Ariane Frey

Voltage Level Shifting Options for the SVT MUX Modules

December 5, 1996

In answer to some questions raised at the BaBar SVT Preliminary Design Review on October 22, 1996, this note is intended to present some more details on the existing options for performing the necessary voltage translation of the clock, command, and data signals sent to and coming from the front-end IC's. The nominal voltage levels for the detector bias are +30 V and -30 V for n– and p–side, respectively. The signals provided by the SVT DAQ system and sent back are on system level which is at the center tap of the detector bias voltage. The front–end electronics runs at the given detector bias voltage in order to minimize the voltage applied across the coupling capacitor on the detector. All digital voltages are +5V and 0V from their respective grounds.

The issues raised addressed mainly the question of safety for the front–end electronics as it is deeply buried in the detector and only accessible in a major access, current separation and reliability. The MUX Module itself is located outside the detector and can be accessed if the beam is turned off.

The MUX Module is described in more detail in the PDR note. In short, one module services 4 HDI's (8 sides), thus 8 clock lines, 8 command lines and 16 data lines, 32 channels in total.

Relevant requirements for the level shifting are:

— Voltage translation of the signals to the nominal $\pm$ 30 V and vice versa.

— Signal transmission at 60 MHz, i.e. a 120 MBaud rate on the clock lines, 60 MBd on command and data lines.

— The 8 command bit streams arrive via the G–Link in parallel and the individual time delays after the whole MUX electronics chain must not be greater than 3–4 ns. The same applies for the 16 data lines coming back from the front–end IC's since they have to be latched into the G–Link simultanously with a window of about 8 ns. The voltage translators should contribute as less as possible to time differences in individual channels.

— For all the considered options it is mandatory that an error in a transmitted bit does not affect the bit information that comes afterwards.

1 Transistor option

This voltage isolation option was the first one considered and is realized on the first prototype MUX Module that is currently at Santa Barbara for the electronics bench test. Thus, it is the only fully prototyped option that is shown to function in talking to the front–end chips. Figures 1 and 2 show the schematics for the circuit. The voltage level shifting is done with discrete bipolar transistors (pnp–transistors: BF569, npn MMBT24LT1) which easily meet the requirements on speed and fast switching. The signal translation is performed by 2 cascodes and a third transistor pair that differentially receives the signals. Much effort was put into a careful biasing of the transistors in order to distribute the total voltage drop as evenly as possible over the 3 transistors.

For testing purposes, a slightly inconvenient feature is that the circuit requires a minimum voltage offset of about 7 V to work properly.

- **Current separation:**
  In this scheme the current paths of the two different voltages (system level and $\pm$ detector bias voltage) are not completely decoupled. A small current of 2.5 – 3 mA is flowing from the voltage source at system level through the detector bias return. Whether this poses a considerable problem or not is not yet fully understood.

- **Safety and reliability:**
  Each of the 3 transistors on one leg of the differential pair is rated to withstand a collector–emitter voltage of 30 V (npn) and 35 V (pnp). Hence, there is a safety margin of a factor of 3. If one of the transistors shorts out, the other two (or even one) is able to withstand the full applied voltage.
  An issue that needs further study is how one will detect that one transistor has failed so that it can be replaced in order not to compromise the HDI electronics safety.

- **Cost:**
  The transistor option is a very cheap one, each of the surface mount transistors just costs $ 0.20 or $ 0.35.

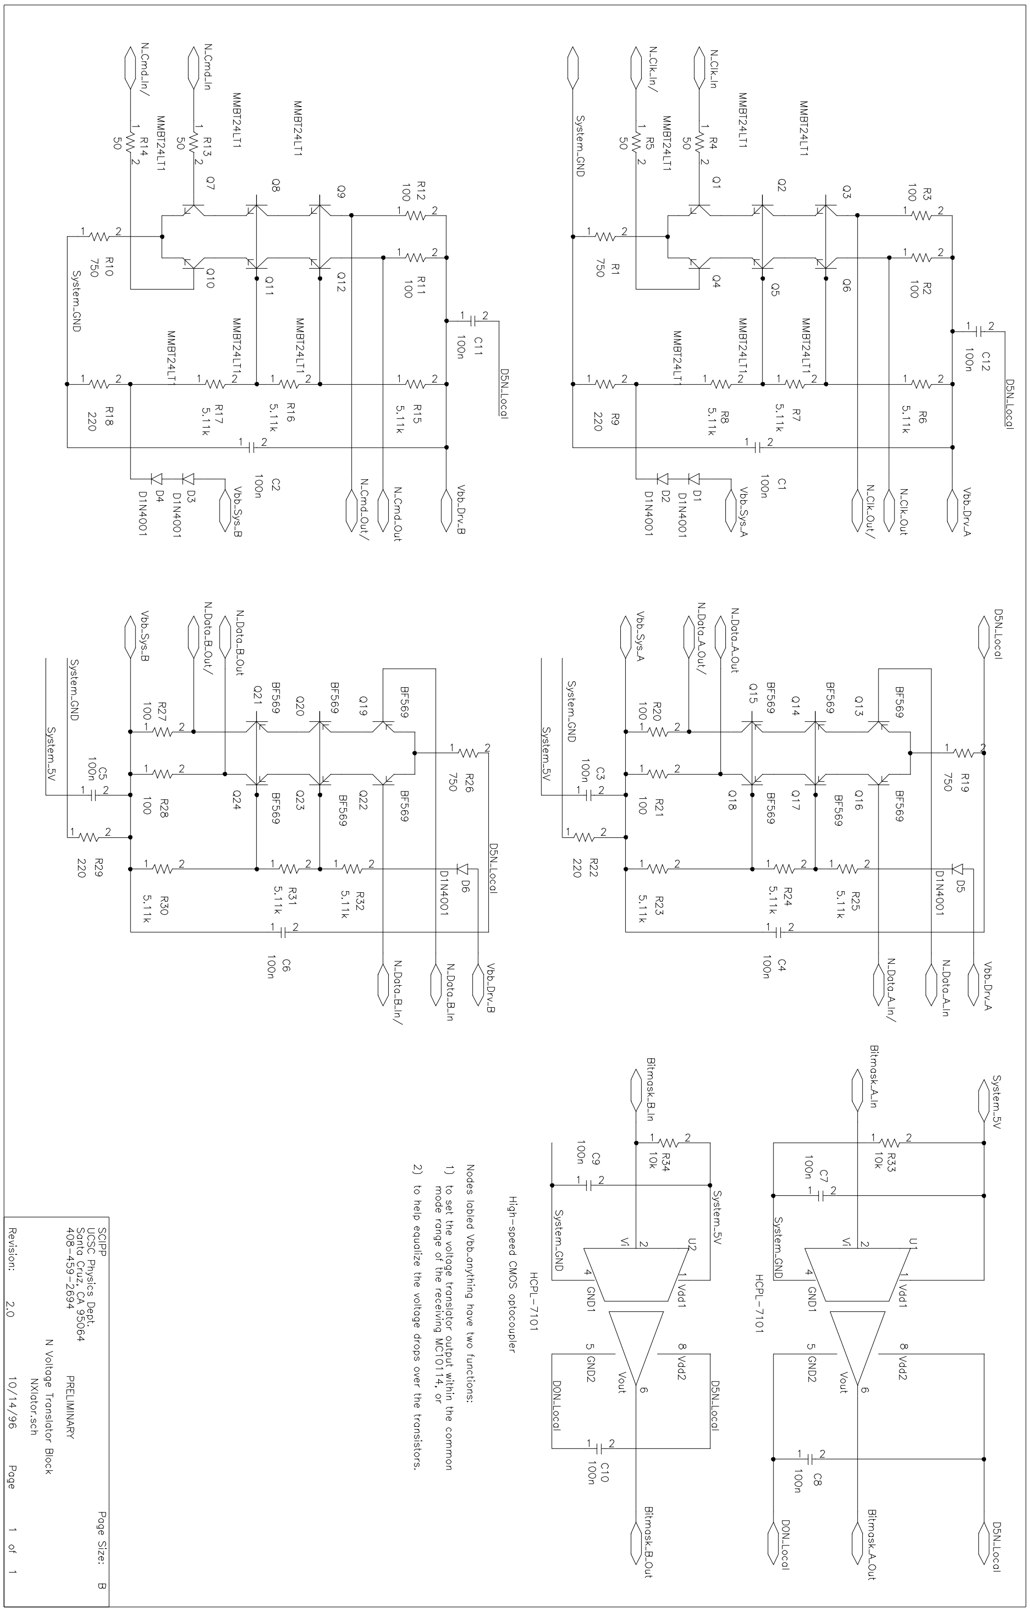

Figure 1: Schematic of the transistor level shifter for the n–side.

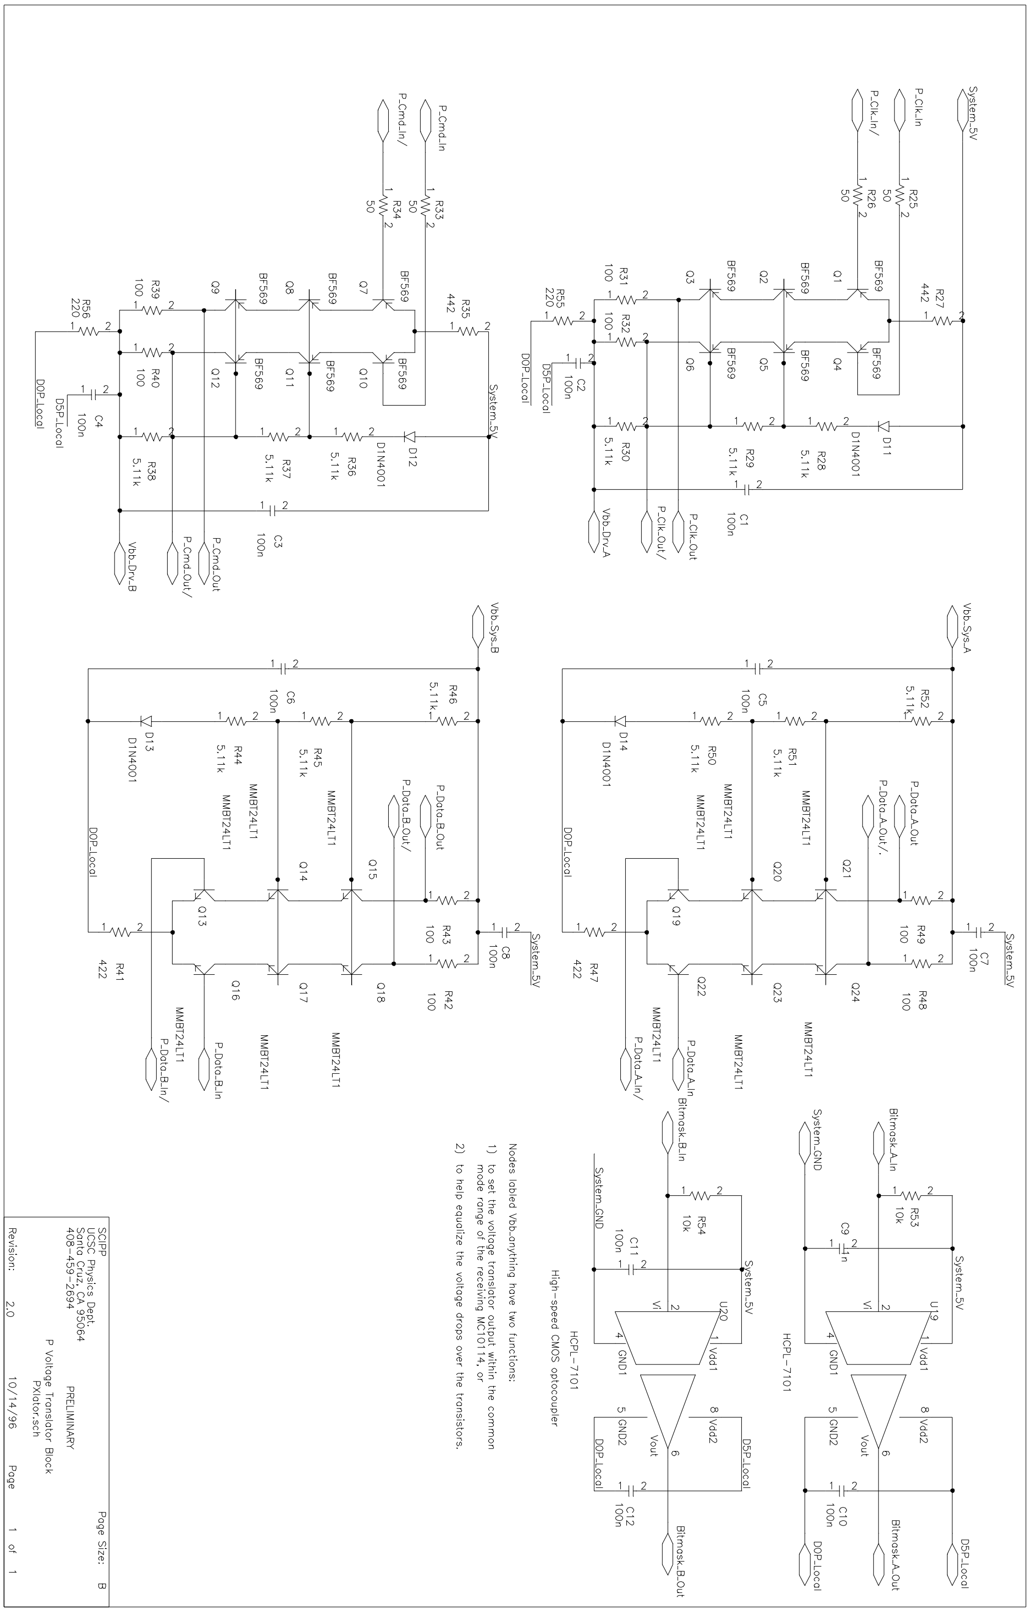

Figure 2: Schematic of the transistor level shifter for the p–side.

2 LED option

The use of optical signal transmission provides a feasible and safe way of voltage isolation. The Hewlett–Packard optical fiber transmitter and receiver, parts number HFBR–1404 and HFBR–2406, respectively, can be used for this purpose. The idea is to use them without the optical fiber in between, but directly mechanically couple the transmitter and receiver instead.

The transmitter/receiver pair is rated at signal transmission rates up to 155 MBd, and during testing we could see signals at frequencies up to 200 MHz getting through, so the bandwith is obviously not a problem here.

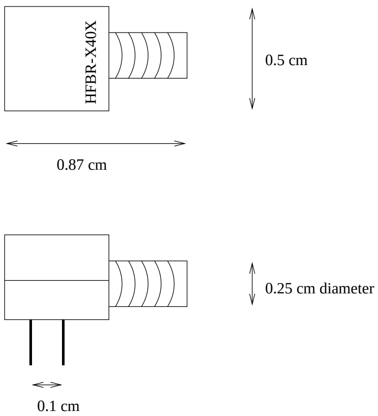

Figure 3: Geometrical dimension of the optical transmitter and receiver.

The geometrical dimensions of the parts are shown in fig. 3. The measurements were performed with the transmitter and receiver soldered onto a testboard facing each other and with the ports aligned w.r.t. each other by means of a aluminum tube. Better ways of optical coupling by cutting the plastic ports off and gluing the parts together in order to bring the optical planes closer together will be studied. The wavelength of the emitted light is 820 nm (GaAlAs). The receiver contains a PIN photodiode and a low noise transimpedance pre–amplifier integrated circuit with a buffered emitter–follower output.

The design of the transmitter driver circuit requires some consideration as the LED is a highly capacitive load ($\sim$ 55 pF at 1 MHz) and the driver must be able to quickly source current into the LED. Tests were performed with the driver circuit shown in fig. 4 where the signals enter as differential PECL. At high speeds a pre–biasing of the LED is recommended and is done via the resistor R3 which provides

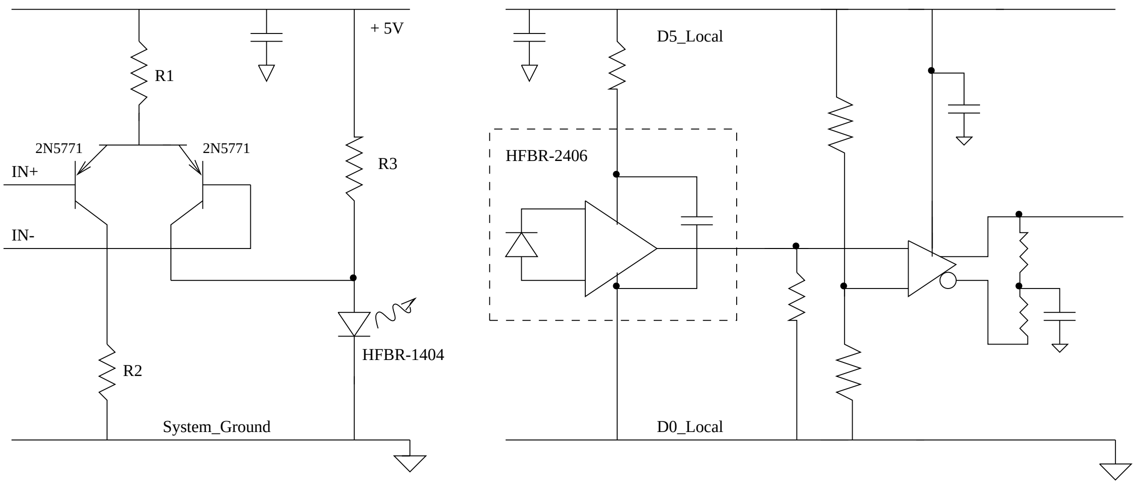

Figure 4: Proposed driver and receiver circuitry for the LED option.

a constant small current through the LED. HP recommends a ratio of 3.91 between the forward current and the pre–bias current for optimum performance. In order to achieve fast on/off-switching of the LED, a forward drive current of $I_F \gtrsim 10$ mA is needed. Our tests were performed with a drive current of $I_F = 30$ mA, we hope to reduce this considerably with an improved optical coupling. It should be noted, however, that even a drive current of 30 mA is very low for these devices, all specs given by HP are measured at 60 or 100 mA (see also the remarks on reliability below). The forward voltage drop over the LED is about 1.5 V at 15 mA and quite stable with temperature (HP measures -0.22 mV/ ^{o}C $I_F = 60$ mA). For $I_F = 30$ mA the total current draw including the pre–bias and the receiver will be about 70–90 mA.

The proposed driver circuit has been prototyped and has been shown to work. The proposed receiver circuit still needs to be breadboarded, so far the signals coming out of the receiver were directly looked at without amplification and digitization. For a drive current of $I_F = 30$ mA we observed signal swings of roughly 100 mV and signal rise times of ~ 3 ns.

- **Current separation:**
  The current paths of the two systems at the different voltage levels are completely separated. Whether the switching of the relatively large currents (compared to the transistor option) from one leg of the differential driver to the other causes a problem for the system due to transients still needs to be addressed.

- **Safety and reliability:**
Safety for the read–out chips on the HDI is not an issue with optical isolation, if a component fails, this channel will be dead, but there is no risk of the full bias voltage being applied to the HDI.
Addressing some concern at the PDR about reliability (need of replacement of a device on the MUX card): HP provides fairly detailed information on reliability of the HFBR—xxxx devices. For the transmitter, failure includes not only an electrical failure, but also when the light output decreases to less than half its original value. At 50^{o}C and I_F = 60 mA their mean–time–to–failure (calculated as device hours/failures) is 2.12 M hours, assuming a 100 % duty cycle. The dependence on the forward current goes as $I_F^{-2.3}$, so for (conservatively taking) $I_F = 30$ mA and a 50 % duty cycle (for the clock it will be 50 %, for data and command lines surely much less) the mean–time–to–failure should be 20.9 M hours, 2383 years. So with 832 channels one expects one failure every 2.8 years for continuous operation. For the receiver the rate at 50^{o}C is 1 failure in 0.6 years continuous operation, for a more realistic temperature estimate of 40^{o}C it is 1 failure/year[1].

- **Cost:**
The cost for the transmitter and receiver are \$ 18.00 and \$ 21.75, respectively. For the driver and receiver circuitry we estimate about \$ 7, summing up to \$ 46.75 per channel (32 channels per MUX Module)

3 Optocouplers

Transmitting signals via Optocouplers is another way of optical isolation. The advantage over the use of LED's is that one doesn't have to deal with the optical coupling of the transmitter and receiver pair, but everything comes quite handy in one DIP chip package without much surrounding circuitry needed. Also the current is lower. Nevertheless, the requirements of high speed poses a problem here, there are no optocouplers on the market that can handle a signal of 60 MHz (corresponding to 120 MBaud) as needed for the clock. The clock signals therefore would have to be

[1] A possibility to check that all the channels are still functional with good light yield in order to avoid unnoticed slow death of channels might be to provide two settings of comparator thresholds, a nominal one and one twice as high which can be selected for example by a switch. During a non–running period every few month one could systematically check whether all channels are still responding with the higher threshold setting.

transmitted by other means (transformers or capacitors as described below). Hence, in the following only the transmission of command– and data–lines are considered, requiring signal transmission at 60 MBaud.

Tests have been performed with the highest speed optocouplers from Hewlett–Packard HCPL–7101 which have a guaranteed 50 MBd data rate (typical 65 MBd, but this number is not guaranteed) and withstand voltage of 600 V continously, 3750 V/1 min.

The HCPL–7101 chips are CMOS technology and incorporate a CMOS LED driver IC, a AlGaAs LED and a CMOS detector IC consisting of an integrated photodiode, a transimpedance amplifier and a comparator with hysteresis. The outputs are 3–state and the chip can be operated with a total current $I_{tot} \gtrsim 10$ mA. So far so good.

The not so nice features of the HCPL–7101 are a quite long propagation delay time, quoted to be typically 27 to 28 ns and propagation time skew spec'ed at 10 ns maximum which would not be acceptable for a good enough synchronization of the 16 parallel data channels. The propagation delay should not change more than about 1 ns over a temperature range from 25 to 40°C. We performed some tests in the lab with 8 different chips (probably all from the same lot, though) and measured propagation delays at 30 MHz from low–to–high of 23.7 ns on average with every chip within the range 23.1–25 ns and delays from high–to–low of 23.4 ns on average, distributed between 22.6 and 24.9 ns. The pulse width distortion was on average 0.5 ns and below 1 ns in all cases. The propagation delay time was strongly dependent on the supplied voltage, and changed by more than 3 ns when changing the supply voltages from 4 V to 5 V.

- **Current separation:**
  As for the LED optical solution, the current paths between the system voltage and the bias voltage are completely separated.

- **Safety and reliability:**
  Given the specifications there should be no danger of a short of any kind between the different voltage rails.
  HP does not provide any data on the reliability of their optocouplers or mean–time–to–failure estimates. From our own experience with these optocouplers which are used on the MUX Module prototype for the (more or less DC) bitmask voltage level shifting, the optocouplers seem to be extremely sensitive

to handling. We actually had to replace 2 out of 4 chips[2]. M. Momayezi from LBL reported similar problems. We therefore are not convinced that this solution will prove to be reliable enough.

- **Cost:**
  The HCPL–7101 chips cost \$ 4.75 each and need little surrounding circuitry.

4 Capacitive coupling

4.1 Clock lines

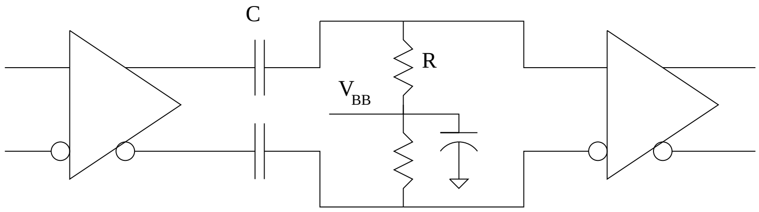

Figure 5: Scheme for the ac–coupling of the clock lines. V_{BB} denotes a convenient voltage to bias the signals to the accepted voltage range of the 10114 ECL buffer chips.

At least for the clock lines (for command and data lines refer to the next section) a simple capacitive ac–coupling is applicable as the clock is continuously toggling at the nominal rate of 60 MHz. A simple ac–coupling for the clock lines is depicted in fig. 5. This circuit has been breadboarded and seems to work fairly well. The signal shapes improve with smaller resistor values at the cost of larger current draw. Tests were performed with $R = 50\ \Omega$ and $C = 0.1\ \mu F$.

- **Current separation:**
  No current is returned through the bias supply.

- **Safety and reliability:**
  Usual surface mount capacitors are rated for 100 V and are available at 500 V ratings, too. To increase the safety margin one could also put 2 caps in series.

[2]For the final design we intend to use a different kind of optocouplers with lower bandwidth and higher reliability to translate the bitmask.

Also, it should be pointed out that the voltage level shifting is not the only place in the system where capacitors are used across the bias voltage, this is also done on the Matching card.

- **Cost:**
  Cheap !

4.2 Command and data lines

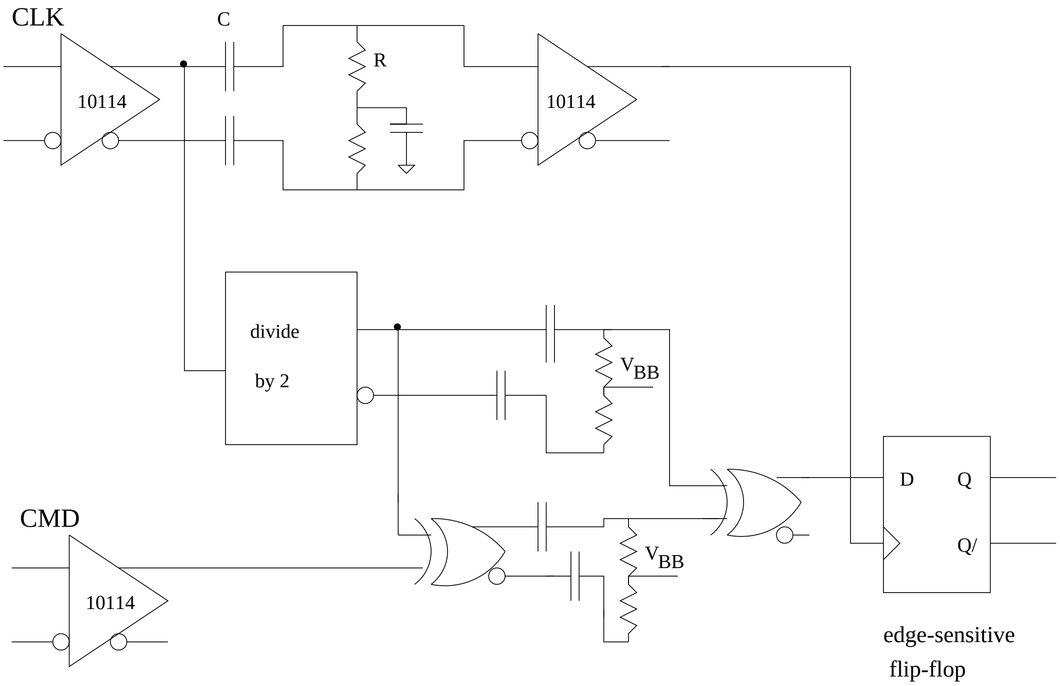

Figure 6: Scheme (not a schematic !) for the chopper circuit to transmit cmd/data.

This is based on an idea of Mark Nyman from LBL. A drawing of the principle is shown in fig. 6. Capacitive coupling of the cmd/data lines faces the problem that the signals can have long streams of just 1's or just 0's. This is overcome by "chopping" the signals with the clock signal divided by 2 (i.e. at the same frequency as the cmds/data). The derived clock/2 is exclusively OR'ed with the commands/data and this signal is then capacitively coupled. The original command/data signal is restored via a second X–OR of the voltage translated clock/2 and cmds/data. To avoid glitches on the output of the voltage translator the signals are put through an edge–sensitive Flip–Flop triggered off the clock at the nominal frequency.

As the clock/2 and the cmds/data have the same period one can think of a worst case scenario where the cmds/data contain a long stream of 101010... If clock/2 and cmds/data are exactly in phase the X–OR would result in a constant output value. Thus, to ensure a signal toggling at a frequency of $\geq$ 30 MHz totally independent of the bit stream on cmd/data lines the chopping clock/2 has to be phase offset w.r.t the cmds/data.

Whether the possible phase window for the clock that triggers the flip–flop is wide enough for reliable performance is currently under study. The rising edge of the clock should be at least 2 ns after the cmd/data signals arrive at the flip–flop inputs, and also it has to come well before the unavoidable glitches that propagates through because of the chopping.

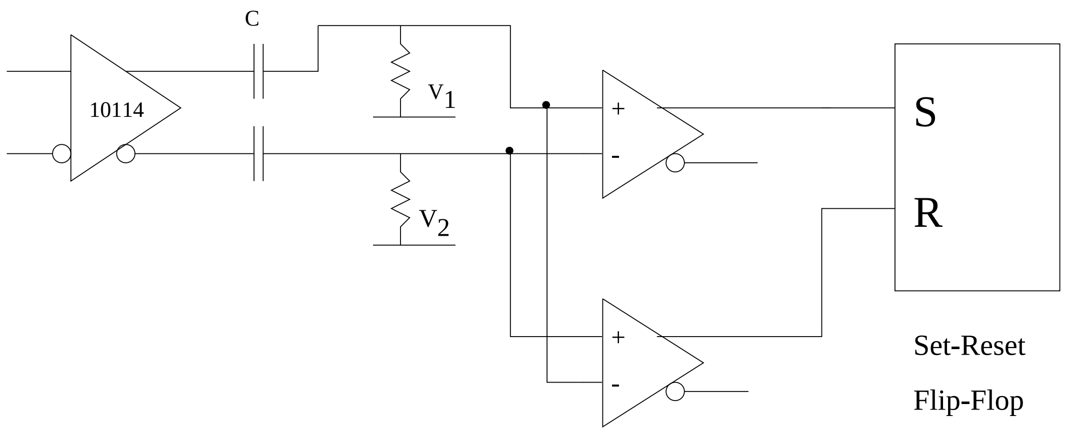

Figure 7: Sketch of the RS–Flip–Flop circuit

A second proposed alternative does not make use of the clock to chop cmds/data, but of the fast signals on the capacitor when there is a transition from 1 to 0 or vice versa in the data stream. A sketch is shown in fig. 7. The time constant of the RC circuit is chosen such that only the fast edges are transmitted through the capacitor, giving a positiv or negativ going signal of roughly 4 ns width and biased with V_1 and V_2 which are offset by a diode drop. The signals are differentially fed into two AD9696 fast comparators which then are connected to the set and reset–inputs of a Flip–Flop. Hence, when there is a positiv going signal on one and a negative going signal on the other capacitor, the comparator outputs may change and consequently the flip–flop may change its output state. A pulse only on one capacitor should not change the outputs in order to suppress noise.

It should be noted that the capacitive coupling options do consume a fair amount of power to achieve nicely shaped signals. An estimate of about 30 mA per chip gives a total current estimate of 100–150 mA for the chopper circuit and $\sim$ 100 mA for the RS–Flip–Flop alternative. The current draw is to first order constant and independent of the logic states.

5 Use of transformers for the clock lines

There are transformer chips available which meet the high speed requirements. However, we see no advantage in a transformer w.r.t. the use of capacitive coupling of the clock lines. Hence, this alternative will not be pursued further.

6 Conclusions

Table 1 summarizes the different options described above.

	Power Consumption	Reliability	Safety for HDI	System Interference	Cost per channel
Transistor	3 mA	2-3	4	small current on bias return	$< \$5$
LED	70–90 mA	2-3	1	$\emptyset$	\$47
Optocouplers	$\gtrsim$ 10 mA	4	2	$\emptyset$	\$5
Capacitors	100–150 mA	1	3	$\emptyset$	few \$

Table 1: Comparison of the different options. Under reliability and safety an estimated rating of the 4 options is given (1 best, 4 worst). System intereference addresses the current path separation issue.

For the clock lines, we think the most preferable option is a capacitive coupling as it is the most simple one, it is cheap and the current draw will be less than needed for the cmd/data lines as in this case no additional IC's are required.

The following discussion therefore only refers to command and data transmission. Our preferences are

1. Capacitive coupling
2. LED's
3. Transistors
4. Optocouplers

The steps to arrive at a final design decision we suggest as follows: If the capacitive coupling via either the chopper or the RS–Flip–Flop method is shown to have good performance we would build a prototype of the whole chain (not only the level shifting) of the MUX Module for 1 channel. If this is successful the prototype printed circuit board would incorporate this option. The only obvious disadvantage over the other options is the higher power consumption, but as the card will be located in a water cooled rack this should cause no problem. Advantages are simplicity and cost.

If the capacitive coupling for cmd/data turns out not to work well enough we will investigate the LED option further which means mainly finalizing the design of and testing the receiver circuitry as well as improving the optical coupling.

In case also the LED option should show serious problems, we have the transistor version as a fall–back solution, it is working in the bench test and does not need further design studies.

The optocoupler solution we prefer the least since the devices would have to work somewhat outside the specifications guaranteed by the manufacturer and we do not have enough experience with them to tend to a design including the optocouplers at the limits of their capability.