A Method for Eliminating Skew Introduced by Non-uniform Buffer Delay and Wire Lengths in Clock Distribution Trees.

Henry M. Wu

Artificial Intelligence Laboratory
and
Department of Electrical Engineering and Computer Science

Massachusetts Institute of Technology

A.I. Memo No. 1422 April, 1993

Abstract

Non-uniformities in buffer delays and wire lengths introduce skew in clock distribution trees. Previous techniques exist for eliminating skew introduced by each of these causes, not both. This method uses a pair of matched variable delay lines to eliminate skew caused both by differing buffer delays and wire lengths.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-92-J-4097.

Introduction

Timing skew in clock distribution trees is caused both by non-uniformities in buffer propagation delay and wire lengths. Many commercial parts exist for solving the problem of non-uniform buffer delays. They work by regenerating the clock signal at the redistribution point with a phase-locked loop (PLL) so that the amplified signal is in phase with the reference signal received. This approach ignores the delay introduced by the wire used to deliver the clock signal to its eventual destination.

[Knight 92][3] describes a method for compensating for wire length delays. While the technique effectively eliminates skew caused by wire delay, it ignores skew caused by the signal buffers.

The following technique eliminates skew caused both by non-uniform buffer delays and difffering wire lengths. As in the technique introduced in [Knight 92], a pair of matched variable delay lines is used to derive the one-way travel time from a measured round-trip delay. The method can be implemented most obviously using two wires from the source to the destination, but can also be modified to require only one wire.

The Technique

Many low-skew clock redistribution buffers work as depicted in Figure 1. A reference clock signal is received by the part. Instead of amplifying the signal, and hence suffering propagation delays that are hard to control due to part-to-part and temperature variations, a local voltage-controlled oscillator (VCO) is employed to generate a new copy of the clock signal which is then amplified. The phase of the amplified copy is measured against that of the received reference, and the VCO adjusted to correct for any errors. In essence, the part functions as a phase-locked loop with the amplified signal tracking the frequency and phase of the reference input.

The problem with this scheme is that the amplified clock must then be transmitted to its eventual receiver over wire with delay. This delay is not compensated for in the control loop of the PLL. This problem can be eliminated if the wire delay is included in the signal that is compared by the phase detector with the reference.

Figure 1: Phase-locked loop redistribution of clock signal. The PLL compares the signal at the edge of the buffer, not the destination.

A circuit as shown in Figure 2 can be used. The amplified clock is fed into a pair of matched variable delay lines. The signal is also sent to the receiver, and brought back along a path with the same electrical length as the forward path. The phase detector PD1 compares these delayed signals, and the delay lines adjusted in tandem until the two delays are the same, namely $2T_{line} + T_{pd1}$ or $2T_{delay} + T_{pd2}$. Assuming T_{pd1} can be made close to T_{pd2}, $T_{delay} = T_{line}$ and the signal at point S represents the same timing as the signal received at point R. Phase detector PD2 compares the signal at S against the reference and adjusts the phase of the amplified clock until the signal at S, and hence the signal at R, is the same in both frequency and phase as the reference.

Figure 2: The two-wire implementation of the technique. The arrival time of the amplified clock signal at the receiver is derived by the matched delay lines and sent to the PLL.

As in [Knight 92], the return path can be eliminated by measuring transmission line effects. If the receiver has high impedance compared to the characteristic impedance of the line, a reflected wave of the same sign as the outgoing wave will appear at the driver after one round-trip delay. Series termination at the driver allows observation of the reflection and prevents further bounces. If the series termination resistance is exactly the impedance of the wire, then the voltage at the wire end of the termination resistor doubles when the reflected wave returns. The circuit is as shown in Figure 3.

Figure 3: Using transmission line reflection to measure round-trip delay. V_{ref} is adjusted so that the comparator triggers on the return bounce of the signal.

Conclusions

The described technique compensates for non-uniform buffer delays and wire lengths to allow skew-free redistribution of clock signals. The technique utilizes a pair of matched variable delay lines to derive the wire delay and incorporates that into the control loop of a PLL clock generator.

References

[1] Hans J. Greub. *Apparatus for Skew Compensating Signals.* United States Patent 4,833,695. 1989.

[2] Mark G. Johnson. A Variable Delay Line Phase Locked Loop for CPU-Coprocessor Synchronization. In *IEEE International Solid-State Circuits Conference*, 1988.

[3] Thomas Knight and Henry Wu. *A Method for Skew-free Distribution of Digital Signals Using Matched Variable Delay Lines.* Memo 1282, MIT Artificial Intelligence Laboratory. 1992.

[4] Ian A. Young, Jeff K. Greason, Jeff E. Smith, and Keng L. Wong. A PLL Clock Generator with 5 to 110MHz Lock Range for Microprocessors. In *IEEE International Solid-State Circuits Conference*, 1992.