Performance and Comparison of Custom Serial Powering Regulators and Architectures for SLHC Silicon Trackers

T. Tic[ab], P. W. Phillips[a], M. Weber[a]

[a] STFC RAL, [b] IoP ASCR

Tomas.Tic@cern.ch

Abstract

Serial powering is an elegant solution to power the SLHC inner trackers with a minimum volume of cables. Previously R&D on the serial powering of silicon strip detector modules had been based on discrete commercial electronics, but with the delivery of the Atlas Binary Chip Next chip in 0.25 micron CMOS technology (ABCN-25) and the Serial Powering Interface chip (SPi), custom elements of shunt regulators and transistors became available. These ASICs can be used to implement three complementary serial powering architectures. The features of these schemes and their performance with 10 and 20 chip ABCN-25 hybrids will be presented.

I. Introduction to Serial Powering

The following subsections will introduce serial powering in the context of experiments on SLHC.

A. A problem and a solution

In the current ATLAS experiment at CERN's Large Hadron Collider (LHC) the SemiConductor Tracker (SCT) comprises 4088 detector modules, each powered by its own power supplies through its own cable. The overall mass of the electrical services is significant and since the path between the detector modules and its radiation intolerant power supplies is long, power losses in the cables are also significant. One of the most important differences between LHC and its upgraded form, super-LHC (sLHC), is the ten times higher projected beam luminosity, resulting in much higher particle hit occupancy. There are not many options how to decrease the number of ghost hits of a micro-strip silicon tracker with binary read out other than making the strips shorter. This in turn results in a much higher number of readout hybrids. With the current power distribution scheme the mass and volume of the electrical services, and the power dissipated in them, would be simply unbearable.

The problem sounds similar to domestic power distribution. Currently we are adjusting the power plant output voltages so that independently powered households receive their 230 volts. Some sort of voltage vs. current trade-off with power management moved closer to the detector readout hybrids will be required with the SLHC in order to keep the services volume and power losses bearable. The comparison with household power distribution is not chosen randomly as an alternative to serial powering would be to employ DC/DC converters, similar to the way in which households employ transformers for AC/AC conversion. Serial powering is less conventional in this respect - and reminds us of occasionally troublesome Christmas tree lights - but in the case of particle detectors it provides a viable solution to the power distribution problem.

B. Serial Powering - System Overview

The voltage vs. current trade-off is very simple with serial powering. A number of detector readout hybrids are connected in series and supplied from a single current source as shown in Figure 1. Several tens of hybrids can easily be connected in this manner. The voltage on each hybrid is regulated locally using one of the shunt regulators, the main topic of this talk. The underlying concept is that the hybrids are electrically similar, drawing similar currents, so the current overhead needed to ensure that the correct voltage may be maintained on each hybrid in the chain can be kept very low. Only one voltage is obtained directly from the shunt and any other required voltages must either be obtained from it or provided separately. High voltage for biasing the sensor shall most likely be provided common to two detector hybrids (as they will sit on the same sensor). Even though the hybrids in a serially powered chain sit at different potentials with respect to ground, communication to and from the outside world does not require especially advanced coupling. Protection circuitry is also required to maintain the integrity of the chain in the event of open loop failures. Just a few '0402'-sized capacitors and perhaps 15 mm^2 of silicon can be the total mass overhead of serial powering.

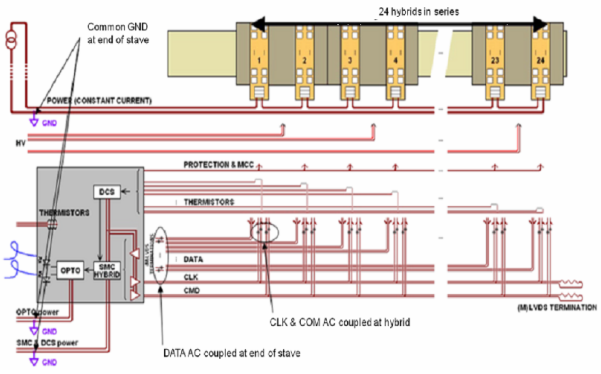

Figure 1: System overview with serial powering [3]

C. Serially powered chain of detector hybrids

In order to look at the chain of serially powered detector hybrids a little bit more analytically, it is convenient to introduce a couple of simplifying conditions like linearity, time invariance

and e.g. Norton's representation of the power supply. These are acceptable conditions in the small signal region and the model of the chain then reduces to its dynamic impedances as illustrated in Figure 2.

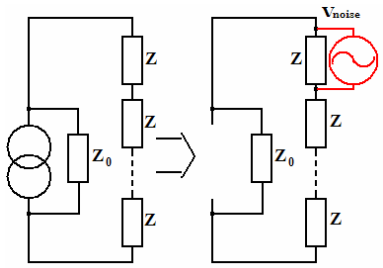

Figure 2: A Generic Chain of Serially Powered Devices

It is useful to quantify how much of a small signal voltage generated on one hybrid in a serially powered chain is transferred onto the others ($A_{V_{noise}}$). This coupling is described by Formula 1.

$$A_{V_{noise}} = \frac{Z(\omega)}{(n-1)\cdot Z(\omega) + Z_0(\omega)} \qquad (1)$$

One can immediately see that the choice of current source not only simplifies the DC conditions of the chain but it also prevents individual detector hybrids from "seeing" each other. The impedance of the detector hybrid should be as low as possible and the output impedance of the current source should be as high as possible. In the low frequency range this is important to prevent oscillation modes in the chain while in the high frequency range this is important to minimize the spreading of any noise to which the detector system may be sensitive. One should make two important remarks at this point. Firstly, no oscillation modes have ever been observed during our studies, even with obsolete shunt regulator designs and current limited laboratory power supplies. Secondly, the circuit designs of the shunt regulators and current source matter only up to a few MHz, above this frequency the impedances are dominated by other factors such as the hybrid layout, decoupling and power cable impedances.

II. Serial Powering and ABCx chips

This section will only present the information useful for further elaboration in the text.

A. Serial Powering and ABCD chip

To give a "proof of principle" result, several serially powered staves were constructed using the ABCD chip (the chip used in the current ATLAS SemiConductor Tracker). The largest of these staves operated a chain of 30 hybrids using serial powering interface boards such as the one shown in Figure 3, based on commerical components. All staves were successful and any initial worries about the concept of serial powering were despatched.

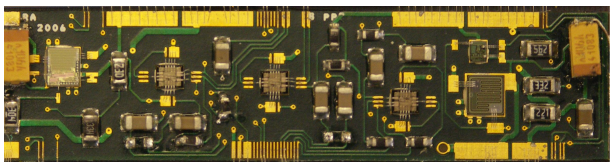

Figure 3: Serial Powering Interface Board, comprising a shunt regulator and AC coupled LVDS buffers for communication.

B. Description of the current ABCn chip/hybrid

The ABCN-25 readout chip incorporates several functional blocks for serial powering which will now be discussed in detail. The chip requires digital and analogue supply voltages. The nominal digital voltage (2.5V), is only slightly higher than the nominal analogue voltage (2.2V), Hence a linear regulator is present on the chip to derive the analogue voltage from the digital voltage. When considering the performance of a shunt regulator it should be taken into account that the overall power supply rejection ratio (the PSRR of the linear regulator convoluted with that of the analogue front end) of the ABCN-25 chip is rather large. If any increase in equivalent noise charge is observed, the disturbance to the digital voltage must be enormous. The likely cause would then be that the digital voltage is too low, with fast changes.

The current consumption of the ABCN-25 chip was expected to be largely dependent on whether the clock signals were present and upon how its internal registers had been configured. Only a minimal, short increase in digital current consumption was expected each time a L1A trigger is received by the chip, corresponding to increased switching activity as the readout cycle begins. Under such circumstances the task of a shunt regulator would be to maintain the correct DC voltage and to provide some small signal filtering, as large changes in current should not occur.

Unfortunately in this version of the chip, the digital current consumption exhibits strong variations coming long after ($\sim$400 μs) L1 accept. The size of this current "bump" depends strongly and monotonically on the discriminator threshold and weakly on many other parameters, but it is always there. If another L1 comes before the previous "bump" has ended the whole process resets and the current consumption drops very quickly to the original level.

This effect produces very sharp, large, time structured peaks in the current consumption of the chip (up to 1.4 A for a 20 chip hybrid, but also observed using single chip ABCN-25 test PCBs).

This is a much more demanding test for a shunt regulator than the expected rare and accidental "worst case" of turning the clock off. Some examples of this are shown in Figure 4. Since L1 accept trigger rate is much higher than 1 per 400 μs this effect does not pose any troubles during standard data taking. Efforts to understand this effect continue. It is likely to disappear with the next version of the chip.

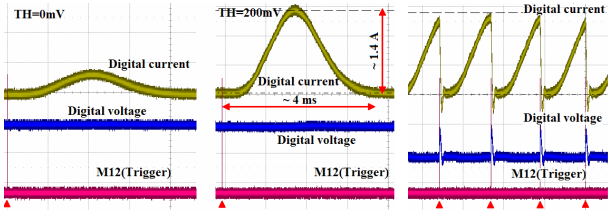

Figure 4: The Current Bump. Red triangles indicate L1A triggers. Left: The bump after single trigger at threshold set to zero; Middle: The bump after single trigger at threshold set to 200; Right: Multiple triggers. The scales are: time - 1 ms/div, digital voltage - 200 mV/div, digital current - 500 mA/div

III. The Three Shunt Regulators

In this section the three main shunt regulator architectures will be presented. Each is now available in fully custom circuitry. The two distributed options are named after their designers. The 'M' scheme designed by Mitch Newcomer is a distributed shunt with external feedback. The 'W' scheme designed by Wladyslaw Dabrowski is a distributed shunt with the internal feedback. The stand-alone shunt regulator option employs the Serial Powering Interface chip (SPi), a chip which provides additional functionality as will be discussed later. The key differences between the schemes are emphasized in Figure 5.

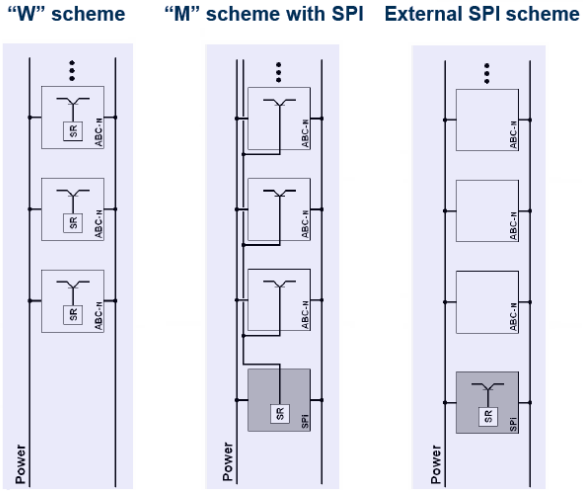

Figure 5: The Three Shunt Regulators Schemes [3]

A. 'M'

As can be seen in Figure 5 this scheme consists of two parts. The shunt transistor-like components are integrated in the ABCN-25 chip while the control scheme is to be somewhere else on the hybrid. First, let's have a look at what is inside of each ABCN-25 chip. There are two sets of current mirrors which can be seen in the schematic shown in Figure 6.

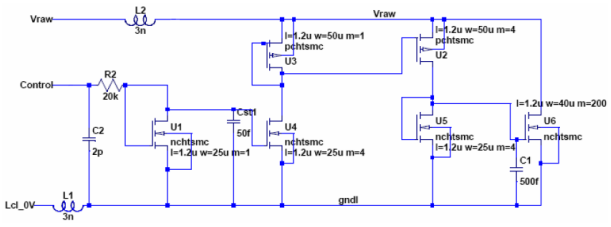

Figure 6: The shunt device present in ABCN-25 (two per chip) - current mirrors [2]

The two sets can be driven separately from a current limited dual output op-amp for improved reliability. This choice of shunting element has many advantages. The input capacitance is small. Charge injected into the control node of the chip is not directly transferred into the digital voltage. Small changes in the control voltage result in small changes in shunting current which deals with the matching problem due to the voltage drops across the hybrid and some small noise on the control bus which can result only in small noise on digital voltage. The linearity is very high when all the transistors are in the strong inversion region. Finally, the current mirrors are very fast and can go from zero to full current (∼140 mA) in less than 50 ns. The shunt is simply perfect and it can be decided later how to use it. The two shunts require very little silicon area and the scheme can be tuned without a new submission of the ABCN-25 chip. The digital voltage on the hybrid with respect to the control voltage exhibits all the properties of a plant (process gain, time constant, dead time) so it is no surprise that the basic control scheme (by M.N.) shown in Figure 7 (left) is reminiscent of an analogue PID regulator.

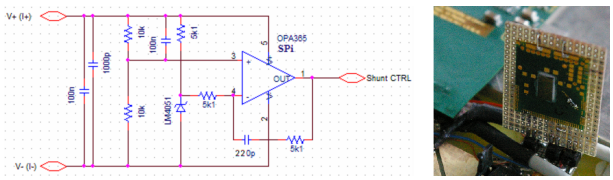

Figure 7: Left: Basic control circuit for the 'M' scheme, Right: Photograph of the current implementation of the control scheme with the hybrid

Intuitively it can be seen that the transfer of the circuit goes to the value given by the two resistors in the negative feed-back for higher frequencies. In the Figure 7 (right) the current implementation of the control scheme with the hybrid is displayed. It makes no sense to fine tune the circuit at this moment as especially the dead time depends on the position with respect to the hybrid. The next iteration of the Liverpool hybrid will have the control scheme incorporated. The 'M' architecture does, however, provide excellent results already. In Figure 8 the transient response of the system and its control voltage to a step in input current is shown. It can be seen that the control voltage reacts almost immediately and there is no overshoot on the digital

voltage. The time constant of the hybrid is $\sim$17 μs.

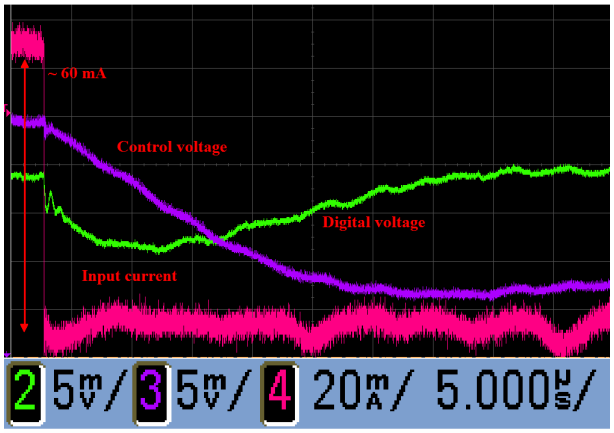

Figure 8: Transient response of the system in 'M' scheme (digital and control voltages as measured at the control circuit) to a step in current

This is extremely encouraging considering the improvised nature of the connection between the control circuit and the hybrid as used for these tests. If there is sufficient current in the chain, the bumps and slopes of the hybrid current consumption no longer appear on the digital voltage rails, as seen by an oscilloscope trace. Even if there is not enough current to cover the bumps, the superior step response of the circuit "softens" the digital voltage time profile so that no increase in ENC (equivalent noise charge) is observed. Figure 9 shows a typical ENC chart for a hybrid operated using this scheme. The ENC value is just below 400 electrons, in good agreement with the design value for the ABCN-25 chip, and the same as is obtained for a hybrid powered from a voltage source. The hybrid was not trimmed.

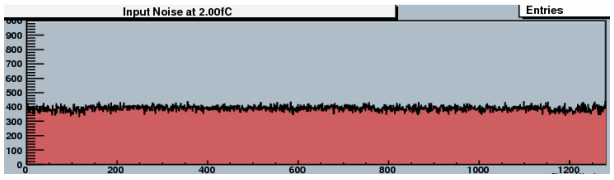

Figure 9: Typical ENC plot obtained using the 'M' scheme. For clarity, only half the channels are shown.

B. 'W'

As suggested in Figure 5 this scheme utilizes one complete shunt regulator within each read-out chip. The scheme is tempting because it does not require any external shunt regulation components. A classical design of a shunt regulator consists of a voltage reference, op-amp and shunt transistor. This design cannot work at all in the case of many shunt regulators connected in parallel. There are some IR drops across the hybrid and also the voltage references cannot be perfectly matched from manufacture which would both result in small number of chips shunting all the shunt current. The special design to overcome this difficulty is best explained by the conceptual diagram shown in Figure 10.

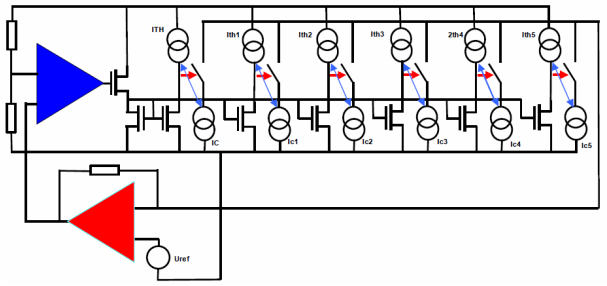

Figure 10: Conceptual diagram of the 'W' scheme shunt regulator present in ABCN-25. [1]

The shunt transistor is a P-MOS and its current is sensed and compared with six different current references. There is a trans-resistance amplifier which adjusts the reference voltage of the shunt op-amp. If the shunt current goes above one of the reference currents, the corresponding correction current source gets connected to the input of the trans-resistance amplifier thus adjusting the set-point voltage of the shunt regulator and the shunt current. One of the reference currents provides over-current functionality while the five others serve for shunt current redistribution within the hybrid during start-up. The 'W' scheme considers huge decline in current consumption of the ABCN-25 to be an accidental situation at which over-current protection should be activated. The shunt transistor is rather large for improved reliability.

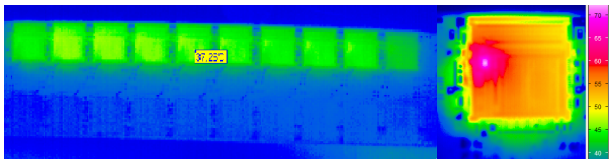

Figure 11: Infrared pictures of the hybrid. Left: Over-current protection activated simultaneously after turning off the clock in the 10 chip hybrid, Right: ABCN-25 shunting extreme current without damage

Figure 11 (left) shows an infrared image of a Liverpool hybrid fitted with 10 ABCN-25 chips with over-current protection activated on all the chips after the clock was turned off. For hybrids fitted with 20 chips, the over-current protection does not work as well as expected. The shunt regulators in this scheme were expected to shunt rather small currents so increasing the supply current in order to cover the ABCN-25's current bumps is not really possible. During the tests with the DAQ system of the hybrid supplied with increased current, the sharp, time dependent peaks of ABCN25's current requirements sometimes make one or a few chips shunt much more current than the other chips. Such a hot-spot is shown in Figure 11 (right). This chip

was likely to be shunting a high current at the time ($\sim$1 A) but the situation has never damaged any chips. As stated earlier the bumps on power consumption cannot be observed at high trigger rates as may be expected during SLHC running, but what about the calibration? The DAQ software can be modified to accommodate the bump by separating the L1 accept triggers in time. In such case the performance of this scheme in terms of ENC plots is as good as with the 'M' scheme and no hot-spot appearance can be observed.

C. 'SPi'

Serial powering interface chip is a versatile chip designed by Marcel Trimpl (FNAL), M.Newcomer and N.Dressnandt (Penn). The idea of the chip is to provide an universal solution for serial powering. It contains linear regulators, LVDS buffers, dual output current limited op-amp for the 'M' control scheme, its own shunt regulator with selectable output voltage etc. Its block diagram is shown in Figure 12.

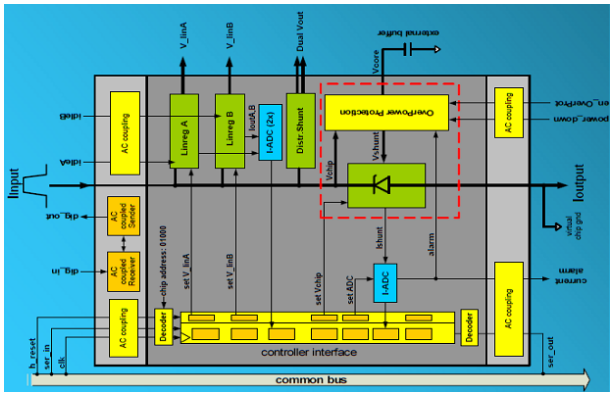

Figure 12: Serial Powering Interface chip block diagram. [3]

The whole talk dedicated to SPi in the Power Working Group session was given by Richard Holt (RAL). Let's just emphasize in this place that SPi has been tested thoroughly in a test stand and used with the hybrid as well.

IV. Current development

Both short and long staves will be constructed with the ABCN-25 chip in the year 2010. These staves will bring together all parts of a serially powered system (including protection etc.) for the first time. The individual shunt regulator architectures will use plug-in boards so that a variety of schemes may be studies. The next iteration of the Liverpool hybrid is specifically designed to accommodate serial powering. The next version of the ABCN chip, ABCN-13, will be built in 130nm technology, and the development of new powering blocks for this ASIC is in progress.

V. Summary

The development of serial powering and its shunt regulators and other powering blocks goes hand in hand with the development of the ABCx ASICs. Previously shunt regulator circuitry based on commercial electronics had been used to build several demonstrator staves based on the ABCD ASIC, and these were seen to perform well. Currently three main shunt regulator options have been implemented in full custom silicon and they are all functional. The characterization of these blocks has provided useful feedback to refine future designs. Several new serially powered stavelets and a full stave will soon be constructed with the ABCN-25 chip. Future ASICs, such as the ABCN-13 and MCC chips, will contain new powering blocks in 130 nm technologies.

References

[1] W.Dabrowski, Design of the distributed shunt regulator integrated in the ABC-N ASIC,
ABC-N Final Design Review (31.1.2008).

[2] M.Newcomer, Distributed Slave Shunt Approach: Description and Simulations (11.12.2007).

[3] The Figures were kindly provided by Richard Holt (RAL).