

VME Breaks Performance Barrier Again! 1000 Mbytes/s Range Possible

Recent developments within the VMEbus community have seen the bus speed doubled three times in a very short period of time. Due to the asynchronous protocol, this increase in performance has been achieved with full compatibility to all earlier boards.

The original theoretical data rate introduced 16 years ago was 40 Mbytes/s with the Block Transfer (BLT) Protocol. For a long time this was sufficient. But with much faster processors, the original data rate becomes more and more a bottleneck. By multiplexing data and address lines and including LWORD, the bus speed (or data rate) was doubled for the first time to 80 Mbytes/s theoretical speed. This was proposed in 1989 and approved in 1994 with ANSI/VITA-1 [1] as the Multiplex Block Transfer Protocol (MBLT).

Currently, the VME64 Extensions Standard [2] is under balloting by ANSI. Reducing the two-fold handshake to a single and using both the rising and falling edges of the handshake signals (strobe and acknowledge). This approach is known as the 2e(dge)VME protocol. Also, by using the new ETL [3] incident wave-switching drivers, the time window has been halved to 50 ns and the speed doubled for a second time to 160 Mbytes/s. The third doubling to 320 Mbytes/s was introduced by Drew Berding (Arizona Digital) and Bustronic, using a patented "star wiring" backplane technology to achieve incident wave switching by an alternative method. Related to a transfer rate of 320 Mbytes/s is a very short time window of 25 ns for a data transaction.

This can only be realised with a source synchronous protocol. This new performance using the 2eSST [4] was presented during last year. It is achieved by omitting the remaining acknowledge signal after each data beat and placing them at the end of a block transfer. This saves all the waiting times needed for the return of the handshake signal. The settling time for a stable digital signal is also omitted by using incident wave switching. In the case of the "VME320 Technology" as a synonym for a star wired backplane, incident wave switching is realised by using very slow edge rates where the reflections merge instead of using stronger drivers. By removing the need for reflections all signals, data and strobe, run in one direction over the backplane, "synchronous to source".

LIMITATION OF DATA THROUGHPUT

Using older protocols, the data speed is limited by the

propagation time through the backplane. The next data beat may start only, when the receiver acknowledges the successful reading of the preceding data. This requires a signal propagating opposite with the direction of the data stream.

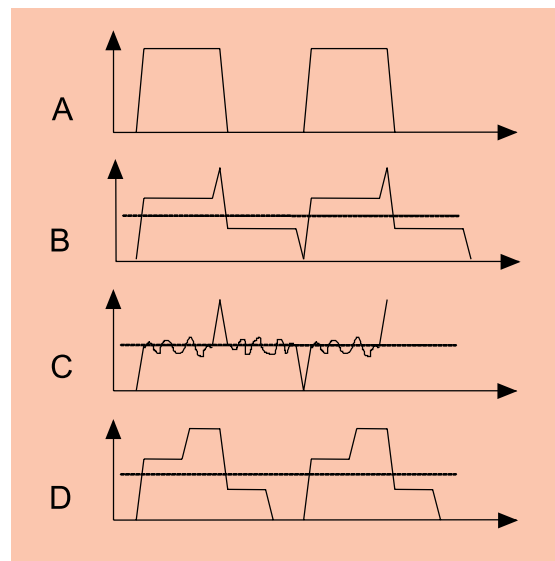


Figure 1. Bus speed is limited by "run time limitation" with increase of speed of the source signal (A), on high impedance traces due to high propagation delay, the reflected wave superimposes with the following edge (B) of the source signal and results in noise (C). Low impedance traces will result in shorter propagation time and can improve performance (D)

A further limitation of speed is originated by the reflection, independent of whether incident wave switching is achieved or not. With increasing data speed or physical frequency, the next edge of the signal arrives at the same time as the reflection and both superimpose on each other. But both have different slope directions: if the reflection comes from a rising edge, the reflection is also rising, but the following signal edge is a falling one. Such superimposed signal is no longer readable as a clear digital signal, it appears like noise voltages centred near the threshold region with small spikes which occur at the time when the edges merge. This

kind of limitation may be called run time limitation. See Fig. 1.

So we have two kinds of limitation of speed:

- a backward running handshake signal and
- backward running reflections.

SOLUTION OF THE PROBLEM

A further increase of bus speed can be achieved by using lower impedance traces to achieve shorter propagation delays [5]. But additionally we need a new quality for getting a much higher increase of speed. This new quality is the source synchronous protocol. Now the limitation of speed will only be caused by the skew of all signals because there is no longer the need for signals to run backward towards the source. The reflections must arrive after a very short time, much earlier than the next signal edge occurs, as indicated or must be removed. The combination of a low impedance backplane with much faster running signals and a source synchronous protocol will open the way into the future.

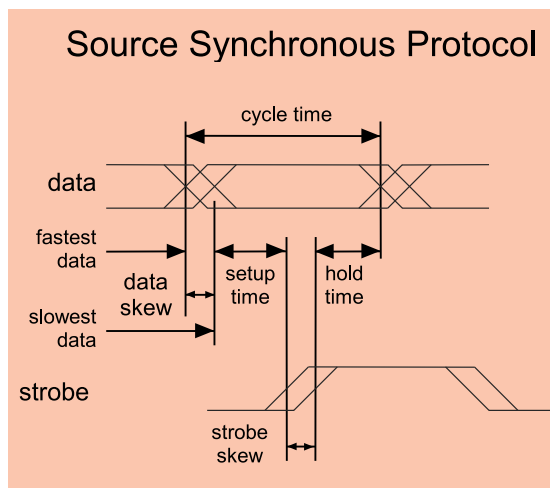


Figure 2. 2 edge Source Synchronous Protocol with timing parameters; depending of write or read cycle, strobe may be DSn^* or $DTACK^*$

The data may be read at the destination, when the latest data signal has reached one set-up time before the fastest strobe has arrived and, vice versa, the fastest data must be still valid one hold time longer after reaching the slowest strobe [4,6], see Fig. 2. The absolute running time, without skew, is identical for all signals and may be subtracted because all signals run in the same direction. So the difference between the fastest and slowest signal, the skew, becomes the bottleneck of performance. It is therefore recommended to use a low skew backplane and transceiver technology. As shown in table 2, the backplane and transition skew of the star wiring backplane contribute greatly to the overall skew. This is due to the very slow edges of this technique. Using low impedance backplanes and the stronger ETL incident wave-switching drivers will reduce the skew considerably as shown in the following paragraph.

CALCULATION OF SKEW

The total skew of the transmission path is determined by the skew of the source, the skew of the transmission line in the backplane and the skew of the destination (Fig. 2) [4].

Driver Skew

The source skew of the ETL [3] driver is taken from the ABTE TI-Data Book [7] as the difference between fastest and slowest gate delay from port "B" to port "A": $t_{pd\ max} - t_{pd\ min} = 5,2\ ns - 1,5\ ns = 3,7\ ns$.

This is the worst case skew over the whole voltage and temperature range. The maximum skew between any signals on a single board will be greatly reduced due to the reduction in temperature and supply voltage variation.

Backplane Skew

An alternative to a star wired backplane, acting as a lumped element of inductance and capacitance, making very smooth signal edges, a backplane using transmission line signals is proposed with this paper. Due to the slow signal edges, generated by the "centred capacitance" of the star wired backplane, higher frequencies over this backplane seems to be limited [8]. This is the reason for using low impedance backplanes traces [5] or, the reflections must be removed. The influence of these signal reflections will be investigated in a future paper.

The backplane skew is calculated by the difference in propagation delay, caused by load capacitance tolerances:

LOADS	CAP. VALUE	TOL
BP-Via	0,8 pF	5%
Connector	0,8 pF	5%
Board-Via	0,6 pF	5%
Trace	4,5 pF	15%
Package	1,0 pF	10%
Transceiver	8,0 pF	25%
Capacitance per slot	15,7 pF	
Weighted Tolerance		18,3 %

Table 1. Backplane Loads and Tolerance of Loads

The trace capacitance is calculated by assuming a 60 W trace with a length of 1,5 inches.

For Calculating the difference of propagation delay, the following basic formula is applied:

$$t_{pd} = t_{pdo} \sqrt{1 + \frac{C_{Load}}{C_o}} \quad (1)$$

Where t_{pd} is the propagation delay of the loaded trace, t_{pdo} is the propagation delay of an infinite trace in a

homogenous dielectric (180 ps/in.) and C_0 is the intrinsic capacitance of a trace, assumed to 3,9 pF/in. for a trace of 46 W. The backplane skew is calculated with 0,4 ns and will be reduced with lower impedances. The backplane skew of a transmission line backplane is slightly lower than of a star wiring backplane. This calculation assumes a fully loaded backplane. With empty slots or boards which do not support A32/D32 or A64/D64 transfers, the backplane skew compared to fully loaded backplanes increases and lower data rates are achieved. For obtaining the highest data rate all unused slots shall be populated, for example with dummy loads of low tolerances of the capacitive load. Boards which do not support the o/m transfers should not be used for highest data rates, or, these boards should also use dummy loads at the non used address- and data lines

Destination Skew

The destination skew is caused by variations in the receiver threshold and may be estimated by the time which is needed for the signal to cross the region between the "LOW" and "HIGH" level. Using the fast and strong ETL drivers will increase the performance compared with the star wiring backplane. Instead of 6 ns skew [6] we will get a value of 0,4 ns: The slope is assumed to approx. 0,5 V/ns and the levels are well known with 1,4V for "LOW" and 1,6V for "HIGH".

CALCULATION OF THE CYCLE TIME

The overall cycle time may be calculated by adding all skews, the setup and hold times [4] (Fig. 2):

Using the same calculation scheme and the same assumptions as it was done for the star wiring backplane [6], we may calculate as follows:

SKREW DESCRIPTION	VME533	TRENEW
Tdata-skew	5,0 ns	3,7 ns
+ Tbackplane-skew	1,5 ns	0,4 ns
+ Ttransition-skew	6,0 ns	0,4 ns
+ Tdata-setup	1,5 ns	1,5 ns
+ Tstrobe-skew	1,0 ns	1,0 ns
+ Tdata-hold	1,0 ns	1,0 ns
Tcycle	16,0 ns	8,0 ns
data rate	500 MB/s	1000 MB/s
data frequency	31,3 MHz	62,5 MHz

Table 2. Skew and data rate calculation; VME533 is the calculation for the star wiring technique of Drew Berding [6].

Theoretically the strobe skew should be the same as the data skew, but in practise, a reduction may be possible. The strobe is only one signal instead of 64 signals and uses the same driver, backplane trace and receiver, all at nearly the same temperature and supply

voltage. That's why the strobe skew is assumed to 1 ns [6].

The data rate is defined as the number of data bytes driven per second along the data transmission path:

Data rate = 1(second) * 8 Bytes / cycle time. The physical data frequency can be calculated by dividing the data rate by 2*8 Bytes. The factor "2" arises from the fact that at the Period T of a square wave signal two data beats will occur.

SIMULATION RESULTS

This timing appraisal is to illustrate how far we are from setting a limit on further increases in the data transfer rate of the "old" VME bus. Computer-assisted simulations based on SPICE models could, in principle, provide clarity concerning the accuracy of the idea discussed in this paper, and determine the parameters of a system's board-backplane before expensive test-boards are manufactured. All simulations and measurements are based on a backplane equipped at all 21 slots. In each case, the "feeder and measurement points" for the simulated signals are the bond pads on the silicon chip inside its housing.

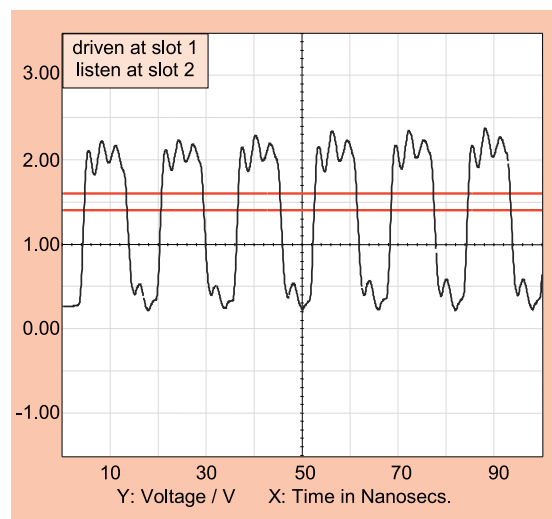


Figure 3. Simulated signal, fed into slot 1 and "measured" at slot 2; A physical data frequency of 62.5 MHz, corresponding to a data rate of 1,000 megabytes/s.

Figure 3 shows the simulated progress of a signal at a data rate of 1,000 M/s on a fully equipped 21 slot VME bus backplane. The physical signal frequency is 62.5 MHz, i.e. a new bit is sent through all 64 data transmission paths every 8 nanoseconds. This then makes 1,000 megabytes per second. The data source is found on the first slot and the recording of the signal follows at slot 2. As a rule, this is one of the two worst case instances. In conventional (VME) bus systems, it is precisely with this configuration that effects on signal transmission time through the backplane, reflections, and also non-homogeneous impedance along the signal path become most noticeable and can cause serious problems with data transmission. But as can be clearly seen in Figure 3, signal quality is good and clearance to the ETL thresholds are adequate. Figures 4 and 5 then show the same signal as

"measured" in the middle of the backplane at slot 11 and at the end of it at slot 21. Signal quality likewise is good, and yet the clearance to the thresholds could be greater at the last slot.




Figure 4. As in Fig. 3, but with the signal "measured" in the middle of the backplane (slot 11)

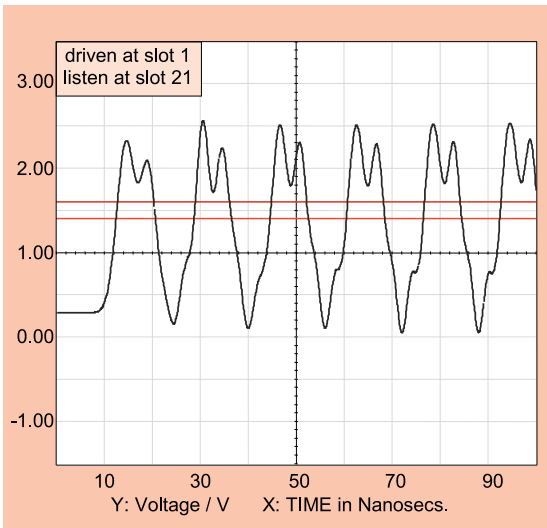


Figure 5. As in Fig. 3, but with the signal "measured" at the end of the backplane (slot 21).

The low pass characteristic of the bussed backplane transmission lines is likewise detectable as the signal edges became flatter. This is caused by the additional capacities at each slot (see Table 1). With an exact routed point-to-point connection, this effect is significantly weaker and the limiting frequency of a backplane transmission line is within the GHz range.

If the shape of the graphs in the figures mentioned is more closely observed, the differences between the first and the following bits become obvious. The first impulse was incited on a quiet path, whereas the pulses that follow have the oscillations generated with the edges of the preceding pulse(s) superimposed on them.

In contrast to the high data rates common in state-of-the-art technology, disturbances incited at the flank because of the substantially shorter cycle time (8 nanoseconds to 100 nanoseconds with a conventional MBLT protocol at VME64) are not yet reduced and must be taken into consideration when simulations and measurements are made. Thus, there are always "several" impulses to be recorded.

The second worst case instance occurs when the operating board is located approximately in the middle of the board. Because of the parallel wiring of both path branches, the wave then "sees" only half the impedance at the "T-junction" while entering the backplane. Figure 6 shows, however, that this is not an issue for Trenew's VME1000 concept.




Figure 6 Simulated signal fed into slot 10 and "measured" at slot 11, the second worst case.

The reader is referred to Figures 7 and 8 for an idea of the many small steps remaining before a solution is found. Whereas the signal at the middle of the backplane certainly has a curve shape that is very unsuit-




Figure 7. Example of where the target of transmitting 1,000 megabytes per second is not entirely reached.

REAL TIME MAGAZINE

SUBSCRIPTION SERVICE

REAL-TIME CONSULT,
Rue de la Justice 23,
1070 Brussels,
Phone. 32-2-520-55-77, Fax. 32-2-520-83-09
E-mail: info@realtime-info.be

BOOKSTORES

EBSCO Subscription Services

P.O. Box 1943, Birmingham,
Alabama 35201-1943,
USA
Phone. 1-205-991-6600
Fax. 1-205-991-1479

Standaard Boekhandel

Industriepark Noord 28 A
9100 St.Niklaas,
Belgium
Phone. 32-3-760.32.11
Fax. 32-2-777.92.63

SWETS Subscription Service

P.B. Box 830,
2160 SZ Lisse,
Netherlands
Phone. 31-25.213-51-11
Fax. 31-25-211-58-88

Dawson France

B.P. 57,91871
Palaiseau Cedex,
France
Phone. 33-1-69-10-47-00
Fax. 33-1-64-54-83-26

Lavoisier Abonnements

14 rue de Provigny
94236 Cachan Cedex,
France
Phone. 33-1-47-40-67-00
Fax. 33-1-47-40-67-02

READMORE

22 Cortlandt Street
New-York 10007-3194
USA
Phone. 1-212-349-5540
Fax. 1-212-233-0746

BLACKWELL'S Periodicals

PO BOX 40
Hythe Bridge Street, OX1 2EU Oxford
Great-Brittain
Phone. 44-1865-79-27-92
Fax. 44-1865-79-14-38

FRANCE Publications

112 Rue Reamur
75002 Paris
France
Phone. 33-1-45-08-55-68
Fax. 33-1-45-08-50-19

Schweitzer Sortiment

Lenbachplatz 1
D 80333 Munchen
Germany
Phone. 49-89-551-340
Fax. 49-89-551-34100

Garmar Libreria Internacional

Fundadores, 5-1. 5
28080 Madrid
Spain
Phone. 34-361-1388
Fax. 34-361-07-50

BTJ AB

Box 1302
SE 11183 Stockholm
Sweden
Phone. 46-8-723-6900
Fax. 46-8-723-0038

Hoepli Ulrico Casa Editrice Libreria SPA

Via Hoepli 5
20121 Milano
Italy
Phone. 39-2-865-446
Fax. 39-2-805-2886

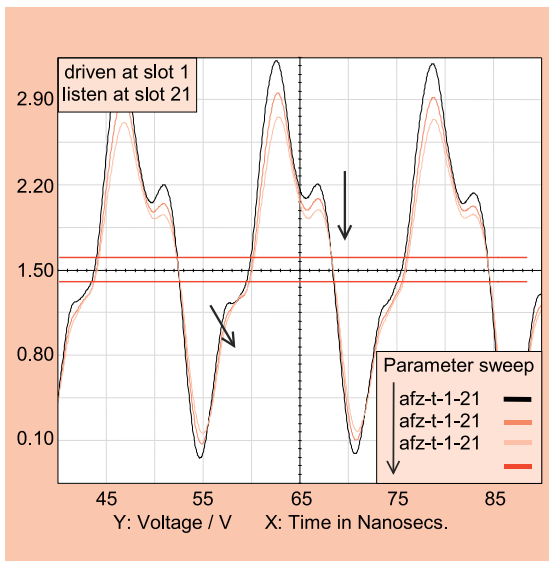


Figure 8. As in Fig 7, but "measured" at the last slot. Here, the signal would already be usable throughout.

able for digital signal transmission, it would already be acceptable throughout as a digital signal were it at slot 21.

In addition, the signal's reaction to a change in a parameter of the system's backplane-daughter card is also represented. At the last slot the effects of a variation in the signal are negligible. At the middle of the backplane, signal distortions are already considerable, causing additional crossing thresholds.

Certainly this very clearly demonstrates the complexity of the assignment: it proves to be substantially simpler to adjust the system parameters so that the signal can, within specified tolerance restrictions, be easily read at only one slot.

But this also very clearly demonstrates the efficiency of a simulation: the latter can save a lot of time, and can avoid the use many expensive prototypes. Using models that are adequately good representations of the

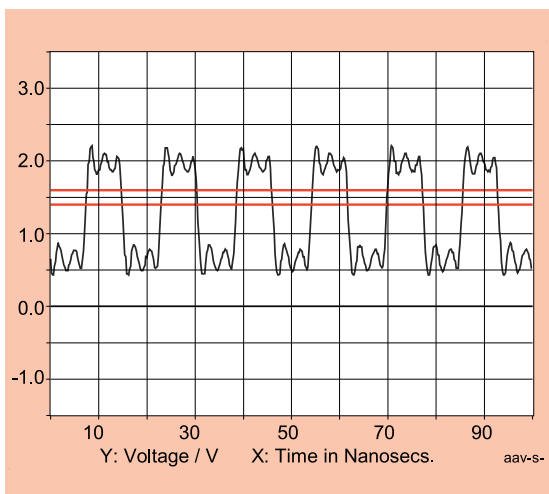


Figure 9. Measured signal of a generator, set at 25 Ohm, operating at slot 1 and measured with an oscilloscope at slot 2; nevertheless fed onto the daughter card at the pin of the transceiver.

real world is the only need. However, as is well known, it is in relation to this that the greatest problems are found in practice.

For this reason, measurements were ultimately taken using a laboratory prototype, built according to the parameters established by the simulations. The graph in Figure 9 accords very well with the corresponding graph in Figure 3.

FINAL CONCLUSIONS

This assignment aims to provide an idea of the order of magnitude of data rates possible with stronger ETL-drivers and a backplane based on low-impedance or reflection-free transmission paths instead of central capacity. The work is being continued with improved models in order to be able to more precisely define the tolerances of the parameters of the system, and also to include other effects such as cross talk.

The open structure of the VME bus, the asynchronous protocol, and the separation from the logical and power interface make continuing adaptation to further developing technologies possible without any need to break the compatibility with legacy boards. It is the opinion of the author, and of others, that this is where the great strength of the VME bus lies. ■

REFERENCES

1. ANSI/VITA 1-1994, VME64 Standard
VMEbus International Trade Association, 7825 East Gelding Drive, Suite 104, Scottsdale, Arizona 85260
Internet: <http://www.vita.com>
2. VITA 1.1-199x, VME64 Extensions Draft Standard
3. VITA 2 - 199x, Draft 0.5; ETL
Enhanced Transceiver Logic Specification
4. VITA1.5 - 199x, Draft 1.3; 2eSST
Source Synchronous Transfer Draft Standard
5. Andreas Lenkisch, "Signal Integrity on VME; Data Rates of 160 Mbytes/s and Higher", SysComms 98, Geneva, CERN, March 1998
6. Drew Berding, "VME Breaks Performance Barrier, Approaches 533 Mbytes/s Range"; RTC Europe, Vol. 2, No. 1, Jan 98, page 33 - 34
7. Advanced Backplane Circuits, Texas Instruments 1995
8. Ulrich Wallenhorst, "Harting-Signal Integrity-VME320 and VME64", VSO Meeting Phoenix May 1997, Meeting Minutes, www.vita.com/vso/minarch

Andreas Lenkisch studied Electronics and Physics at Technical University Ilmenau from 1974 - 1978. He joined Trenrew in 1990 to develop standard and custom backplane solutions for high speed computer applications. Introduced "Classic" VMEbus product range with high power distribution capability and enhanced high frequency behaviour. In 1995, he headed Research department to extend the companies design capabilities and investigate emerging bus standards and related technologies.