

**“Overclocking Data Storage Subsystems:
A SATA-IV Proposal”**

**Paul A. Mitchell, B.A., M.S.
Systems Development Consultant**

All Rights Reserved

Author's Brief Bio

began using computers in 1971, U.C. Irvine, grad school

published author, computer graphics, Harvard 1977

frequent contact with system performance issues

biased for singular enthusiasts, DIY builders,
low-end budgets, human productivity

high-performance should NOT
require top dollar -or- special expertise
be limited to large organizations

U.S. patent pending on very high-speed storage device

Human Productivity: Small Changes Add Up

Example: **save 1 ½ seconds every working minute**

FTE = 2,000 hours x 60 minutes = 120,000 minutes/year

120,000 @ 1.5 = 180,000 seconds saved

180,000 / 3,600 = 50 human hours saved

100 people @ 50 hours = 5,000 human hours / year

5,000 / 2,000 = **2 ½ FTE**

The Legacy PCI Slot

33 MHz clock rate

x 32 bits per cycle =

1,056 Megabits per second (“1G”)

/ 8 bits per byte transmitted =

132 Megabytes per second (one direction)

NOTE equivalence with **ATA-133** (aka “PATA”)

Overclocking Data Storage Subsystems: A SATA-IV Proposal

SATA-III

6.0 GHz clock (“6G”)

/ 10 bits per byte =

600 MB per second (one direction)

SATA-II

3.0 GHz clock (“3G”)

/ 10 bits per byte =

300 MB per second (one direction)

SATA-I

1.5 GHz clock

/ 10 bits per byte =

150 MB per second (one direction)

PCI Express 3.0

8.0 GHz clock

/ ~8 bits per byte = **(note change to “8” here)**

1.0 GB per second (one direction)

PCI Express 2.0

5.0 GHz clock

/ 10 bits per byte =

500 MB per second (one direction)

PCI Express 1.0


2.5 GHz clock

/ 10 bits per byte =


250 MB per second (one direction)

Why change
from
10 bits / frame
to
130 bits / frame?

8b/10b “Legacy Frame”:



128b/130b “Jumbo Frame”:



Western Digital Corporation: “Advanced Format Technology”

Figure 1

Legacy Architecture





Figure 2

Advanced Format Architecture



SATA-IV Proposal

“Sync” with PCIe 3.0 “jumbo frame”:

1 start bit + 16 data bytes + 1 stop bit (“128b/130b”)

6.0 G / 8 = **750.0 MB per second**

3.0 G / 8 = **375.0 MB per second**

1.5 G / 8 = **187.5 MB per second**

Overclocking Data Storage Subsystems: A SATA-IV Proposal

Exact Divisor is:

$$130 \text{ bits} / 16 \text{ bytes} = \underline{\underline{8.125}}$$

$$16.0 \text{ G} / 8.125 = 1,969 \text{ MB/s}$$

$$12.0 \text{ G} / 8.125 = 1,477 \text{ MB/s}$$

$$8.0 \text{ G} / 8.125 = 985 \text{ MB/s} \quad \text{SATA-IV default}$$

$$6.0 \text{ G} / 8.125 = 738 \text{ MB/s}$$

Visible Present Workstation

5.25" 4-in-1 enclosures proliferate for 2.5" devices
(e.g. Icy Dock, Thermaltake, Enhance Technology, etc.)

4 x SATA-IV channels in RAID 0 mode

@ 984.6 MB/s = 3,938 MB/s

controller efficiency x (500 / 600) (assumed)

3,282 MB/s **3.2 GB/s**

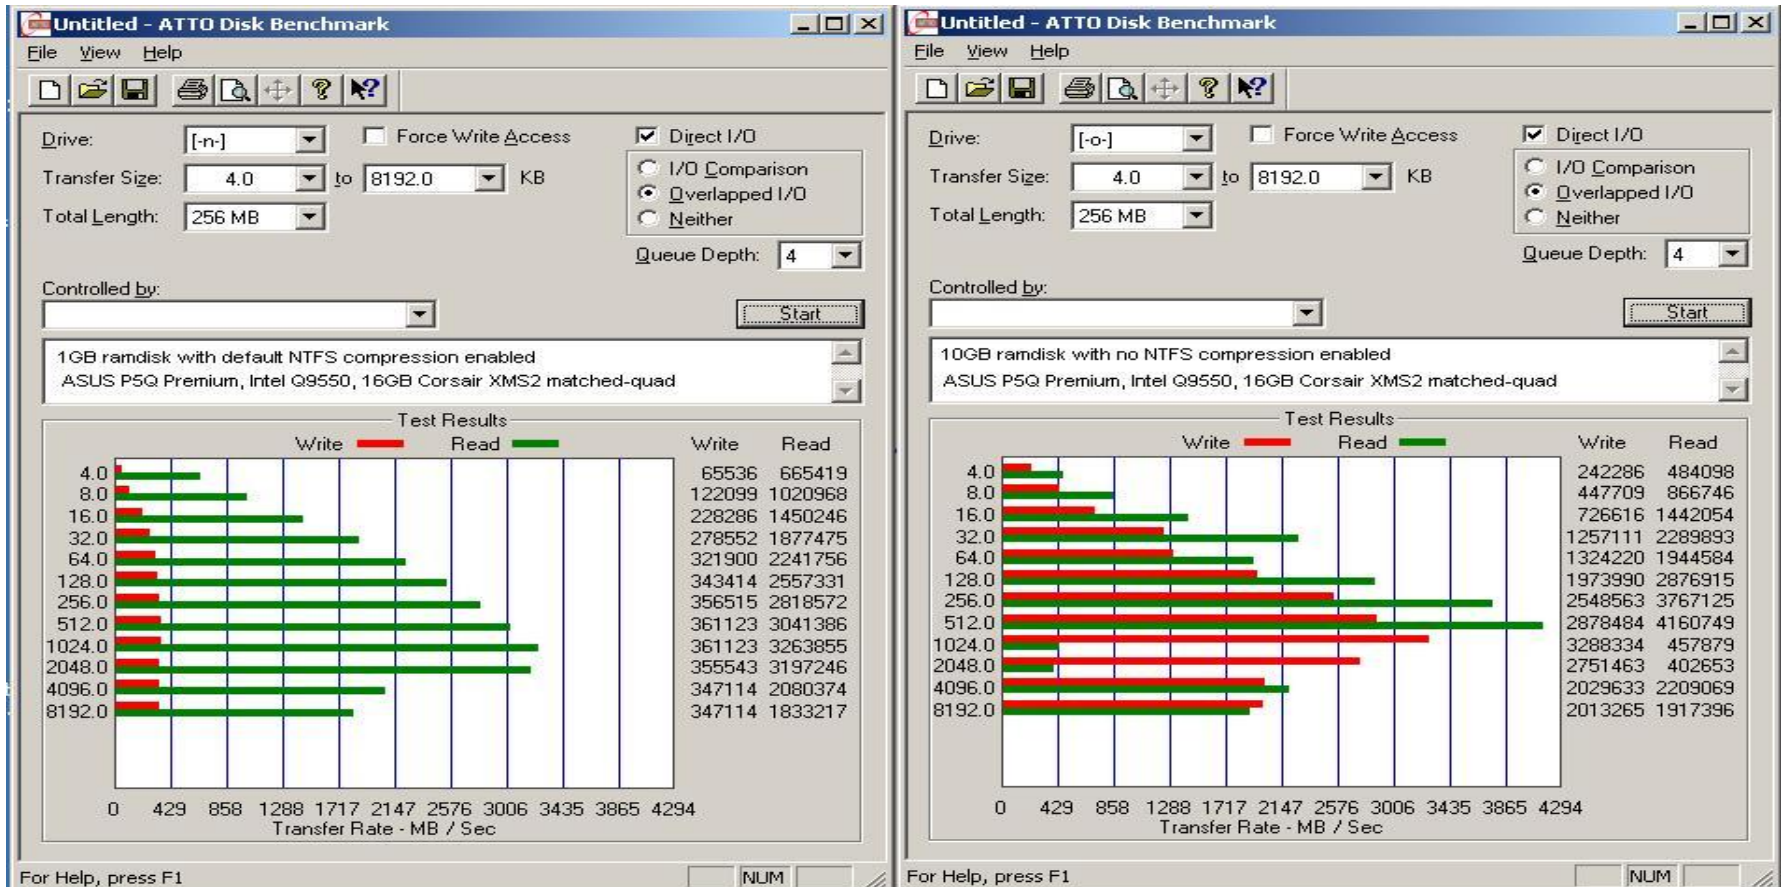
Overclocking Data Storage Subsystems: A SATA-IV Proposal

Compare “Raw Read” Results with *RamDisk Plus*
10GB ramdisk, 16GB dual-channel DDR2-800 matched quad
ASUS P5Q Deluxe motherboard, Intel Q9550 quad-core CPU



Overclocking Data Storage Subsystems: A SATA-IV Proposal

ATTO Results with *RamDisk Plus*:
compare NTFS compressed -and- uncompressed ramdisks



Deluxe Future Workstation

2 x 5.25" 4-in-1 enclosures =
8 x SATA-IV channels in RAID 0 mode
@ 984.6 MB/s = 7,877 MB/s
controller efficiency x (500 / 600) (assumed)

6,564 MB/s **6.5 GB/s**

controller efficiency x (0.760) (measured)

5,986 MB/s **5.9 GB/s**

Existing Hardware Examples Highpoint RocketRAID 2720SGL



- x8 PCIe 2.0 edge connector
- 2 x SFF-8087 ports
- 8 x 6G SATA/SAS ports
- 4.0 GB/s upstream bandwidth

Overclocking Data Storage Subsystems: A SATA-IV Proposal

ICY DOCK model MB994SP-4SB-1
4 x 2.5" SSD/HDD in 1 x 5.25" Bay
SATA Hot-Swap Backplane Cage



Overclocking Data Storage Subsystems: A SATA-IV Proposal

Adaptec 2236600-R mini SAS x 4 (SFF-8087)
to (4) x 1 Serial ATA fan-out Cable - 0.5M



Conclusion

The SATA-IV Standard should include support for:

variable channel speeds, perhaps with pre-sets:

6, 8, 12 and 16 GHz (and so on, beyond the horizon)

optional “jumbo frames” identical to PCIe 3.0 spec:

1 start bit + 16 data bytes + 1 stop bit (“128b/130b”)

Overclocking Data Storage Subsystems: A SATA-IV Proposal

Further Reading:

“Overclocking Data Storage Subsystems: One Approach to Variable Channel Bandwidth,”
by Paul A. Mitchell, July 27, 2010

http://benchmarkreviews.com/index.php?option=com_content&task=view&id=11178&Itemid=21

“Visible Computer Futures,” by Paul A. Mitchell, October 31, 2010

http://benchmarkreviews.com/index.php?option=com_content&task=view&id=12508&Itemid=21

“Technical Review and Evaluation of RamDisk Plus Software,” October 1, 2009

http://benchmarkreviews.com/index.php?option=com_content&task=view&id=8479&Itemid=23

“A Major Performance Enhancement to the Industry-Standard SATA Protocol,” May 26, 2010

http://benchmarkreviews.com/index.php?option=com_content&task=view&id=10291&Itemid=22

“Advanced Format Technology White Paper,” by Western Digital Corporation

<http://www.wdc.com/wdproducts/library/WhitePapers/ENG/2579-771430.pdf>

“HighPoint 2720SGL RocketRAID Controller Review – Amazing 3GB/s Recorded With 8 Crucial C400 SSDs,”
by Paul Acorn, February 14, 2012

<http://thesdreview.com/our-reviews/highpoint-2720sgl-rocketraid-controller-review-amazing-3gbs-recorded-with-8-crucial-c400-ssds/>

PCI Express® 3.0 Frequently Asked Questions

http://www.pcisig.com/news_room/faqs/pcie3.0_faq/

Highpoint website: <http://www.highpoint-tech.com>

Icy Dock website: <http://www.icydock.com>

Adaptec website: <http://www.adaptec.com>

Reader's Notes: