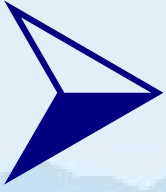# Metsähovi four-Gbps recorder and it's applications

Jouko Ritakari, Jan Wagner,
Guifré Molera
Metsähovi Radio Observatory

Background

➢ VLBI: High-speed, high-volume data acquisition

➢ Correlated in special HW correlators, until now computers have been too slow.

➢ Five years ago hard disks replaced the tape recorders

The Metsähovi correlator engine

➤ Developed for the Cell processor (Sony Playstation 3) in 1H 2007

➤ About ten times more powerful than Intel-based solutions

➤ Less platform-dependent: Only a few pages of code need to be rewritten if platform changes

➤ Needs high-speed input data

Problem(s):

- Getting data in and out of specialised disk controller and old PATA disk packs is slow

- Data input is moving to 10 Gbps Ethernet when iBob or DBBC are used

- Testing four-Gbps streaming is easier with four-Gbps recorder

- The seven-year-old VLBI data recorder de-signs would benefit from speedup

Solution:

➢ Feasibility study: a four-Gbps recorder made from only commercial off-the-shelf compo-nents

➢ 10 Gbps Ethernet for input/output

➢ Fast international data streaming using the Tsunami protocol

➢ Fallback plan: If it does not work at four Gbps, use two at two Gbps each
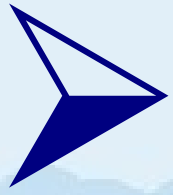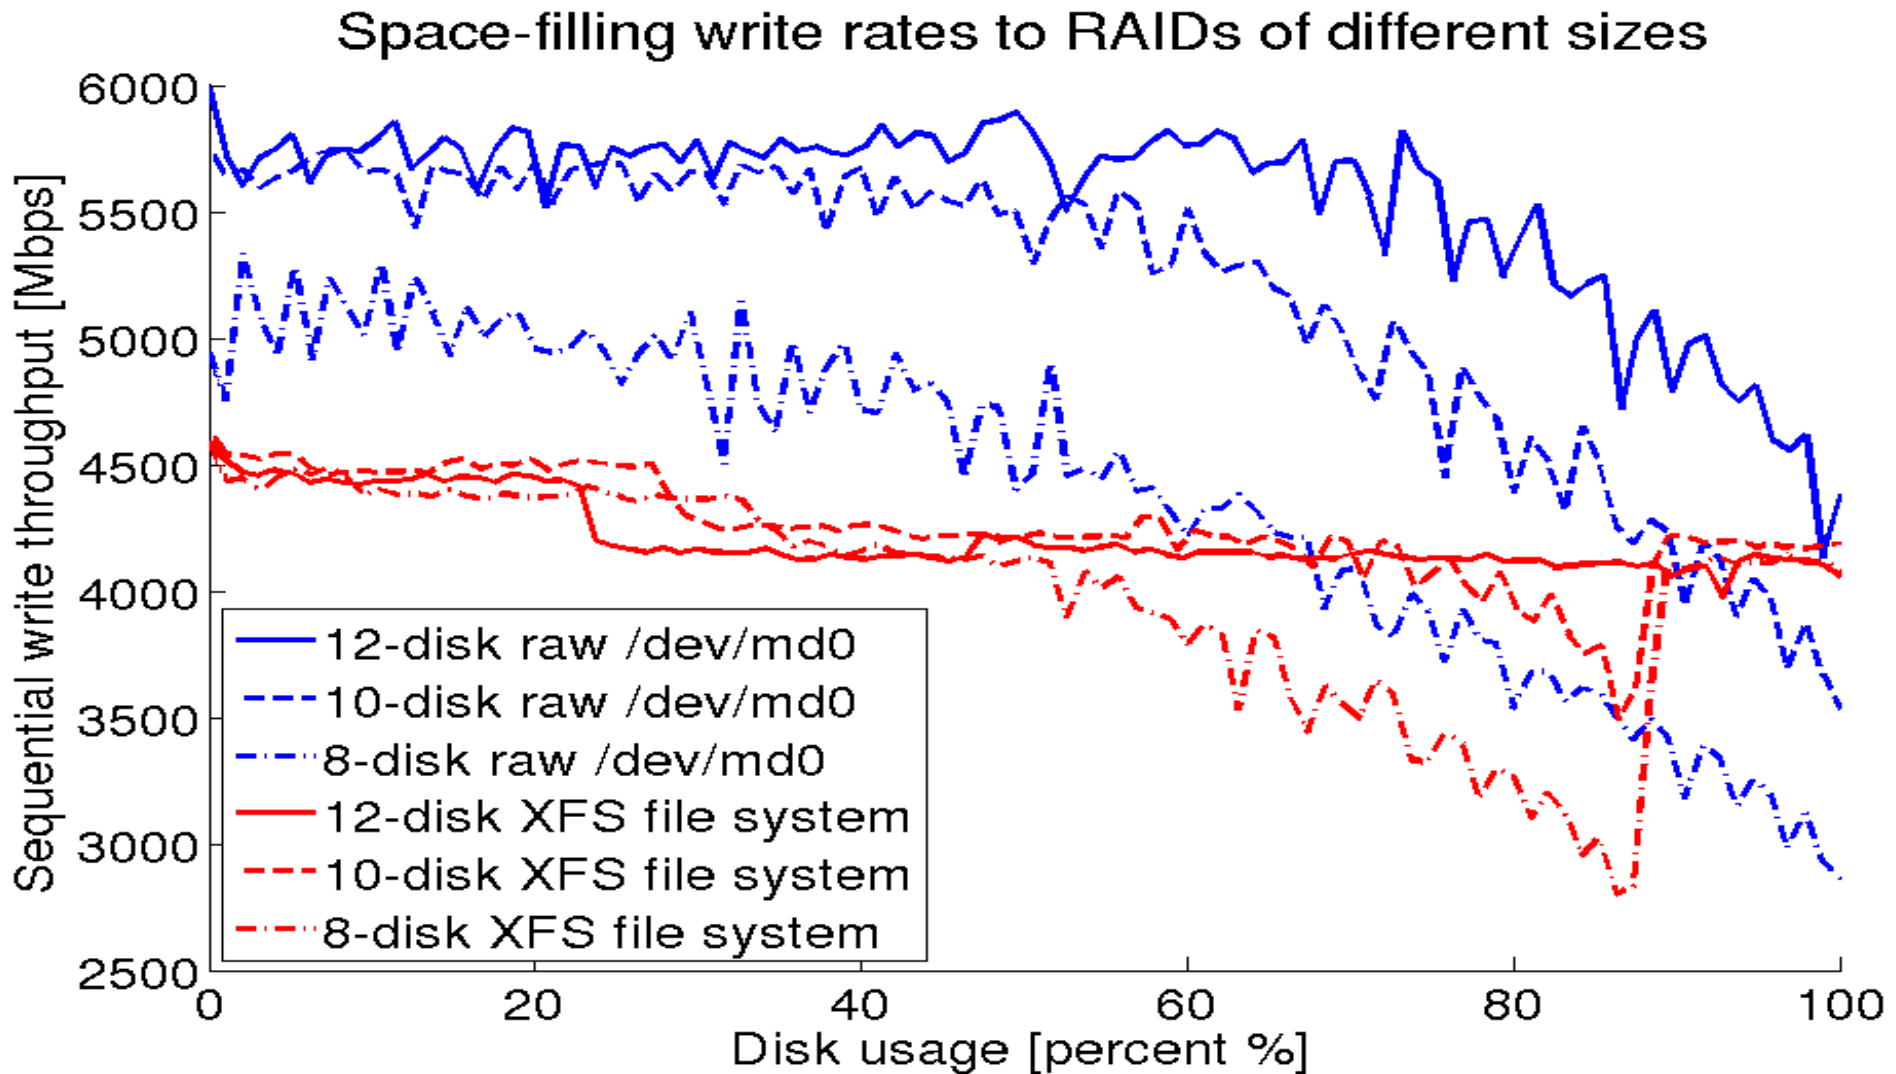
Two sampler subsystems

- Primary sampler: Berkeley iBob/iADC with two 1 Gsamples/s 8-bit ADCs, dual 10Gbps Ethernet output

- Secondary sampler (for old system): Four-channel 8-bit or two-channel 12-bit ADC, max. 66 Msamples/s sampling speed, VSI output
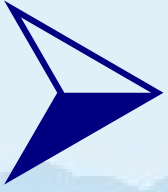
Disk subsystem tests

➢ Original tests with 12-disk raid array, one-year-old Seagate 250 GB disks

➢ Best write speed with ext2 fs 5.5Gbps, de-creases to half of that on inner tracks

➢ Replaced the disks to newest and fastest Samsung Spinpoint F1 750 GB

➢ Consistent > 4Gbps performance

Samsung 750GB disk performance

Space-filling write rates to RAIDs of different sizes

Legend:
- 12-disk raw /dev/md0
- 10-disk raw /dev/md0
- 8-disk raw /dev/md0
- 12-disk XFS file system
- 10-disk XFS file system
- 8-disk XFS file system

10 Gbps Ethernet tests

➤ Equipment: Myrinet and Chelsio NICs

➤ Initial iperf tests: 2-3 Gbps, 100% CPU

➤ Improved to 8 Gbps with jumbo frames

➤ Turned off udp checksum, Ethernet has already one.

➤ Final iperf UDP results:

　➤ MTU 1500: 4.88 Gbps

　➤ MTU 4470: 8.55 Gbps

　➤ MTU 9000: 9.04 Gbps

Combined 10GBE + disks

➢ Both tasks are CPU-limited

➢ Used two dual-core AMD 2212 processors and threads to divide the load

➢ Optimum performance with two six-disk raid arrays and two udp streams
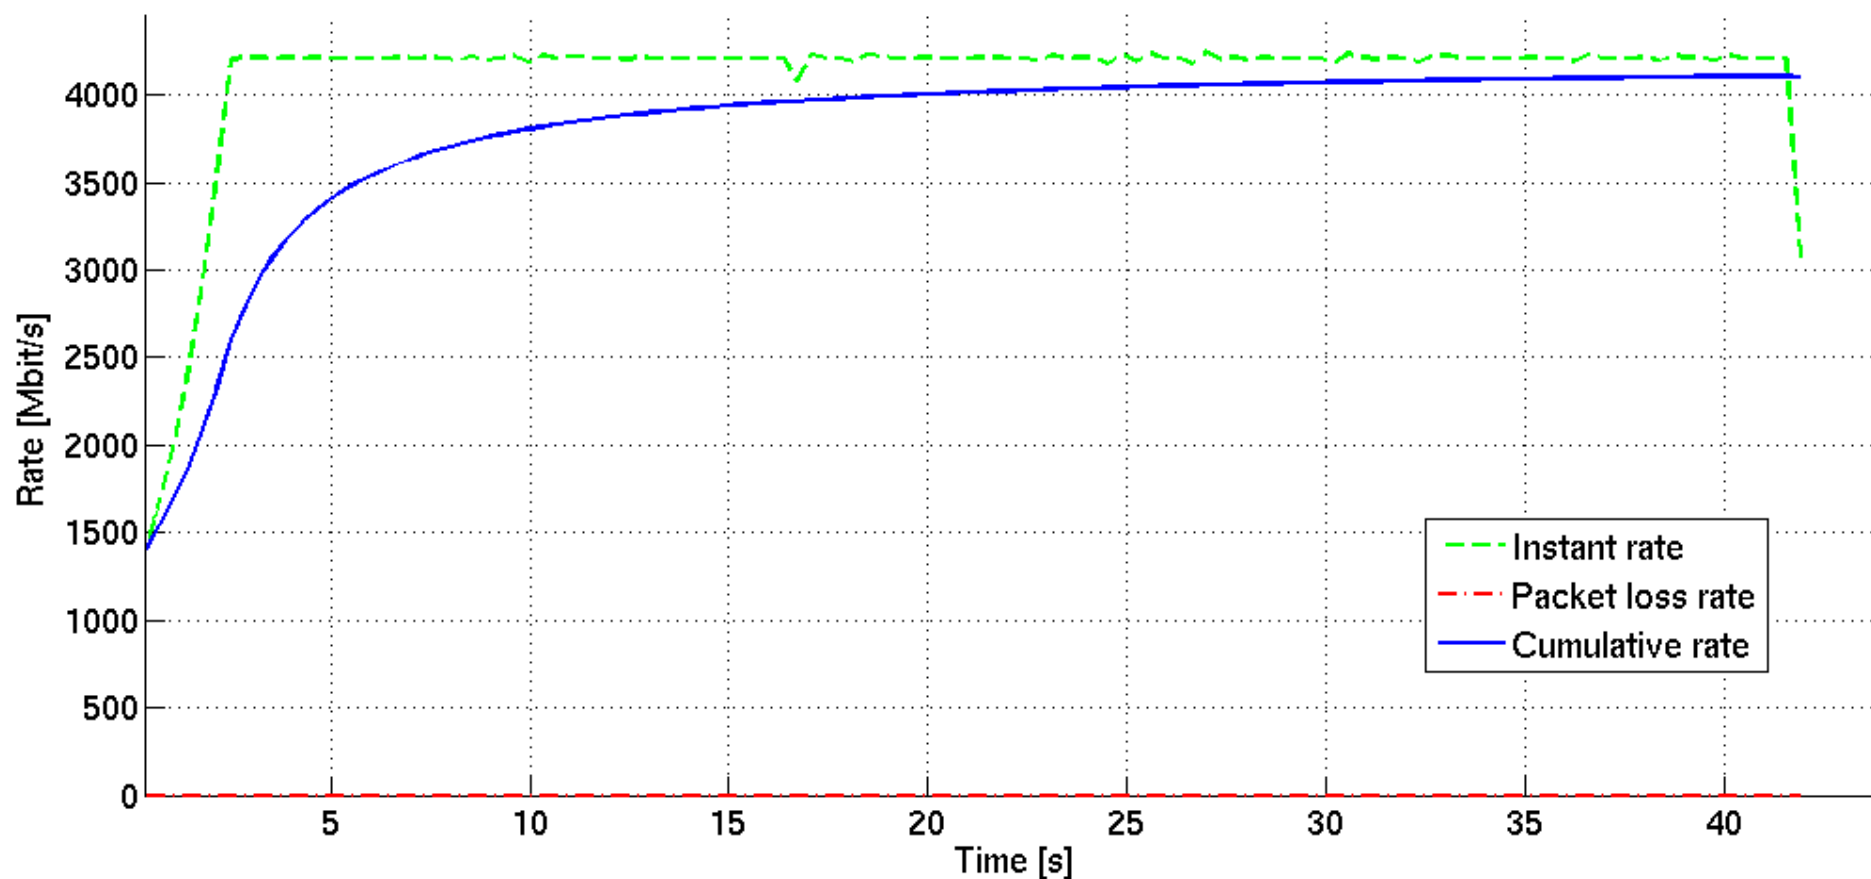
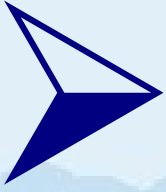➢ Sustained 4 Gbps transfer from 10GBE to disks
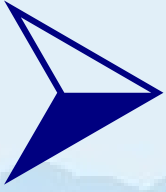
Tsunami tests

➢ Four-Gbps tsunami would be very useful

➢ ...so would be direct transfer from iBob udp/ip packetizer to correlator...

➢ A single tsunami transfer runs at 4 Gbps, two parallel transfers are a little faster

➢ And: Zero packet loss!!!
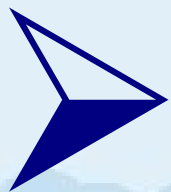
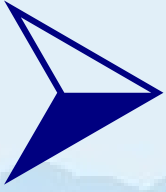# Lossless 4 Gbps Tsunami transfer

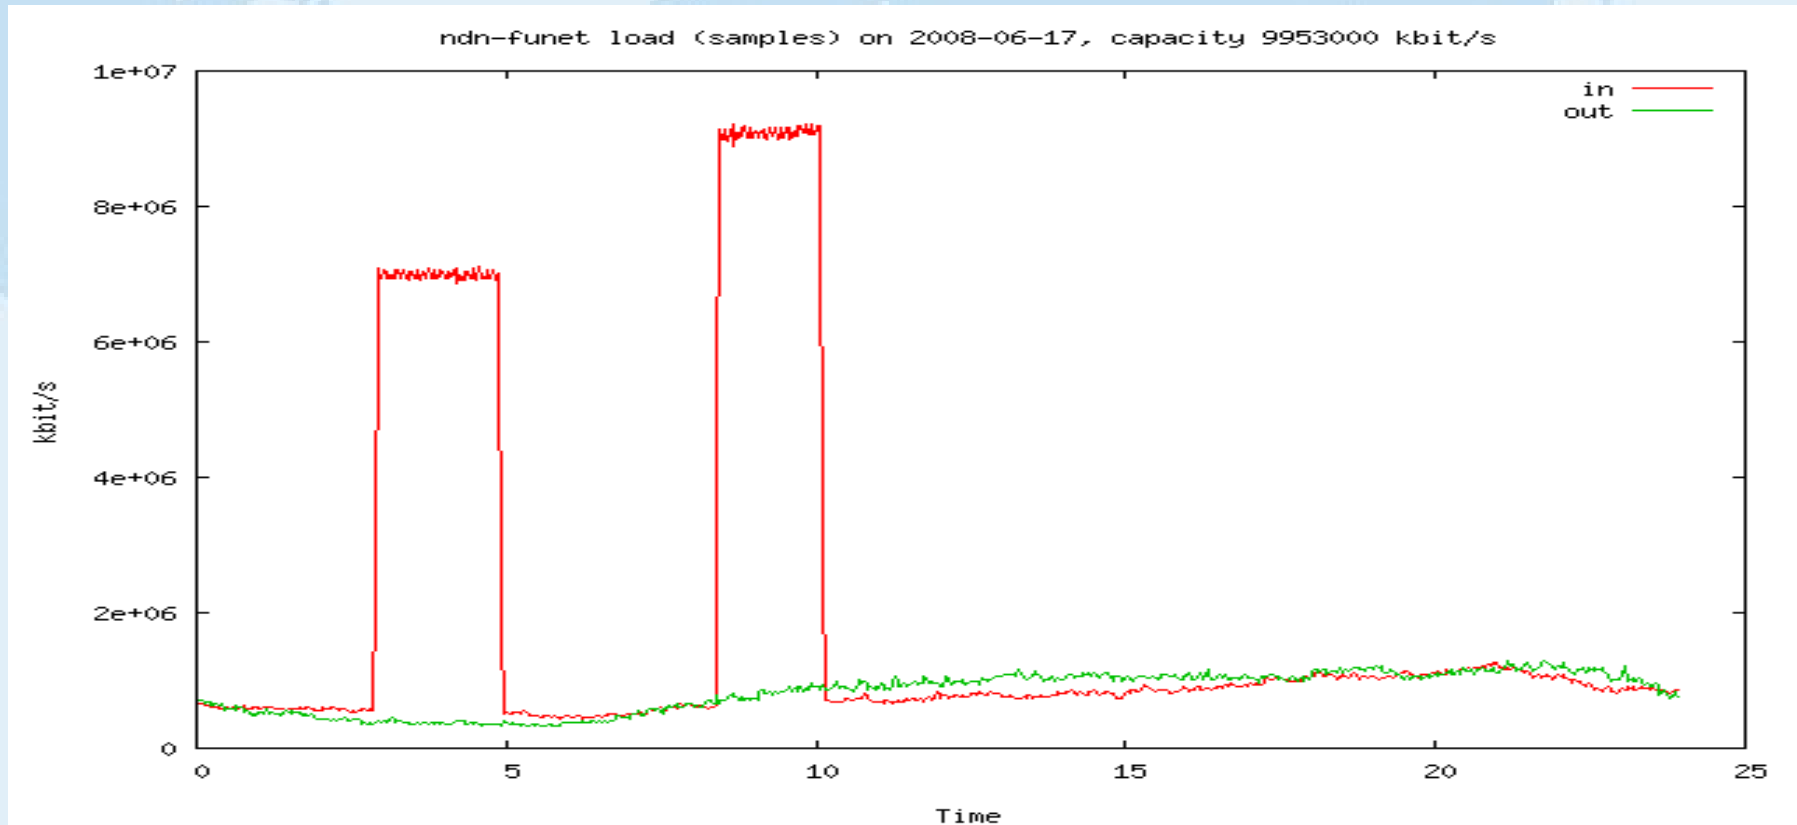4 Gbps File Transfer of a 20GB File to RAID-0

Scientific results

➢ Maser and Kronian water observations (Pogrebenko & al)

➢ Venus Express and Mars Express spacecraft tracking

➢ 8 Gbps eVLBI demonstration in the 7[th] eVLBI workshop, Shanghai

8 Gbps demo, Shanghai

ndn-funet load (samples) on 2008-06-17, capacity 9953000 kbit/s

New improved disk-packs

➢ Two prototypes with different lifting mechanisms

➢ COTS parts, 30TB capacity (11/2008)