Optically Isolating an I²C Interface

Using Clare's CPC5902 and CPC5903

Technical Brief TB-101

Abstract: The ideal model of an open-drain bus as a simple, shared pull-up resistor is no longer instantaneously valid when using bus isolators with real-world propagation delays. If the delays are not considered in the design of the hardware, then there may be differences in the number of rising or falling edges seen at each isolated bus.

1 Introduction to the I²C Bus

The I²C bus is a set of hardware and software rules that allows communication between multiple devices over a shared, two-wire interface. In operation the bus uses one line (SDA) for data and the other line (SCL) for clock. In standard mode (100kbps) or fast mode (400kbps), each of the two wires can be pulled high through an external pull-up resistor. Thus if any of the multiple devices asserts a logic low onto either line of the bus, the logic low is seen at all devices. Rules defining master and slave devices determine who is allowed to drive the bus lines, and when. The bus specifications were defined and are maintained by Philips Semiconductors (now known as NXP Inc.)[1]

The device that is acting as the bus master drives the clock line (SCL) when communicating with slave devices. Any addressed slaves respond at a time defined by this clock by asserting the data (SDA) line using a specified protocol. The simplest system contains only one master device with all other devices always responding as slaves. In more complex systems multiple devices can take turns being bus master, and the SCL line will be driven by whichever device is master at that time.

There are many applications requiring ground isolation or logic-level translation between devices where the simple I²C protocol is very useful. The small number of wires (2) minimizes the required number of isolators and keeps down the cost of the isolated systems, although each direction of the wire has generally required its own isolator.

Systems in which the bus can be isolated such that only slave devices exist on one side of the isolation barrier do not need bidirectional isolation of the SCL line. More specifically, a bus containing only I²C slaves that do not need to implement clock stretching (in order to slow down the bus to the speed that they can use) does not need bidirectional isolation of the SCL line. The CPC5903 optical isolator is configured for use in these systems (see **Figure 1**): It has one bidirectional path (for SDA) and one unidirectional path (for SCL).

Systems in which devices on either side of the isolation barrier may function as bus master need bidirectional isolation of both lines; the CPC5902 is configured for this application (see **Figure 1**). The CPC5902 also functions very well when one channel is only used in a unidirectional manner. It supports these applications:

- Slave-only-on-SideB
- Clock-stretching-slave-on-SideB
- Bus-master-on-SideB

and should be used for slave-only systems in which yet-unknown future devices might be attached to the isolated bus at some later time.

The remainder of this Technical Brief discusses optical, bidirectional bus isolators, of which the CPC5903 has one, and the CPC5902 has two.

Figure 1 CPC5902 & CPC5903 Application Diagrams

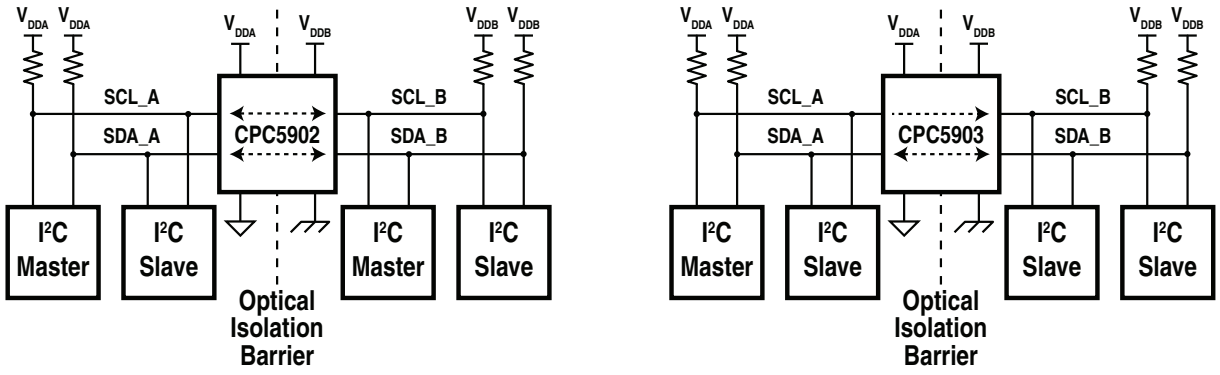

[1] NXP Semiconductor, "I2C-bus Specification and User manual" UM10204_3, June 2007

2 Isolating a Bidirectional Open-Drain Bus

The I²C bus, which contains at least one bidirectional open-drain bus, does not easily lend itself to optical or other isolation. The ideal bus model of the SDA line is a shared pull-up resistor connected to bidirectional I/O pins at multiple devices. The bidirectionality of the bus in an isolated bus system causes a departure from this ideal model: If a real-world active bus repeater or bus isolator is allowed to drive its own input low at both its I/O pins, then the bus will latch the first logic low asserted, and allow no further activity!

Figure 2 This Bus Repeater Latches Up

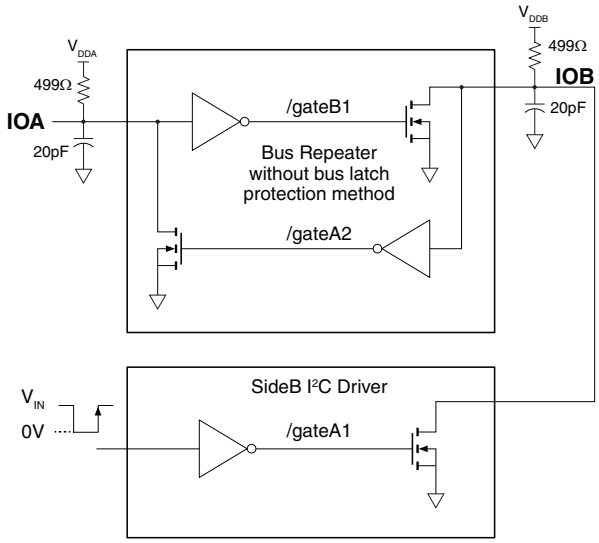

Figure 3 Bus A and Bus B Both Stuck Low

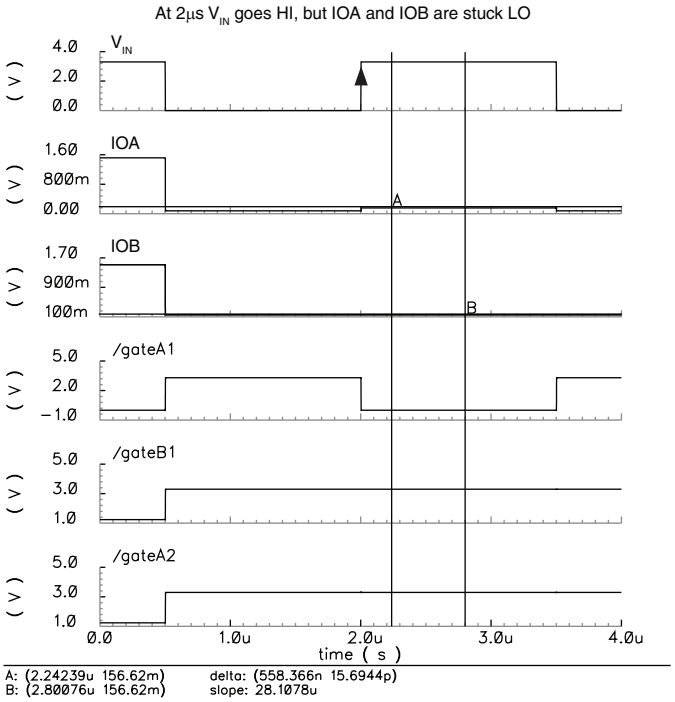

One way to avoid the latch-up problem is to use a separate input and output pin on all devices at SideB (so that when the repeater output is driven low it does not drive the repeater input low). There are some peripherals, which have been intentionally designed to run with standard optoisolators, that do utilize this extra wire to provide separate data input and output wires at devices at SideB. This does require an extra package pin, and prohibits the use of most I^2C devices (without the extra pin) at the isolated SideB bus. This also complicates debugging of the bus because most I^2C bus analyzers do not have a provision for the third wire, and thus can only be used on SideA of the isolated bus.

Figure 4 Separate In and Separate Out at SideB

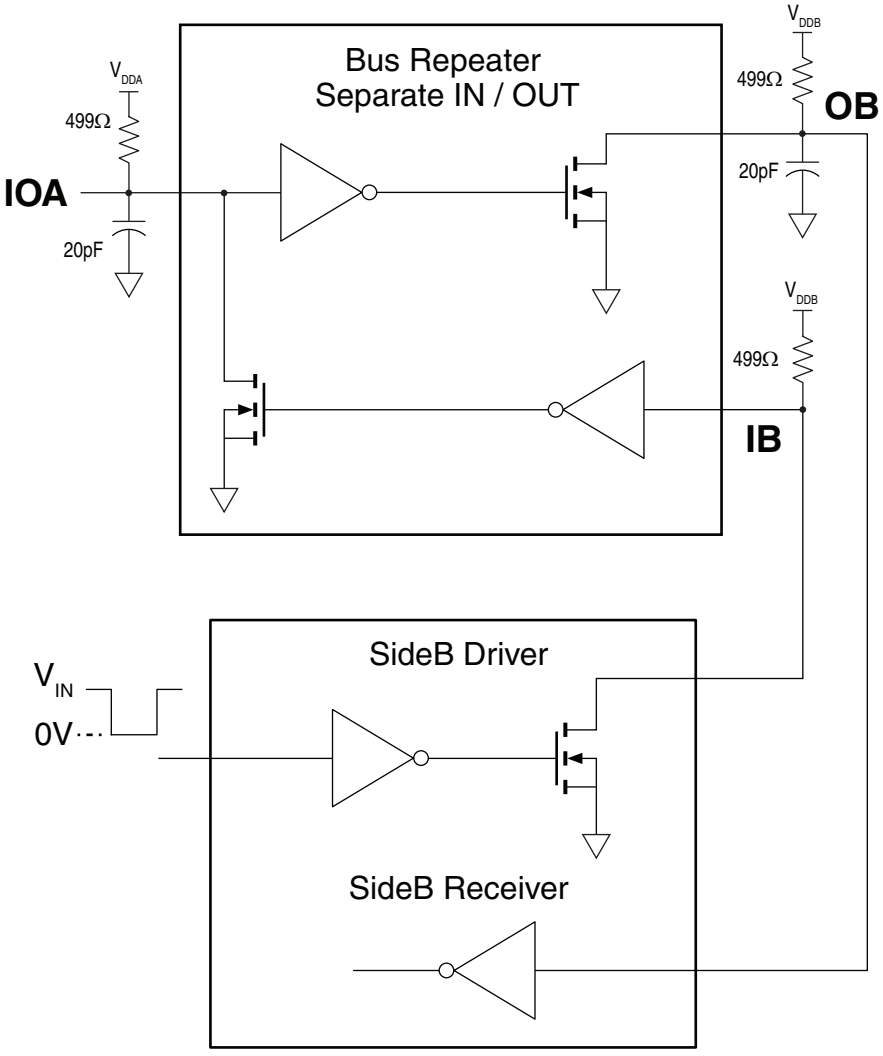

A different method of bus logic low latch-up prevention is used in the CPC5902 and CPC5903, and by some other bus repeaters and isolators. This method uses non-standard three-level logic at the isolated SideB. This method only requires 2 wires, and works with all I^2C devices and bus analyzers. It also works with the separate data-in/data-out pin devices if the data-in and data-out pins are shorted together at SideB.

Figure 5 Three-Level System at Side B

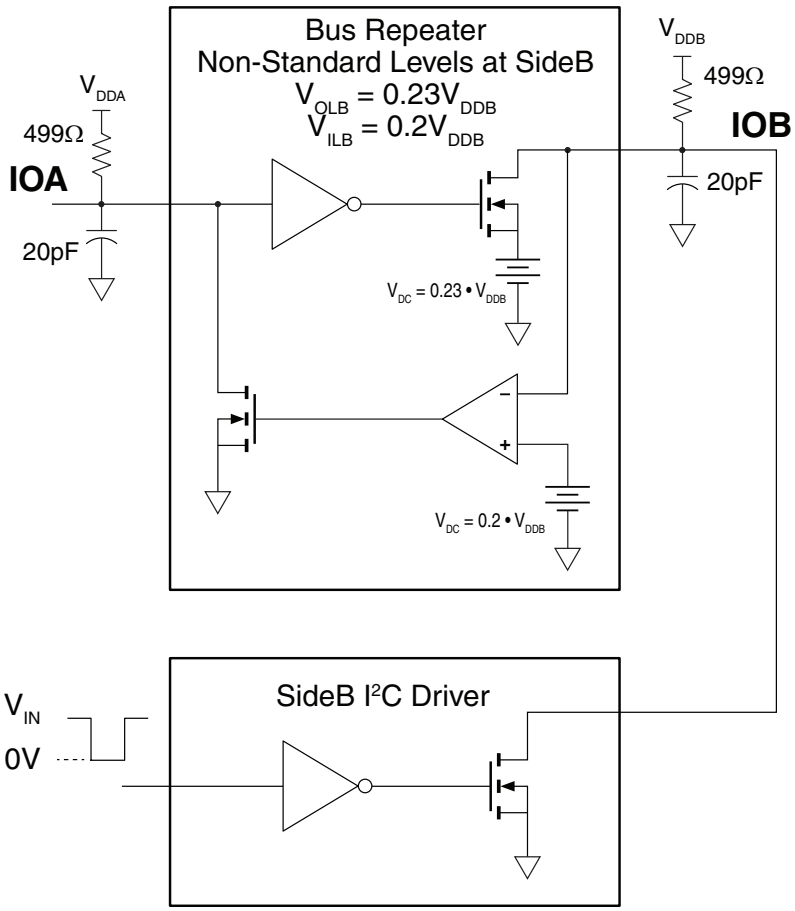

If the repeater output driver is designed such that it can only drive V_{OLB} down to $0.23*V_{DD}$, and the receiver input threshold, V_{ILB}, is designed to switch near $0.2*_{VDD}$, then the repeater output cannot drive its own input low. The I^2C specification defines V_{IL} to be any input voltage below $0.3*V_{DD}$. All devices attached to the bus that meet this spec will correctly see the non-standard output level, V_{OLB}, as an asserted logic low. When using this method, a careful designer must pick the value of pull-up resistors on SideB of the bus such that all devices on that bus can drive lower than $V_{ILB}=0.2*V_{DD}$. This may require picking a slightly higher value pull-up resistor.

As an example, consider a system in which there are both 3mA-drive I^2C standard-mode devices and 6mA-drive I^2C fast-mode devices on the isolated SideB of the bus. The pull-up resistor value in a usual system must be picked so that the standard-mode (weakest) device will pull the bus lower than $0.3*V_{DD}$ when asserted low. However, the pull-up resistor must be picked so that the weakest device will pull lower than $0.2*V_{DD}$ when used at the SideB of the CPC5902 and CPC5903 isolators, as described below.

2.1 Three-Level System Pull-Up Resistor Selection

If the V_{DDB} supply voltage is specified to be minimum 3V to maximum 3.6V, then the minimum value of the pull-down resistance for usual, standard levels would be that value that pulled the bus down to or below $0.3*V_{DD}$ when V_{DDB}=3.6V. At $3.6V_{DD}$, 0.3*3.6=1.08V, and 3.6V-1.08V=2.52V, which must drop across the pull-up resistor when the minimum guaranteed output drive of 3mA is applied. Thus the minimum value of resistor for this standard-level, standard-mode system is 2.52V/3mA = 840Ω. The designer would need to pick a standard value which is slightly greater than this; how much greater depends on the guaranteed resistor tolerance, perhaps picking 866Ω for 1% tolerance resistors.

For the CPC5902 and CPC5903 non-standard level SideB, for V_{DDB} between 3V and 3.6V, use the same method, but require that the bus pull down to $0.2*V_{DDB}$ when V_{DDB}=3.6V. At $3.6V_{DD}$, 0.2*3.6=0.72V and 3.6V-0.72V=2.88V, which must drop across the resistor when the 3mA minimum drive is applied. Thus the minimum value of resistance for the CPC5902 and CPC5903 family is 2.88V/3mA=960Ω. A standard value of 1% tolerance resistor that will stay above this minimum is 976Ω. Thus, in this example, the use of CPC5902 or CPC5903 requires a 976Ω R_{PULLUP} to guarantee operation under all conditions, while a standard level implementation would have required 866Ω.

Note that I^2C fast-mode-capable devices guarantee 6mA of output drive, and therefore could use smaller resistors, but the presence of the 3mA drive standard mode devices on SideB of the bus make the larger resistance mandatory to insure their operation. In practice, most manufacturers design their standard mode output drivers to supply more than the required 3mA of the I^2C-specification, and so the values used above prove to be quite conservative.

Also note that SideB of one CPC5902 or CPC5903 cannot be connected to SideB of another CPC5902 or CPC5903 if communication of one through the other is required. This is because the output logic low of one will not be seen as a valid logic low at the input of the other. SideB of isolators can be connected to SideA of more isolators if cascaded operation is desired.

3 Time Delays in Bus Repeaters and Isolators

Another way in which an isolated bus system departs from the ideal pull-up resistor model is in the timing of transmitted versus received data. For an ideal, non-isolated system of multiple devices tied to the same pull-up resistor, there is very little time delay between when each device on the bus sees the level go to an asserted logic low or return high. Unfortunately, delays of tens to hundreds of nanoseconds can exist within isolators, and such delays must be carefully accounted for to insure legal values of specifications such as data setup and data hold times.

The LED-photodetector cascade within an optoisolator is inherently a lowpass filter. The primary bandwidth limiter is usually capacitance at the photodetector, which must be charged and discharged by the small photo-generated current. In a logic optoisolator this electro-optical filter is then followed by a comparator. This low-pass filter suppresses very short logic pulses, and adds a time delay to any logic signal applied. Later sections of this paper will show that the effects of this filter on pulse-widths and delays can be modified by using different latch-up suppression methods.

In this Tech Brief, simple, buffered RC filters are used to model the low-pass characteristic, and a pair of logic inverters used to model the high-gain comparator (see **Figure 6)**. This circuit is used elsewhere in this Tech Brief as a circuit block named “Filtered 60ns Delay.”

Figure 6 “Filtered 60ns Delay” Circuit

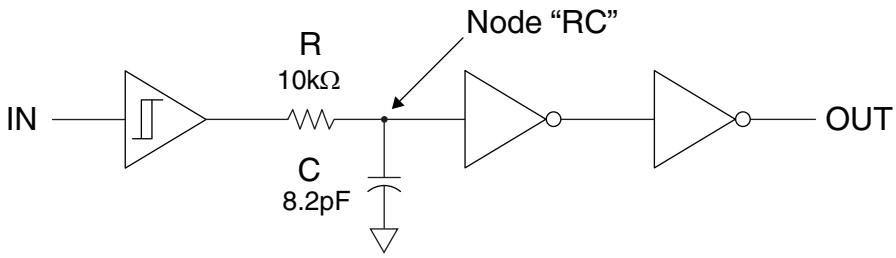

For applied pulse-widths greater than 5 times the RC time constant (5RC) at Node “RC,” this circuit exhibits nearly the same filtering for going-low and going-high signals. For example, the “Filtered 60ns Delay” circuit delays the rising edge of a 400ns pulse by 60ns and the falling edge by 78ns. The small asymmetry is because the comparator switching threshold voltage is lower than $V_{DD}/2$. An asymmetry in this direction generally also exists in real logic optoisolators when the capacitance at the photodetector can be charged quickly by photocurrent from an over-driven LED, but can only self-discharge at a lower rate.

However, a low-pass filtered comparator is a non-linear circuit even for digital signals. When applied after a “long” (greater than 5RC) wait, even a short pulse will ramp Node RC up from zero to the comparator switching threshold voltage in the expected amount of time (see **Figure 7**). But if the pulse width is not greater than or equal to 5RC, then the Node RC does not ramp all the way to V_{DD} after the rising edge is applied (see **Figure 8**). Thus when the falling edge arrives and the Node RC starts to go negative, it does not need to drop all the way from V_{DD} down to the comparator threshold.

Figure 7 below shows the Node RC charging all the way to V_{DD} for a nominal input pulse. The going-low delay is similar to the going-high delay, and the output pulse is only slightly stretched compared to the pulse at its input.

Figure 7 Applied Pulse Greater Than 5RC

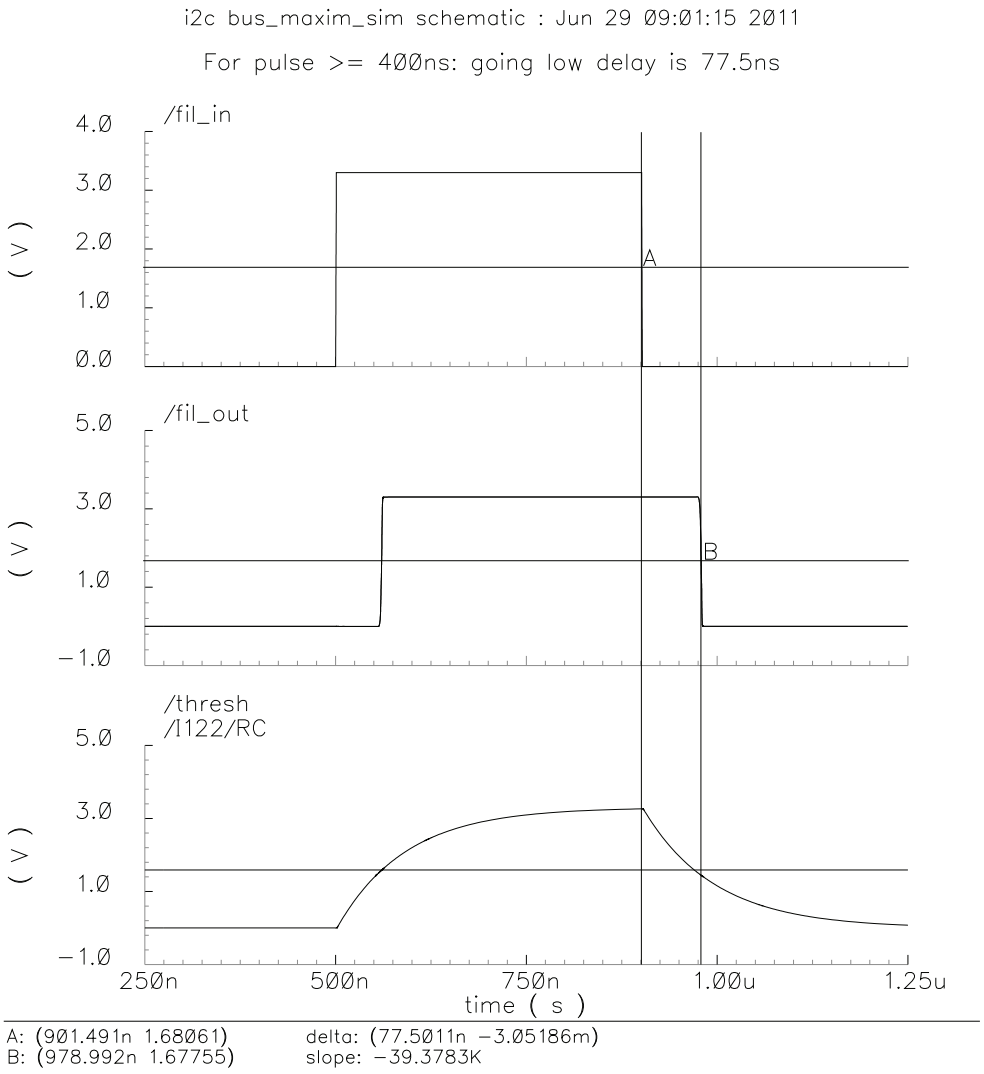

Figure 8 below shows the Node RC barely charging up to the comparator threshold voltage, and very quickly going below it when the falling edge starts to pull Node RC negative. It shows the two problems that should be addressed in a robust isolator design: (1) when very short pulses are applied, the t_{PHL} delay becomes much shorter than the delay for longer pulses; (2) the output pulse-width, which is lengthened for long pulse-widths, is severely shortened for short applied pulses.

Figure 8 Applied Pulse Not Greater Than or Equal to 5RC

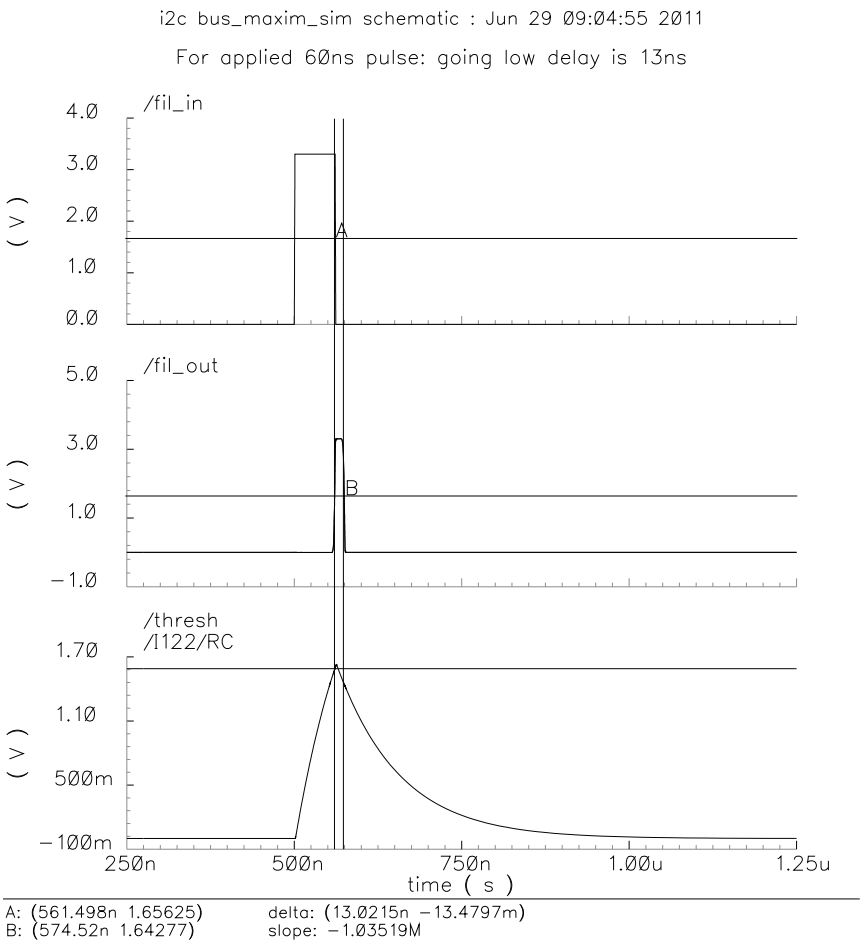

Applied Pulse	t_{PLH}	t_{PHL}	Pulse Out
60ns	59.8ns	13ns	14ns
150ns	59.8ns	59.6ns	150ns
400ns	59.8ns	77.5ns	419ns

This table shows the t_{PLH}, t_{PHL} and output pulse-width vs input pulse-width at “Filtered 60ns Delay” for a pulse applied after 5RC or longer of stable setup.

Note that in the optoisolator designs to be shown in later sections, there is a logical inversion before the input to the Filtered 60ns Delay and another after its output. This alters the effect of the pulse distortion as the t_{PHL} (not t_{PLH}) at the IO pins remains constant and the t_{PLH} (not t_{PHL}) at the IO pins is the delay which is greatly affected by applied pulse duration after sufficient setup. For I^2C applications this is generally a preferred topology as valid “long” asserted low pulses are slightly stretched instead of being slightly shaved.

The next sections will show the effects of adding the Filtered 60ns Delay at all the optical interfaces of various optoisolator topologies. In order to maximize simulated bus bandwidth assume only 20pF of total load capacitance and fast-mode compatible 499Ω pull-up resistors to 3.3V at all IO pins. An asserted low pulse of varying duration will be applied at SideB and the resulting pulses at both SideA and SideB will be compared.

4 Separate SDAin and SDAout

Devices which use the separate SDAin and SDAout pins topology for latch-up avoidance are meant to be used with standard unidirectional logic optoisolators. When the optics are replaced with Filtered 60ns Delay blocks, the effects on pulses of various lengths can be examined.

Figure 9 Separate Input & Output with Delays

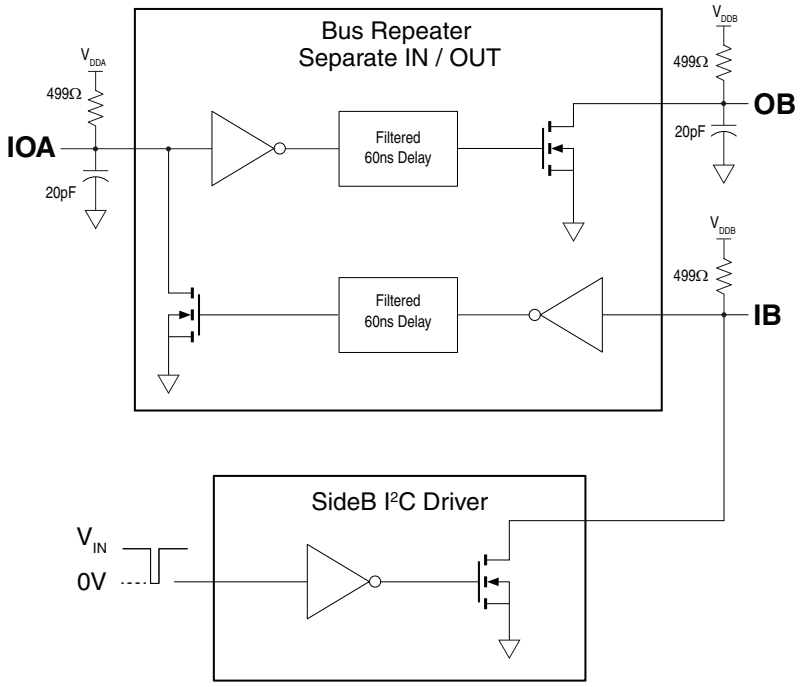

4.1 Separate SDAin / SDAout Method:

IB	IOA	OB	Comment
<55ns	none	none	fine, suppresses short pulses
57ns	5.5ns	none	glitches IOA
80ns	54ns	none	2 extra edges at IOA vs. OB
82ns	57.5ns	8ns	glitches OB (See Figure 8)
100ns	86ns	64ns	fine, same number of edges IOA, IB, OB

The I^2C specification for fast-mode operation requests that internal filtering at devices on the bus suppress pulses of less than 50ns. Note that the isolators themselves are not “on” the bus, they “are” the bus, and thus they do not have to perform this suppression. The devices attached to the isolators are supposed to suppress the glitches. However, the isolator can either help or hurt glitch suppression. In the table above, glitches less than 55ns applied at IB do not propagate at all to IOA or OB; so for very short pulses this isolator aids in glitch suppression.

However, when the applied pulse length gets to 57ns, less useful effects begin to appear. The optics have shaved the width of the applied pulse: a 57ns pulse applied at IB yields only a 5.5ns pulse at IOA. Instead of helping in glitch suppression, the filtering within the optics is now causing glitching at IOA when a non-glitch pulse has been applied at IB.

Figure 10 Separate Input & Output Shaves Pulse

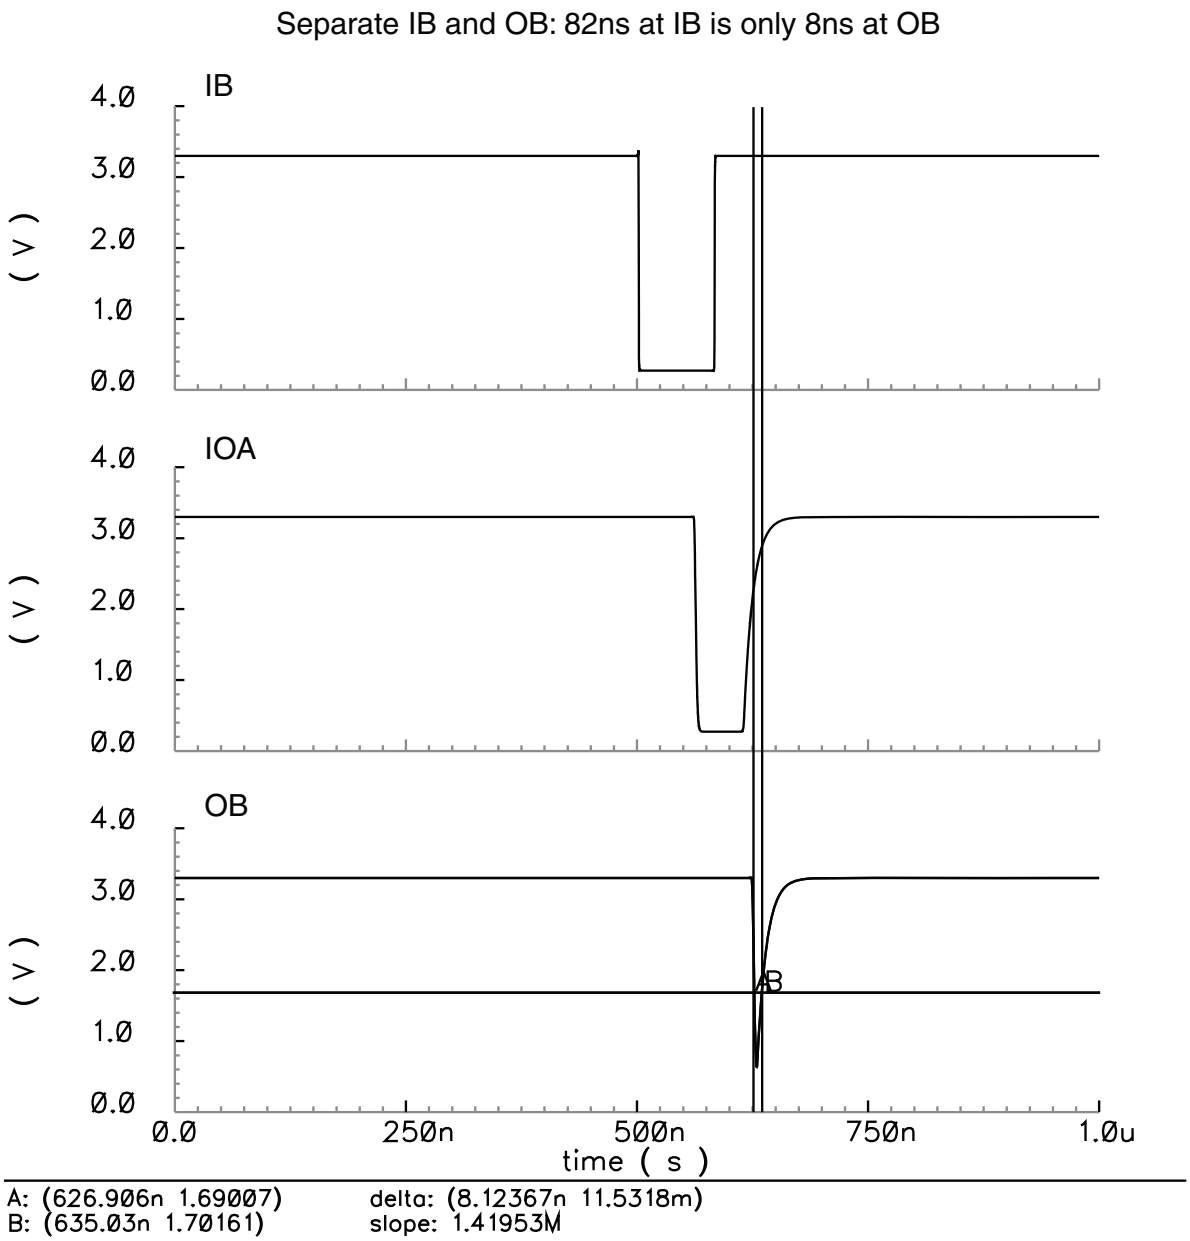

As the pulse width is increased to 80ns, the pulse at IOA grows to 57.4ns in length, but the additional filtering of the A-to-B optics has so far suppressed all indication of pulsing at OB. Thus the bus on the IOA side has seen one more negative edge and one more positive edge than the bus at OB. If this channel is being used for clock, then this can be a very serious problem.

At 82ns width, OB finally begins to get a pulse, but only an 8ns glitch. Again the optics has taken a more-legal 82ns pulse and made a glitch from it, this time at the OB port. Finally, at applied pulses near 100ns, the OB pulse grows to 64ns in length, and the same number of edges that were generated at IB will be seen both at IOA and at OB.

5 Simple Three-Level Logic

Simple non-standard three-level logic as a latch-up prevention method also suffers from glitch generation mechanisms. The worst path is when IOB is driven. This is because, unlike SideB, the isolator at IOA is allowed to drive its own input at IOA. The signal thus gets optically filtered going from B to A and then again going from A back to B

Figure 11 Simple Three-Level Method with Delays

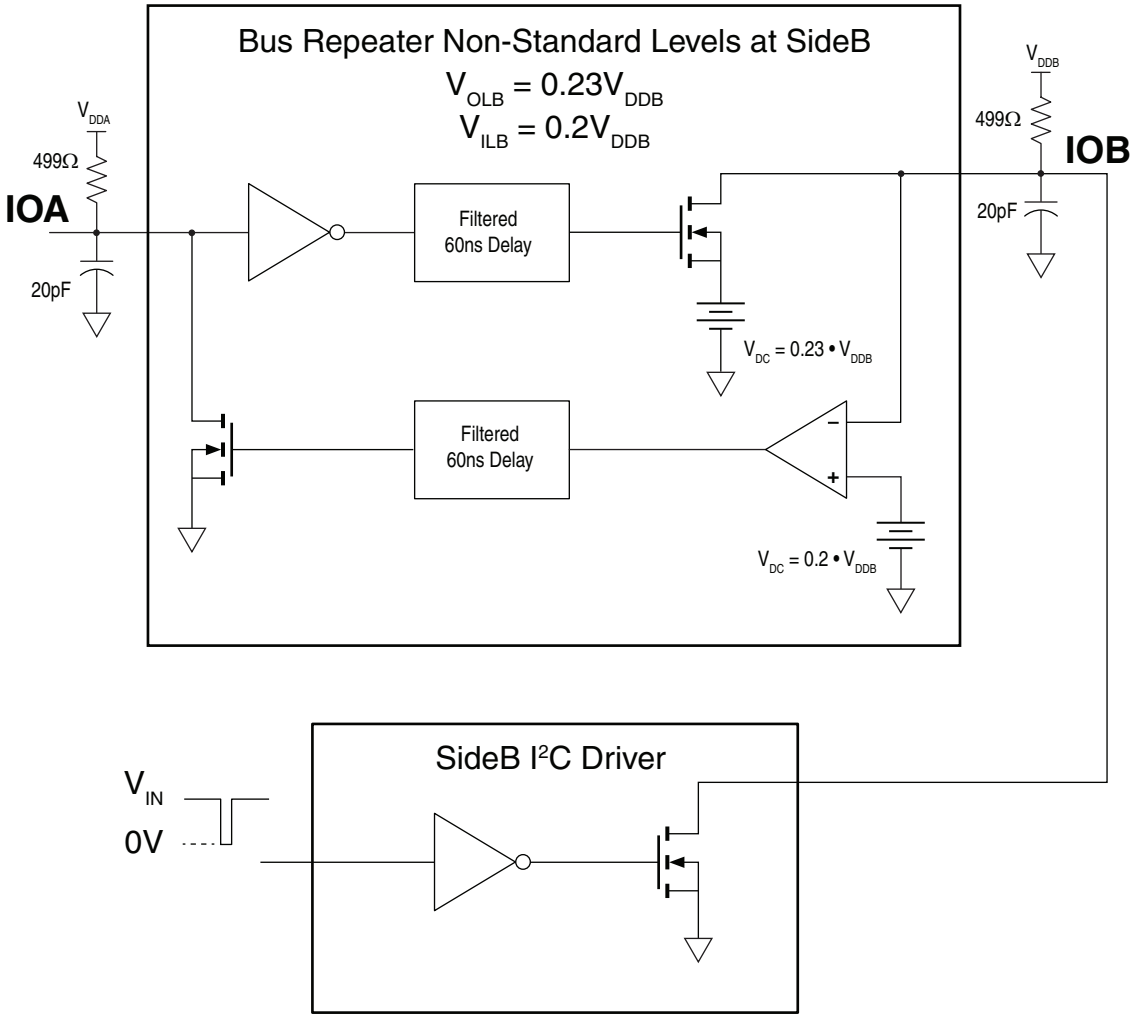

5.1 Simple Three-Level Method:

V_{IN}	IOB	IOA	Comment
<55ns	<61ns	none	fine, suppresses short pulses
57.5ns	64ns	8.7ns	glitches IOA
82ns	87ns	57ns	IOA is better, getting longer
82.5ns	88ns, 39ns, 6nsLow	8ns	double pulse at IOB! (See Figure 12)
118ns	120ns, 2nsHi, 99nsLow	110ns	double pulse at IOB!
>120ns	122ns, 227nsLow	113ns	fine, same number of edges

In the table above, glitches less than 55ns applied at V_{IN} do not propagate at all to IOA; so for very short pulses this isolator aids in glitch suppression. Note that the V_{IN}-driven open-drain driver is applied to IOB, but that capacitance at IOB causes the pulse width applied at V_{IN} to grow by up to 6ns.

However, when the applied pulse length gets to 57.5ns the optics again have shaved the width of the applied pulse - a 64ns pulse applied at IOB yields only an 8.7ns pulse at IOA. As the pulse-width grows the output at IOA grows. However, when the output pulse-width at A is long enough to pass through the optical filter at A to B, a very disturbing phenomenon appears. The return signal from A to B causes an extra “echo” pulse at IOB! This extra set of edges at IOB continues until the pulse width applied at IOB is longer than two (2) optical turning-on delays. After that the signal at IOB is stretched by the delayed echo, but this is not a problem, moderate amounts of pulse stretch are generally allowed by I^2C timing specifications.

Figure 12 Simple Non-Standard Levels Double-Pulses at Side B

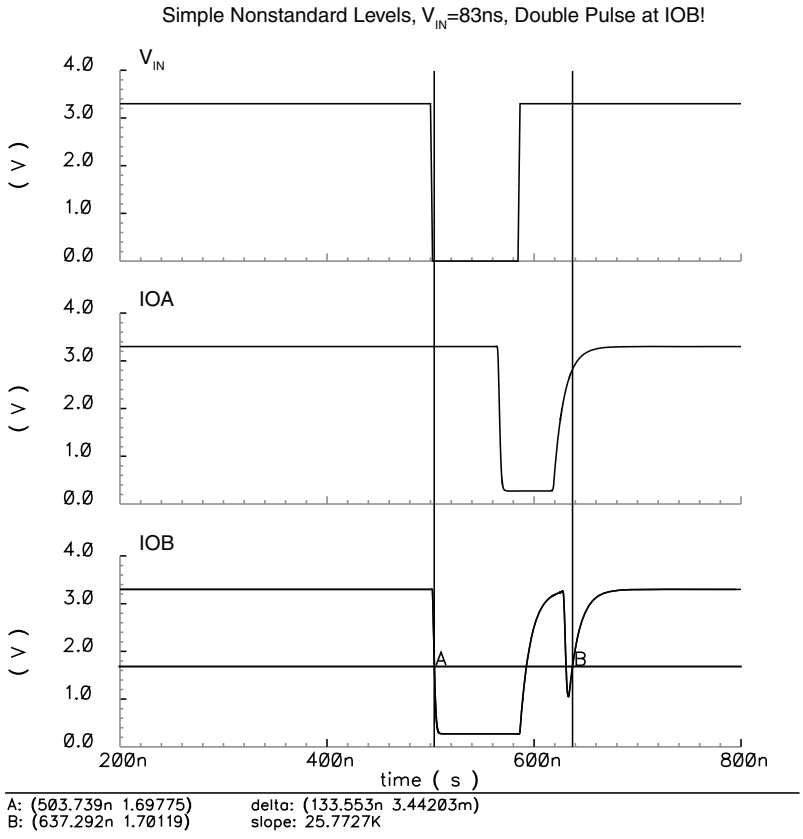

6 CPC5902 & CPC5903 Methodology

In order to overcome the pulse-shaving and double-pulsing problems, the CPC5902 and CPC5903 use a filtered, pulse-stretched, self-driven system with the three-level method on their bidirectional channels in a manner similar to that shown in **Figure 13**.

Figure 13 Three-Level with Pulse-Stretch and Self-Drive

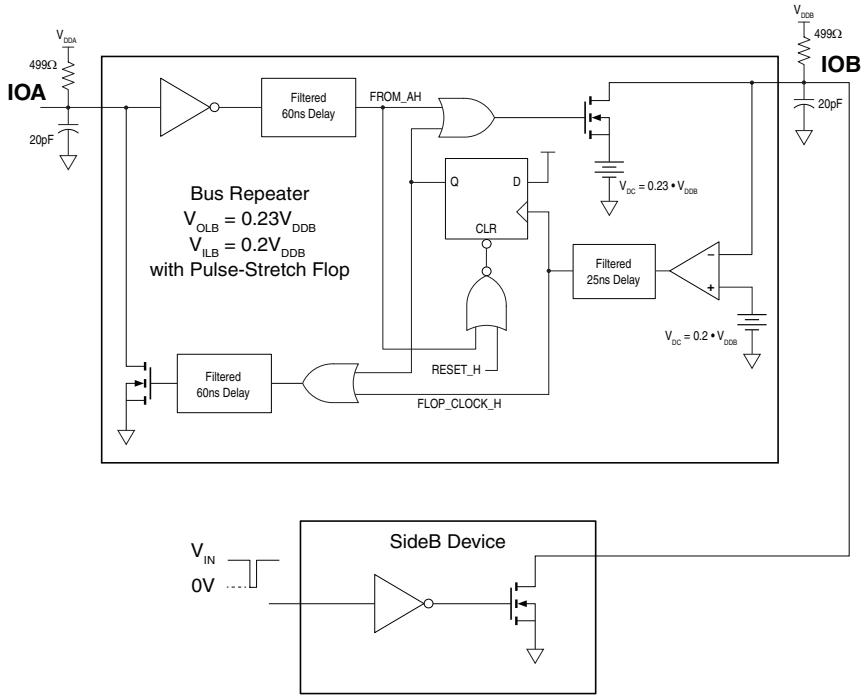

6.1 Filtered Pulse-Stretch Self-Drive Three-Level Method:

V_{IN}	IOB	IOA	Comment
<25ns	<27ns	none	fine, suppresses short pulses
25ns	271ns	123ns	pulse-stretches IOA, IOB
90ns	271ns	123ns	pulse-stretches IOA, IOB
120ns	271ns	123ns	pulse-stretches IOA, IOB
150ns	310ns	154ns	pulse-stretches IOB
1000ns	1198ns	1022ns	pulse-stretches IOB

Signals asserted by devices on the bus at IOB first encounter an input filter which rejects very short pulses. In the example circuit this is shown as delay_25ns in series with the input at IOB. When using the CPC5902 and CPC5903, the effective low-pass filtering of this block can be increased by increasing C_{LOAD} at the input. A delay of 25ns is used here to differentiate this filter from the optical filter. Pulses applied to I/O at SideB (IOB), which are long enough to pass through the input filter, assert the signal FLOP_CLOCK_H. Assuming that a short pulse just long enough to generate a brief logic high at FLOP_CLOCK_H has been applied, then that brief logic high is routed through the OR gate to the B_to_A optical channel for transmission. The presence of this "long enough to be legal" signal is also stored by the flip-flop.

Figure 14 Non-Standard Levels with Flip-Flop

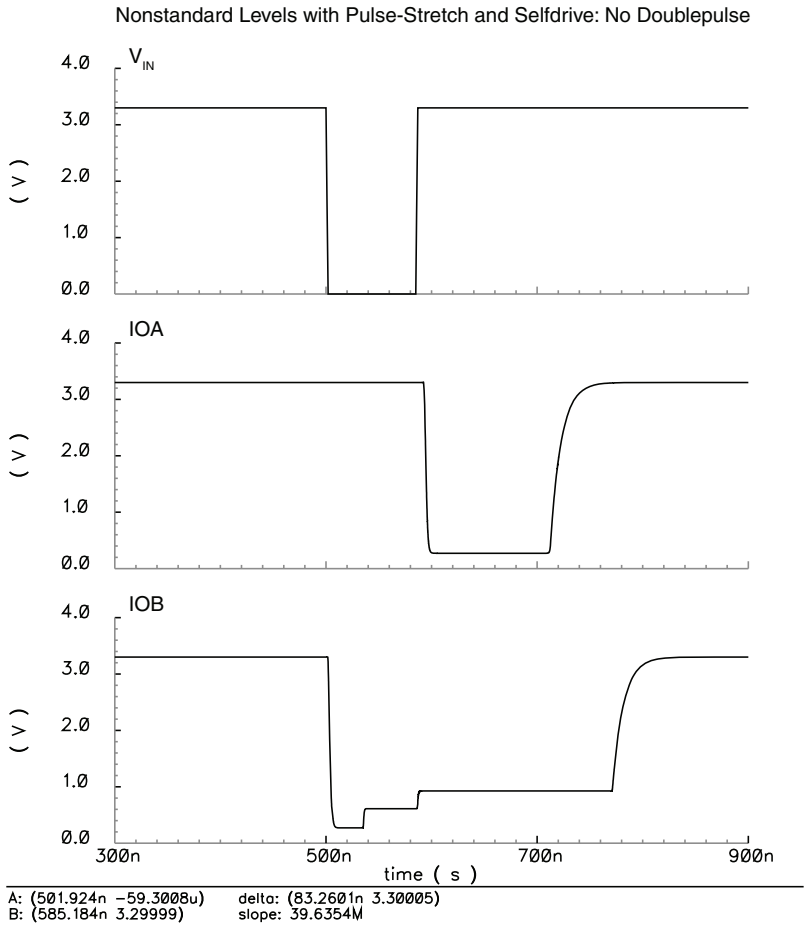

The stored signal is then applied to the non-standard output level driver at IOB, which normally drives the I/O to $0.23V_{DDB}$. Thus the part self-drives its own input while it is already being driven by a device on the bus. Unlike the simplified drive circuit shown here, the circuit used in the CPC5902 and CPC5903 can only sink current, having little effect on the signal while it is still being driven below $0.2*V_{DDB}$ by a device on the bus.

While self-driving the IOB port, the stored signal is also applied through the OR gate to the B_to_A optical channel. A short FLOP_CLOCK_H signal can return low, but the stored signal will remain high. The OR gate at the input to the B_to_A optical filter insures that the optics will continue to be driven until an edge is returned though the A_to_B optical filter to clear the flip-flop. Thus the drive into the B_to_A optics is stretched to a minimum length of two (2) optical turning-on delays. As was shown in the simple three level case above, the minimum length of the stretched pulse is now always the minimum required to avoid double pulsing. Moreover, the stretched drive-pulse-width tracks the actual delays at the two optic filters, so that it is always of optimal length. This feedback can also be interpreted as an error reduction loop. The B_to_A optics continues to be driven until it has been verified by an A_to_B transmission that the data successfully arrived.

Upon a turning-on edge being received at FROM_AH (see **Figure 13**), the flip-flop is cleared. Assuming that the device driving IOB has ceased externally driving it, then the drive into B_to_A is de-asserted. The de-asserted edge propagates through a B_to_A turning-off delay, then a delay determined by the $R_{PULLUP}*C_{LOAD}$ at IOA. After an additional A_to_B turning-off delay, the deasserted edge becomes available at FROM_AH, and is used to deassert the drive at IOB. Thus the minimum pulse at IOB is:

(delay_on_BtoA) + (delay_on_AtoB) + (delay_off_BtoA) + (RCdelayA) + (delay_off_AtoB)

From the table, note that for pulses long enough not to be suppressed by the input filter, there are always the same number of edges asserted at both sides of the isolator. As the applied pulse-width becomes longer than delay_on_BtoA + delay_on_AtoB, the pulse-width at IOA is no longer stretched by the flip-flop circuit (although the pulse will be stretched slightly because the turning-on delay is shorter than the turning-off delay).

The pulse at IOB is then only stretched by a time determined by:

(delay_off_BtoA) + (RCdelayA) + (delay_off_AtoB)

Thus for legal length I^2C fast mode pulses (logic low longer than 1.3µs) the flip-flop circuit is not used to stretch the drive to the optics. The flip-flop circuit and the self-drive circuit at SideB will stretch an applied legal pulse at IOB by an amount primarily determined by the turning-off optical delays as above. This pulse stretch is similar to what happens to pulses applied to a resistor-pulled-up bus with a large capacitive load. The I^2C fast-mode specifications will tolerate stretches of a few hundreds of nanoseconds even at 400kbps rates, and will tolerate even more for smaller values of C_{LOAD}.

7 Conclusion

The ideal model of an open-drain bus as a simple, shared pull-up resistor is no longer instantaneously valid when using bus isolators exhibiting real-world propagation delays. If the delays are not carefully considered in the design of the hardware, then there are potential differences in the number of rising or falling edges seen at each isolated bus.

The CPC5902 and CPC5903 family of isolated bus repeaters uses verification feedback from the standard level SideA back to the non-standard level SideB and self-drive of the SideB I/O in order to insure that the same number of clock edges are seen at both busses. This feature greatly reduces the probability of an undetected error in data transmission across the bus.

For additional information please visit www.clare.com

Specifications: TB-101-July 1, 2011