Clock Distribution in High-Performance PCs

SCAA030A
February 1997

Texas Instruments

IMPORTANT NOTICE

Texas Instruments (TI) reserves the right to make changes to its products or to discontinue any semiconductor product or service without notice, and advises its customers to obtain the latest version of relevant information to verify, before placing orders, that the information being relied on is current.

TI warrants performance of its semiconductor products and related software to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Certain applications using semiconductor products may involve potential risks of death, personal injury, or severe property or environmental damage ("Critical Applications").

TI SEMICONDUCTOR PRODUCTS ARE NOT DESIGNED, INTENDED, AUTHORIZED, OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT APPLICATIONS, DEVICES OR SYSTEMS OR OTHER CRITICAL APPLICATIONS.

Inclusion of TI products in such applications is understood to be fully at the risk of the customer. Use of TI products in such applications requires the written approval of an appropriate TI officer. Questions concerning potential risk applications should be directed to TI through a local SC sales office.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards should be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance, customer product design, software performance, or infringement of patents or services described herein. Nor does TI warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used.

Contents

Introduction

Personal computer and workstation designers are pushing the operating speeds of new equipment to ever-higher frequencies with technological advances in areas such as RISC/CISC microprocessors, high-speed SRAMs, and cache memories. At higher frequencies, the timing delays and uncertainties associated with clock-signal generation and distribution in a system become critical factors in determining the system's overall performance and reliability. System performance is optimized by carefully considering the attributes of the components used in designing the clock circuit. The clocking network is the heart of any system. There are two main aspects to this network: clock generation and clock distribution. Clock generation is accomplished by taking the output of some source (a crystal oscillator, for example) and manipulating it to obtain pulses with a specific frequency, duty cycle, and amplitude. These signals are then fanned out to the various system components by a clock-distribution network (see Figure 1). As system speeds rise to 33, 40, or 50 MHz and clock periods grow shorter, the uncertainties of meeting setup, hold, and pulse duration requirements become critical due to a narrowing time window. A clocking system that does not fully consider these uncertainties will suffer degraded performance and reliability.

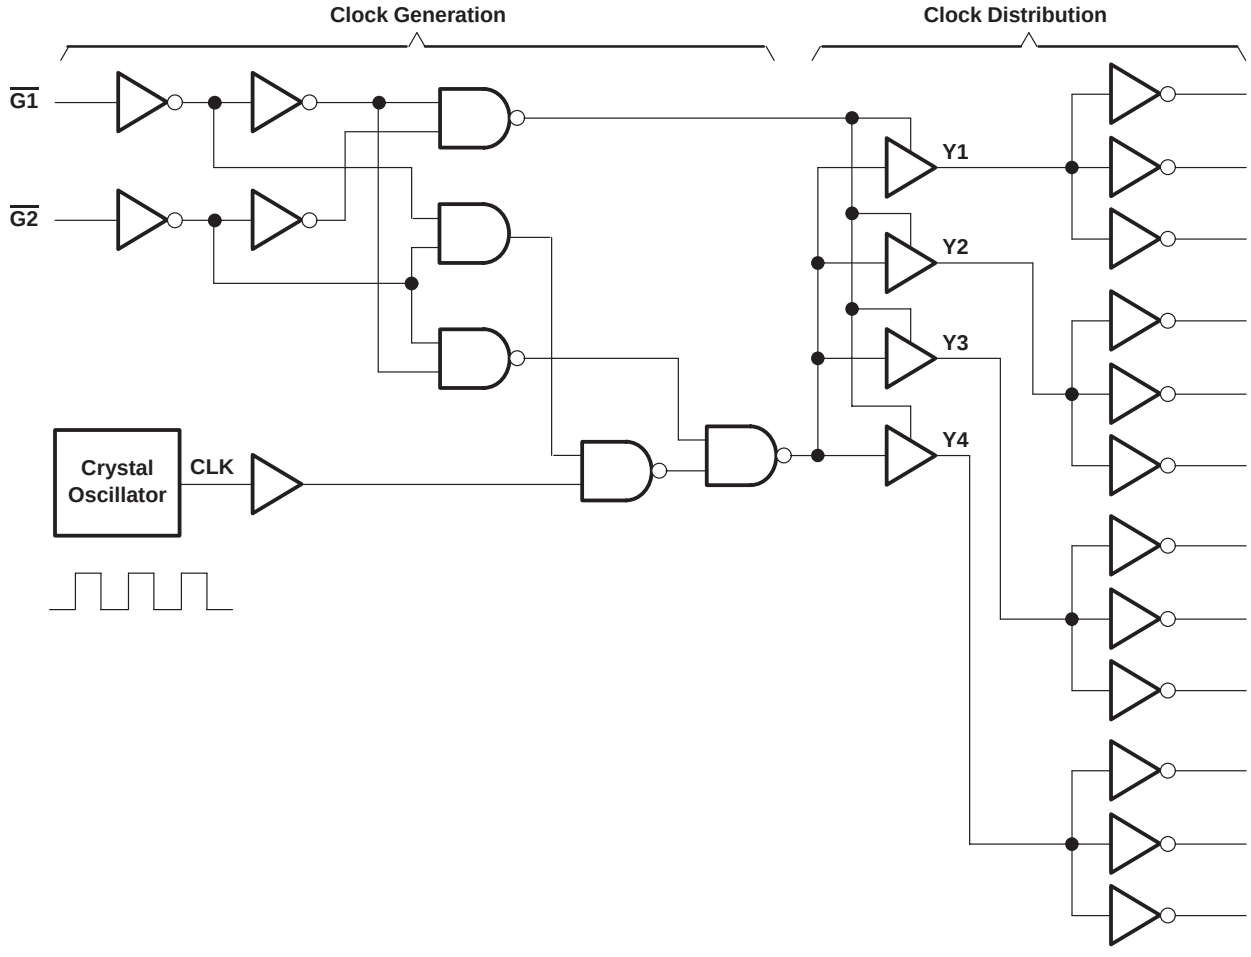

Figure 1. Clock Network

Texas Instruments (TI), as well as several of the other major integrated-circuit (IC) vendors, offers the designer a choice of clock generation and distribution circuits commonly referred to as clock drivers. Clock-driver ICs are able to provide clock generation functions (frequency multiplication, division, and duty-cycle control) as well as clock-distribution functions (buffering and fanout) with timing specifications unavailable on older CMOS and TTL logic families. Advances in process technology have allowed IC vendors to offer circuits with tight specifications on switching speeds and skew parameters. The high speed, fast edge rates, and tight skew specifications offered on clock-driver data sheets give the designer additional flexibility, but component selections must be closely tied to proper board-layout techniques. The purpose of this application report is to discuss considerations in the design of clocking networks for high-performance systems wherein proper clock generation and distribution are essential.

What Is Skew?

Any discussion of clock-driver attributes ultimately centers around skew. Simply defined, skew is the difference between the expected and actual arrival time of a clock pulse. In an ideal clock circuit, propagation delays remain fixed and equal for high-to-low and low-to-high transitions over the entire ranges of supply voltage, operating temperature, and output loading and are independent of the number of outputs switching. However, the world is not ideal, and definitions have evolved to help the designer deal with the various types of skew that can be encountered. Clock-driver performance can be described in terms of five types of skew as defined in JEDEC Standard 99, clause 2.3.5.

Output Skew

Output skew, $t_{sk(o)}$, is the difference between any two propagation-delay times when a single switching input or multiple inputs switching simultaneously cause multiple outputs to switch, as observed across all switching outputs. This parameter is used to describe the fanout capability of a clock driver and is of concern when making decisions on clock buffering and distribution networks.

Input Skew

Input skew, $t_{sk(i)}$, is the difference between any two propagation delay times that originate at different inputs and terminate at a single output (see Figure 2). Input skew describes the ability of a device to manipulate (stretch, shrink, or chop) a clock signal. This is typically accomplished with a multiple-input gate wherein one of the inputs acts as a controlling signal to pass the clock through. $t_{sk(i)}$ describes the ability of the gate to shape the pulse to the same duration, regardless of the input used as the controlling input.

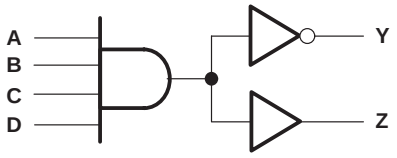

Figure 2. Example of a Gate With Input Skew

Pulse Skew

Pulse skew, $t_{sk(p)}$, is the difference between propagation delay times t_{PHL} and t_{PLH} when a single switching input causes one or more outputs to switch. $t_{sk(p)}$ quantifies the duty-cycle characteristic of a clock driver. Certain applications require a fixed duty cycle for proper operation. As an example, the CLK2 input of an MC68020 processor operating at 40 MHz requires a duty cycle of 50 ±5%. $t_{sk(p)}$ is a measure of a clock driver's ability to supply such a precisely controlled pulse.

Figure 3 defines output, input, and pulse skews.

Process Skew

Process skew, $t_{sk(pr)}$, is the difference between identically specified propagation-delay times on any two like ICs operating under identical conditions. $t_{sk(pr)}$ quantifies the skew induced by variations in the IC manufacturing process but not by variations in supply voltage, operating temperature, output loading, input edge rate, input frequency, etc. Process skew is commonly specified and production tested under fixed conditions (e.g., V_{CC} = 5.25 V, T_A = 70°C, C_L = 50 pF, all inputs switching simultaneously).

Limit Skew

Limit skew, $t_{sk(l)}$, is the difference between the greater of the maximum specified values of t_{PLH} and t_{PHL} and the lesser of the minimum specified values of t_{PLH} and t_{PHL}. Limit skew is not observed directly on a device, but is calculated from the data sheet limits for t_{PLH} and t_{PHL}. $t_{sk(l)}$ quantifies for the designer how much variation in propagation delay time will be induced by operation over the entire ranges of supply voltage, temperature, output load, and other specified operating conditions. Specified as such, $t_{sk(l)}$ also accounts for process variation. In fact, all other skew specifications ($t_{sk(o)}$, $t_{sk(i)}$, $t_{sk(p)}$, and $t_{sk(pr)}$) are subsets of $t_{sk(l)}$; they will never be greater than $t_{sk(l)}$.

In general, not all skew parameters of a device are of interest, but their discussion is included for illustration. The designer's goal is to minimize skew to an acceptably small fraction of the system clock period. A design rule of thumb is that clock skew should be less than 10% of the system clock period.

The desired operating frequency determines the designer's *skew budget*, or the maximum amount of skew allowed. For example, a system operating at 33 MHz has a period of 30.3 ns, and using the 10% rule, the allowable skew budget is 3 ns. At 50 MHz, the period is reduced to 20 ns and the permissible skew is now a scant 2 ns. Components in the clock network must be carefully selected to meet the shrinking skew budget as operating frequencies increase.

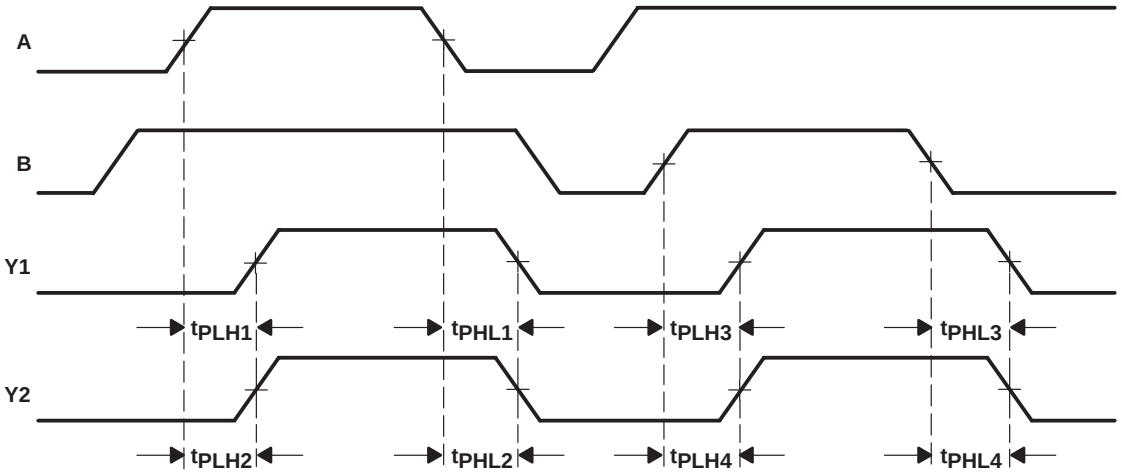

$t_{sk(o)}$ is calculated from the following expressions: $|t_{PLH1} - t_{PLH2}|$, $|t_{PHL1} - t_{PHL2}|$, $|t_{PLH3} - t_{PLH4}|$, $|t_{PHL3} - t_{PHL4}|$

$t_{sk(l)}$ is calculated from the following expressions: $|t_{PLH1} - t_{PLH3}|$, $|t_{PHL1} - t_{PHL3}|$, $|t_{PLH2} - t_{PLH4}|$, $|t_{PHL2} - t_{PHL4}|$

$t_{sk(p)}$ is calculated from the following expressions: $|t_{PLH1} - t_{PHL1}|$, $|t_{PLH2} - t_{PHL2}|$, $|t_{PLH3} - t_{PHL3}|$, $|t_{PLH4} - t_{PHL4}|$

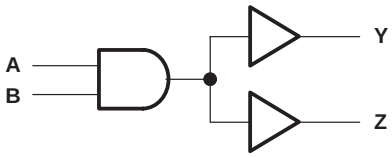

Figure 3. Skew Definitions

Power Dissipation

Power consumption becomes an important consideration as operating frequencies rise, but is often overlooked as a designer tackles other issues. A much-touted aspect of devices fabricated in CMOS and BiCMOS technologies is low power consumption, especially when compared to equivalent devices fabricated in a purely bipolar process. At lower frequencies, this generalization holds true, but power consumption at higher frequencies becomes less a function of process technology and more a function of output loading. To illustrate this point, the dynamic power (P_d) consumed by a CMOS device will be examined. The dynamic power consumed consists of two components:

- Power used by the device as it switches states
- Power required to charge any load capacitance

P_d is easily calculated using the following expression:

$$P_d = \left(C_{pd} \times V_{CC}{}^2 \times f_i\right) + n(C_L \times V_{CC} \times f_o) \quad \textbf{(1)}$$

Where:

C_{pd} = power dissipation capacitance of the device as specified on the data sheet
f_i = input switching frequency
f_o = output switching frequency
C_L = load capacitance on each output
n = number of outputs switching

This example assumes all outputs are equally loaded. Power consumed by the device switching logic states occurs because fabricated transistors on a chip are not ideal. Parasitic capacitances are present, and they must be charged and discharged. C_{pd} quantifies the magnitude of these parasitic capacitances and is in the range of 24–30 pF for clock-driver devices fabricated using TI's EPIC™ Advanced CMOS process. Power needed to charge the load capacitance, C_L, makes up the second half of the equation. The designer has some control over C_L, while C_{pd} is strictly a property of the device chosen. Power consumed by both is a direct function of the frequency at which the system operates and is usually fixed by the processor speed. The designer's goal is to reduce and evenly distribute the load that a device must drive. The SPICE data shown in Figure 4 for the CDC337 clock driver graphically illustrates the power-consumption tradeoffs that can be made between the switching frequency and output loading.

The CDC337 is fabricated in TI's EPIC-IIB™ BiCMOS process and contains four buffered outputs that switch at the clock frequency and four divide-by-two outputs that toggle at one-half the clock frequency. It is specifically designed for applications requiring synchronized output signals at both the clock frequency and one-half the clock frequency. Power consumption also has implications for packaging at both the component and system levels. The general trend is that the system box shrinks with each advance in performance. This, in turn, requires smaller power supplies, closer board-to-board spacing, and reduced capacity for free air flow, all of which are not compatible with power-hungry designs. Increased system packaging density usually requires the use of surface-mount components that do not have the higher heat dissipation properties of larger DIP devices but do allow closer board-to-board spacing and component placement on both sides of the circuit board. All of these factors can drive up system operating temperature and cost. Power consumption can make or break a system design and should not be treated lightly.

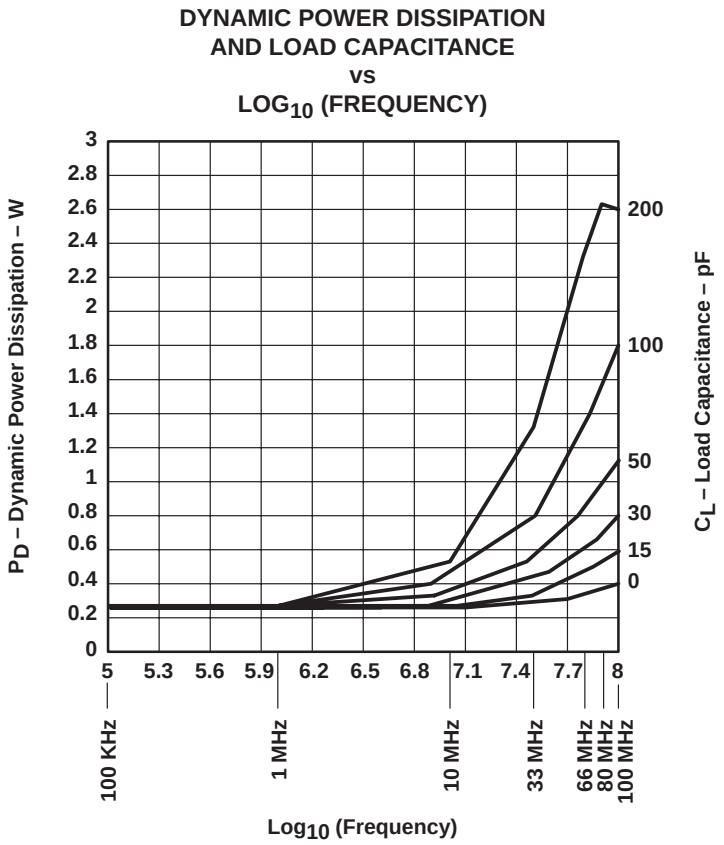

Figure 4. CDC337 Power Dissipation

High-Speed Design Considerations

A number of tightly specified, high-speed, high-drive clock-driver circuits are available to aid the designer in developing a system-clocking network. By carefully following established high-speed circuit design techniques, the designer can achieve superior performance from standard components and not incur the high cost of custom components. Transmission-line effects take over in high-speed, high-drive designs, and attention to detail during board layout is critical. A sound high-speed design uses all of the following techniques and rules of thumb:

- Keep output loading as light as possible. This reduces power consumption, allows switching at higher frequencies, and reduces skew.
- Equally load all outputs where possible
- Use short, equal-length etch runs
- Avoid sharp corners that may induce unwanted reflections, ringing, and overshoot due to discontinuities
- Properly decouple all device V_{CC} pins as close to the pins as possible. The best high-frequency filtering is often accomplished with a combination of capacitors. An effective combination is 0.1 μF in parallel with 0.01 μF or 0.005 μF. Use RF-quality (low-inductance) capacitors.
- Use a multilayer board with low-impedance power and ground planes to minimize circuit noise
- Select components with low-noise characteristics. Components optimized for low noise usually have multiple V_{CC} and GND pins interspersed among the device outputs to help reduce noise.

Summary

High-performance systems demand a carefully designed clock-generation and distribution network. The designer has the challenge of outlining a design that meets tighter timing, lower power consumption, smaller space, lower operating temperature, and lower cost requirements. Timing performance is optimized by reducing clock skew. Skew can be minimized by selecting components optimized for low-skew applications and by tightening power-supply requirements to ±5% instead of ±10%. Power consumption is largely a function of operating frequency but can be reduced by making output loading as light as possible. The use of surface-mount components saves board and system box space but requires analyzing tradeoffs in power consumption, air flow, and operating temperature. High-performance, low-cost clock-driver components are now available that can help the designer with this performance juggling.

Acknowledgment

The author of this report is A. R. Austin.