WHITE PAPER

Hardware Design Considerations for Teledyne e2v's Space-Grade DDR4

June 2021

By Dr. Rajan Bedi

TELEDYNE e2v | Semiconductors
Everywhereyoulook™

Dr. Rajan Bedi is the CEO and founder of Spacechips, which designs and builds a range of advanced, L to K-band, ultra high-throughput on-board processors, transponders and OBCs for telecommunication, Earth-Observation, navigation, internet and M2M/IoT satellites. The company also offers Space-Electronics Design-Consultancy, Avionics Testing, Technical-Marketing, Business-Intelligence and Training Services. (www.spacechips.co.uk). Rajan can also be contacted on Twitter to discuss your space-electronics' needs: https://twitter.com/DrRajanBedi

Spacechips' Design-Consultancy Services develop bespoke satellite and spacecraft sub-systems, as well as advising customers how to use and select the right components, how to design, test, assemble and manufacture space electronics. We teach semiconductor memories on our FPGAs for Space Applications Training Course.

ABSTRACT

Fast, highly reliable, and radiation tolerant memories are a Must Have feature for complex Space Edge Computing systems : DDR4 will allow the Space industry to offer **higher-throughput on-board processing** and **increased acquisition times enabling new Earth-Observation**, **space-science** and **telecommunication** applications, e.g. ultra high-resolution imagery, live streaming video and **on-board AI**.

I previously introduced **Teledyne e2v Space Radiation Tolerant DDR4 for space applications (DDR4T04G72)**, offering 4 GB of volatile storage at a clock frequency up to 1.2 GHz and a data rate of 2.4 GT/s (bandwidth of 172.8 Gb/s), and **this white paper will explore deeper areas of such a Space DDR4 memory**.

First part of this paper will go into **technical considerations on SDRAM**, before extending to the **DDR4 architecture itself and its organization**, since DDR4 contains new architectural and hardware features which need to be considered to ensure design is right-first-time.

As examples, **Point-to-point connection between a PolarFire FPGA** and Teledyne e2v DDR4T04G72 will then be highlighted, as well as **connection of multiple DDR4 devices to a Xilinx KU060 FPGA**.
Finally, I will conclude with the **benefits of DDR4 over DDR3, and single & multi-mode DDR4 configuration storage and capacity will be given**.

Previously I introduced DDR4 for space applications offering 4 GB of volatile storage at a clock frequency up to 1.2 GHz and a data rate of 2.4 GT/s (bandwidth of 172.8 Gb/s). Compared to previous generations of SDRAM, DDR4 contains new architectural and hardware features which improve capacity, performance, scalability, system-level reliability and power efficiency. In this post, I introduce these, discuss timing and signal-integrity considerations, and the connectivity of this memory with FPGAs to ensure your avionics design is right-first-time.

WHITE PAPER

Hardware Design Considerations for Teledyne e2v's Space-Grade DDR4

June 2021

By Dr. Rajan Bedi

TELEDYNE e2v | Semiconductors
Everywhereyoulook™

An SDRAM architecture comprises memory cells organised in a two-dimensional array of rows and columns as illustrated in Figure 1. To select a particular bit, it is first necessary to address the required row and then the specific column. Once the desired row is open, it is possible to access multiple columns, hence improve speed and reduce latency through successive read/write bursts.

To increase word size, the memory has multiple arrays which means when a read/write access is requested, the memory only requires one address to access 1 bit from each array.
To increase overall memory capacity, banks are added to the internal structure of SDRAM as shown below. Bank interleaving further increases performance and each can be addressed individually.

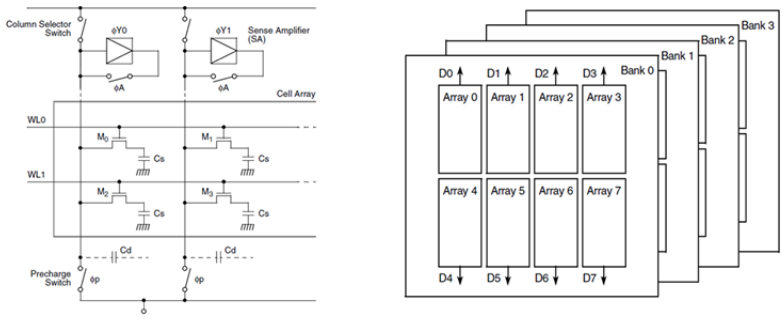

Figure 1 : SDRAM bit cells and the organisation of a DDR chip.

The core speed of SDRAM is slower than its I/O rate and multiple words of data are accessed during every column command which are then serialised to/from the interface. DDR4 is based on an 8n-prefetch architecture which transfers two n-bit wide data words per clock cycle at the I/O. A read or write operation comprises a single 8n-bit wide, four cycle burst transfer at the internal DRAM core and eight corresponding n-bit, one-half clock cycle transfers at the I/O pins.

DDR4 extends the above SDRAM architecture by introducing bank groups allowing a prefetch of eight in one group, and a second to be executed independently in another. Effectively, DDR4 time-division multiplexes its internal bank groups to hide the fact that the internal core requires more time than that required by a burst of eight words at the I/O interface. Compared to DDR3, DDR4 improves performance by offering more banks with significantly smaller row sizes meaning devices can cycle through different banks at a higher rate. The organisation of DDR4 memory is illustrated below: to support higher storage capacities without adding extra address pins, DDR4 uses a newly-defined ***ACT_n*** input to multiplex addresses on the command pins, ***RAS***, ***CAS*** and ***WE***. If ***ACT_n*** is low, these inputs are used as address ***A16***, ***A15*** and ***A14*** pins respectively. When ***ACT_n*** is high, they resume their normal functions as specified in the SDRAM command truth table.

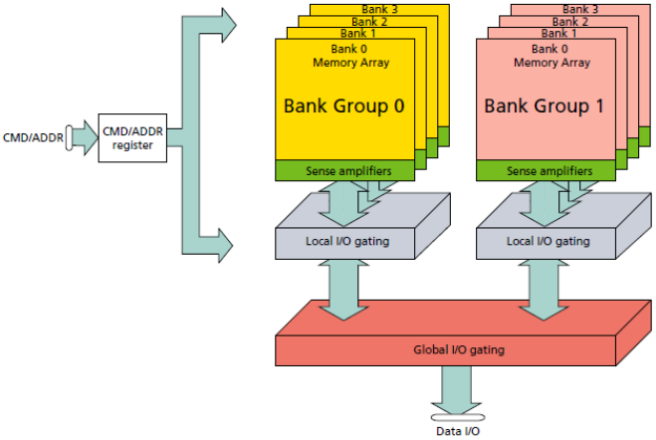

Figure 2 : DDR4 Bank Groups.

Hardware Design Considerations for Teledyne e2v's Space-Grade DDR4

June 2021

By Dr. Rajan Bedi

TELEDYNE e2v | Semiconductors
Everywhereyoulook™

Teledyne e2v's, 4 GB, rad-tolerant DDR4T04G72 is an MCP containing five die, four of which offer 1GB (8 Gb) of storage each, 512 Mb x 16 bits, organised in two groups with four banks in each as shown above. To bolster reliability, a 72-bit data bus is created comprising 64 data and 8 bits for error detection and correction. This ECC function is realised within the fifth die. The device uses an internal 8n-prefetch buffer to maximise high-speed operation and offers programmable read, write and additive latencies.

DDR4 has introduced a number of hardware features to reduce power consumption: firstly, the I/O supply (VDDQ) has been reduced to 1.2 V from the 1.35 V rail used by DDR3. A separate 2V5 voltage, Vpp, has been added to activate the internal word line lowering power dissipation by 10%. The I/O electrical interface for the data bus has changed from push-pull, series-stub terminated logic (SSTL) to pseudo open drain (POD) signalling as illustrated below. By terminating to VDDQ instead of 1/2 of VDDQ, the amplitude and centre of the signal swing can be tailored to each design's need. POD I/O reduces switching current when driving data since only 0's consume power. DDR4 also offers Data Bus Inversion which assigns fewer bits low dissipating less power. Reduced switching results in less noise and a cleaner data eye.

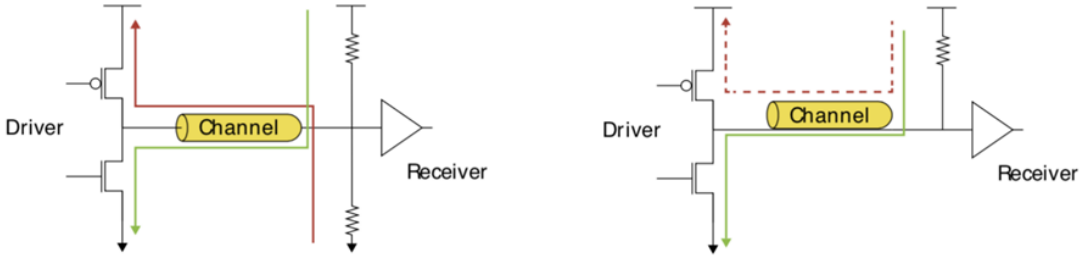

Figure 3 : DDR3 push-pull I/O signalling (left) vs. DDR4 POD (right).

Collectively, the reduced VDDQ voltage, the use of an external Vpp supply to boost the word line, the change to POD signalling and VDDQ termination, as well as the previously discussed smaller row size with lower activation currents, have reduced overall power consumption compared to DDR3 SDRAM. At similar data rates, a DDR4 device has a 30% advantage in power efficiency. This improvement can be used to operate the SDRAM device at higher speeds or lower dissipation for the same performance. A power-prediction spreadsheet and ICEPAK/ECXML thermal models are available for the DDR4T04G72.

At a system level, DDR4 offers improved reliability, availability and serviceability (RAS). Real-time, CRC error detection of the data bus during write operations as well as parity checking of the command and address busses are available as shown below. Unlike DDR3, DDR4 can be configured to block commands upon detection of a parity error.

Hardware Design Considerations for Teledyne e2v's Space-Grade DDR4

June 2021

By Dr. Rajan Bedi

TELEDYNE e2v | Semiconductors
Everywhereyoulook™

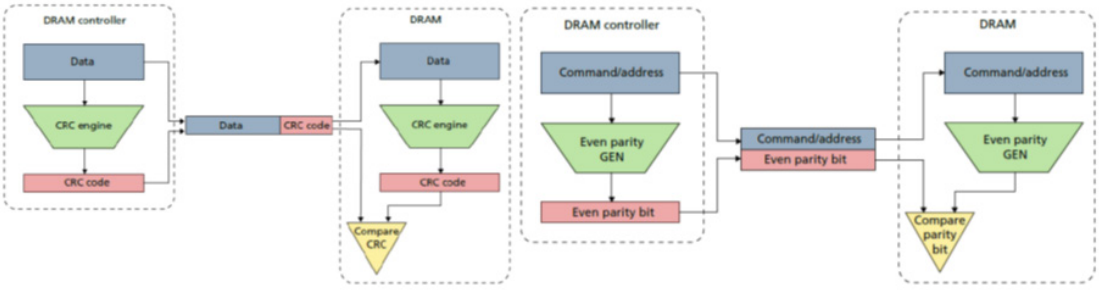

Figure 4 : System-level error detection offered by DDR4.

For those soft errors that cannot be fixed using ECC during the lifetime of the memory, DDR4 offers a post-package repair function to correct rows which have become faulty. Not only does this increase the reliability and longevity of systems, but it also provides a further mechanism to protect against single-events upsets.

DDR4 also offers a Connectivity Test Mode (CT) to check the continuity of the PCB traces between the memory and the controller without invoking the SDRAM's initialisation sequence. Unlike conventional boundary-scan testing where test patterns are shifted in and out of devices serially during each clock, CT mode uses a faster, parallel interface.

The DDR4 I/O interface is a true source-synchronous design, where the data is captured twice per clock cycle using a bidirectional data strobe, ***DQS***. During a READ operation, ***DQS*** is output by the memory, co-incident with the data, and for WRITEs, the strobe is provided by the controller centred with respect to the data, providing a synchronous reference. To improve signal integrity as data transfer rates increase and amplitudes decrease, the clock and strobe signals are differential to cancel out common-mode noise. At a PCB level, ***DQS*** has identical loading to the data bus and should be routed similarly. The other address, command, control and data signals still operate in single-ended mode, which makes them more susceptible to noise, crosstalk and interference.

Prior to PCB layout, it is important to decide how much of the available timing budget to allocate to routing mismatch. This can be determined by thinking in terms of time or as a percentage of the overall period, e.g. with a clock frequency of 1.2 GHz, the period is 833 ps. Typical flight time for FR4 is 6.6 ps/mm, so length matching traces to 1 mm consumes around 1.6% of the total period for track tuning. If your design does not push the performance limits, you can allocate a larger percentage of the overall timing budget to length mismatch to provide more routing flexibility and ease the layout effort.

When calculating PCB propagation delays, note that these vary for inner (stripline) and outer (microstrip) layers because their effective dielectric constants are different. Vias represent additional length in the Z direction and the number of vias in matched lines should be the same with identical span to ignore their impact on the overall timing budget.

Before PCB fabrication, post-layout simulation is recommended to confirm timing margins and signal integrity. IBIS and Spice models are available for the DDR4T04G72 to allow you to confirm electrical and timing compliance early in the design cycle. I use Mentor Graphics' (now Siemens) Hyperlynx Linesim and Boardsim to verify pre and post-layout signal integrity respectively, to optimise termination and drive strength, and to validate timing margins to allow sign-off before manufacture. An EBD model is currently being developed.

Hardware Design Considerations for Teledyne e2v's Space-Grade DDR4

June 2021

By Dr. Rajan Bedi

TELEDYNE e2v | Semiconductors
Everywhereyoulook™

To verify the signal integrity between the controller and the memory using an internal routing layer, Figure 5 illustrates the eye diagram predicted by Linesim of a PolarFire rad-tolerant FPGA connected to a data line of a single DDR4T04G72. Multiple DDR4 devices can also be connected to a single FPGA each with its own IP controller: Figure 5.

To increase overall storage capacity, the same soft IP can also command multiple DDR4 devices placed in either fly-by or clamshell topologies, i.e. common clock, address, control and data signals, with each SDRAM having its own chip-select input as illustrated below. In this case, the transmission lines are longer and the capacitive load higher, so simulation is necessary to confirm the required driver's current strength. Each KU060 DDR4 controller has a maximum data bus width of 80 bits, can access up to five external memories and the FPGA can instantiate two of these IPs.

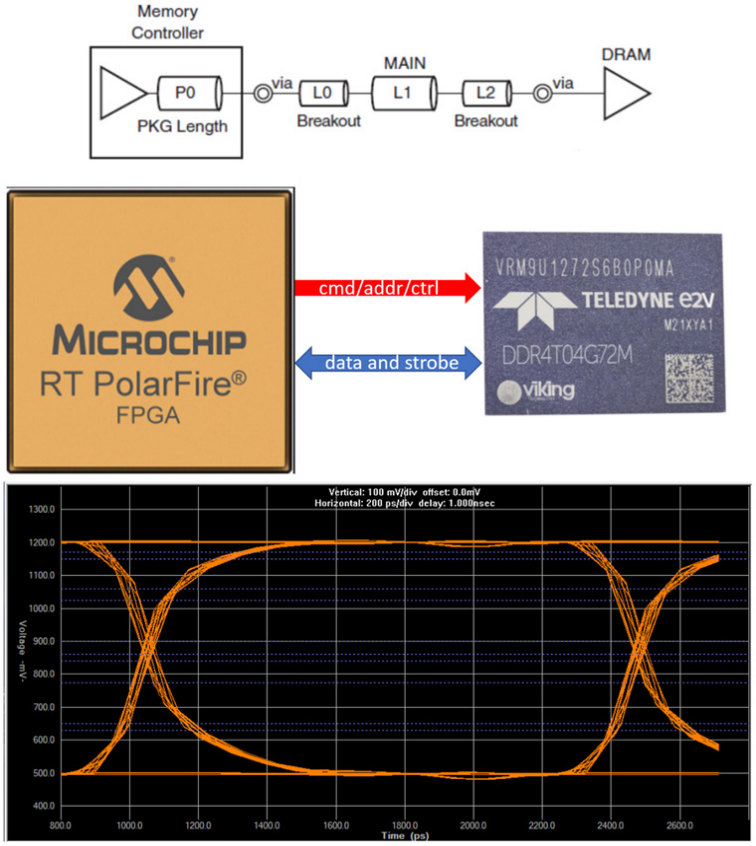

Figure 5 : Point-to-point connection between PolarFire and the DDR4T04G72.

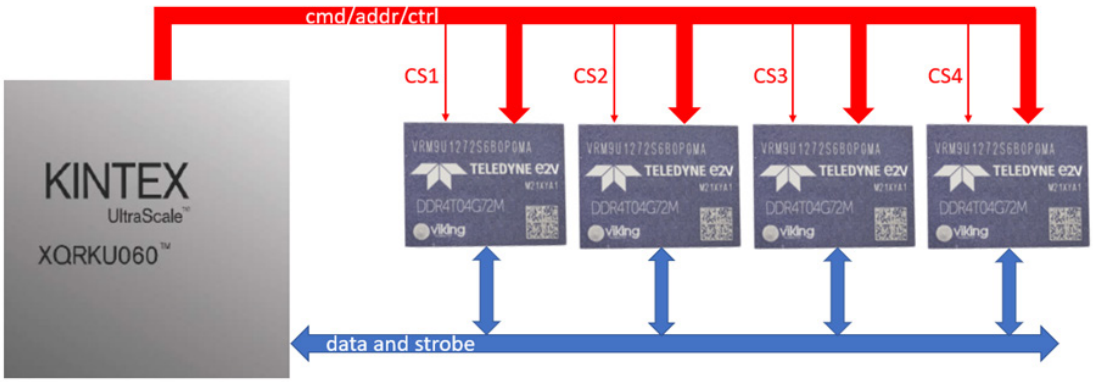

Figure 6 : Connection of multiple DDR4 devices to a Xilinx KU060 FPGA.

Xilinx offers a video which demonstrates how to instantiate a DDR controller as well as resources to calculate the maximum rate and the number of external SDRAM devices that can be connected to their FPGAs. PolarFire's DDR4 IP offers a data-bus width of 72 bits allowing the connection of four DDR4T04G72s as shown above.

When the DDR4T04G72 is connected to Xilinx's KU060 or Microchip's PolarFire rad-tolerant FPGAs, Table 1 summarises the resulting storage capacities and bandwidths assuming a data rates of 1.33 and 1.86 GT/s respectively. The total number of DDR4 IPs which can be instantiated within either FPGA depends on your specific I/O usage so please confirm your configuration using either the Vivado® Design Suite or Libero® SoC. NanoXplore's NG-Ultra will also support DDR4 SDRAM.

	KU060 Single	PolarFire Single	KU060 Multi-mode	PolarFire Multi-mode	KU060 Multi-IP	PolarFire Multi-IP
No. of DDR4 memories	1	1	5	4	10	24
No. of FPGA DDR4 IP controllers	1	1	1	1	2	6
Total Storage (GB) (Excluding ECC)	4	4	20	16	40	96
DDR4 bus width (bits) (Excluding ECC)	72	72	72	72	(2 x72)	(6 x 72)
Maximum Data Bandwidth (Gb/s)	134	96	134	96	268	575.8

Table 1 : System storage capacity and bandwidths.

DDR4's data signals, ***DQ***, ***DQS*** and ***DM_n***, have dynamic on-die termination (ODT) built into the FPGA controller and SDRAM, and ordinarily, external termination resistors would need to be placed at the far-end of the address, command, control and clock nets as shown in Figures 7 and 8. However, these are not required for the DDR4T04G72 which contains ODT for all its high-frequency interface signals.

Conventional tree-topology routing creates a stub whose length increases with the number of receivers reducing the bandwidth of this transmission line. This attenuates the high-frequency components that form the rising and falling edges of the signal, shrinking the eye opening at the SDRAM. Fly-by routing reduces the number of stubs and their lengths.

Hardware Design Considerations for Teledyne e2v's Space-Grade DDR4

June 2021

By Dr. Rajan Bedi

TELEDYNE e2v | Semiconductors
Everywhereyoulook™

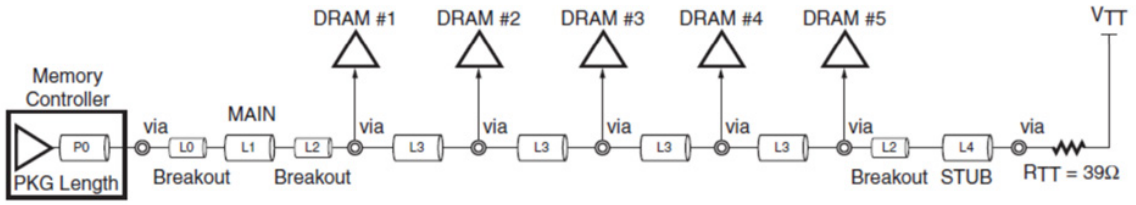

Figure 7 : Fly-by termination of DDR4 command, address and control signals.

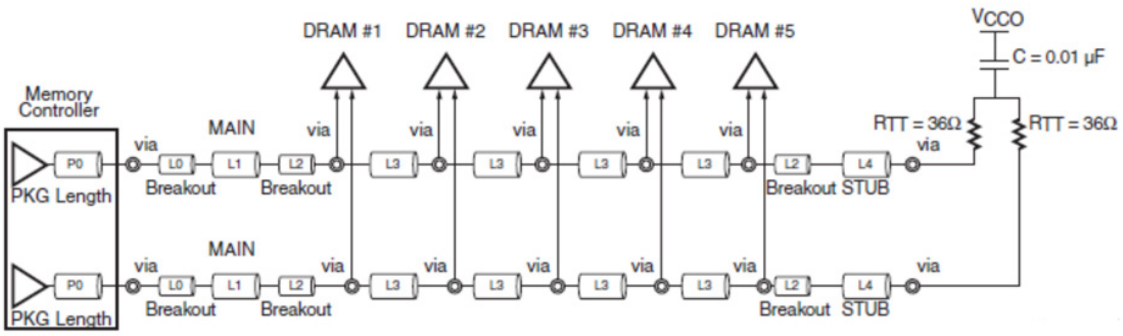

Figure 8 : Fly-by termination of DDR4 differential clock input.

DDR4 has on-die capacitance for the core as well as the I/O and, therefore, it is not necessary to allocate external capacitors for every power-pin pair. However, a minimum amount of PCB decoupling is specified for the DDR4T04G72 to prevent the supply from drooping when the SDRAM core requires current for refresh, read and write operations. Decoupling also provides current during reads for the output drivers. The core requirements are lower frequency requiring larger capacitance values, whereas, the drivers switch at higher rates necessitating low inductance and less capacitance.

You have completed your schematic design, layout, pre-fabrication timing and signal-integrity checks, sub-contracted the assembly of the PCB and verified that the new board powers up as expected. You are now ready to start using the memory, however, prior to operation, DDR4 has to be initialised so the SDRAM understands its operating frequency and delay parameters. DDR3 uses a voltage divider to create Vdd/2 as a reference to decide if the ***DQ*** signals are 0 or 1 as shown in Figure 3. DDR4 uses an internal voltage reference, ***VrefDQ***, whose value must be set by the memory controller during the initialisation phase. Furthermore, SDRAM requires periodic calibration of the output driver impedance and ODT values to minimize variations in voltage and temperature, a process known as ZQ calibration.
The final step before DDR4 can be used is known as Memory Training, which calculates the read/write delays between the SDRAM and its controller. As shown in Figure 6, for multiple DDR4 chips connected to an FPGA, each device may be physically located at a different distance from the controller resulting in individual, flight-time skews between the clock, strobe and the data. Write Levelling compensates for these differences! In a fly-back topology, each chip receives the command, address and control at a different time and Read/Write Centering ensures data can be reliably read from or written to the SDRAM by always capturing in the middle of the data eye. Memory Training initially calibrates the interfaces to ensure adequate margin prior to operation.

Hardware Design Considerations for Teledyne e2v's Space-Grade DDR4

June 2021

By Dr. Rajan Bedi

TELEDYNE e2v | Semiconductors
Everywhereyoulook™

To conclude, DDR4 will allow the satellite industry to offer higher-throughput on-board processing and increased acquisition times enabling new Earth-Observation, space-science and telecommunication applications, e.g. ultra high-resolution imagery, live streaming video and on-board AI. As discussed, DDR4 contains new architectural and hardware features which need to be considered to ensure your design is right-first-time. In addition to the device datasheet, a user guide is also available for the DDR4T04G72.

For the first time, DDR4 will allow satellite and spacecraft manufacturers to avail of the large memory bandwidths that have been exploited by our commercial cousins for the last six years. Compared to existing, qualified DDR3 SDRAM, the DDR4T04G72 can be used with the latest space-grade FPGAs and microprocessors providing:

- A 62% increase in memory bandwidth (0.172 Tb/s with a data rate of 2.4 GT/s), doubling current transfer speeds
- A 25% increase in storage capacity
- 76% reduction in physical size
- 30% reduction in power consumption

TELEDYNE e2v Semiconductors

For further information, please contact:

Thomas Guillemain,
Marketing & Business Development,
Data Processing Solutions
thomas.guillemain@teledyne.com

TELEDYNE e2v Semiconductors

For further information, please contact:

Thomas Porchez,
Application Support,
Data Processing Solutions
thomas.porchez@teledyne.com

TELEDYNE e2v Semiconductors

For further information, please contact:

Jane Rohou,
MarCom Manager.
jane.rohou@teledyne.com