

Allowing server-class storage in embedded applications

20x better power efficiency, 20x lower cost

More performance with limited resources: local data storage in embedded applications can be a real design challenge. That comes from the huge amount of fast data coming from various types of sensors, such as high resolution camera, testbench recorder for industrial analytics, or data acquisition in a physics experiment. An embedded application has a limited budget in term of BOM cost, power consumption and space. In this article, a new way to implement high performance data storage is presented, allowing the use of server-class storage technology in an embedded environment.

Existing storage solutions

Data storage solutions for embedded applications can be built in the device, like serial flash memory or eMMC, or on a removable support, like SDCard, USB media or mSATA SSD. The table below provides an overview on the performance and capacity range.

	eMMC	SD Card	USB stick	mSATA SSD
Capacity	64GB	512GB	512GB	1TB
Throughput	400MB/s	90MB/s	200MB/s	600MB/s

Some embedded applications require higher performance. As an example, a high end video system based on 4k sensor at 60 frames per second requires $4096 \times 2160 \times 3 \times 60 = 2\text{GB/s}$. This is the raw data rate (obviously, the throughput is lower for compressed data). It is easy to understand that standard storage media can't be used for such performances. The bottleneck is on the media physical interface. A new media is needed: PCI Express (PCIe). This is a well-known physical interface used in server, computer and embedded applications. It is based on a high speed serial interface (up to 8Gbit/s in PCIe Gen3) combine in one or multiple lanes.

The PCIe interface has entered the storage market few years ago, through PCIe SSD products, in order to accelerate intensive applications such as big data analytics and data base. It has been followed by the introduction of the NVM Express (NVMe) specification, in order to optimize the protocol, and to leverage the market acceptance. NVMe SSDs were first available as a PCI card, then in a 2.5inch form factor. Recently, M.2 products, and even

as BGA chip, have been introduced, making it easier to be integrated in a consumer application, such as ultrabooks and tablets. In term of performance, a PCIe Gen3 x4 based interface can provide near to 4GB/s in a M.2 form factor, 6 times higher than mSATA.

	eMMC	SD Card	USB stick	mSATA SSD	NVMe SSD
Capacity	64GB	512GB	512GB	1TB	4TB
Throughput	400MB/s	90MB/s	200MB/s	600MB/s	4GB/s

Another performance indicator is the IOPS number: IO per second, where an IO is typically 4kB. Coming back to the high end video system example, and assuming that the video data are stored in a 4kB block on the storage drive, that leads to about 500kIOPS. This is a very standard number for server application, but very high for an embedded application. In the server, on the host side, the NVMe protocol is managed by the CPU as a software driver. In order to reach 500kIOPS, 1.8 cores at 3.3GHz are required, running at 100% of utilization. In an embedded context, that leads to about 4 x 2GHz cores processing capabilities, only for the NVMe management for 500kIOPS. Such computing system is more from the server domain than from the embedded domain, leading in a high BOM cost (about \$400), and a high power consumption (about 50W).

Here is the dilemma, NVMe SSD seems to be the only solution to support high performance storage, but it requires strong computing capabilities on the host side, which is at the opposite of the embedded requirements.


IP-Maker solved this issue with the introduction of the NVMe host IP.

IP-Maker has developed an optimized NVMe management for embedded applications, without using any CPU. The NVMe host driver has been implemented as a full hardware IP to be integrated in a FPGA or ASIC. This IP is integrated between the PCIe root port and the cache memory. The IP fully controls the data flow based on the NVMe protocol. Thanks to its optimized architecture, it can be easily used with a low cost FPGA, making it applicable to embedded applications.

In term of performance, it can be linked with up to a PCIe Gen3 x8 interface, delivering a sub-microsecond latency. Because it uses the NVM Express standard interface, it can be connected to any commercial NVMe SSD available on the market. Compared to a CPU-based system, it comes with 20x better power efficiency, and 20x lower cost.

Theory of operations

The IP can be used for both read and write operations. The following description focuses on a write system, such as a high end video recorder. The camera sensor is connected to the FPGA, the data transmitted from the sensor to the cache memory (external DDR). As soon as the buffer in the cache is ready, the NVMe host manager is configured with the necessary information regarding the data to store from the buffer to the NVMe SSD. That includes the buffer start address and the data size. Then the host manager sets the NVMe commands and manages the data transfer. As soon as the NVMe data transfer is done, the buffer becomes available for new data.



Implementation block diagram

Applications

Many embedded applications can benefit from this high performance storage technology. That provides the possibility to store more and faster data without adding any expensive BOM cost. Therefore that allows embedded system companies to design a new generation of applications with additional storage capabilities, providing more value to their customers.

High resolution camera	Testbench	Military	Embedded vision	Medical imaging	Gateway
Raw data for 4k and 8k sensors	Industrial testbench, monitoring system	Rugged embedded system, avionics recorder for playback sequence	High quality video with limited system space: robotics, drone, video surveillance	Raw data for portable ultrasound application	For IoT embedded server, with space and power limitation, for edge analytics computing

Application examples

Ready for the future of storage

With the increasing demand for high performance analytics application in cloud computing, many new technologies have been developed for the server market. IP-Maker succeeded in transferring the use of these technologies in the embedded world. The NVMe host IP from IP-Maker will leverage a new generation of embedded applications, with server-class high performance storage, and with embedded-class power requirement and BOM cost.

In a very near future, new non-volatile memories will emerge, such as MRAM, RRAM or 3DXP. That will dramatically reduce the SSD latency, by one or two orders of magnitude. The NVMe host IP, with its sub-microsecond latency, is ready for this new generation of storage technology.

About IP-Maker

IP-Maker is a leader in Intellectual Properties (IP) for high performance storage applications. IP-Maker's NVM Express (NVMe) technology provides a unique hardware accelerated solution that leverages the PCIe SSD performances, including ultra-low latency and high throughput. IP-Maker is a contributor to the NVMe specification. The ASIC and FPGA IP portfolio includes NVMe, Universal NandFlash Controller and ECC IP cores. The combination of the IP-Maker technology and its associate services dramatically cuts time-to-market.

Contact information

Les Néréides
55 rue Pythagore
13290 Aix en Provence
France

www.ip-maker.com
contact@ip-maker.com
+33 972 366 513