Generation of Synchronized Processor Clocks Using the CY7B991 or CY7B992

Introduction

Many modern systems use multiple processors operating simultaneously in order to increase performance and improve throughput. Timing analyses and interprocessor communications are significantly simplified if the clocks to the processors occur at exactly the same time. This application note explains the problem and presents a technique for generating synchronous clocks to two Intel 80960CA processors using the Cypress CY7B991 Programmable Skew Clock Buffer (PSCB), also known as RoboClock. The technique is then extended to "n" processors.

Design Requirements

The processors require 33-MHz clocks and are operated in the x1 mode (i.e., CLKMODE = HIGH). In this mode the output clock, PCLK1, PCLK2, are also 33-MHz and can be phase shifted plus or minus two nanoseconds from the input clock, CLKIN. This is due to the internal (2X) Phase-Locked Loop (PLL) in the processor. The minimum CLKIN LOW duration is 10 ns and the minimum CLKIN HIGH duration is also 10 ns. In addition, the maximum cycle-to-cycle CLKIN period variation is plus or minus 0.1%. Another requirement is that the RESET input to the processor be held LOW for at least 10,000 CLKIN cycles after V_{CC} and CLKIN have stabilized (are within their specifications) before it is allowed to go from LOW to HIGH.

Clock Interconnections

Figure 1 illustrates the interconnections required for clock synchronization. The connections are the same if the CLKIN frequency is 66 MHz. However, the CLKMODE input must be tied to ground. The PCLK1 output is then 33 MHz.

Theory of Operation

During power turn-on, the $\overline{\text{RESET}}$ input of each processor must be held LOW by external logic (not shown). Typical power supply turn-on times are in the 50 ms to 500 ms range. After the power supply and 33-MHz oscillator have stabilized, it takes 10,000 cycles of the 33-MHz CLKIN input to processor 1 before its PCLK1 output is within its specification. This is 300 microseconds.

During this time, the TEST input to the CY7B991 must be held HIGH. When this is done, the internal Phase-Locked Loop of the CY7B991 is disabled and the signal at the REF (reference) input is passed through to the 1Q0 output, and then to the CLKIN input of processor 2. Again, 300 microseconds must pass before the output PCLK1 of processor 2 is stable. Both processors are now running at 33 MHz, under control of the oscillator. The PCLK1 clocks of the processors, however, are not synchronized.

Synchronization of the Processor Clocks

The next step is to cause the TEST input of the CY7B991 to go from HIGH to LOW. This causes the Phase-Locked Loop within the CY7B991 to ad-

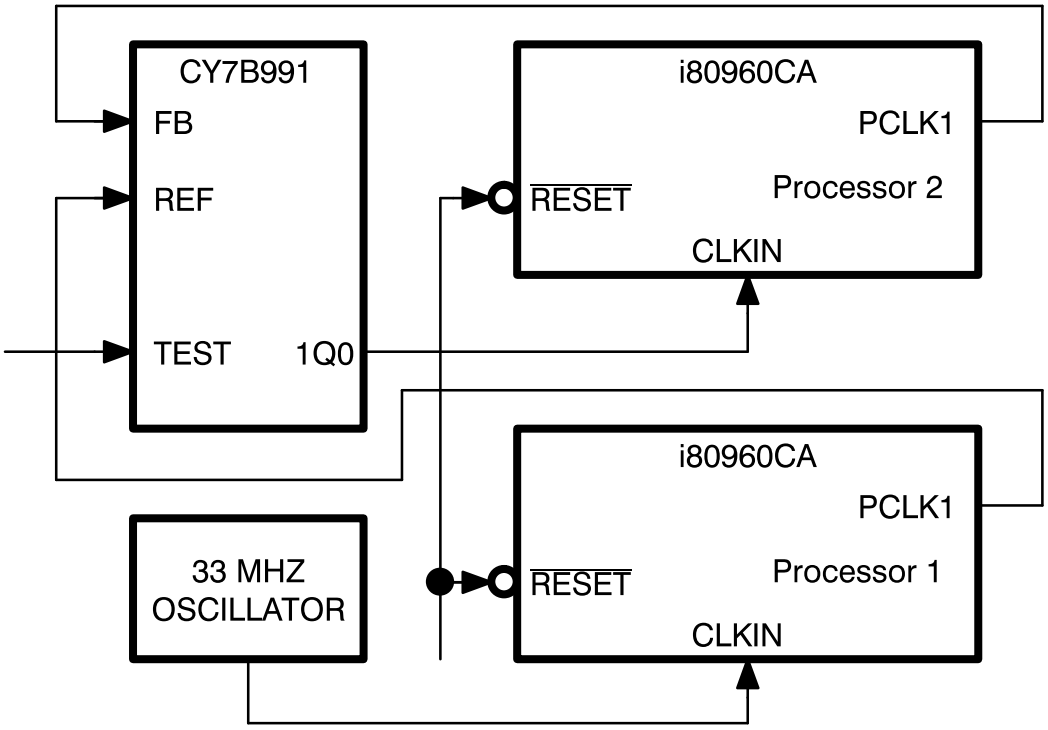

Figure 1. Clock Connections for Synchronization

just the phase and frequency of the 1Q0 output, which is driving the CLKIN input of processor 2, such that the rising edges of the signals on its FB and REF inputs are aligned. Because the PCLK1 output of processor 2 is a function of its CLKIN input, when the CY7B991 adjusts its 1Q0 output, the PCLK1 output follows. The result is that the PCLK1 outputs and, therefore, the CLKIN inputs of the two processors are synchronized. What this means is, that for all practical purposes, there is "zero delay" between the rising edge of the signal on the REF input and the signal on the FB input. However, what is more important is that this alignment is adaptive and dynamic because it occurs on a cycle-by-cycle basis, and, therefore, is not influenced by variations in power supply voltage or temperature. After the processor clocks are synchronized, the $\overline{\text{RESET}}$ lines to the processors can transition from LOW to HIGH.

CLKIN Cycle-to-Cycle Variation

The next step is to calculate the maximum cycle-to-cycle variation of the CLKIN input to processor 2 and make sure that it is within the 0.1% specification on the 80960CA data sheet. At 33 MHz the clock period is 30 ns, so 0.001 x 30 x 10^{-9} = 30 picoseconds per cycle.

The Phase-Locked Loop of the CY7B991 requires approximately 50 microseconds to lock. This corresponds to 50 microseconds divided by 30 ns per cycle, or 1,667 clock cycles. The worst-case condition is that the two processor clocks are 180 degrees out of phase when the signal at the TEST input of the CY7B991 transitions from HIGH to LOW. One-half a cycle of a 30 ns period clock is 15 ns. Fifteen nanoseconds divided by 1,667 cycles is 9 picoseconds per cycle. This is much less than the plus or minus 30 picoseconds (60 picoseconds total) specified on the 80960CA data sheet

Synchronization of Many Processors to a Single Clock

Figure 2 illustrates how three processors can be synchronized. The first runs off of the oscillator and the other two are synchronized to the first by using two CY7B991s. The PCLK1 output of processor 1 is the

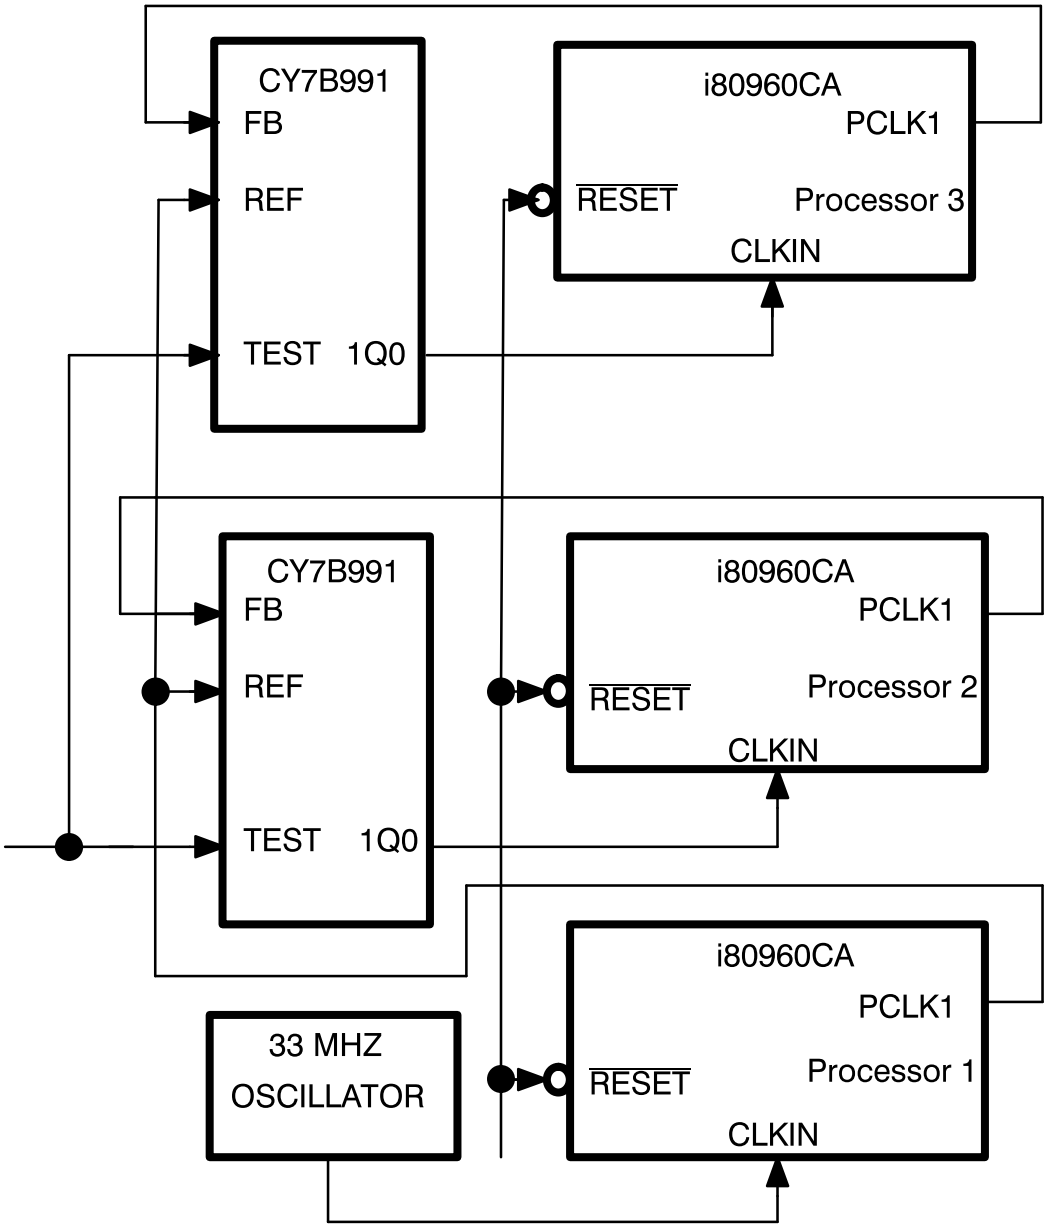

Figure 2. Clock Connections for Synchronization of Three Processors

reference for the two CY7B991s. Each controls the clock to a processor. Thus n−1 CY7B991s are required to synchronize n processors.

The advantage of using separate RoboClocks for processor 2 and processor 3 is that, because of the analog nature of the internal RoboClock PLL, the PCLK1 output of each is independently and dynamically adjusted, on a cycle-by-cycle basis, with the PCLK1 output of processor 1. This is accomplished by applying the PCLK1 output of processor 1 to the REF input of the two RoboClocks and tying the PCLK1 outputs of processors 2 and 3 to the FB inputs of two separate RoboClocks.

One CY7B991 can be used to control many processors if they do not have on-chip Phase-Locked Loops. Or, the system designer may choose to not use the processor clock output.

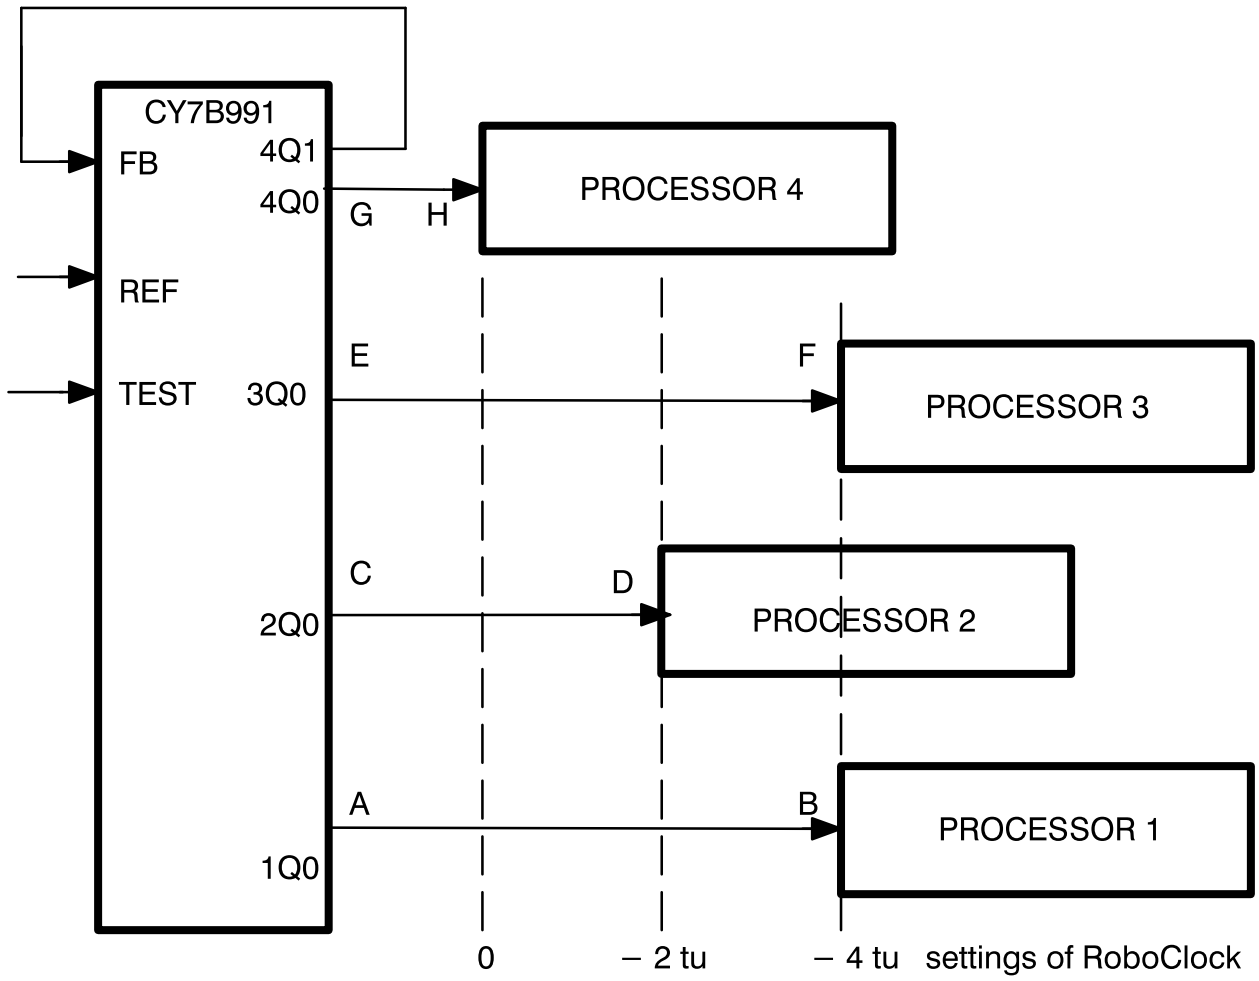

Figure 3. One RoboClock Driving Multiple Processors

Driving Multiple Processors From One RoboClock

Figure 3 illustrates one CY7B991 driving multiple processors that are located at different distances.

Advantages

The advantages of the configuration illustrated in *Figure 3* are (1), that one CY7B991 can drive up to seven processors using seven of the eight RoboClock outputs and (2), that the select inputs can be used to adjust the timing of the Q outputs to compensate for variations in trace length, so that the clocks to the processors arrive at exactly the same time.

For example, the propagation delay of trace G H is two timing units, that of trace C D is four timing units, and those of traces E F and A B are six timing units. It is required that the clocks to all of the processors arrive to each at the same time.

The first step is to select a “zero” point as a timing reference. This is the clock at processor 4, which is point H. However, in real life, the propagation delay of trace G H is two timing units. What RoboClock does is precisely align the rising edge of the signal at its FB input with the rising edge of the signal at its REF input. The length of the fed back output trace (4Q1 to FB) should be as short as possible. It not only simplifies the timing analysis, but also reduces the noise introduced into the PLL.

Limitations

There is no feedback from the clock outputs of the processors, so they cannot be individually, dynamically aligned with the REF clock, as is done in *Figure 2.* A second limitation is that the eight outputs of the CY7B991 are grouped in four pairs and are adjustable only as pairs. This means that a maximum of three (pairs of) processors can be aligned independently if each is a different distance away from RoboClock. By matching the trace length within pairs (i.e., 1Q0, 1Q1) up to seven processors can be driven, each with its own, dedicated output.

Determination of Delay Settings

For purposes of explanation, the "zero" is chosen at point H, which is the closest physical point to RoboClock. Point F is four timing units farther away than point H, so its signal must precede (timewise) that at point H, so that the signals at points F and H occur at the same time. Therefore, the select input controlling the 3Q outputs is set at −4 timing units. In a similar manner, the select input controlling the 2Q outputs is set at −2 timing units and that controlling the 1Q outputs is set at −4 timing units. When this is done, the signals at points H, F, D, and B occur at the same time, which is the zero point shown. The timing is shown in in *Figure 4* below.

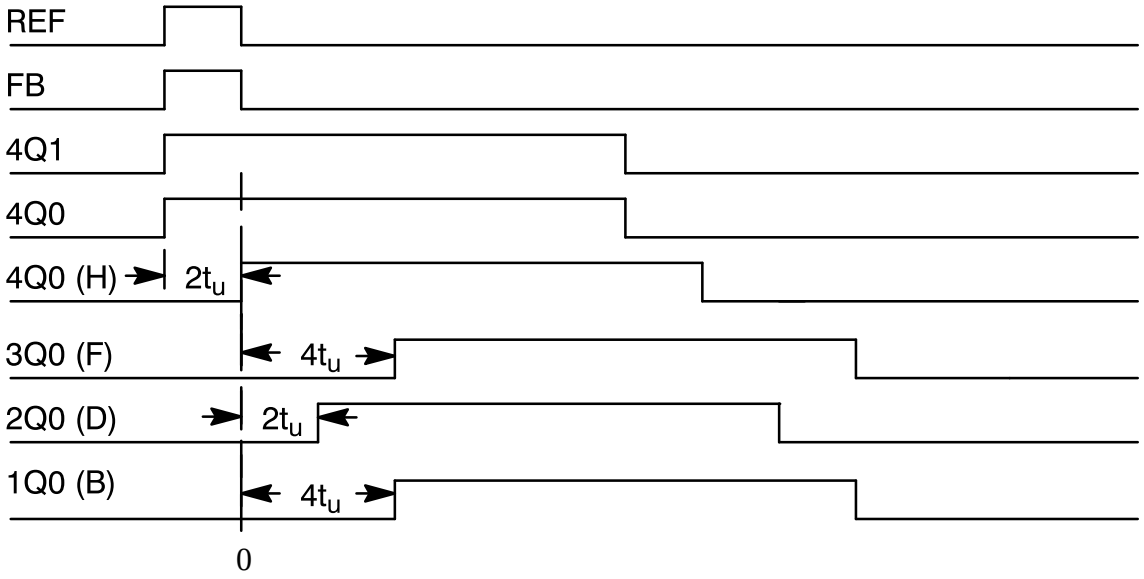

Figure 4. Timing Diagram for Figure 3 (before select control settings)