silicon compiler [12]. EPOCH receives its input in a synthesizable subset of Verilog HDL: behavioral datapath elements were manually converted from the behavioral model into netlists of SSI and MSI




Fig. 9. Layout plot of the Ravel-XL chip. It is implemented in a $0.8\text{-}\mu$ three-metal CMOS process, and the final dimensions of the chip are approximately 1.18×1.18 centimeters on a side.

macro cells defined in the EPOCH library, while behavioral control modules were input directly from the architectural models. EPOCH performed standard-cell logic synthesis for the behavioral controllers, and provided technology mapping for the library cells, as well as timing-driven placement, routing, and buffer and power-rail sizing. The EPOCH static timing analyzer, TACTIC, was used in the determination of the critical path. The longest sensitizable path in the design was found to lie in the datapath, and results in a clock frequency of 33 MHz. The chip contains 900 000 transistors, dissipates 1.06 Watts and occupies 1.4 cm^2 of die area in a $0.8\text{ }\mu$ three-metal CMOS process. It will be packaged in a 256 pin PGA package. Because of the large pin count the chip is pad-limited: without the pad frame the chip core is only 0.75 cm^2 . A layout plot of the chip is shown in Fig. 9.

A stylized chip floorplan showing its functional units and their major interconnections is depicted in Fig. 10. In this figure, the relative size of each functional unit roughly corresponds to the area it occupies on the chip; for clarity, however, the position of each unit may not correspond exactly to its actual chip placement. This is particularly true for the control logic: shown as a single unit on the floorplan, it is actually partitioned by the physical design tools into blocks of standard cells that are used to fill the gaps created during the placement of the datapath components. The largest block on the chip is the $2\text{ K} \times 54\text{ b}$ data cache (32 b words + 22 b tags). The other functional units identified on the floorplan—most of which have been described already—can be divided into the following four groups.

- 1) *Chip Interface* which includes the host interface (HI), code memory interface (CMI), and data memory interface (DMI).
- 2) *CF Pipeline Stage* which is the code fetch and decode (CFD) unit.




Fig. 10. Stylized chip floorplan showing major functional units and their address and data interconnections. The relative sizes of the functional units are approximately correct, though for clarity the placement of the components have little relation to that on the chip layout shown in Fig. 9.

- 3) *DF Pipeline Stage* which includes the data fetch (DF) unit and the operand Base Registers (BR).
- 4) *EW Pipeline Stage* which includes the gate evaluation (GE) unit, the gate evaluation register files (RF) and the violation queue (VQ).

The physical interface to the interleaved code memory is achieved by maintaining a 32 entry circular prefetch queue in the CMI. A controller in the CMI attempts to keep the queue full by continuously issuing read requests to the memory to prefetch instruction words. Concurrently, the CFD unit removes entries from this queue and performs the necessary instruction decoding and operand routing. Immediate operands are routed to the appropriate register: gate delays and synchronizer setup/hold parameters are written to the RF in the EW stage; base addresses in LDB instructions are written to the specified BR. Operand address displacements are posted to a 16 entry queue that is accessed by the DF unit. The DF unit removes these displacements and pairs each with an appropriate BR before issuing a read request to data memory through the DMI. The GE unit and its associated register files implement the custom datapath described in Section IV-B and shown in Fig. 3. Dual-bank registers, shown shaded in that figure, allow the CFD unit and the DMI to write data to one bank while the GE unit operates on data in the other bank, as required by the structure of the pipeline (see Fig. 6). The DMI processes reads and writes to the write-through data cache and

to the external data memory. It accepts requests from four sources: 1) operand reads from the DF unit, 2) operand writes from the GE unit, 3) violation writes from the VQ, and 4) reads/writes from the HI. Priority for access is given first to operand read requests, second to operand write requests, and last to violation write requests. Requests from the host occur only when the pipeline is stopped, so no notion of priority is needed in this case.

VI. PERFORMANCE MEASUREMENT AND COMPARISON

In this section we compare the performance of Ravel-XL to that of several other representative logic simulators. Both ED as well as LC simulators, implemented both in hardware and in software, are represented. Since the algorithms and system architectures used by the different simulators and accelerators are quite diverse we use the M-EGPS metric introduced in (9) to insure consistency. In addition, since many of the hardware accelerators achieve their speed using multiple boards—each consisting of a single processor pipeline and local storage—in parallel, we consider the board to be the atomic unit for performance comparisons. Where appropriate, we discuss multi-board system performance, and note which systems are scalable.

A. Benchmark Results

We benchmarked several software simulators including Verilog-XL, a Verilog interpreter from Cadence Design Systems [11], VCS, a Verilog compiler from Chronologic Simulation [13], and the software implementation of Ravel [31], [32], [37]. For these simulators the EGPS figures are computed directly from experimental run-times using the ISCAS-89 sequential benchmark circuit suite [7] with sequences of randomly-generated input patterns. Experiments performed with the Verilog-HDL model of Ravel-XL allow a direct comparison to be made between Ravel-XL and the other software simulators. The performance of Ravel-XL is compared with several ED hardware accelerators: MARS [1], [2], the XP product family from Zycad Corp. [43], and the Fujitsu SP [33]. It is also compared against several LC accelerators: an unnamed system by Zasio *et al.* [42] and the family of IBM simulation engines (LSM [9], YSE [15], and EVE [17]). For these systems the peak performance figures are estimated from published simulation data. Since the activity levels in these simulations are not given, the EGPS figures for ED simulators are estimated assuming a 10% activity level, which is typical for circuits we have tested. Performance estimates at a 100% activity level are also derived in an attempt to show where the trade-off between the ED and LC methods lies. A summary of the performance study is given in Table I.

1) *Ravel-XL Performance Measurements:* Assuming a circuit composed of 3-input gates and a 100% data cache hit rate, (14) predicts a 4 CPG peak performance for Ravel-XL. At 33 MHz this yields a speed of 8.25 M-EGPS which is 40 (respectively, 20) times faster than Ravel in its full long/short path (respectively long-path-only) simulation mode. However, this estimate does not take into account the structure of the

TABLE I
SIMULATION BENCHMARK RESULTS

System	Algorithm	Timing Model	Peak Speed (10^6 EGPS)		Capacity (gates)	scalable?
			activity = 100%	activity = 10%		
Verilog-XL	ED	1 value	.004	.04	n/a	N
VCS	ED	1 value	.04	.40	n/a	N
MARS (one board)	ED	rise/fall	.065	.65	64K	Y
Ravel (long & short)	LC	min/max	.20	.20	n/a	N
Ravel (long only)	LC	1 value	.40	.40	n/a	N
Zycad XP (one board)	ED	rise/fall	2.5	25	256K	Y
Zaslo <i>et al.</i>	LC	unit delay	5.0	5.0	256K	N
Ravel-XL (one board)	LC	min/max	6.6	6.6	$< 2^{30}$	Y
IBM EVE (one board)	LC	unit delay	12.5	12.5	4K	Y
Fujitsu SP (one board)	ED	unit delay	12.5	125	64K gates 5MB mem	Y




Fig. 11. Experimental results obtained with the Verilog-HDL model of Ravel-XL using the ISCAS89 suite of synchronous sequential benchmark circuits.

test circuits or the number of cycles lost to cache misses or pipeline hazard stalls.

Fig. 11 shows experimental results measured with the Verilog-HDL model of Ravel-XL. The figure shows how the number of cycles required to simulate each gate changes with circuit size. Since the average number of gate inputs may not be constant across the various circuits in the benchmark suite, we also graph the average number of cycles to process each gate input. The results show a high simulation cost for small circuits—this is due to the difficulty of scheduling gates without read-after-write (RAW) pipeline hazards. After this initial spike, the simulation cost increases slowly due to increasing cache miss rates. Finally, in circuits larger than about 10,000 gates, the cost appears to taper gradually off to a near constant value as the code scheduler is able to partition the circuit into strongly connected cache-resident blocks.




Fig. 12. The fraction of cycles spent by Ravel-XL waiting for RAW hazards and cache misses to be resolved, as a function of circuit size.

According to Fig. 11 we should expect a simulation rate closer to 5 CPG for large circuits, which will reduce our predicted performance to about 6.6 M-EGPS, or about 33 times faster than the software version of Ravel.

It is instructive to examine the fraction of clock cycles that are wasted while waiting for RAW hazards and cache misses to be cleared. As shown in Fig. 12, almost 40% of the processor cycles for the largest circuits are spent servicing RAW and cache-miss stalls. We expect this percentage to drop significantly with better compilation of the circuit equations (see Section VII-B). Fig. 13 shows how the performance of Ravel-XL, measured as the average number of cycles to simulate each gate-input, varies with cache miss rate. These numbers were generated using the ISCAS-89 s38584.1 benchmark circuit by artificially forcing cache misses at the desired rate. As can be seen in the figure, performance drops off linearly with an increase in miss rate.

It is worth pointing out that the overhead of communicating with the host will be negligible in most cases. Asynchronous host writes to Ravel-XL cost 16 clock cycles, and reads between 15 and 18. As an example, it will require 10 ms to download the 20,705 gate ISCAS89 benchmark circuit s38584 to the code memory at the beginning of a simulation, and the cost of writing/reading the 290 primary input/output values each cycle represents only about 5.3% of total simulation time.

2) *Software Simulators*: In its current implementation Ravel generates a simulation program in the MIPS R3000 instruction set [20]. Table II lists the number of machine instructions generated for a typical gate. At an ideal CPI of one on the benchmark workstation, and assuming an average of three-inputs/gate, Ravel runs at about 100 CPG. This ideal CPI rate is rarely achieved, however, because of the lack of locality in the instruction stream produced by Ravel. Experiments indicated a dramatic increase in the cache miss rate as soon as the size of the simulation loop exceeded the size of the instruction cache [32]. As we mentioned in Section IV-E, it has been observed that memory system performance degradation

TABLE II
THE NUMBER OF MACHINE INSTRUCTIONS GENERATED BY THE RAVEL COMPILER A TYPICAL GATE

Delay Model	2-input	3-input	4-input	n -input
long & short path	71	100	129	$71 + 29(n - 2)$
long path only	33	46	59	$33 + 13(n - 2)$
zero delay	8	12	16	$8 + 4(n - 2)$




Fig. 13. The variation in Ravel-XL performance, measured as the average number of cycles to simulate each gate-input, as the miss rate increases. The test circuit is S38584.1.

due to lack of locality is a problem common to LC simulators in general [22], [23]. Even for moderately sized circuits of several thousand gates the observed CPI was 2 or larger, yielding a minimum CPG of 200 for a typical three-input gate. The benchmark workstation, a DECstation 5000/240 running at 40 MHz, can be expected to achieve 0.20 M-EGPS with the full simulation model and 0.40 M-EGPS with long-path-only delays. This agrees with the simulation data gathered in [31], which observed a long-path-only simulation speed of 0.355 M-EGPS for the ISCAS-89 S1196 circuit, a typical circuit with a 13% activity level, and which is large enough to cause the CPI to be around 2.

Experiments using the ISCAS-89 sequential circuit suite have shown the software implementation of Ravel to operate about ten times faster than Verilog-XL, and at about the same speed as VCS, for circuits with activity levels near 10% [31]. In these experiments Ravel was run in long-path-only mode to more closely match the single-delay model of Verilog. Based on this data, Ravel-XL is expected to run 165 times faster than Verilog-XL and 16.5 times faster than VCS, and at a 100% activity level Ravel-XL would achieve speedups of 1650 and 165 respectively.

3) *Event Driven Hardware Accelerators*: The MARS hardware accelerator is a micro-programmable system that can be

programmed to simulate at many abstraction levels. Numbers reported here, for a single-board system programmed for a two-phase multiple-delay algorithm and running at 10 MHz [1, p. 35], are about an order of magnitude slower than Ravel-XL. This system is easily scalable using multiple boards in parallel, and the authors expect an almost linear increase in speed and capacity with a multiboard system.

Zycad Corporation markets a hardware accelerator using an ED algorithm which supports multiple-value rise/fall delays that achieves a speed of 25 M-EGPS at a 10% activity level. Circuit activity levels must be above 50% before Ravel-XL is faster. This system is scalable with up to 16 boards, obtaining a linear speedup as boards are added. Arbitrary delay models based on function calls can be used at a performance penalty.

Perhaps the fastest simulator reported is the multiprocessing SP system from Fujitsu. The only reported run times are given relative to an internal software simulator, complicating performance estimation. However, they report a maximum of 800 million *event* evaluations per second for a 64 processor system. Extrapolating back, we estimate 12.5 M-EGPS per processor at a 100% activity level, though the conditions required for peak performance are not given. In addition, the SP only supports a unit-delay model, perhaps accounting for its high performance relative to the others ED approaches.

4) *Levelized Code Hardware Accelerators*: Several other hardware accelerators use the LC technique. One by Zasio *et al.* obtains 5 M-EGPS, though it is limited to a unit-delay model for timing. The most successful LC systems were those designed by IBM, the logic simulation machine (LSM), Yorktown simulation engine (YSE) and engineering verification engine (EVE), with EVE being the most recent. All share a common architecture, which also bears some resemblance to that of Ravel-XL: it is a multi-processor system, each board made up of a single gate processing pipeline and local instruction and data memories. A CISC-style instruction is used, but theirs is limited to a constant five inputs. Boards can be scaled in parallel using a large crossbar switch, up to a maximum of 512 boards. They claim a peak throughput of over three billion EGPS and a capacity of two million gates for a full EVE system—a 500 K gate benchmark ran at 490 M-GEPS. The IBM systems are also limited to a unit-delay model for timing.

VII. FUTURE WORK

In this section we make some retrospective observations about the implementation and state the goals for a second-

generation chip. We also discuss ongoing work with several code optimization problems in the Ravel-XL compiler. We conclude with some observations on the use of Ravel-XL in a multiprocessor configuration.

A. Architectural Improvements

As shown in (14), the speed of Ravel-XL is currently limited by memory throughput. With a higher bandwidth to memory and more parallel hardware in the gate/synchronizer evaluation datapath we could conceivably obtain a simulation rate of 1 CPG. This will require the use of technology such as a multi-port cache in the data memory and a faster interface to the code memory, such as a Rambus RDRAM [18]. As the simulation speed increases the write-through cache will quickly limit performance, requiring a more complex caching scheme, in conjunction with deeper write-buffers, to limit the frequency of off-chip writes.

Other improvements that are planned include the ability to model gated-clocks and tri-state busses. The support for gated clocking may require a notion of conditional execution (i.e., branching) in the algorithm, and could introduce significant complication in the hardware. The modeling of tri-state busses will require a representation for impedance values and a new wired-logic primitive. Tri-state busses can currently be modeled by collapsing them into equivalent OR or AND gates, though CMOS bus contention will not be correctly modeled.

B. Compiler Issues

In the design of Ravel-XL we made an effort to create a flexible system in which the compiler would not be required to perform expensive optimizations to achieve reasonable performance. The only problems in the code generation process that require potentially expensive optimizations are: 1) the ordering of the gate evaluations in the instruction stream to maximize temporal locality, and 2) the ordering of the waveform variables in data memory to control spatial locality. To obtain a 100% data-cache hit rate we must guarantee that each gate waveform value stays resident in the cache from the time that it is written until the last of its fanout gates have been evaluated. The traditional level-order (breadth-first) traversal of the circuit graph identified with LC simulators may, for this reason, lead to poor data cache performance. This will be particularly noticeable if the width of the circuit at any given topological level (number of gates per level) exceeds the size of the cache.

A preliminary version of a compiler for Ravel-XL has been implemented to obtain the data shown in Section VI-A1). To address the problem of improving the temporal locality in the code, and thus the cache miss rate, we have explored several traversal techniques as alternatives to the strict level-order traversal. Basically, the compiler attempts to broadcast a gate output to its fanout gates as soon as possible after it has been evaluated to maximize the likelihood of its presence in the data cache, while at the same time minimizing the average lifetime of all cache entries. In general, this problem is NP-complete




Fig. 14. The effects of two different code ordering strategies on the cache miss rate (2 K word cache).

[26], but we have obtained good results with simple heuristics using a recursive depth-first traversal [38] of the circuit. The algorithm starts at a primary-output and recursively expands its fanin-cone, generating code for each gate (if it has not already been evaluated) as it returns from the recursion. Since cache misses are likely to result on any signals that fanout from this block of gates to other blocks, we attempt to choose the next primary-output from a set of candidates that uses some of the current set of unresolved fanout paths.

The only problem with this technique is that the recursive traversal encourages the scheduling of gates followed immediately by gates they fan out to, resulting in read-after-write (RAW) pipeline hazards that cause frequent stalls. To correct this problem the compiler tries to ensure that there is at least one unrelated gate scheduled between two gates connected by a common signal. The effectiveness of these two optimizations over the simple level-order traversal is shown in Fig. 14.

C. Multiprocessor Systems and System Scalability

With careful partitioning, large digital circuits can be simulated in parallel with minimal interprocessor communication and synchronization. Indeed, many of the faster logic simulation hardware accelerators use parallel techniques [1], [2], [9], [15], [17], [33], [43]. Ravel-XL was designed with support for multi-processing in mind. Multiple Ravel-XL boards can be placed on a single backplane and one design partitioned among these boards. Synchronization can be handled in one of two ways. If circuits are partitioned only at synchronizer boundaries, communication among the different boards is necessary only at the beginning of each clock phase when new input vectors are loaded. If a circuit must be broken between synchronizers, however, WAIT instructions can be placed in the code at required synchronization points. In these configurations communication occurs only through the backplane and is managed by the host. This creates an obvious bottleneck, but is a cheaper alternative to the complex crossbar interconnection system found in many other multiprocessor systems.

VIII. CONCLUSIONS

In this paper we described Ravel-XL, a hardware accelerator for levelized-code (LC) digital logic gate simulation. An architecture was developed to implement the Ravel LC simulation algorithm in hardware, and a single-chip VLSI implementation was presented. The Ravel algorithm adopted a unique waveform model that allows timing information to be calculated during the levelized traversal of the traditional LC simulation process. This eliminates one of the serious limitations of LC techniques when compared with event-driven (ED) algorithms, namely, the inability to perform accurate timing simulation. Ravel-XL, by implementing the Ravel algorithm in hardware, is able to pipeline the gate simulation process and take advantage of the parallelism available in the code to provide a significant speedup over Ravel running on a general-purpose computer. Further efficiency is gained by customizing the design of the memory system to prevent simulation speed degradation when simulating large circuits.

This implementation is capable of executing an order of magnitude faster than the Ravel algorithm running in software on a general purpose computer, and two orders of magnitude faster than Verilog-XL, an ED simulator, when simulating large circuits with high event-activities. In a single-board configuration, the simulation speed of Ravel-XL is also competitive with those of several other commercial and research hardware accelerators, and its simple highly-integrated implementation should give it a significant price/performance advantage. Ravel-XL is also easily scalable to multi-board parallel simulation configurations, and should be capable of simulating at speeds comparable to those of other parallel simulation accelerators such as YSE, EVE, and the Zycad XP.

Work is still in progress on the simulation front-end software and code compilation and optimization software. An important goal in the project is to prevent the need for expensive code pre-processing. Large circuits will require some optimizations in scheduling the code to prevent data-cache misses, but preliminary work suggests that simple algorithms will be sufficient in most cases. Work is also continuing on the problem of circuit partitioning to minimize the connectivity of circuit blocks split over different processors in a multiboard parallel Ravel-XL configuration. In addition, we are examining improvements to the architecture based on experience gained during the current implementation.

ACKNOWLEDGMENT

The authors would like to thank J. Bell for his work on the Ravel-XL compiler, and also the anonymous reviewers for their helpful suggestions and constructive criticism.

REFERENCES

- [1] P. Agrawal, W. J. Dally, W. C. Fischer, H. V. Jagadish, A. S. Krishnakumar, and R. Tutundjain, "MARS: A multiprocessor-based programmable accelerator," *IEEE Design & Test of Computers*, pp. 28-37, Feb. 1987.
- [2] P. Agrawal and W. J. Dally, "A hardware logic simulation system," *IEEE Trans. Computer-Aided Design*, pp. 19-29, Jan. 1990.
- [3] Z. Barzilai, J. L. Carter, B. K. Rosen, and J. D. Rutledge, "HSS—A high-speed simulator," *IEEE Trans. Computer-Aided Design*, pp. 601-617, July 1987.
- [4] T. Blank, "A survey of hardware accelerators used in computer-aided design," *IEEE Design Test*, pp. 21-38, Aug. 1984.
- [5] R. K. Brayton, R. Rudell, A. Sangiovanni-Vincentelli, and A. R. Wang, "MIS: A multiple-level logic optimization system," *IEEE Trans. Computer-Aided Design*, pp. 1062-1081, Nov. 1987.
- [6] M. Breuer and A. Friedman, *Diagnosis and Reliable Design of Digital Systems*. Woodland Hills, CA: Computer Science, 1976.
- [7] F. Brglez, D. Bryan, and K. Kozminski, "Combinational profiles of sequential benchmark circuits," in *Proc. ISCAS89*, 1989.
- [8] M. C. Browne, E. M. Clarke, D. L. Dill, and B. Mishra, "Automatic verification of sequential circuits using temporal logic," *IEEE Trans. Comput.*, vol. C-35, pp. 1035-1044, Dec. 1986.
- [9] T. Burggraaf, A. Love, R. Malm, and A. Rudy, "The IBM Los Gatos logic simulation machine hardware," in *Proc. IEEE Int. Conf. Comput. Design*, Oct. 1983, pp. 584-587.
- [10] T. M. Burks and K. A. Sakallah, "Min-max linear programming and the timing analysis of digital circuits," in *Proc. Int. Conf. Computer-Aided Design*, 1993, pp. 152-155.
- [11] *Verilog-XL Reference Manual*, Cadence Design Systems Inc., version 1.6, 1991.
- [12] *EPOCH Designers Handbook*, Cascade Design Automation Corp., EDH-1.0Beta, 1992.
- [13] *VCS Reference Manual*, Chronologic Simulation, version 2.0, 1993.
- [14] J. Crapuchettes, "TURBOchannel interface asic functional specification, revision 0.6 (preliminary)," Digital Equipment Corp., TRI/ADD Program, Aug. 31, 1992.
- [15] M. Denneau, "The Yorktown simulation engine," in *Proc. 19th ACM/IEEE Design Automation Conf.*, June 1992, pp. 55-59.
- [16] Digital Equipment Corporation, "TURBOchannel specifications-version 2C," Digital Equipment Corp., TRI/ADD Program, EK-TCDEV-DK-004, Sept. 1991.
- [17] L. N. Dunn, "IBM's engineering design system support for VLSI design and verification," *IEEE Design & Test Comput.*, pp. 30-40, Feb. 1984.
- [18] M. Farnwald and D. Mooring, "A fast path to one memory," *IEEE Spectrum*, pp. 50-51, Oct. 1992.
- [19] J. L. Hennessy and D. A. Patterson, *Computer Architecture, A Quantitative Approach*. San Mateo, CA: Morgan Kaufmann, 1990.
- [20] G. Kane and J. Heinrich, *MIPS RISC Architecture*. Englewood Cliffs, NJ: Prentice Hall, 1992.
- [21] Y. S. Lee and P. M. Maurer, "Two new techniques for compiled multi-delay logic simulation," in *Proc. 29th Design Automation Conf.*, 1992, pp. 420-423.
- [22] D. M. Lewis, "A hierarchical compiled-code event-driven logic simulator," *IEEE Trans. Computer-Aided Design*, June 1991, pp. 726-737.
- [23] —, "Performance issues in a compiled-code hardware accelerator," in *CAD Accelerators*. New York: Elsevier Science, 1991, pp. 47-59.
- [24] —, "A compiled-code hardware accelerator for circuit simulation," *IEEE Trans. Computer-Aided Design*, pp. 555-565, May 1992.
- [25] D. M. Lewis, M. H. van Ierssel, and D. H. Wong, "A field programmable accelerator for compiled-code applications," in *Proc. Int. Conf. Comput. Design (ICCD)*, 1993, pp. 491-496.
- [26] B. A. Mailoy, E. L. Lloyd, and M. L. Soffa, "Scheduling DAG's for asynchronous multiprocessor execution," *IEEE Trans. Parallel Distrib. Syst.*, vol. 5, May 1994, pp. 498-508.
- [27] P. M. Maurer, "Two new techniques for unit-delay compiled simulation," *IEEE Trans. Computer-Aided Design*, vol. 11, pp. 1120-1130, Sept. 1992.
- [28] P. M. Maurer and Y. S. Lee, "Gateways: A technique for adding event-driven behavior to compiled simulations," *IEEE Trans. Computer-Aided Design*, pp. 338-352, Mar. 1994.
- [29] A. N. Parlakbilek and D. M. Lewis, "A multiple-strength multiple-delay compiled-code logic simulator," *IEEE Trans. Computer-Aided Design Integr. Circuits and Syst.*, vol. 12, pp. 1937-1946, Dec. 1993.
- [30] *Enterprise Emulation System User's Guide*, Quickturn Syst., Inc., 1991.
- [31] M. Riepe and K. Sakallah, "Delay accurate compiled-code synchronous gate-level Verilog simulation," in *Proc. 2nd International Verilog HDL Conference*, March 1993, pp. 121-127.
- [32] M. A. Riepe, J. L. Bell, E. J. Shriver, and K. A. Sakallah, "Assigned-delay compiled-code multiphase synchronous logic simulation," (in preparation).
- [33] M. Saitoh, K. Iwata, A. Nakamura, M. Kakegawa, J. Masuda, H. Hamamura, F. Hirose, and N. Kawato, "Logic simulation system using simulation processor (SP)," in *Proc. 25th ACM/IEEE Design Automation Conf.*, 1988, pp. 225-230.
- [34] K. A. Sakallah, T. N. Mudge, and O. A. Olukotun, "check T_c and min T_c : Timing verification and optimal clocking of synchronous digital circuits," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 1990, pp. 552-555.

- [35] —, "Analysis and design of latch-controlled synchronous digital circuits," *IEEE Trans. Computer-Aided Design Integr. Circuits Syst.*, vol. 11, pp. 322–333, Mar. 1992.
- [36] K. A. Sakallah, T. N. Mudge, T. M. Burks, and E. S. Davidson, "Synchronization of pipelines," *IEEE Trans. Computer-Aided Design Integr. Circuits Syst.*, vol. 12, pp. 1132–1146, Aug. 1993.
- [37] E. Shriver and K. Sakallah, "Ravel: Assigned-delay compiled code logic simulation," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 1992, pp. 364–368.
- [38] R. Tarjan, "Depth-first search and linear graph algorithms," in *Proc. SIAM J. Comput.* vol. 1, no. 2, June 1972, pp. 146–160.
- [39] E. Ulrich, "Exclusive simulation of activity in digital networks," *Commun. ACM*, vol. 12, no. 2, pp. 102–110, Feb. 1969.
- [40] L. Wang, N. E. Hoover, E. H. Porter, and J. Zasio, "SSIM: A software leveled compiled-code simulator," in *Proc. 24th ACM/IEEE Design Automation Conf.*, 1987, pp. 2–8.
- [41] Z. Wang and P. M. Maurer, "LECSIM: A leveled event driven compiled logic simulator," in *Proc. 27th ACM/IEEE Design Automation Conf.*, 1990, pp. 491–496.
- [42] J. Zasio and P. Hwang, "A low-cost high-performance leveled compiled-code simulation accelerator," in *Hardware Accelerators for Electrical CAD*. New York: IOP Publishing, 1988, pp. 46–56.
- [43] Zycad Corp., "The XP product family," marketing literature.



Michael A. Riepe received the B.S. degree in computer engineering (with highest honors) from the University of California, Santa Cruz in 1991, and the M.S. degree in electrical engineering from the University of Michigan, Ann Arbor in 1993. He is currently working towards the Ph.D. degree in electrical engineering from the University of Michigan.

His research interests include digital GaAs VLSI integrated circuit design, CAD for high-performance VLSI layout and simulation, and computer architecture.



João P. Marques Silva received the Eng. and Master degrees in electrical and computer engineering in 1988 and 1991, respectively, from the Instituto Superior Técnico at the Technical University of Lisbon. He received the Ph.D. degree in electrical engineering from the University of Michigan, Ann Arbor in 1995.

He is an Assistant Professor at the Instituto Superior Técnico (IST), Lisbon, Portugal, and a Researcher at the Instituto de Engenharia de Sistemas e Computadores (INESC). His research interests

include design and analysis of algorithms and CAD for integrated circuits and systems.



Karem A. Sakallah (S'78–M'80–SM'92) received the B.E. degree (with distinction) in electrical engineering from the American University of Beirut, Beirut, Lebanon, in 1975, and the M.S. and Ph.D. degrees electrical and computer engineering from Carnegie Mellon University (CMU), Pittsburgh, PA, in 1977 and 1981, respectively.

In 1981, he joined the Department of Electrical Engineering at CMU as a Visiting Assistant Professor. From 1982 to 1988, he was with the Semiconducting Engineering Computer-Aided Design Group at Digital Equipment Corporation, Hudson, MA, where he headed the Analysis and Simulation Advanced Development team. Since September 1988, he has been an Associate Professor of Electrical Engineering and Computer Science with the University of Michigan, Ann Arbor. His research interests are primarily in the area of computer-aided design of integrated circuits and systems, with particular emphasis on numerical analysis, timing verification and optimal clocking, multilevel simulation, modeling, knowledge abstraction, and design environments.

Dr. Sakallah has served on the technical program committees of all major CAD conferences and is currently an Associate Editor of the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS. He is a member of the Association for Computing Machinery.



Richard B. Brown (S'74–M'76–SM'91) received the B.S. (with highest honors) and M.S. degrees in electrical engineering (computer emphasis) from Brigham Young University in 1976. He received the Ph.D. degree in electrical engineering (with concentration in solid state and VLSI) from the University of Utah in 1985. His dissertation work included development of a custom MOS fabrication process and integration of digital and analog circuitry to form a novel solid-state chemical sensor.

From 1976 to 1981, he worked in computer design as Vice-President of Engineering at Holman Industries, Oakdale, CA, and then as Manager of Computer Development at Cardinal Industries, Webb City, MO. In September 1985, he joined the faculty of the Department of Electrical Engineering and Computer Science at University of Michigan. He has been involved in establishing the VLSI graduate program of the University of Michigan and introducing a uniform set of electronic CAD tools into the curriculum. He has taught Introduction to Semiconductor Device Theory, Solid-State Devices, and Digital Electronics, and both introductory and advanced VLSI design courses. Since 1987, he has done research in VLSI digital GaAs circuits and high-performance computing systems. His group has optimized processor architectures for GaAs SRAM's and high-performance GaAs microprocessors, and developed CAD tools which provide general support of high-speed circuit design and automatic hardware description language to layout compilation of GaAs DCFL circuits. He holds five patents and consults in the areas of solid-state sensors and circuits, and electronic design automation tools.