Taking Advantage of SSDs for Real-Time Recorders

Embedded Tech Trends 2013

Rodger H. Hosking
Pentek, Inc.

Real Time Recorder Objectives

- Digitize, Store and Playback Wideband Analog Signals
  - Signal frequencies > 3 GHz
  - Bandwidths > 1 GHz
- Multi-channel Synchronous Operation
  - Continuous or Triggered Capture
- Store High-speed Digital Signals
  - 1 GbE, 10 GbE, LVDS, and SerialFPDP
- Time and GPS Location Stamping
- Guaranteed Zero Data Loss
  - Disk Fault-tolerance for High-reliability
- Operation in Lab, Field and Severe Environments
- Tens of Terabytes of Storage
- Exploit Server Class PC Hardware Technology and Costs
  - Removable Storage
- Easy to Use
- Flexible Architecture

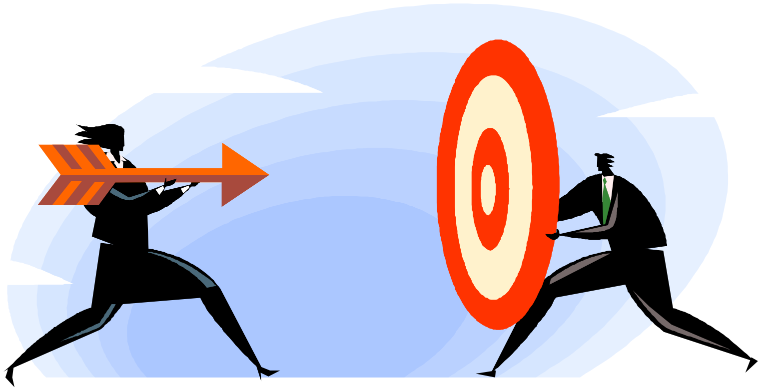

Challenges for High-Speed Recorders

- Signal acquisition and digitization
- Maintain signal integrity
- Synchronization across multiple channels
- Triggering and capturing accuracy
- Hardware data flow bottlenecks
- Disk drive speeds
- Optimize disk storage space
- RAID controller limitations
- Operating system overhead
- Buffering and latency issues
- Shock, vibration and temperature
- Ensuring zero data loss

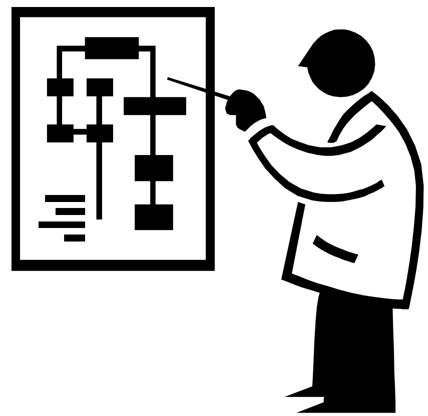

Advanced Technology for Recorders

- High-Performance Data Acquisition & Software Radio Modules
  - PCIe, XMC and VPX form factors
  - 90 dB SFDR signal integrity
  - Multi-channel synchronization, triggering and gating
  - Hardware DMA controllers - up to 8 GB/sec across PCIe
- Server Class PCs with PCIe Links
  - Multi-core CPUs with clock speeds to 4 GHz
  - Advanced PCIe Chipsets, Switch and Bridges
  - Fast DDR3 SDRAMs operating at 10 GB/sec

- RAID controllers delivering over 2.8 GB/sec
  - PCIe system interface and SATA III disk interfaces
- Solid State Drives
  - Immune to vibration and shock
  - SATA III transfer rates to 500 MB/sec or more

Magnetic Drives vs. Solid State Drives

- Magnetic drives use a spinning platter and numerous moving mechanical parts
- SSDs use silicon devices with no moving parts

Inside a Magnetic 2.5” Drive

Inside an SSD

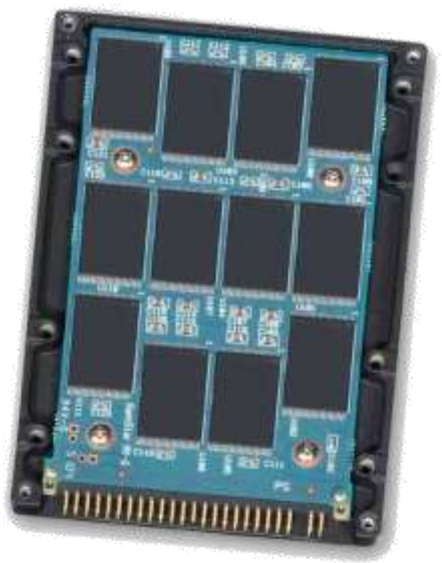

Comparing Disk Drive Technologies

- Typical disk drives used in high speed recorders

Feature	3.5" Magnetic	2.5" Magnetic	2.5" MLC Solid State	2.5" SLC Solid State
Cap (high end)	2 TB	1 TB	500 GB	250 GB
Read Rate	138 MB/sec	150 MB/sec	400 MB/sec	500 MB/sec
Write Rate	138 MB/sec	150 MB/sec	400 MB/sec	500 MB/sec
Weight	750 gm	111 gm	75 gm	75 gm
Power	10.7 W	1.75 W	0.28 W	0.25 W
Oper. Shock	30 G	350 G	1500 G	1500 G
Oper. Temp	5 to 55 ^{O}C	0 to 60 ^{O}C	0 to 65 ^{O}C	0 to 70 ^{O}C
Cost	$200	$180	$1,000	$3,000
Cost / GB	$0.10	$0.18	$2.00	$12.00
Write Cycles	∞	∞	10,000	100,000

SSD Types

- Same basic technology used in all SSD types:
  - Charge level (# of electrons) is stored on the oxide layer within each cell

SSD Type	Description	Bits Per Cell	Charge Levels	Max Write Cycles
SLC	Single Level Cell	1	2	100,000
eMLC	Enterprise MLC	2	4	30,000
MLC	Multi Level Cell	2 or 3	4 or 8	10,000
TLC	Three Level Cells	3	8	1,000

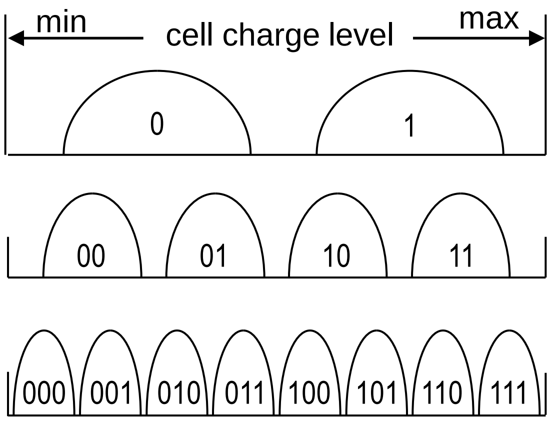

- Writing Operations
  - Cell must be erased before a new charge level is written
  - Oxide charge storage layer wears out each time a cell is erased & written
- Fewer charge levels per cell improves reliability & reduces error rate
  - Determining the stored charge level is easier with fewer levels to resolve
- eMLC
  - Improved MLC reliability through controller enhancement techniques

SSD Enhancement Techniques

- SSD NAND Memory Cell Operation
  - Writes musts be performed by first erasing a cell - not overwriting a non-empty cell
  - Cells are organized in memory blocks, each block containing many pages
  - Writes can only be performed on all the cells within a page
  - Erasing can only be performed only at the block level

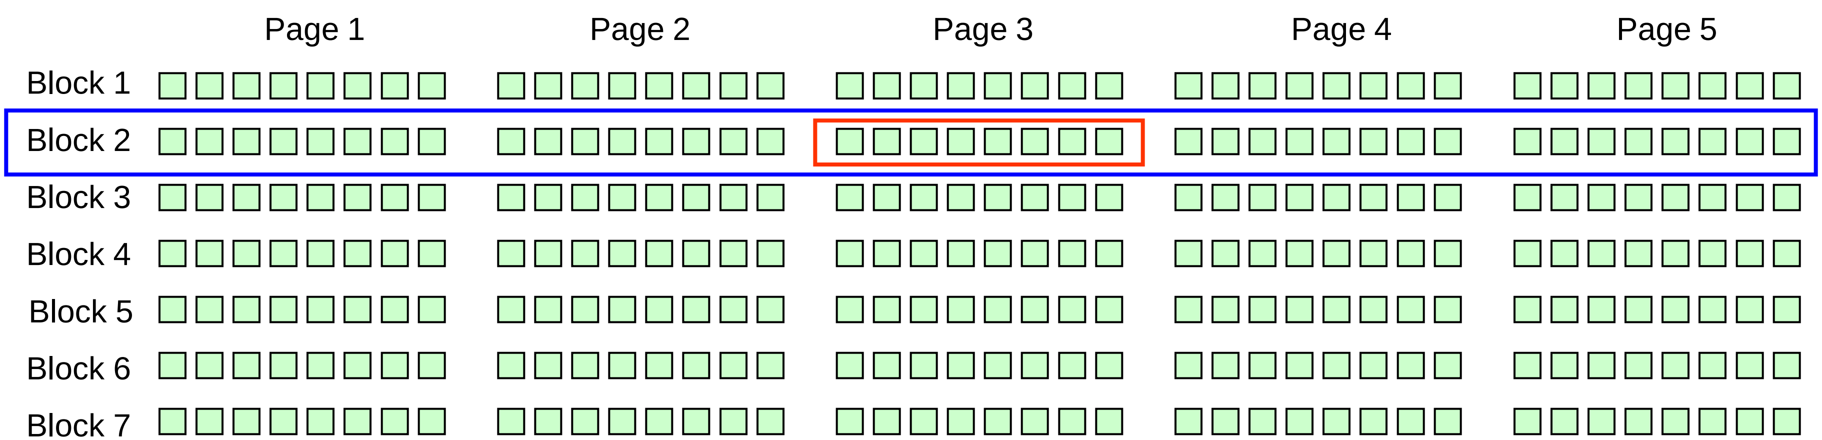

SSD Enhancement Techniques

- Wear Leveling and Garbage Collection
  - Microcontroller inside the SSD minimizes block erasures
  - New data is cached
  - A new erased NAND page must be allocated for each new write operation
  - The new page address is mapped into the old page address
  - The old page is then marked for erasure
- Over Provisioning
  - A portion of physical memory is reserved as swap space for wear leveling
  - Typically 7% is used: a 480 GB SSD really has 512 GB physical memory
  - Increasing the over-provisioning portion improves speed, endurance and reliability

Magnetic Drive R/W Speed

- An empty magnetic drive starts writing data on the outer tracks of the disk platter and then works towards the center track
- The outer tracks are longer (larger diameters), and the disk rotates at a constant speed
- So, the linear speed of the head against the platter decreases as the disk begins to fill
- The read/write transfer rate is proportional to the linear head speed
- This characteristic must be considered when designing for guaranteed read/write rates

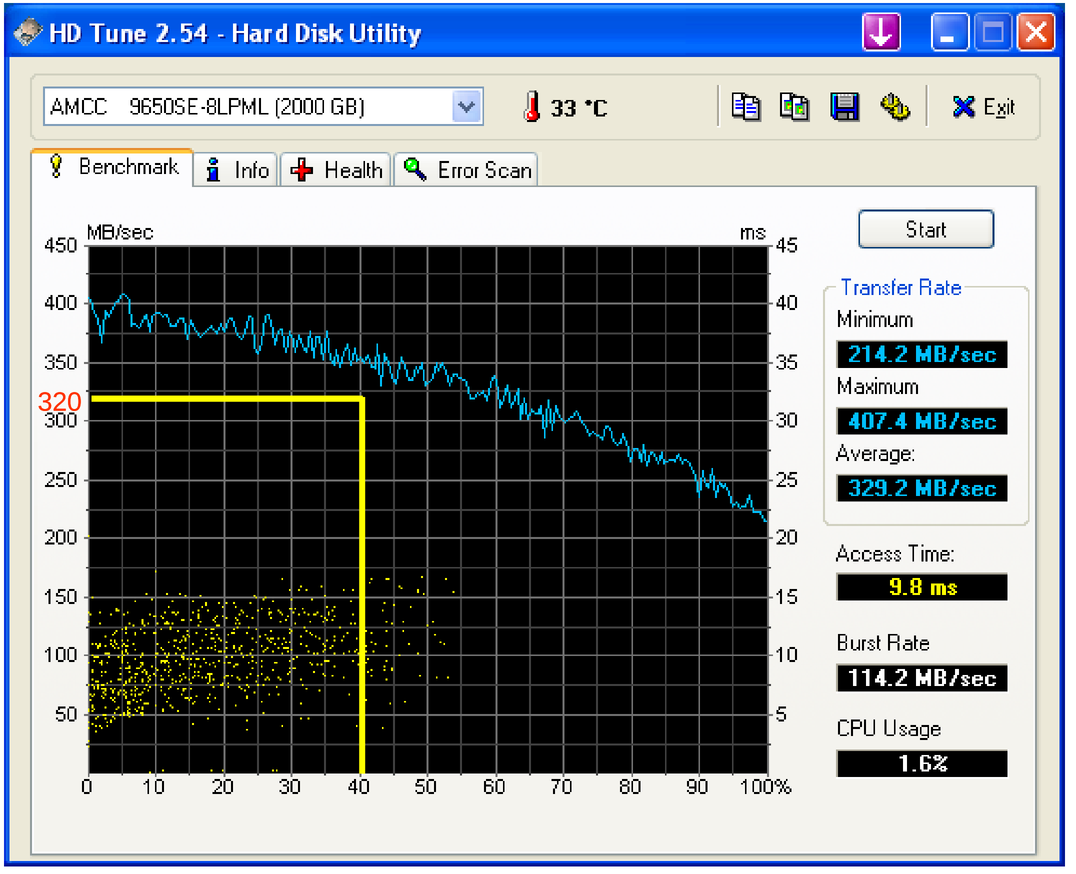

5 Drive RAID 0 – 500 GB WD RE3 – Magnetic 3.5”

Solid State Drive R/W Speed

- Solid state drives have no physical dimensional or rotational effects
- Array of non-volatile solid-state memories
- So, the read/write transfer speeds are uniform across the drive capacity
- Wear-leveling controller processing can affect the R/W speeds
- Some wear-leveling controllers impose unpredictable – completely unacceptable for real-time recorders
- SSDs must be very carefully characterized !!

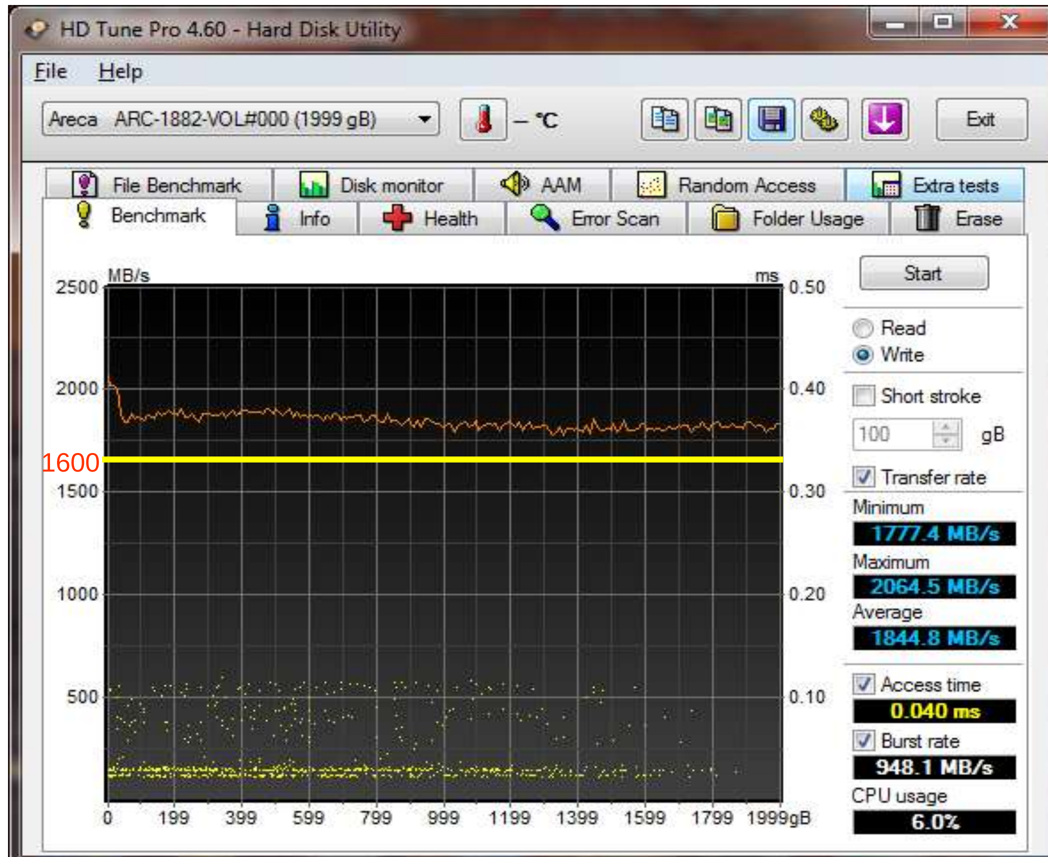

8-Drive RAID 0 – Crucial 256 GB M4 - SSD 2.5"

RAID Controllers

- **R**edundant **A**rray of **I**ndependent **D**rives
- Aggregates storage capacity of multiple disk drives
- Aggregates transfer speed of multiple disk drives – now to 2.8 GB/sec
- Provides levels of fault tolerance for critical, high-reliability applications

16-port PCIe Gen2 SATA-2 RAID Controller

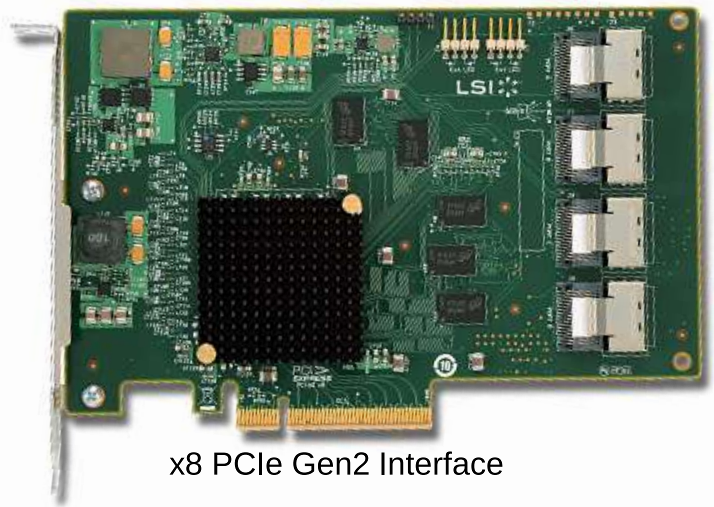

LEVEL	Description	Number of Drives*	Capacity Efficiency	Fault Tolerance	Write Speed Improvement
RAID 0	Block level striping	2	100%	none	200%
RAID 1	Mirroring (duplicating)	2	50%	1 drive	100%
RAID 5	Block striping + parity	3	66%	1 drive	200%
RAID 6	Block striping + 2x parity	4	33%	2 drives	100%
RAID 10	Block striping + double mirrors	4	25%	2 drives	200%

* Minimum number of drives shown. Benefits shown for that minimum number of drives.

Recording System Signal Flow

- Hardware DMA controllers handle all critical, real-time data transfers
  - PCIe Links eliminate data flow bottlenecks
  - Between I/O module and system memory
  - Between system memory and RAID controller
- System CPU & Operating System
  - Initializes the I/O modules, RAID controller, and DMA engines
  - Manages performance through interrupts
  - Neither one "touches" the real-time data !

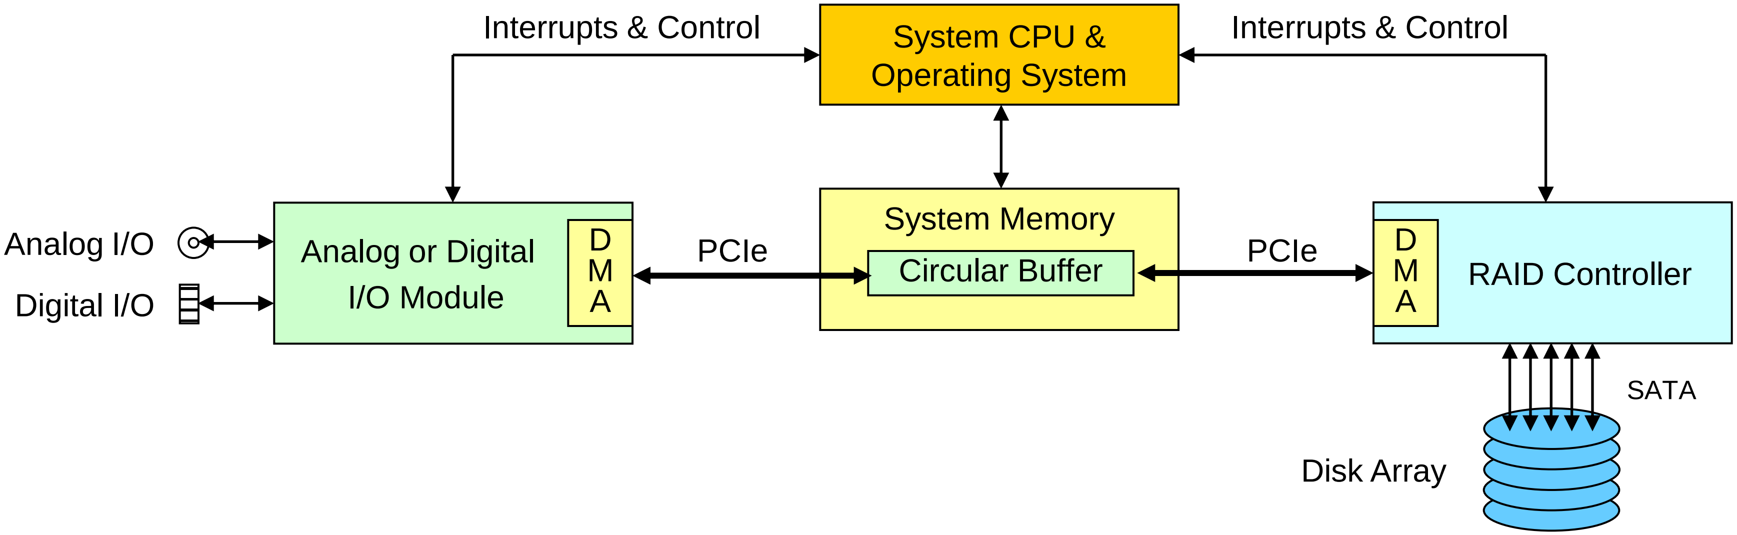

Pentek Talon™ Recorder Products

- Wideband analog recording and playback
  - Sustained recording rates up to 3.2 GB/sec in a 4U chassis
  - Real time signal bandwidths up to 1.5 GHz
  - Synchronous and asynchronous multi-channel operation
- Playback Capabilities
  - Recorded files can be reproduced through D/A
- Data Transport Capabilities
  - Removable Drive Bays/Drive Modules
- Ready-to-Use Instrument
  - Intuitive user-friendly GUI
- Integrated Signal Viewer
  - Virtual oscilloscope and spectrum analyzer
- Multiple form factors
  - Lab, portable and rugged versions
- Built upon a Windows 7 Workstation
- NTFS file system
  - Allows the user to work with files immediately after capture
  - Supports all Windows signal analysis and processing tools (Matlab)

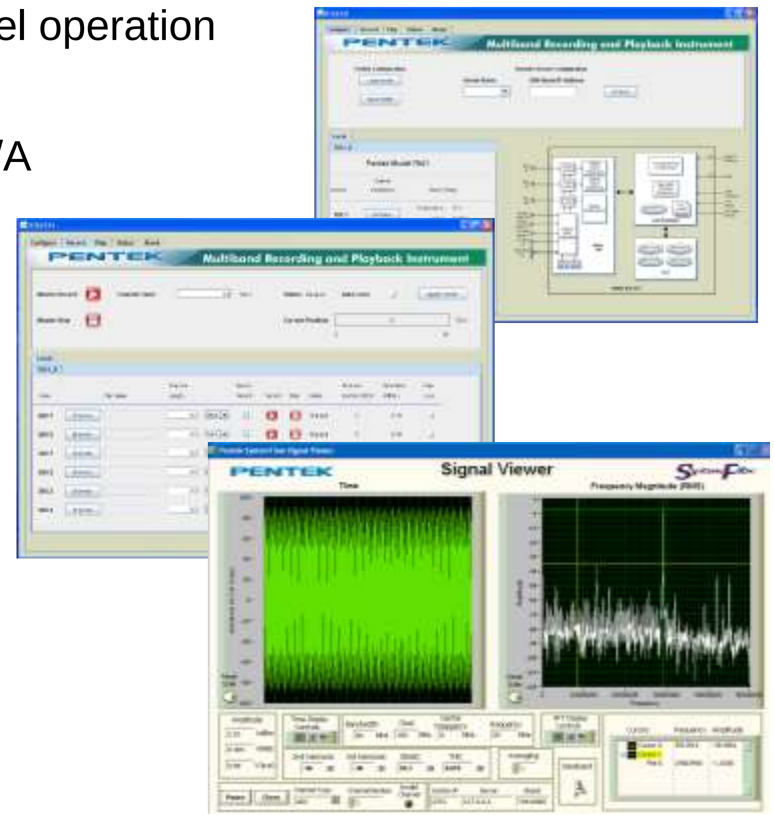

Talon 200 MHz 16-bit A/D Systems

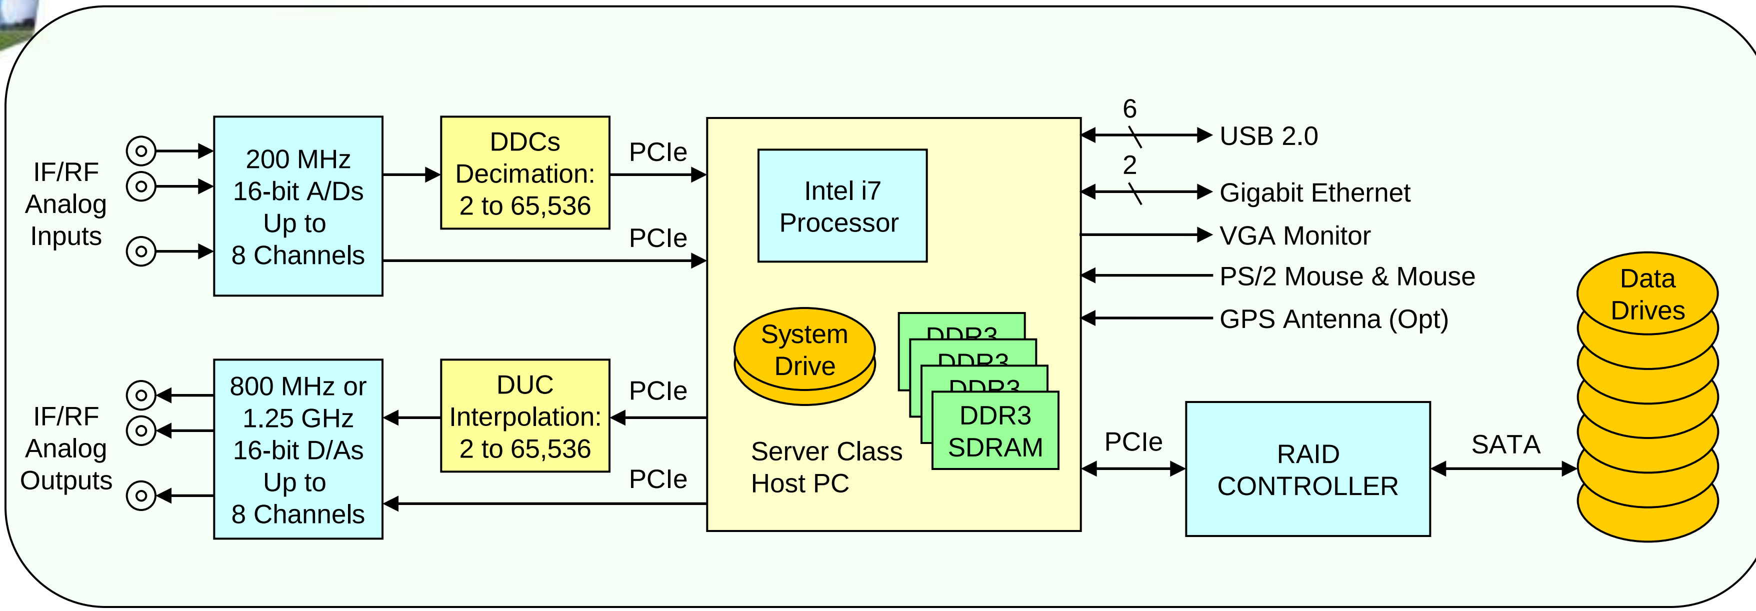

RTS 2726 Portable
Solid State Drives

RTS 2706 Lab System
3.5” High Capacity Drives
24 TB Storage

RTR 2746 Rugged
Solid State Drives
Harsh Environment

RTX 2786 3U VPX
Solid State Drives
Conduction Cooled

SSDs for Wideband Recorders

- SSDs offer read/write rates 3 times faster than magnetic
  - Uniform data rates across entire disk
  - No moving parts or mechanical structures
- Fewer SSDs required for high-speed RAID arrays
  - Example: 2.8 GB/sec RAID array can use 8 SSDs instead of 24 magnetic drives
- SSDs are well suited for extreme military environments
  - Virtually immune to shock & vibration
  - Equal or better than hard drives at extended temperatures
- SSDs Power Consumption
  - Less than half the Watts/GB compared to magnetic
- Endurance Issues
  - SLC offers reliability rivaling the best magnetic drives
  - MLC controller strategies are improving dramatically
  - Sequential writes in recorders minimize endurance issues
- Costs for SSDs
  - MLC cost per GB is 12-20 times higher than magnetic
  - SSD cost per GB is dropping much faster than magnetic