

Using Evolutionary Algorithms for Signal Integrity Checks of High-Speed Data Buses

Maynard C. Falconer
Kiran Kamisetty
Adam Norman
Konika Ganguly
Intel Corporation

Platform & System Technology Group
Hillsboro, OR 97124
Email: maynard.c.falconer@intel.com

Kristina Morgan
Department of Electrical and
Computer Engineering
University of Texas
Austin, TX 78712-0240

Email: kdmorgan10@mail.utexas.edu

Garrison W. Greenwood
Department of Electrical and
Computer Engineering
Portland State University
Portland, OR 97207-0751
Email: greenwd@ece.pdx.edu

Abstract—Today’s high performance computer systems must have fast, reliable access to memory and I/O devices. Unfortunately, inter-symbol interference, transmission line effects and other noise sources can distort data transfers. Engineers must therefore determine if bus designs have signal integrity—i.e., the bus can transfer data with minimal amplitude or timing distortion. One method of determining signal quality on buses is to conduct a set of data transfers and measure various signal parameters at the receiver end. But the tests must be conducted with stressful test patterns that maximize inter-symbol interference to help identify any potential problems. In this paper we describe how an evolutionary algorithm was used to evolve such test patterns. All test results were obtained intrinsically.

I. INTRODUCTION

Evolvable hardware uses evolutionary algorithms to perform physical design. During each generation of the evolutionary algorithm hardware designs are created, tested and either kept or discarded depending on their demonstrated behavior. This process continues, generation after generation, until an acceptable design is produced. The behavior of each design can be found intrinsically, where all tests are conducted on a physical implementation of the design, or extrinsically where all tests are conducted with simulators.

A survey of evolvable hardware literature shows most researchers usually concentrate only on creating designs that exhibit a desired functional behavior. Typically the researcher defines a set of test inputs and the actual hardware response for each test input is compared against the expected response. All “errors” are recorded and the hardware design without any errors is the optimal design. But correct functional behavior is not the only design criteria. With higher operating speeds, lower operating voltages and smaller noise margins, today’s digital designers must worry about *signal integrity* (SI) on the data buses connecting computer system components. In other words, designers must verify that all data bus signals are of sufficient quality to guarantee reliable data transfers. Note that SI is independent of whether or not the computer system processed the data correctly; it only considers whether or not the data transfer itself was correct.

SI must never be taken lightly. Degraded signals usually only corrupt data, although in pathological cases they can cause a system crash [2]. Engineers must run extensive SI checks to stress the bus design before manufacturing starts making printed circuit boards. Stress tests are conducted intrinsically. They involve transferring data signals, called *stressful test patterns*, down the bus at full speed and then monitoring the signal quality at the receiving end. In an ideal case these stressful test patterns will induce SI problems engineers can analyze and then forestall by modifying the bus design and/or providing specific layout guidelines. But therein lies the problem: no one knows exactly what a stressful test pattern should look like and there is no deterministic method for generating them. We have successfully used evolutionary algorithms to evolve stressful test patterns for a high-speed data bus. In this paper we describe how that was done.

Our results were obtained on external data buses between chips, but they could be used for buses inside chips just as easily. This is important because system-on-a-chip designs must also worry about SI [7]. The next section provides background material on SI and some needed definitions. (A thorough treatment of SI can be found in [6].) In Section III we present preliminary results from intrinsic SI testing of an Intel chipset.

II. BACKGROUND

A. Front Side Bus

The front side bus (FSB) is a physical bi-directional bus that connects a CPU to the chipset for communication with the rest of the computer system devices such as DRAM, video cards and so forth. In particular, the FSB serves as a backbone between a chipset and the CPU. This chipset serves as a connection point for all other buses in the computer system.

Normally the CPU operates at some multiple of the FSB speed. For example, a 100 MHz FSB might have the processor running at 550 MHz, which corresponds to a 5.5 bus/core ratio. High data transfer rates are achieved by transferring data multiple times per clock cycle. For instance, by transferring

two bytes at a time, on both the rising and falling clock edges, effectively gives a four times throughput over a 100 MHz FSB.

B. Signal Integrity

The bandwidth demanded from busses, such as the FSB, has been roughly doubling every two years. Increased switching speeds, combined with ever increasing chip density, has made the characterization of interconnects crucial to the proper operation of computer systems. For example, variations in transistor behavior become more acute as devices get smaller. Circuitry can be added to receivers to compensate for these variations, but the additional circuitry increases the receiver input capacitance which, in turn, affects the interconnect behavior. Greater bandwidth requirements are often met by a combination of transferring data in parallel and using higher transfer rates. Unfortunately, transmitting data in parallel introduces crosstalk between data lines, which can corrupt the data transfer. In fact, crosstalk delay is a major contributor to timing uncertainties [1]. Additionally, operating voltages are steadily moving downward. While lowering voltage levels may yield lower crosstalk, it also makes it more difficult to maintain adequate noise margins.

Every trace on a printed circuit board distorts pulses sent down it, sometimes to a point where the receiver cannot unambiguously determine what information was transmitted. The degradation depends on a number of factors including trace length, the type of material used, the connectors used, transmission line effects (reflections and so forth) and the limited bandwidth of drivers and receivers. Both amplitude reduction and stretching of the pulse's rising and falling edges can be expected.




Fig. 1. Diagram showing how a pulse can degrade during transfer down a printed circuit board trace. The pulse on the right, sometimes called a “runt pulse” has a lower amplitude and extended rise and fall times.

Voltage margin at the receiver end is an important measure of signal quality. It is defined as the difference (in mV) between the amplitude of a signal at the receiver at the sample time and the receiver's threshold voltage. Extremely small margins are sufficient for reliable data transfers in ideal systems where there is no noise. However, in the real-world systems are not ideal; the margin must be large enough to counter crosstalk, any system noise, measurement errors and the limited receiver bandwidth.

Engineers often use eye diagrams to get a qualitative picture of SI in a communications channel. (See Figure 2.) These type of diagrams are created by monitoring the receiver input with a sampling oscilloscope. Bits sent over the channel are aligned to a common timing reference and then recorded.

The persistence of the display forms a composite image. The diagram is fuzzy because noise and transmission line effects cause data distortions; the vertical thickness is caused by amplitude fluctuations whereas the horizontal thickness comes from timing variations. The size of the opening gives a visual representation of the timing and voltage margins.

An eye diagram provides a quick check if the data meets the margin specifications. A rectangle, with the width equal to the timing margin and the height equal the voltage margin can be superimposed over the eye diagram. The signal has acceptable margins if the rectangle fits inside the eye opening.

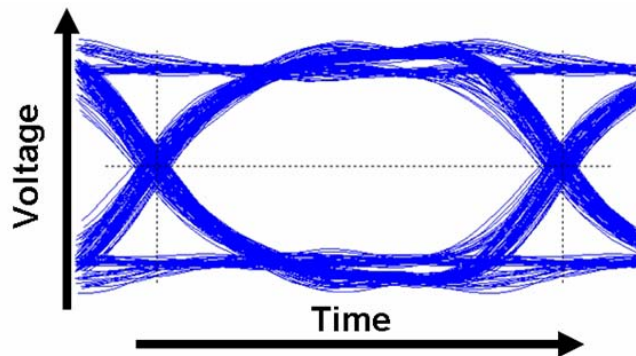


Fig. 2. A typical eye diagram.

SI testing and analysis determines whether noise, transmission line effects or limited driver and receiver bandwidths are likely to cause data transfer errors in a given bus design. The end product of SI testing is a set of guidelines engineering can provide to manufacturing to help prevent potential problems.

C. The Role of Evolutionary Algorithms in SI testing

Today's high demands on interconnect bandwidth force individual bus lines to lie within extremely close proximity of each other with very small rise times to achieve the increased transfer rate. Transmission line effects, primarily reflections, can actually change the operating margins on bus lines. A compact bus layout makes *inter-symbol interference* (ISI) very problematic. For example, ISI can cause a bit transmitted at time t_0 to affect a bit transmitted at some later time $t_0 + \tau$. The number of bits that can be transmitted, and therefore affected, between time t_0 and time $t_0 + \tau$ is referred to as the *ISI depth*. If we assume the ISI depth is 10, then each test pattern sent down a bus line is ten bits long. Hence, there are $2^{10} = 1024$ possible test patterns per line. At nominal FSB transmission rates of 1 Giga-transfers/s (1 GT/s = 1 Gbyte/s/pin), an exhaustive test takes less than 1 ms.

If there is no interline ISI then all 64 bus lines can be validated in parallel. However, in most cases assuming there is no interline ISI is naive; even moderate interaction levels lead to a combinatorial explosion. For example, suppose a given bus line is susceptible to ISI effects from four other neighboring bus lines. For this set of five lines an exhaustive test now requires 32^{10} (over one trillion) 10-bit patterns! If that isn't bad enough, power supply lines can also introduce

interference and in-die process variations further complicate matters. Don't forget we only considered five bus lines in the above analysis when the typical FSB has 64 data lines.

Clearly exhaustive testing is impractical, so only a subset of possible test patterns can be used. Nevertheless, it is important for interconnect validation to exercise the system under test with stimulus patterns that stress the operating margin of the interconnect to the greatest possible extent. The difference in margin between a benign and a stressful stimulus pattern for a FSB is typically greater than 100 mV, or in terms of errors, the difference between an interconnect that creates errors on the order of seconds versus a system that has a time between errors measured in the tens of thousands of years. Deduction and historical results are the most common ways of choosing stressful stimulus patterns, but those methods are tedious and do not always produce good test sets. This is where evolutionary computation methods are beneficial; we can evolve the test patterns needed to conduct SI testing.

III. INTRINSIC SI TESTING

In this section we describe how an EA was used to evolve stressful test patterns for a FSB. The system under test consisted of a desktop PC system with an Intel 965 CPU and an Intel[®] 975X Express Chipset [4] with the FSB operating at 800 MT/s (= 800 Mbyte/s/pin). The 975X chipset has a memory controller hub (MCH) backbone architecture to support faster memory accesses and increased graphics requirements [5]. Figure 3 shows the 975X chipset block diagram. Note the FSB connecting the MCH with processor.




Fig. 3. Intel[®] 975X Express Chipset Block Diagram

One way to validate a bus design is to show it won't fail after conducting extensive testing by transferring test patterns at full bus speeds. In our method an EA tries to create

stressful test patterns which will maximize potential ISI effects and thereby maximize the probability of a FSB data transfer failure. The failure probability is tied to the signal's voltage margin, which is defined as the difference between the receiver threshold voltage level and the signal voltage level at the time of sampling. The FSB design is considered good if the most stressful pattern possible was unable to induce a bus failure (with a guardband incorporated into the analysis). Of course the difficulty lies in determining what the most stressful possible pattern looks like.

The communication channel itself is a single ended transmission line, where the voltage margin to the high signaling state is designated HIGHSIDE and to the low signaling state designated LOWSIDE. The test begins with a test pattern sent down the communications channel from the driver to the receiver. At the receiver the received pattern is compared against the expected pattern and any errors are recorded.

Margin tests are conducted for both the HIGHSIDE and the LOWSIDE, but for brevity only the LOWSIDE test is described. The comparisons between received and expected signals are conducted with respect to a reference voltage (V_{ref}), which helps to determine if a given signal is a 1 or 0. (V_{ref} is a voltage threshold at the receiver. Any signal with a voltage $V > V_{ref}$ is high, whereas any voltage $V \leq V_{ref}$ is low.)

Let V_{nom} be a nominal voltage, which is typically 0.66% of the bus supply voltage V_{tt} . Initially V_{ref} is set equal to V_{nom} and a test pattern is transferred. If no transfer error occurred at that voltage, V_{ref} is decreased by some small ΔV and another data transfer and comparison takes place. This process repeats until at some V_{ref} the pattern comparison fails. The margin is the difference between V_{nom} and the V_{ref} at which the comparison failed. Hence, each decrease moves V_{ref} further from V_{nom} , which increases the likelihood of a transmission error¹. Fitness is based on the millivolts of margin at the receiver, and since we are searching for the most stressful pattern, lower margins denote higher fitness. The HIGHSIDE tests are conducted in the same manner, except V_{ref} starts close to V_{nom} and increases towards V_{tt} .

At first glance making fitness inversely proportional to the voltage margin may seem backwards. Indeed, under this definition the lowest fit patterns are those that didn't cause any failures at any V_{ref} value. But remember our objective is to evolve stressful test patterns that could cause FSB data transfer errors. We want to find patterns with low voltage margins because they are the ones most likely to have ISI problems.

The EA uses a population of 320 individuals with the initial population constructed from historical information (see below). Each test pattern is 64 bits wide, to span the FSB data bus, and is 25 bits deep for an ISI depth. Hence, each individual's size is $64 \times 25 = 1600$ bits. Only the 160 best individuals from the previous generation are used for parents. Half of the offspring in the next are created by just

¹Under normal operating conditions $V_{ref} = V_{nom}$. We only move V_{ref} away from V_{nom} to determine the voltage margin.

copying the parents. The other half are created by copying the parents and then applying 4-pt crossover. All offspring are then subjected to mutation, which is bit-flipping applied with a 5% probability. Fitness, as stated above, is the voltage margin.

The initial population is not randomly generated, as is done in most EAs, but rather was constructed from historical data. These initial patterns were designed to stimulate different resonant frequencies and exercise odd and even crosstalk events. A small number were selected based on previous simulation runs or intuition resulting from an inspection of the physical bus layout.

Table I shows the tests conducted. All test runs were intrinsic, conducted on physical hardware. Tests are typically conducted for both the HIGHSIDE and the LOWSIDE, even for the same driver/receiver pair, because SI effects are not necessarily symmetrical. That is, it may turn out that the voltage margins for a HIGHSIDE test are not the same as the LOWSIDE test. (As a side note, V_{ref} is usually set with a resistor divider network. Engineering may adjust the resistor values to change V_{ref} to a value that equalizes the voltage margins. This is an example of how SI testing influences the physical design of a bus.)

The EA typically converged in 40 to 50 generations, which required around 12 hours of test time. In all cases a stressful input pattern with a significantly lower margin was found. For example, from the first row in the table the worst-case V_{ref} in the initial population was 560 mV, which for $V_{nom} = 800$ mV corresponds to a 240 mV margin. The EA found a test pattern with a $V_{ref} = 592$ mV. This evolved test pattern yields a voltage margin of 208 mV, which is a 32 mV reduction. Not only does this evolved test pattern exhibit ISI effects, but the associated lower voltage margin leaves the system more susceptible to other noise sources; it therefore is a good stimulus to use for investigating a variety of potential SI problems.

It is important readers fully appreciate what has been accomplished here. Engineers had estimated the amount of noise inherent in a high-speed bus design using a set of test patterns derived from historical data and previous experience. However, an EA found stressful test patterns intrinsic testing showed would generate even higher noise levels than originally anticipated. In the worst case this higher noise level will require printed circuit board layout changes or possibly even changes to the circuit design itself [3]. It is also worth noting a 32 mV reduction in the voltage margin is a significant amount, roughly corresponding to the difference between a 1 second mean time to errors and a 1.6 week mean time to errors.

Figures 4 and 5 show the best EA run for one of the LOWSIDE and HIGHSIDE tests respectively. In the LOWSIDE plot $V_{nom} = 800$ mV so the difference between the plotted points and V_{nom} —i.e., the voltage margin—is also directly proportional to the fitness. Hence, the plot also depicts fitness vs generation number. The plot does not rise monotonically because the EA was generational with no elitism. Of course the best fit solution found in any generation could be recorded off-line and kept from participating in any future reproduction

DVR/RCVR	Test	Worst-case V_{ref} Init. Pop. (mV)	Worst-case V_{ref} Final Pop. (mV)	ΔV_{ref} (mV)
MCH/CPU	L	560	592	32
CPU/MCH	L	544	584	40
CPU/MCH	H	976	936	40

TABLE I

SI TESTS ON INTEL CHIPSETS WHERE H (L) REPRESENTS A HIGHSIDE (LOWSIDE) TEST. IN ALL CASES $V_{TT} = 1.2V$. ΔV_{REF} IS THE CHANGE IN VOLTAGE MARGIN BETWEEN THE BEST FIT INDIVIDUAL IN THE INITIAL POPULATION AND THE BEST FIT INDIVIDUAL IN THE FINAL POPULATION.

operations. This follows what is typically done in the (μ, λ) evolution strategy.

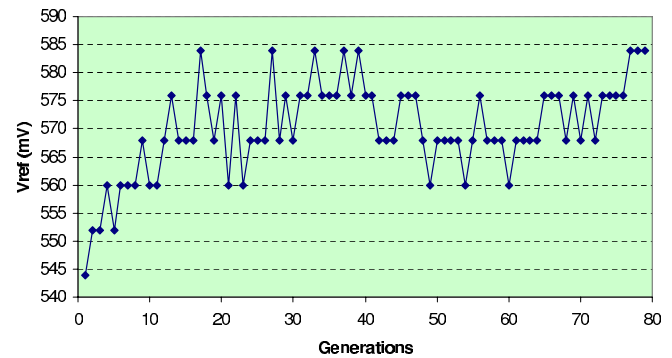


Fig. 4. V_{ref} mV vs generations for LOWSIDE tests with $V_{tt} = 1.2V$. Plot is equivalently fitness vs generations (see text)

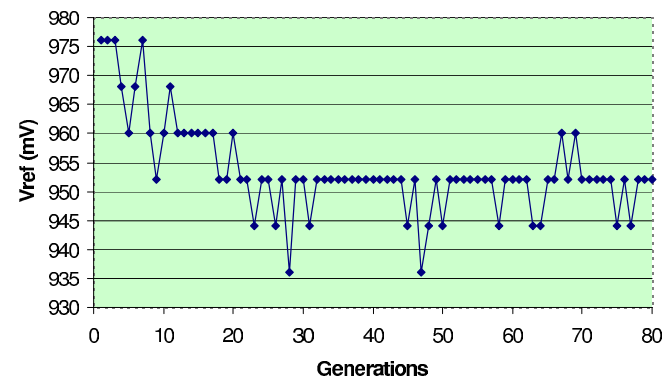


Fig. 5. V_{ref} mV vs generations for HIGHSIDE tests with $V_{tt} = 1.2V$.

IV. FINAL REMARKS

In this paper we described how an EA could be used to test hardware. Up to this point the evolvable hardware community has largely ignored fitness evaluation for items other than functional behavior. Today's complex designs—to say nothing about tomorrow's designs—cannot be incorporated into product lines without having undergone a full suite of parametric testing. Noise margins, voltage margins and other factors must be quantified with rigorous tests to certify SI in the design. Indeed, the design cannot be considered finished until

engineering has completed the SI analysis. Moreover, these test results ultimately influence the physical bus design itself. EA techniques are a powerful tool for design verification. The evolvable hardware community must start developing evolutionary computation methods to support hardware testing that goes beyond just functionality checks.

REFERENCES

- [1] F. Caignet, S. Bendhia, and E. Sicard. The challenge of signal integrity in deep-submicrometer CMOS technology. *Proceedings of the IEEE*, 89(4):556–573, 2001.
- [2] G. Edlund. Anatomy of a signal integrity failure. In *DesignCon 2003 Proceedings*. Int'l Engineering Consortium, 2003.
- [3] L. Green. Understanding the importance of signal integrity. *IEEE Circuits & Devices Magazine*, 15(6):7–10, 1999.
- [4] Intel Corporation. *Intel[®] 975X Express Chipset Data Sheet*, November 2005. Document Number: 310158-001.
- [5] Intel Corporation. *Intel[®] Pentium[®] D Processor Data Sheet*, September 2006. Document Number: 310306-006.
- [6] H. Johnson and M. Graham. *High Speed Signal Propagation: Advanced Black Magic*. Prentice-Hall, 2003.
- [7] Mehrdad Nourani and Amir Attarha. Built-in self-test for signal integrity. In *Design Automation Conference*, pages 792–797, 2001.