

Performance of new Run IIb EVB SCPU's

Markus Klute, Ron Rechenmacher, Bruce Knuteson, Steve Tether
MIT / FNAL

April 23, 2004

VMIC 7805 boards¹ will be used to read event fragments out of CDF DAQ and Silicon VRBs in the Run IIb Event Builder. This short note documents the tests performed in preparation for a production order of these boards.

Two VMIC 7805 prototype boards have been installed in a Run IIb Event Builder test stand located on the first floor of B0. The VMIC 7805 boards are a later model of the VMIC 7750 boards that have been used successfully in D0's Run II data acquisition system. Our VMIC 7805 boards run at 1.7 GHz, have 256 MB of RAM, and have a 100 Mb/s and a 1 Gb/s Ethernet port.²

The VMIC 7805 boards are 6U boards. CDF Adapter cards, with 5A fuses replaced by 15A fuses, are used to place the VMIC 7805 boards into the 9U VME crate. Four CDF Adapter cards are available in our test stand; the CDF Adapter cards holding the SCPU's in the current Run IIa EVB system will be used in the Run IIb system.

The VMIC 7805 boards require +5V (necessary). In order to power two micro-D serial ports on the card (optional), +12 V and -12 V are needed.

The VMIC 7805 boards have been used to read out CDF DAQ and CDF Silicon VRBs. The measured data rate when reading a single direct memory access (DMA) is 40 MB/s. A secret register bit on the Tundra Universe chip can be set to enable fast readout. With this bit set, a data rate over the VME backplane with a single DMA of 50 MB/s is obtained. These represent the maximum possible rates achievable, without processing or interrupts.

Code on the VMIC 7805 boards has been written with one thread performing a VME read of a VRB (with the VRB in emulation mode), and copying the data to shared memory; and a second thread sending the data from shared memory over GbE through our Cisco 6509 switch³ to a PC in our test stand. The VRB read is realistic, consisting of a byte count read followed by a DMA chain of 32 kB spread over 7 VRB reads.

We have derived the following formula for the speed of a VRB read into the memory of the VMIC 7805 board under realistic conditions:

$$\text{speed} = \text{totalEventSizePerCrate} * N_{\text{crates}} / \text{VmeReadTimePerEvent}$$

¹ Data sheet available at <http://mit.fnal.gov/~knuteson/cdfevb2/doc/VMIVME-7805%20Spec.pdf>.

² The VMIC 7805 board can also be obtained with a 2.2 GHz clock rate, and/or with 512 MB or 1 GB of memory.

³ Data sheet available at <http://www.cisco.com/en/US/products/hw/switches/ps708/ps711/index.html>.

$$\text{VmeReadTimePerEvent} = a + N_{\text{VRB}} * (b_1 + b_2) + \text{totalEventSizePerCrate} / (50 \text{ MB/s}),$$

where

$a = 5 \mu\text{s}$ = time between the interrupt that signals the end of the DMA and the resuming of the thread execution,

$b_1 = 5 \mu\text{s}$ = time in between DMAs within a chain,

$b_2 = 5 \mu\text{s}$ = time for reading byte count (same origin as b_1).

For example, with $N_{\text{VRB}} = 7$ and $\text{totalEventSizePerCrate} = 500 \text{ kB}/16 \text{ crates} = 32 \text{ kB}$, we obtain $\text{VmeReadTimePerEvent} = 5 + 7*10 + 640 = 715 \text{ us}$, and $\text{speed} = 500 \text{ kB}/715 \text{ us} = 700 \text{ MB/s}$. Our actual measurement under these conditions is $\text{speed} = 44.1 \text{ MB/s} * 16 = 705.6 \text{ MB/s}$, in agreement with prediction. This rate is for VME VRB read only, not sending to Level 3.

The total data transfer rate from the VRB through the GbE switch into the PC representing Level 3 is somewhat less than this: we observe 36 MB/s. The cause of the drop in rate from over 44 MB/s to 36 MB/s is believed to be due to PCI bus conflict between the thread sending the event data out over GbE and the thread reading data in from the VRB. Details are under investigation.

The Run IIb Event Builder specifications require a total throughput of 500 MB/s at 1kHz rate. At a rate per crate of 36 MB/s, with 16 crates, 500 MB/s is achieved.

In the new system, we expect most of the increased event size (from 250 kB in Run IIa to 500 kB in Run IIb) to come from increased data read out by the COT. In the Run IIb system, the COT data path will go through PCs, which will house software emulating VRBs and the same SCPU code as in the 7805 boards. The rate specification for the VME SCPUs is thus really $250 \text{ kB/s} * 1 \text{ kHz} = 250 \text{ MB/s}$. This spec has been cleared by a factor of two.

The current Run IIa Event Builder system reads VRB data out of 15 VME crates. Allowing for possible increase in the number of VME crates used, use of two VMIC boards in a continuing test stand for the prototyping of software changes after the final system is in place, and the necessity of having spares on hand, the number of VMIC 7805 boards needed is 20. With two in hand, an order will be placed in the next few weeks for 18 VMIC 7805 boards.