Document Mark_6_Recorder_Pa10			Description
Description Discussion Points	Document ID TBD	Revision PA10	
Created by / Signature M. Taveniku	Title Mark 6 Recorder (proposed)		Date 2011-01-22

Discussion Points and Proposal for:

XCube Logger for Next Generation VLBI Data System

MARK 6 Recorder (proposed)

XCube Logger/Recorder and Disk Pack

Table of contents

1. INTRODUCTION 3

1.1. WHAT IS THE BASE FOR THE XCUBE RECORDER 3
1.2. XCUBE RECORDER DEVELOPMENT STATUS 5
1.3. XCUBE RECORDER TO MARK 6 SPECIFICS DISCUSSION 6
1.4. RELATION TO OTHER WORK 8

2. GENERAL SPECIFICATIONS 10

2.1. HIGH ALTITUDE OPERATION 12

3. REQUIREMENTS FOR VLBI OPERATION 13

3.1. OPERATIONAL CONCEPT 13
3.2. DETAILED REQUIREMENTS 13

4. QUESTIONS AND ANSWERS 15

5. DEVELOPMENT PLAN 23

5.1. RESOURCES 23
5.2. SCHEDULE (TBD AND AGREED ON) 24
5.3. RECORDER PERFORMANCE AND RELIABILITY VERIFICATION 24
5.4. HARDWARE DEVELOPMENT 25
5.5. SOFTWARE DEVELOPMENT 26

1. INTRODUCTION

This document has been prepared to assist in the consideration of the XCube real-time data logger system as the basis of a next-generation VLBI data system. Unlike the current-generation Mark 5x series of VLBI data systems, the XCube system is built entirely of high-end commercial off-the-shelf (COTS) hardware components; the XCube system's real-time capabilities are realized through special software that manages the data flow to disks in innovative ways. Though originally built for a different market, the XCube system is built around the same basic real-time high-data-rate requirements that VLBI's demands, particularly with respect to management of slow and/or failing disks. Adapting the XCube system to VLBI special needs appears to be a straightforward task.

In December 2010, an XCube data-recording system was demonstrated at Haystack Observatory to record at a sustained rate of 8Gbps with data provided by two simultaneous 4Gbps 10GigE data links to an RDBE unit. Current expectations are that a system built with more recently available COTS hardware will readily sustain at least 12Gbps to a disk array, possibly as high as 24Gbps.

The transition from the current Mark 5x series of data systems to an XCube-based system appears to be straightforward and can be done on a gradual basis. There is also potential for using the XCube system with modified Mark 5 disk modules, allowing significant cost recovery of expensive module components.

The XCube system is designed to be flexible in terms of configuration to support a multitude of operational scenarios including high-altitude operation, burst mode operation, ultra-high speed recording (using Solid State Disks) as well as synchronous operation across multiple recorders. This flexibility will allow it to adapt well to the varying demands of diverse VLBI observations.

1.1. What is the base for the XCube Recorder

The XCube recorder concept is based on the observation that commercial hardware and software components are becoming exponentially faster and cheaper at the same time, and that applications such as large data bases and multimedia streaming have similar requirements to the needs in the VLBI community. The systems and techniques developed for the commercial applications are similar, but not exactly coincident with requirements for high-performance sensor systems. However, with slight (careful) modifications to the commercial components and system they can meet VLBI requirements. The task in this case, is to, reliably, continuously, and in real time capture and record to removable media multiple GByte/s from one or more sources simultaneously.

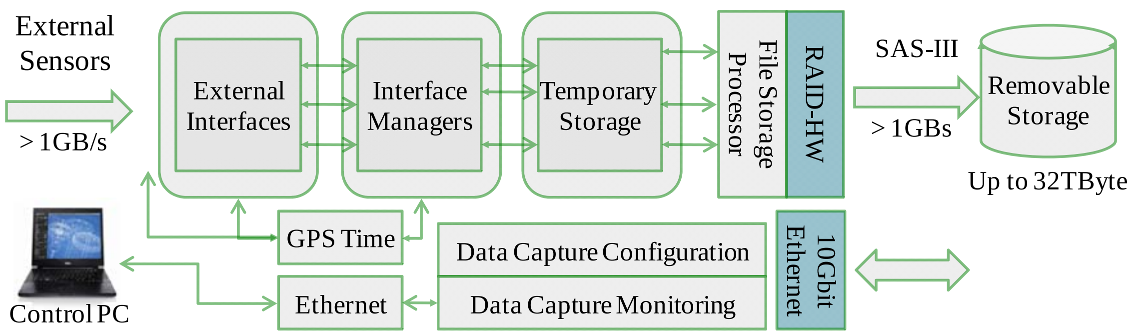

Figure 1: Multiple external data sources provide real-time data streams to be stored on disk.

The XCube recorder uses no special hardware, just carefully chosen components and then software modifications to drivers, careful management of data-streams and system memory to provide reliable performance.

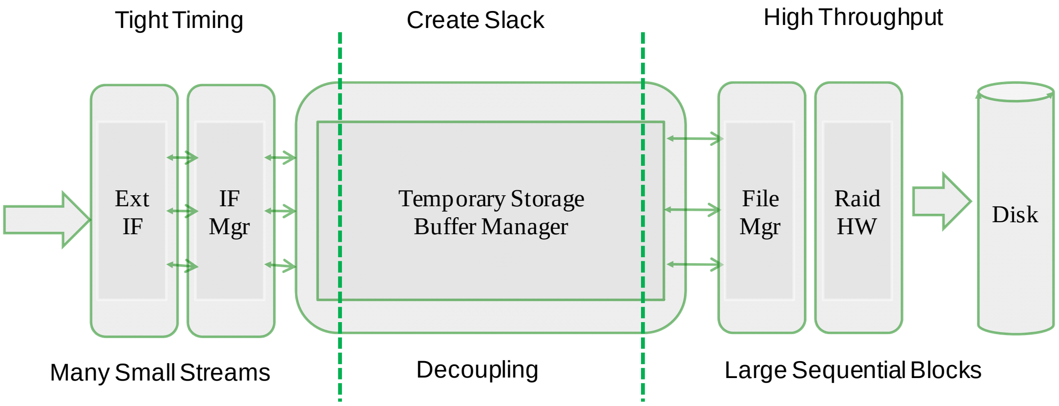

Figure 2: Data-path with two halves, left side handles streams of small packets with tight timing, while the right halve maximize throughput with long latency large block transfers.

The basic concept for the software defined data capture is to have a "custom made" driver for a specific interface (say 10Gbit Ethernet) and have that tuned to reliably handle the real-time small-packet-based data stream from that device. On the other side, writing data to disk needs to be done in large continuous chunks, with sometimes long delays between writes. This is done by the File Writer software module that takes care of scheduling disks buffers. There is a discrepancy between the line-speed packet-processing and the large-buffer long-latency block processing of the disk subsystem. To address this difference, a set of elastic stream buffers connect the input and output streams, decoupling them from each other.

In order to create a flexible data capture environment for the generic product the software is designed so that each hardware interface has an associated specialized "Interface Manager" (IF Mgr). These interface managers are specifically written to handle either a specific hardware interface in a generic fashion (Ethernet for example) or a specific interface and device combination (for example a digital backend, such as the RDBE using UDP/VDIF connected to 10Gbit Ethernet).

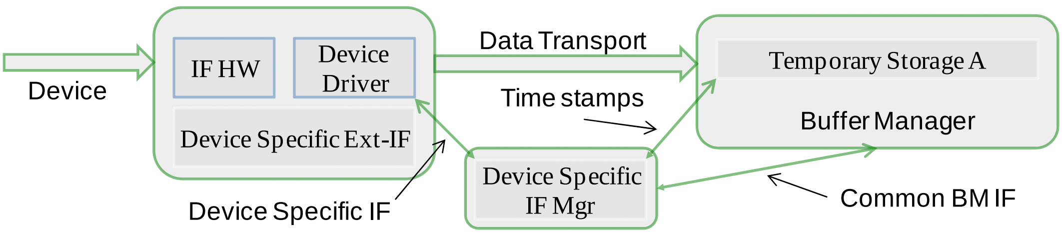

Figure 3: An interface manager provides device-specific handling of a hardware port, or a port and external device combination.

The Interface managers are standalone programs that can provide device specific handling, such as IP-filtering, or Image reconstruction for a camera. There will be an SDK developed enabling skilled users to write their own interface managers.

In order to make this process efficient and reliable, a large amount of software development work goes into providing efficient data-transfers between the blocks, careful management of process scheduling and interrupt management. On top of the modules that do actual work, there must be higher-level software modules for control, configuration, diagnostic, and other management tasks.

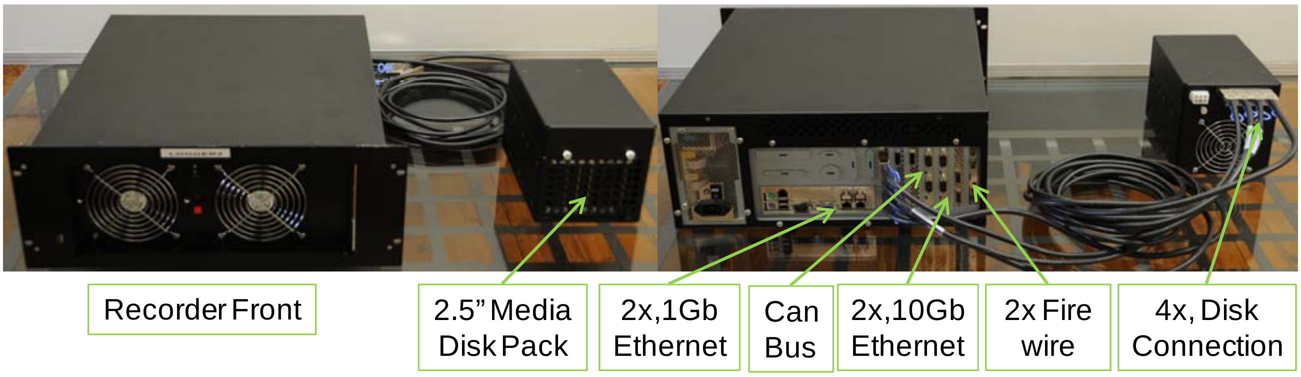

Figure 4: The recorder front and 2.5" media 16-disk pack (left); the back of the recorder with IO and disk pack connections (right).

You can compare the XCube system to a highly tuned race car, where all components, both hardware and software, are chosen or written with the single goal of high-performance real-time data recording. This means that, although this system is "Software Defined", tuning work needs to be done for each platform and combination of hardware used to ensure proper performance.

1.2. XCube Recorder Development Status

Currently, the recorder is used for data capture in an automotive application. In that application the XCube recorder simultaneously accepts, and keeps track of, 4 CAN bus streams (target data from 3 radar, 3 cameras, vehicle sensors and more), a 1394B fire-wire camera (raw frames), a high-speed USB (radar, raw-samples), 6 GigE-Vision cameras, a generic 1Gbit Ethernet input stream, and command and control information from an operator console. In all a continuous data stream of 800MB/s (6.4Gbps) is presented to the recorder.

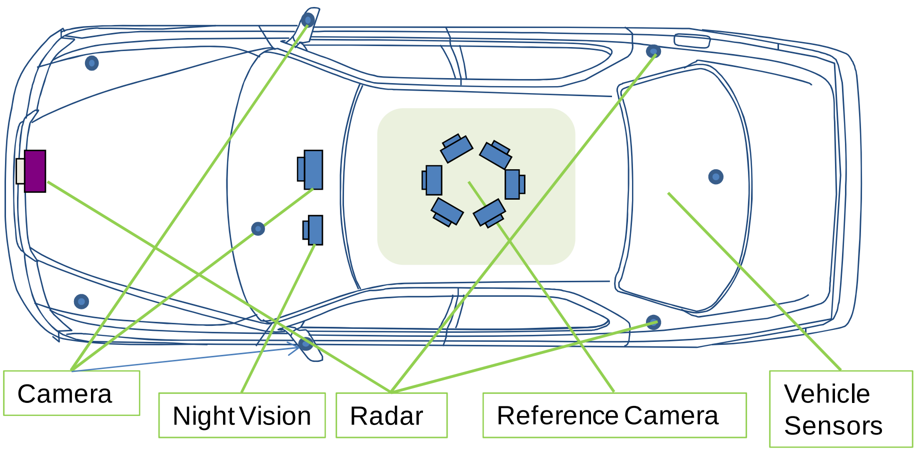

Figure 5: Sensor setup for the European test vehicles. Vehicle system and Reference systems produce over 800MB/s continuous data.

The Ethernets operate in a mix of RAW mode and Image capture mode (UDP transport and image frame reconstruction on the fly). The loggers have been operating in the field since October 2010, and currently there are 3 vehicles equipped. In this project new functionality for diagnostics, configuration management as well as additional capabilities such as a night vision system and additional diagnostics channels from the application computers are being developed.

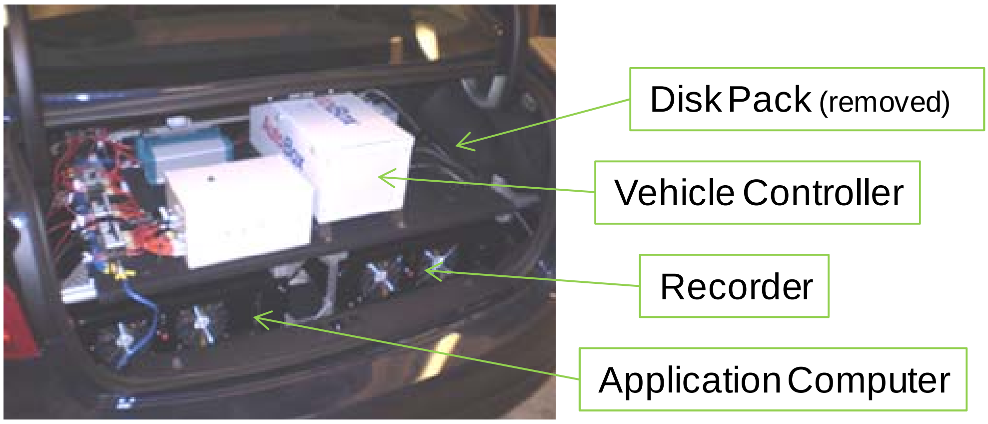

Figure 6: Recorder and Application Computer mounted in trunk of Volvo C60.

The Disk subsystem in the field works well up to 1100MB/s (8.8Gbps) of continuous capture. We have recently changed disk controllers and are rolling out a more reliable and faster version that runs reliably up to 1710MB/s (13.6Gbps) continuous recording. We are verifying performance with dual controllers and expect to see up to 3400MB/s (27Gbps) throughput to disk.

The chassis, processors, cooling and infrastructure components are debugged and working as expected. We have so far not experienced any problems with the chassis or disks operating in the cars, but we only have a few hundred hours of operation data so far.

A data-tap mechanism is under development that will enable users to get a small trickle stream of data from the streams being captured. One specific function is to grab a few video frames from the camera streams to verify exposure and focus while the system is running. This capability might be useful to extract phase-cal information from VLBI data.

Based on the field experience, the XCube team is working on a training class, and is responding to feedback from users by adjusting operational procedures to make the systems more user-friendly. The changes have mostly to do with providing better feedback to the operator on system status plus better diagnostics of the disc subsystems.

1.3. XCube Recorder to Mark 6 Specifics discussion

The XCube logger accepts multiple simultaneous 10GigE data streams for capture to a (removable and replaceable) set of disks. The capture chain from 10GigE to disk is currently working reliably with streams up to 7Gbit effective data rate per Ethernet channel.

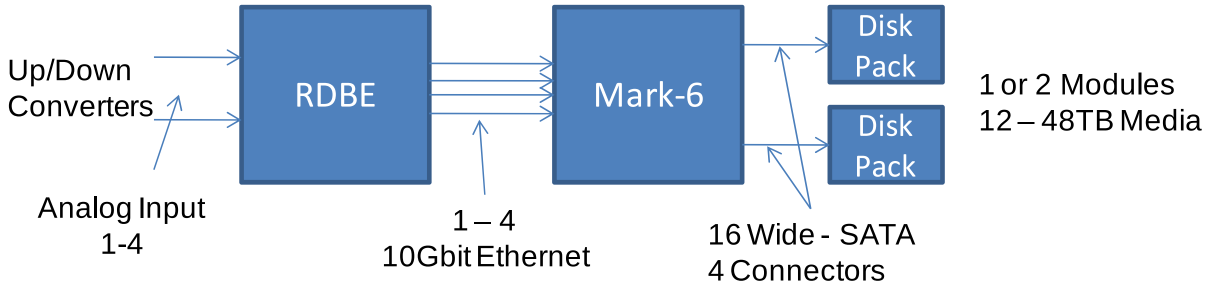

Figure 7: Dataflow from the analog frontend through the RDBE via 10GigE through the Mark 6 to removable disk packs via 2 sets of 16 wide SATA/SAS channels.

We believe acceptance of up to 4 parallel 10GigE channels is possible, though this configuration has not yet been tested. This capability will allow the simultaneous capture of data from multiple RDBEs connected to a single XCube system, up to an aggregate data rate of at least 12Gbps with the current capabilities.

We have done an initial investigation of the Mark 5C command set and the match to the XCube system. Of the 32 Mark 5C commands, most require very little work to transfer to the equivalent XCube Recorder command. Table 1 details the current compatibility status.

bank_info?	Get bank information (query only)	Ok, but needs extension
bank_set	Select active bank for recording	Support can be added
data_check?	Check data starting at position of start-scan pointer (query only)	Support can be added
dir_info?	Get directory information (query only)	Ok
disk_model?	Get disk model numbers (query only)	Ok
disk_serial?	Get disk serial numbers (query only)	Ok
disk_size?	Get disk sizes (query only)	Ok
disk_state	Set/get Disk Module Status (DMS): last significant disk operation	Support can be added
disk_state_mask	Set mask to enable changes in DMS	Not applicable
DTS_id?	Get system information (query only)	Ok
error?	Get error number/message (query only)	Ok
fill_pattern	Set the fill pattern for missing packets in PSN Mode 1 or 2	Not supported
get_stats?	Get disk-performance statistics (query only)	Ok
MAC_list	Acceptable MAC source addresses that will be filtered ()	Ok
mode	Set data recording mode	Ok, similar in setup
net_protocol	Set network data-transfer protocol	Ok (raw/filtered)
OS_rev?	Get details of operating system (query only)	Ok
packet	Set the packet acceptance criteria	Similar in setup
personality	Set the personality of the application	XCube
pointers?	Get current byte values of pointers (query only)	Not Applicable
protect	Set/remove erase protection for active module	Ok

record	Turn recording on | off; assign scan label	Correspond to Start/Stop
recover	Recover record pointer which was reset abnormally during recording	Not Applicable
reset	Reset Mark 5 StreamStor card (command only)	Ok (resets recorder)
rtime?	Get remaining record time on current disk set (query only)	Ok
scan_check?	Check data between start-scan stop-scan pointers (query only)	Not Applicable
scan_set	Set start-scan and stop-scan pointers	Not Applicable
SS_ifconfig	Set the StreamStor 10G daughter card interface (not supported)	Ok, configures actual 10GbE
SS_rev?	Get StreamStor firmware/software revision levels (query only)	Not applicable
start_stats	Start gathering disk-performance statistics	Ok
status?	Get system status (query only)	Ok
VSN	Read/write extended-VSN	Similar system command exists

Table 1: Mark 5C Command set with notes on potential implementation on the XCube recorder. “Not Applicable” means that the command is specific to the “StreamStor” card and not needed. The items marked with Ok, have an equivalent function available in the XCube recorder (needs translation code to be compatible with Mark5C).

The Mark 5C control code (not specific to StreamStor) has been developed in python and is called DRS*. When discussing the DRS implementation with Chester Ruszczyk the code will easily run on the XCube Recorder.* The implementation of the DRS allows addition of an XCube personality without re-architecting the code. With the well designed DRS code, it should not be more than a few weeks of effort to port and integrate the code.

The XCube disk module for sixteen 2.5” media is ready to go and can be used directly to at least 8Gbps. A disk module for 3.5” media is in the XCube development plan as a standard product (legacy SATA media can be used in new boxes). Additionally, an initial concept and feasibility study has been done to develop an inexpensive kit that will allow existing Mark 5 modules to be modified for use with the XCube system, though re-use of Mark 5 ‘PATA’ modules will require installation of SATA disks.

1.4. Relation to other work

As indicated earlier, the XCube system was originally developed for use in an automotive application as part of a larger suite of capabilities and tools, some of which might be of interest to the VLBI community.

Data may be simultaneously captured from a number of different sources (e.g. Ethernet data sources, USB devices, high-resolution cameras, etc.). Though the various data streams will be mixed on the recording disk array, the XCube system automatically keeps track of the source, along with corresponding high-resolution time tags, so that each frame of may be uniquely identified on playback. Such a capability might be useful in the VLBI application for recording auxiliary data streams (for example: system monitoring, auxiliary instrument data capture, video streams), though it is understood that these sorts of tasks are normally done with separate subsystems in the VLBI application.

For the automotive application, playback data are captured and stored in a distributed simulation and storage cluster (DSSC). The DSSC provides facilities to store multiple 100's of PByte of data as well as to run "simulations" or algorithms on that data. DSSC also provides connections to Mathworks tools as well as intelligent-vehicle simulators such as "Pre-Scan" from TNO (prescan,TNO).

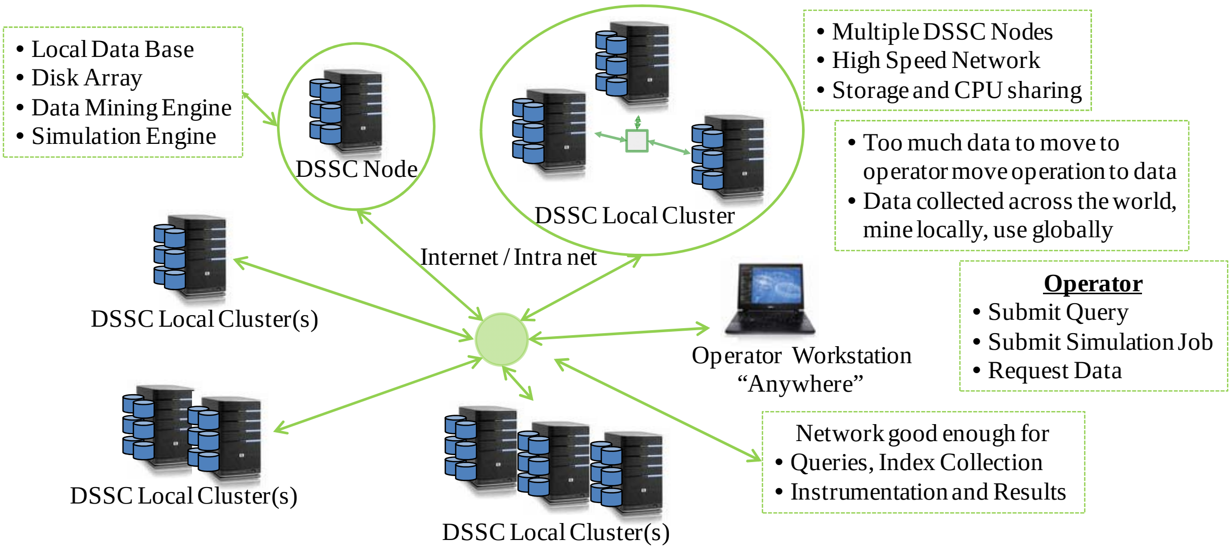

Figure 8: Conceptual view of XCube DSSC. Data are distributed across a large number of servers that provide users anywhere access to data and simulation tools.

In order to find specific data of interest, DSSC implements a search engine "Google for measurement data" and a data mining harness for "tagging" the stored measurement data. This facility consists of a data-mining harness and a facility to automatically run algorithms searching through and updating the search databases with information found.

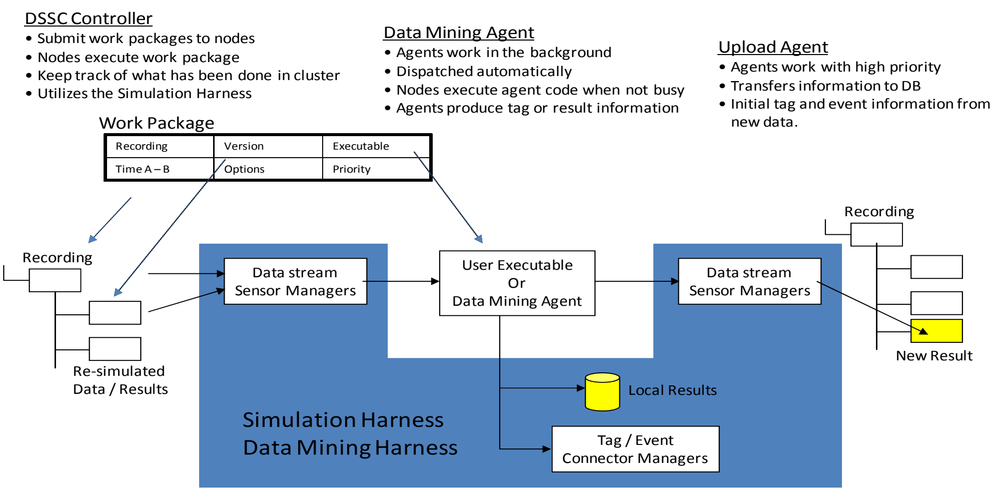

Figure 9: Data mining and simulation harness enables automatic search through data in the system.

The applications that use these types of systems are growing, and specifically applications in automotive active safety, unmanned vehicles, robotics, and signal intelligence.

The XCube system and tools are expected to benefit from the continued development that supports multiple users and applications. In turn, customers benefit from development costs being spread across a large base of users.

2. GENERAL SPECIFICATIONS

The Recorder and Disk Packs are designed to be operated in the "field" meaning that they in general are "rugged". The designs are made so that they can be operated in cars, or in environments that experience shock and vibrations as well as large variations in temperature. They are however, not designed to be operated in wet or outdoors environments. In general these are specifications apply to the recorder and disk packs, with additional reservation for individual disk models that will conform to the manufacturers' specification.

1. Recorder platform
    a. Linux operating system, available for user applications
    b. 12GByte standard, configurable main memory up to 128GByte (motherboard dependent)
    c. 1 or 2 available 1Gb Ethernet interfaces for control and data capture
    d. 4 USB ports for control or capture
    e. VGA port for external display.
2. Recording capability
    a. 4Gbps continuous recording from a single 10Gb Ethernet
    b. 8Gbps continuous recording from dual 10Gb Ethernet
        i. Requires disk pack capable of sustaining 8Gbps recording (typically 16 disks, depending on disk performance characteristics)
    c. Scalable to 12Gbps continuous recording (16Gbps possible, with four 10GigE)
        i. Roughly 7Gbps is achievable on a single 10GigE data line
        ii. Projections indicate that a sustained rate of 32Gbit/s (4GByte/s) should be possible within the next couple of years.
3. Recording Interfaces
    a. Input Dual Channel 10Gb Ethernet (CX4)
        i. Raw mode or IP filtered mode supported
        ii. Optional Quad Channel 10Gb Ethernet (CX4) matching RDBE output capability
        iii. Optical 10GigE data links can be used if needed.
    b. Output 16 Channel SAS/SATA to disk modules
        i. External connections to disk module 4 SATA cables (each supporting 4 disks) plus a separate power cable.
        ii. Optional additional disk controllers.
        iii. Additional disks can be added (in units of four) for increased performance.
4. Output Interface
    a. Dual or Quad Channel 10Gb Ethernet (Cx4)
        i. Optional Infiniband or other network.
    b. Disk Packs can be NFS mounted through the recorder (any IP interface)

5. Operator Interface(s)
    a. Command line operator interface
    b. UI based configuration interface.
    c. VSI-S Compatible Machine to machine interface for connecting to VLBI infrastructure. This interface conforms to Mk-5C command set as much as possible.
6. Recorder Mechanical
    a. 4U, 7"(H), 19"(W), 16"(D) rugged industrial chassis.
    b. Operational temperature: 0 – 45ºC
    c. Altitude: -100 – 3050m (5300m with reduced temperature range 0-30ºC)
    d. Humidity: 5% – 95% non-condensing
    e. Shock and Vibration (TBD)
7. Disk Pack Mechanical and environmental
    a. Self contained with power and cooling
    b. Operating ambient temperature: 0-45ºC
    c. Altitude: -50m – 3,050m
    d. Humidity: 8% – 95% non-condensing
    e. Vibration defined by disk manufacturer specification, media dependent
    f. 2.5" Media Data Pack
        i. Capacity 12TB or 16TB
           (Seagate Momentus 750G, Western Digital WD10TPVT, others on request)
        ii. Altitude: -300m – 3,050m
        iii. Size: 6"(H) x 6"(W) x 9"(D)
        iv. Weight: 2.5 – 3.5kg depending on media
        v. Power: External supplied by recorder (5V, 16A, 12V, 1A)
        vi. Power consumption: ~ 45W (read/write)
    g. 3.5" Media Data Pack (16 individual disks)
        i. Capacity (16TB, 32TB, (48TB*))
           *Requires Seagate 3TByte drive qualified when available
        ii. Altitude: -60m – 3050m
        iii. Size 3U (H) by 16" (D) by ½ Rack Width
        iv. Weight: 12-14kg
        v. Power: External DC, (12V, 16A, 5V, 8A)
        vi. Power Consumption: ~160W (read/write)
    h. Optional 8 disk Data Pack
        i. Same specification as 16-pack
        ii. Size 3U(h) by 8" by ½ Rack(w)
        iii. Power Consumption: ~80W (read/write)
    i. A Mark 5 disk module converted for use with the XCube system will have the same specification as the XCube 8-disk data pack. The Mark 5 disk module does not have built-in cooling so it will require external cooling infrastructure from a Mark 5 or an external enclosure.

2.1. High altitude operation

The XCube recorder will be able to operate up to 5300m to accommodate operational requirements at the highest-altitude radio observatories, including the Atacama, Chile, observatory sites (ALMA). The table below summarizes the highest altitude (from Wikipedia) observatory sites in the world.

Name	Elevation	Type	Pressure
University of Tokyo Atacama Observatory	5640 m (18,500 ft	Optical	52kPa / 51%
Chacaltaya Astrophysical Observatory	5230 m (17,160 ft	Cosmic Ray	54kPa / 54%
Atacama Large Millimeter Array(ALMA)	5190 m (17,030 ft)	Microwave	54kPa / 54%
APEX	5104 m (16,745 ft)	Sub-millimeter	55kPa / 55%
Large Millimeter Telescope(LMT)	4580 m (15,030 ft)	Microwave	59kPa / 58%
Mauna Kea Observatory	4190 m (13,750 ft	Optical/mm	62kPa / 61%
Very Long Baseline Array (VLBA), Mauna Kea Site	3730 m (12,240 ft)	Radio Telescope	66kPa / 66%

Though the ALMA telescope is above 5000m, and the XCube system could be used at that altitude data from the telescope will normally be transmitted to a control facility at ~3000m altitude for recording. The altitudes of APEX, LMT and Mauna Kea present a challenge for rotational-disk recording without pressurization, though solid-state disks are not affected by altitude.

Note: The recorder is cooled by forced air and consequently will be affected by a decrease in air pressure. The cooling fans are less effective resulting in lower cooling capacity. Similarly lower mass-air-flow through heat sinks decreases heat transfer. At 5300m both fans and heat sinks lose about 45% of efficiency compared to sea level. XCube compensates for this loss by decreasing the allowed ambient operating temperature to 0- 35ºC.

Disk operation at high altitude is mostly affected by partial loss of the air cushion that disks use to keep heads flying just above the magnetic surface. Too much loss of this air cushion can cause the disk heads to ‘crash’ onto the magnetic surface and damage the disk. Cooling can also be somewhat of an issue, but it is not the primary factor in this environment.

For high-altitude operation XCube proposes to develop a small pressure container that will house disk-packs used at sites above 10,000ft (12,000ft with selected disks). The (low) pressure container will house the standard disk packs and have the cables connect through a bulkhead connector providing the connections. Cooling will be provided by conductive heat transfer through the pressure housing (cooling fins on the outside of the box) and internal heat transfer is provided by the air-movers contained in the disk packs. The boxes will be designed to be hermetically sealed, but leakage can always occur (much less expensive box if small leaks are ok), therefore the pressure containers include a small micro controller and a miniature compressor to monitor air-pressure and temperature and to maintain appropriate operating pressure inside the container.

3. REQUIREMENTS FOR VLBI OPERATION

This section lists a set of requirements for VLBI operation that are a result of discussion with Haystack and NRAO personnel over the past several months.

3.1. Operational Concept

There are two modes of operation that are important from the recorder (logger) point of view, recording and playback. The first is in the field where it is set to record data. Here the recorder is typically connected to a digital back end (DBE) system, a remote control system, and one or more disk packs. The DBE connects to the Recorder over one to four, 10Gbit Ethernet links, each providing a UDP packet stream to the recorder. In addition to the 10Gbit Ethernet links the DBE and the recorder communicate over a 1Gbit Ethernet for control and status information.

The data received in the recorder are written to removable Disc Packs, each disk pack containing 8 or 16 physical discs. There can be one or more disk packs attached to a recorder at any time. The recorder can selectively (by configuration) store data across any number of disks available. When a disk pack is full it is detached from the recorder and sent to another location for correlation.

In the correlation step the Recorder is used in playback mode. In this mode the recorded data contained on the disk-packs are either sent out over the 10Gbit Ethernet ports in the same way as they were received, or the disks are mounted (NFS or Local) and files are directly read by the correlation software. If the files are kept in the parallel file systems as they were recorded, replay speeds can be fast, potentially faster than when they were recorded.

3.2. Detailed Requirements

The XCube recorder shall meet the following requirements for VLBI:

Data-recording requirements include:

- Ethernet data-transmission protocol to recorder will normally be UDP.
- Recording to a selected subset of connected disks.
- Continuous data capture and recording of a least 4Gbps from one 10GigE data source; recording on 8 disks (2.5”, 3.5” or SSD); burst packet rates may exceed 4Gbps for very short periods of time, but average will be no more than 4Gbps.
- Continuous data capture and recording of a least 8Gbps from two 10GigE data sources; recording on 16 disks (2.5”, 3.5” or SSD); burst packet rates on each 10GigE data link may occasionally exceed 4Gbps.
- No data loss is expected in normal operation with good disks.
- Slow disks shall not affect the recording performance provided the aggregate data-rate capacity of all disks (including the slow ones) remains above the target recording data rate.
- If a disk fails, full-rate recording will continue (with minimal interruption) to the remaining disks, if possible, or at a lower rate (with some lost data) if the aggregate recording data-rate capacity of the remaining good disks in insufficient for full-rate recording.
- The user may specify that either full Ethernet packets, or only the data-payload, will be recorded.

Data-playback requirements include:

- Playback of data shall be in the same order that it was received.
- Sustained playback at the rates approaching disk-playback data rates where possible.
- No data loss is expected in normal operation with good disks.
- In the event of bad disk(s), the system will recover all possible data from other disks. If a disk became bad during recording, it is recognized that all data recorded on that disk prior to failure will be lost, but all remaining recoverable data should be available to the user.
- In the event of a slow disk(s) during playback, the user will normally accept the performance degradation, but should be able to declare the disk bad and continue playback (but without data from the bad disk) at higher speed.
- In some cases recorded data will be out-of-order, normally causing it to be out-of-order on playback as well. The current playback concept is that such out-of-order packets will be managed by the correlator system and will not be the responsibility of the playback system (though this can be further discussed).

Operational requirements:

- Monitor and control through a standard port (IP-socket) as if connected to local terminal.
- Record in either Promiscuous mode or from specified IP/Ethernet addresses
- Start/stop recording
- Query available disk space
- Query list of recorded files/data
- Maintenance of a 'permanent' unique Volume Serial Number (VSN) written to disk that is assigned to each disk module for identification purpose. The VSN must survive all normal recording, playback and delete operations and should be readable if <u>any</u> disk in the module is functioning normally.
- During recording or playback operation, inactive modules may be removed and replaced without affecting the operation of active disk modules. Furthermore, simple operations may be performed on inactive modules (such as mounting, dismounting, reading module/disk information, etc) without affecting active disk modules.
    - Any operations affecting disks in the system while a recording is active may result in data loss.
- Meaningful disk-performance statistics should be kept during recording/playback operations and made available to the user in order to identify underperforming or failed disks. During active operation, the user should be notified when any individual disk performance falls below a user-set threshold.

Environmental requirements:

- Normal operation at altitudes 0-10000ft; high altitude operation (up to 18000ft) will be possible with specially designed pressure enclosure for disk modules
- Normal operation in ambient temperatures 0ºC - 40ºC

Other requirements (at least strong desires)

- Ability to read disks (recorded on the XCube system) on a suitable standard Linux workstation; may require a compatible hardware disk controller.

- Option to write a "file-reader" module that can run on Linux operating system and provide a socket- based interface to read data from the disks. This software should be able to take commands to read packets, write packets, open and close file-sets.

4. QUESTIONS AND ANSWERS

During meetings and conversations over the past months there has been a multitude of questions that I tried to collect and answer here:

1. What is the reason for the XCube/Mark 6 recorder in the first place?
   a. XCube is developing a toolset for rapid development of complex multi-sensor systems applications.
      i. This includes collection of "real world" data from platforms (at GByte/s data rate from multiple sensors and platforms concurrently).
      ii. The second step is to store reliably store hundreds of petabytes of data
      iii. Third step is to find interesting parts of the data. XCube is developing a "Google" for measurement data containing a distributed search engine and data mining agents to manage immense amounts of real world data.
      iv. Fourth step is to be able to process data and run algorithms on that data. XCube develops a platform with integrated simulators and MathWorks tools for this purpose.
   b. The background of the recorder comes from the above mentioned work, as well as work I have done with the RDBE. It seemed that a lot of effort has been put into making the Mark 5 recorders work with the new system requirements. It also seemed that we were using a system that although good in its design was driven past its design limits and that there were easier and less expensive ways due to technology evolution to provide recording capability for the future.
   c. XCube as a company is working with similar recording (data capture and logging) for the automotive industry and the signal intelligence community. We think that the logger developed for those markets using COTS components could easily be adapted for the VLBI community
2. Is the logger software open source?
   a. No, the core of the recorder is proprietary software.
   b. The system is designed to allow a skilled user to write their own modules to capture data from different sources and connect the data-streams to the system. Great care needs to be taken when writing these modules since they directly affect performance.
      i. XCube will provide training and other services to help develop such interface managers if needed.
   c. There is a specification for users to write their own interface managers if needed
   d. There is a specification for how to control and configure the system from an external program. The system is setup to be remotely controlled by another application.
3. Can I run other software on the Recorder?
   a. Yes, the recorder is based on a Linux operating system and the system is available to run other software.
      i. There is an intricate interplay between the data recording and other applications running on the machine, the interface managers are running at

high priority, and they are locked into memory, but there are still possibilities to negatively affect capture performance from other software programs running concurrently with data capture.

   b. In fact we expect to run most of the current Mark 5 code and utilities on the XCube recorder. It should also be possible to run the software correlator code directly on the recorder if desired.

4. Is it possible to license the software?
   a. There are currently no plans to license the software.
   b. There is an intricate interplay between the hardware and software. The hardware although off-the-shelf needs to be carefully selected and the software tuned to fit the particular hardware configuration. XCube selects, qualifies, tests and guarantees the function of the systems.
5. What Linux distribution is supported?
   a. The logger is an “embedded system” with the capability to run other software. Currently the recorder uses Ubuntu 10.x LTS distribution.
6. Will the XCube system support ‘burst mode’ recording?
   a. The architecture of the XCube system inherently supports ‘burst mode’ recording up to the full rate that data can be written to memory. A more efficient way to use ‘burst mode’ with the XCube system is to start disk recording at the same time as the start of data-capture, then continue recording until the RAM memory buffer fills up; this will maximize the duty-cycle of the data capture at speeds greater than the fully sustainable capture rate.
   b. Maximum “burst-mode” capture speed depends on multiple factors, such as interface boards used, limits on the PCI-express bus, available memory bandwidth etc.
7. Does the XCube system support a ‘permanent’ Volume Serial Number (VSN) written to disk that is normally assigned to each disk module for bookkeeping purposes and is not affected by normal recording/playback operation?
   a. The XCube system does not currently support a VSN, but we are planning to make the necessary modifications to support a VSN as required.
   b. The XCube system stores disk information in a set of “permanent” files on the disks in a disk pack. These files contain information about the disks usage, performance and identity. These files are part of the system and will not be erased.
8. What disk media are supported?
   a. In theory (and so far in practice) any SATA-II or higher drive should work with the system. In reality there are some differences in interpretations of the SATA specifications between different vendors, as well as differences in interface (electrical and low level) implementations that make some drive and controller combinations work better than others. We have found that some controllers (Adaptec and 3ware products) have difficulty working reliably with certain versions of Western Digital and Seagate disks, resulting in performance degradation.
   b. There are performance differences between disks, and interfaces to those disks. In general most disks will work. Every disk model is different and there are large performance differences between them.
   c. In order to be able to (with reasonable effort) qualify hardware, software, and be able to support the systems over time, XCube will work with customers to qualify a selected set of drives that are suitable for their needs.

   d. XCube has a policy to use commercial volume production disks to optimize for price/performance or price/volume, or the specific needs of our customer.
   e. Today there are Seagate Barracuda 2TB, Western Digital Caviar 1.5TB, Seagate Momentus 750GB, and Western Digital Caviar blue 1TB, disks in the field. As soon as it is released we will support Seagate 3TB drives (Q1/Q2 2011)

9. What should we do with our old media?
   a. There are a couple of options for legacy media. For relatively new disks XCube will work together with the customer to qualify (electrical and performance) disks, then there are two options
      i. Full support, which means that the disks will go into the regression tests for future hardware (controllers, cables, disk packs, etc.). These disks will have full support from XCube.
      ii. Support for current systems, which means that they will be certified to run with the current generation systems, but for new generations of the hardware the disks will have to be recertified. This is the same procedure as older disks when they will be phased out over time.
   b. XCube will not support older non-SATA/SAS disks, they are both too slow and we do not have interfaces capable of providing PATA externally. In theory it would be possible to provide PATA to SATA bridges and use older disks, but we would highly discourage such an option.
10. Can the XCube system be used with existing Mark 5 disk modules?
   a. Obviously, existing Mark 5 modules cannot be directly used with the XCube system. However, a plan is being developed to provide a kit, consisting of a new module backplane and a new front panel that will allow both Mark 5 PATA and SATA modules to be modified to work with the XCube system. Of course, all PATA disks will need to be replaced with SATA disks. The new front panel will have two 4-disk SATA connectors for cables from the XCube box, as well as one connector for a power cable. Since the Mark 5 modules have no built-in cooling, external cooling must be provided, either by inserting the modules into an existing Mark 5 chassis (there is no electrical connection to the Mark 5) or with an enclosure that provides suitable cooling (basically just requires a fan tray under the disk modules).
   b. An adapter will be made available to fit the 2.5” and 3.5” XCube-developed disk modules into the Mark 5 disk shipping containers. New shipping containers specially designed for XCube-developed modules will also be available
11. What about disk performance? What should I expect from my disks?
   a. There are many factors that affect the performance of disks and disk packs. Rotational speed, bit-density on the disk, seek time, cache structure, interface speed and interface capability (how writes and reads are managed on the disks) and mechanical components on the drive, to name a few.
   b. For rotational media a 5400 rpm drive in general will sustain speeds up to 70GByte/s, while individual drive parameters may be as much as 10% off in a real test. Manufacturer specifications are generally somewhat better than performance achievable in the real world.
   c. For 7200 rpm drives, the sustained speed is about 105MB/s per drive, with similar spread between individual drives and manufacturers.
   d. For solid state drives (SSD) the sustained read and write rates are different and depend on the type and construction of the drive as well as the interface to the disks.

It is possible to sustain well over 200MB/s per drive. Today the cost may be prohibitive for situations where the ultimate performance is not needed.

   e. For Disk Packs the performance depends on drive configurations and the number of drives that run in parallel. The drive scheduling software and hardware interfaces allow the disks performance to scale with the number of disks used in the disk packs with minimal degradation in performance.
      i. Eight, 7200rpm drives sustain 800MB/s, 16 drives sustain 1600MB/s
      ii. Eight, 5400rpm drives sustain 540MB/s, 16 drives sustain 1080MB/s
      iii. Eight, 200MB/s SSD sustain 1600MB/s, 16 drives sustain 3200MB/s (dual controllers are needed due to PCIe limitation of 2.5GB/s)
   f. Note: With careful management, some disks may perform a little better than the above mentioned numbers. We have successfully run sixteen 7200rpm drives in the lab at 1711MB/s (13.6Gbps) continuous recording of entire disk packs. We expect that we can get even more performance out of specific disks but would not recommend that for production systems yet.

12. The Mark 5 systems go to considerable length to minimize disk-head motion during recording in order to eliminate as much rotational and head-seek latency as possible. How does the XCube system manage/minimize these issues?
   a. XCube uses a different scheme compared to the Mark 5 recorder that doesn’t rely on tight low level control of the disks. This scheme allows the disks to work at their average rate instead of peak rate.
   b. XCube does minimize the head movements as well as order writes, by ordering writes and pre-allocating files, although it is not necessary for 8Gbps or less operation.
13. It is not that hard to write code to do this, why should I not just do it myself?
   a. It is possible for a good software team with proper guidance and knowledge of XCube approach to replicate the function of each part of the system. A reasonable investment estimate of about 2 man years of coding and testing of software and hardware combinations is reasonable. In addition to this effort, the systems will have to be supported, new hardware and media will have to be qualified over time.
      i. Even though the software is done and working there is still a significant investment in making it a supportable product.
      ii. In order to properly support the systems there is a significant investment in man hours and infrastructure that needs to be sustained for the life of the product.
   b. XCube has to date invested over 2 man years in the software and 1 man year in qualification and testing of hardware and software components, as well as custom design of chassis and disk-packs to be suitable in an industrial environment
   c. XCube has multiple customers for the systems and will continue to develop and support the Recorders over the long term.
14. Is the XCube recorder COTS (commercial off the shelf)
   a. The hardware is constructed out of commercial off the shelf components. There are no “custom” components developed only for the recorder itself. All components are bought on the open market and have commercial use.
   b. The components used are carefully selected and qualified to work with each other and other components in the system to provide reliability and performance.

   c. As with all complex systems not just "any" component will work, great care will have to be taken to compose the right set to form the system
15. COTS doesn't have a long lifespan typically18 months or less, how do you ensure that systems work over time, and that we still get the benefits of COTS
   a. The XCube recorder is defined by software and is based on a Linux operating system platform. However, the software and hardware is highly dependent on each other. In order to get predictable performance XCube carefully selects components with "long" commercial life-span when possible.
   b. We manage hardware obsolescence by having a limited number of configurations in the field, and qualify new hardware before the old boards become obsolete. In most cases we can keep a particular set of components for about 3 years (long term support). Each year we re-evaluate the choice of base platform and re-qualify a new board to be the production version. We also use rigorous regression tests to ensure that the new and old systems work the same (new systems will in most cases be more capable, but are backwards compatible).
16. XCube claims that the recorder is software defined and standards based; does this mean that it can run on anything?
   a. Our data capture mechanism is defined by software, and it resides on a standard Linux based software platform. However, in order to achieve the performance necessary and determinism the software is tuned to the hardware platform. In other words, all pieces need to fit together for the system to perform well.
   b. The recorder is not a standalone software package it is a complete system with hardware and software components that are designed to work together. The components are selected and qualified so that they will provide reliable performance over time.
17. Mechanical disks are inherently unreliable, what measures are taken to detect and prevent failures?
   a. The disks are kept running as cool as possible with plenty of cooling to keep heat related problems to a minimum.
   b. We will use the SMART data on disks that are equipped with that capability to detect early indications of failure. Research by Google has shown that this data is not very reliable, but if monitored over time increases in search time and seek errors can predict drive failure in some cases.
      i. Smart data is currently available through an external utility
      ii. Further diagnostic software should be written to better take advantage of these features
   c. Each disk in a disk-pack is individually monitored by the software and a "permanent" log is written to the disk.
   d. The disks are written in parallel, if a disk fails, only the data on that disk are lost, not the entire capture. The system will continue to capture data for the reminder of the run, excluding the failed disk from the storage array.
18. Can we run multiple disk packs on a single machine and switch between them?
   a. Yes, the recorder as a default has connections to 16 disks, these disks can be configured as 2 sets of 8 if needed. The recorder can be outfitted with a second controller and then it is possible to run disks sets as 1 by 32, 2 by 16, or 4 by 8.

i. In fact any combination is possible (1, 3, 7 and so on) but for human management to be reasonable we should limit to the above mentioned disk sets.

b. In the current implementation the recorder will not switch disks in the middle of a capture, and disks can't be replaced while a capture is running.

19. What about multiple insertions and wear on connectors?
    a. The recorder uses commercial external cables between itself and the disk packs; each cable supports four disks in a disk pack. By leaving the cables attached to the recorder and using the connectors on the disk-packs, the total number of insertions seen at disk-packs and recorders is minimized. Replacement cables are available off-the-shelf for less than $70 each.
    b. The Amphenol SFF-8088 connectors used have a nominal insertion durability of 500 cycles, and a guaranteed minimum of 250. At many typical VLBI stations, a disk-pack connection cycle occurs once every couple of days on average, so that a cable at these stations should last at least a year or somewhat longer; some extra cable life might be obtained by switching cables end-for-end somewhat before 250 cycles on the heavy-use end; some heavy-use stations will require more frequent cable renewal. Correlators will generally need to replace cables more often due to the rapid flow of disk modules through the correlator system. Disk-pack connectors, on the other hand, typically experience no more than 25-30 connection cycles a year and should not be a concern during the normal life of a disk pack.
    c. Note on cable/connector usage: As a service procedure, when changing cables or connectors on records or disk packs, a note should added to the usage log in the recorder or disk pack, respectively, in order to properly track ongoing usage.

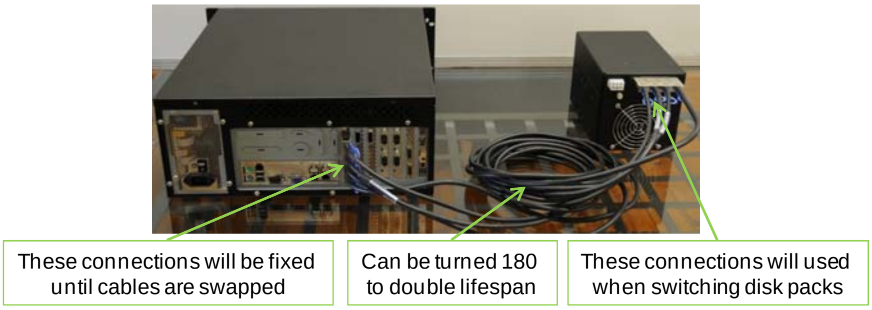

Figure 10: cables connected to recorder will not be frequently inserted, while the ones connected to disk packs will. Cables can be rotated 180 degrees to double lifespan.

20. How do I know when it is time to replace the cables?
    a. Cables are inherently passive components this means that there needs to be an operational procedure to manage cables.
    b. When cables start to go bad, errors will be detected during recording, but cables should have been replaced much before that point is reached.
21. What kind of file system is used? Is the file system proprietary?
    a. The disk packs use an EXT4 file system, which is a Linux standard. The disks are easily readable from any system that can read EXT4.
    b. The disk packs do not use any special formatting so an appropriately configured workstation can read the files off the disks.

   c. The data are stored in an XCube log file format that has an open specification. For the Mark 6 recorder it will be possible to store data directly in VDIF format if that is more appropriate.
   d. The output file format is easily modified to be suitable for any specific application. The recorder uses a parallel writing scheme to the disks which requires the files to be self-describing in terms of packet order (there needs to be a way to order packets, by sequence numbers or timestamps).

22. How are “slow” disks managed?
    a. XCube has developed a disk scheduler that balances workload on the disks based on each individual disks relative performance in the disk pack. The scheduler will also shut down faulty disks, as well as manage retry and timeout occurrences.
    b. The reason for the custom scheduler is to get around problems such as slow or faulty disks that severely affect performance of traditional RAID controllers and make them unsuitable for high performance real-time data capture.
23. What if XCube goes out of business
    a. All hardware and software is under version control, and in the unlikely event that XCube no longer can support the systems all necessary software and hardware design will be released to our customers.
    b. XCube has been in business since 2003 and is growing its customer base with a healthy growth potential.
24. Is this a custom design for VLBI only?
    a. The recorder is developed together with the VLBI community. However, the recorder is part of the generic high speed logger developed by XCube for other purposes. This means that development will continue and benefit the VLBI community.
    b. The components of the recorder are readily available on the open market. Sheet metal and packaging are modified or developed to fit the application requirements.
25. What is the roadmap for the recorder?
    a. The software and hardware will be continuously developed for the VLBI and other markets driven by customer demands.
       i. There is a miniaturized version requested from the automotive community.
       ii. There is a high altitude version requested from the signal intelligence community.
       iii. There are customer requests for increased recording speed and increased diagnostic utilities which will be developed.
    b. The current plan for recording capability is to immediately release a 4Gbps version that will follow by an 8Gbps software upgrade. The same system will be able to support 12Gbps and 16Gbps recording (with additional hardware)
    c. The Data packs will include new 3TB 3.5” drives and there is a high-altitude data pack planned.
26. How is replay managed?
    a. There are many ways that files can be retrieved from the recorder.
       i. Replay through the Ethernet interfaces.
       ii. Mount the disks over Infiniband or Ethernet (single or dual 10Gbit) or any other network.
    b. Correlator code can be run on the Recorder, and data can be retrieved from other recorders to do correlation “in-place”

i. The recorder has a powerful server class computer with plenty of compute and memory capabilities to run other applications too.

ii. It might be an interesting point of investigation to look at a parallel version of correlation across multiple recorder nodes.

c. The playback operation uses a parallel file system scheme that require the files to be re-assembled from the disks (in chunks of 32MByte or larger), and that as a side effect enables playback at GByte(s)/s speeds.

27. How much RAM is supported? Can I add more RAM?

a. The amount of supported RAM depends on the particular system board used. Currently the boards support up to 128GByte RAM.

b. The base configuration comes with 12GByte, if needed the machines can be equipped with additional RAM.

28. How do we control the recorder?

a. The recorder has a recorder server process as the main interface to the system. This process listens to an IP-socket and understands VSI-like commands from external programs.

b. In the integration project we will port the Markk-5 code so that it will run on the recorder and interface to the underlying XCube control code.

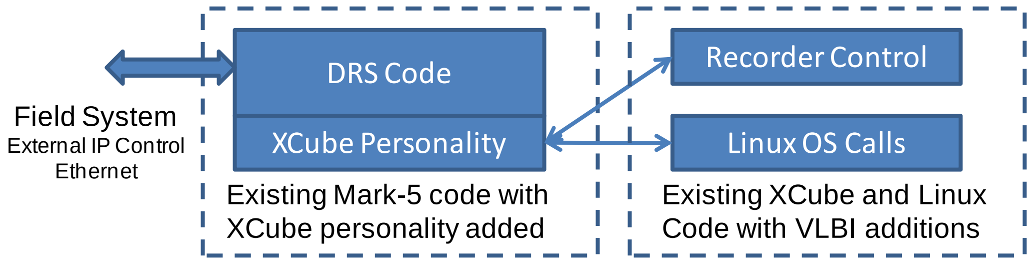

Figure 11: Code running inside the recorder

29. How about identification of components?

a. Disk packs are identified with a mark-plate in the front of the disk-pack. Each disk pack has a model code that identifies the capacity and configuration of it. DP_XXTB_YYP, a serial number, and manufacturing date.

i. Inside the disk pack there is a permanent file that has the hardware and software versions of all disks.

ii. In addition to the identification file, there is a usage log and an error log.

b. Recorders are identified with a Mark plate in the front panel. The recorder is marked with model number, version, serial number, and manufacturing date, In addition the memory, input interface and output interface models displayed

i. There is a permanent file in the recorder with complete hardware and software configuration information.

ii. There is a complete usage log, as well as a complete error log.

c. The files containing hardware and software configuration information are version controlled and each file can be traced back its original configuration.

5. DEVELOPMENT PLAN

Based on information from discussions during the month of December, it will be a tight schedule to make the deadlines. XCube will dedicate 1 full time resource to support the development effort. We will also support with hardware development of the 3.5" media packs.

Since the recorder already operates at well above the 4Gbps recording speed needed for the initial deployment and it has been demonstrated to work at 8Gbps, it seems that the main questions would be to integrate software, and to verify reliability of the system. Dividing the effort into a performance and reliability track and a software development and integrations track can minimize risk and ensure project success. The performance and reliability evaluation can start immediately using available hardware and software components. The software development and integration can run in parallel and be tested separately without negatively affecting the other pieces. The hardware development is limited to a backplane for use of the legacy Mk-V disk packs, and the sheet metal for the use of 3.5" disk packs. The media that we want to use can be verified, tested and certified using different packaging (available) concurrently with hardware development.

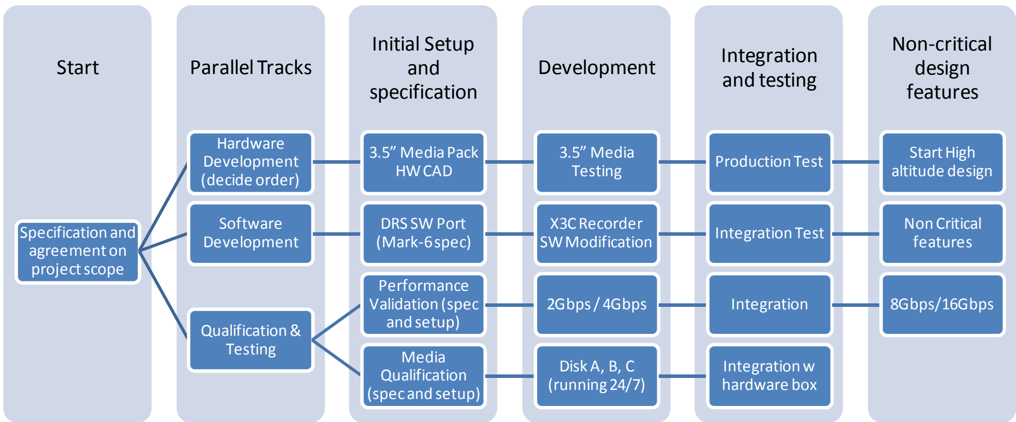

Figure 12: Development flowchart showing the major development tracks and task.

There is a design concept for high altitude operation that involves a low pressure shell casing for the data packs. This shell contains a backup mini compressor and a micro controller with a pressure and temperature sensor that ensures appropriate pressure and temperature for reliable disk operation. The box will be large enough to encompass any of the data packs.

5.1. Resources

XCube can make available resources to do the necessary hardware and software development. These resources are available immediately (software developer, test engineer, hardware engineer).

For the specification parts as well as the integration of correlator code we would need help from Haystack and NRAO personnel.

Note: After discussion with Chester Ruszczyk and a walk through the DRS code it is obvious that this port is straightforward and can easily be done with XCube personnel.

Recorder hardware and Disk Packs are available for testing and development purposes. It would be good to have access to an RDBE for data generation (alternatively, internal XCube hardware could be used, but testing would be more realistic with an RDBE connected).

For later stages in integration testing (February timeframe) we should run real world tests and connect the recorder(s) to a complete setup including a Field System.

5.2. Schedule (TBD and agreed on)

Start of project can be done immediately scoping and initial decisions on direction need to be made.

Performance and reliability testing, hardware development and software development can be done in parallel. The initial parts of the work can be completed in less than 3 months.

We currently project a start date of January 10 and an end date of March 31.

- First demonstrable system with VSI-S commands and integrated commands (5 weeks)
- First tested system (8 weeks)
- Full integration testing (10 weeks)

We need to work out details on the schedule, but this is definitely possible to execute within a 3 months schedule.

5.3. Recorder performance and reliability verification

The performance and reliability testing can be done with a simple setup of a RDBE frontend, a Mark 6 recorder, and potentially a 10Gbit Ethernet switch in between. The setup could be done with available hardware at XCube facilities, provided we can borrow a functional RDBE frontend unit. The Recorder can run with multiple disk-packs at both 4Gbps and 8Gbps on a 24/7 schedule from the start of project for as long as needed.

A reasonable use would be to have an alternating schedule qualifying different disk media with the available 3.5” and 2.5” media carriers. By using the system for development during daytime and testing during nighttime we can get 16hours of recording data per day from day 3 after start of project.

Major tasks:

1. Develop test schedule including what parameters to record, and which environment conditions to test. We would probably run the systems on the warm side. We have no pressure chamber, so we might want to do some testing at high altitude (although high-altitude testing can be delayed until all else is done).
2. Develop test script to execute test, including script to verify that the captured data are correct. (1 week)
3. Setup test hardware for test (hours)
4. Write test report (automated with a few hours of manual work)
5. Steps 3 and 4 will be repeated for each disk setup and capture speed. We would probably run alternating 4Gbps and 8Gbps captures depending on the workload the team wants to generate.

The team needs to decide what media we want to qualify during the test period. XCube currently has, 4 different media qualified for use with the recorder. It would probably be a good idea to run through testing with some of the legacy drives in use today. It would be possible to run double disk packs on a recorder (to get more data capture time on the disk, but it would require a little more software development).

5.4. Hardware development

The hardware development needed is outside the critical path due to equivalent electrical characteristics of the carriers used in performance testing. A decision needs to be made which specific media carriers we want to use first. The 2.5" media carriers are ready to go, so no development is needed for them. Regarding the 3.5" media carriers there are options for using the legacy Mark 5 carriers, a 16 disk carrier and an 8 disk carrier. The 16 disk 3.5" media carrier is in XCube development plan, but we could make either of the 8-disk carriers first.

Based on previous carrier designs, we expect about 1 month of lead time from start of project to first working hardware, and 3 more weeks to adjust for production. This is possible since the design is mainly passive and a repeat of one already made.

We assume that the pressure container will be the last hardware piece to be developed. The main design decision there is to decide if the legacy Mark 5 data packs need to be supported. If that is the case, then there needs to be cooling infrastructure developed for those data packs.

Major tasks:

1. Decide which order to develop the disk packs and write specifications. (1 week)
2. If Mark 5 disk pack
   a. Hardware cad of new front panel (2-3 days, 1 week)
   b. Hardware cad of new carrier backplane (1 week)
   c. Lead time for prototype manufacture (3 weeks)
   d. Testing of new backplane (2 weeks)
   e. Second spin (if needed) (1+3 weeks)
   f. Qualification of media with disk pack (2 weeks each)
3. If XCube 3.5" media disk pack
   a. Finish hardware cad of design (1 week)
   b. Prototype production (3 weeks)
   c. Testing of new box (1 week)
   d. Qualification of media (2 weeks per media)
4. Shipping container for disk packs
   a. Write specification
   b. Select outer box (from set of standard shells)
   c. Hardware cad of protective foam (1 week)
   d. Lead time for production (3 weeks)
   e. Packing crate testing (3 weeks)
5. Pressure container development.
   The concept for the low-pressure container is that of a low-pressure housing that encloses a regular data pack. The pressure housing has a backup miniature compressor and a

microcontroller with a pressure sensor that ensures proper pressure and temperature in the enclosure.

a. Finalize specifications (1 week)
    i. There is a conceptual design on the table that needs to be finalized.
b. Hardware CAD of box (2 weeks)
c. Software development for controller (4 weeks)
d. Bulkhead connector development (outsourced 3 weeks lead time)
e. Production lead time (4 weeks)
f. Testing of pressure container and disk pack. (4 weeks)
    i. Most of the testing can be made at sea-level by running the pressure at +1 bar. This will fully replace actual high-altitude testing, due to low pressure inefficiencies of cooling and air moving units.
g. Operations manual (2 weeks)

XCube has resources identified and available to do the work with the hardware design as well as production.

5.5. Software development

There are several levels of software that need to be developed. There is a critical path of software that is needed to make the system work in the VLBI system together with the RDBE front-ends. This layer of software includes a port of the current Mark 5C software to the new platform and the basic control interface to the underlying recorder software, and finally a playback utility for software correlation. Then there is a layer of new functionality that will improve the ease of use of the system, and increase the diagnostic abilities. The third set includes new functionality not previously available in the system. For this set, the initial systems will use “manual” configuration until the new functions are implemented.

Initial thoughts on major tasks

1. VSI-S interface Implementation to the Recorder
    a. XCube has a generic command parser already implemented for internal use. This is the code that was used as a base to develop the command parser for the RDBE. There is a recorder manager process listening for commands on an internal socket that the Mark 5C code can connect to.
        i. Write a specification for command transfers (1 week,)
    b. Most of the Mark 5C commands are generic in nature, and should directly be transferrable from the current to the new generation.
        i. Port Generic Mark 5C interface code to XCube recorder. The code is written in Python and assumes a Linux platform, so the code should directly port. There are slight differences between the platforms and testing need to be done. (2 weeks)
    c. The specific commands relevant to new implementation are start/stop and media commands and diagnostics. These commands are relatively compatible with the current system as is. This is initial basic level recorder commands to load media, unload media, start and stop capture etc. We expect no major hurdles to implement these functions however there should be a reasonable amount of testing done here.
        i. Implementation of commands (2 – 3 weeks)

        ii. Testing of commands (2 weeks)

2. Operational procedure
   I think this is mostly a management and scientist work package. With this I think we need to sit down and decide how we will be using the recorders and then write the operational procedure (XCube will definitely need help with this package)
    a. Operator manuals for the new equipment. A how to guide for installing media, booting up machines and inspect status etc. (1 week)
        i. This guide will have to highlight operational differences between the Mark 5 and XCube systems. It will also have to contain “debug” procedures to ensure that the new system is working correctly.
    b. Handling procedure for disk packs. (1 week)
        i. The disk packs (depending on selection order) will have slightly different handling requirements than the old ones. These need to be properly explained and operators trained on the procedure.
    c. Identification of parts, serial numbers and configuration data.
3. Command Software Implementation
    a. Interface current Mark 5C software to XCube recorder. This code will mostly run unaffected by the platform change. However, the media commands and the start/stop commands need to be modified.
    b. New recorder commands: for remote operation of new functionality additional commands will be needed (This need to be better defined).
        i. These commands manage the specifics of network setup and how to mount media specific to the XCube recorder. (4 weeks)
        ii. New diagnostic functionality. These functions can be incorporated over time, to enhance system abilities. (initial 4 week effort)
4. Interface Manager implementation, this is the code that defines exactly what gets stored in the files on disk:
    a. By default the Recorder will store the Ethernet frames from the network, either from all addresses or filtered by one specific address. This package is already implemented by default. (0 time)
    b. It may be desired to extract the payload directly from the input stream. By only storing the payload stream to disk the data-rate will be slightly lower and improve overall performance. This is a relatively simple modification of the current Ethernet interface manager (1 week implementation, 1 week testing)
5. Replay software. This software package is needed for replay and correlation functionality. There are multiple variations of this package. We can implement something relatively straightforward first for first-cut version, then make improvements and additional features.
    a. The recorded data are by default stored in an X3C-log file format, in regular files on the disks. When reading back for replay, the data need to be appropriately ordered from the files. This can be done in many different ways, one being just reading the files and consuming the packets from the stream in the right order. Another is to have a small utility that pushes the packets back out to a socket interface that the correlation code can read.
    b. Create interface code for reading the data files directly in the correlator code (Need better information here, should about 2 weeks of effort plus testing, need input from Walter and Chester). It might make sense to directly run correlation code on the recorders.

   c. Create interface code to push data packets from the recorder to an appropriate socket interface (IB, Ethernet, or local). This code is essentially done, but we need to modify it to work with VDIF data, and agree
6. Software integration testing.
   a. Write test specification for integration testing of system. This will define integration test for RDBE / XCube Recorder as well as XCube recorder to software correlation.
   b. Implementation of integration tests (2 – 3 weeks TBD).
   c. Test (3 – 4 weeks?) We might want to run a few tests in parallel with the current Mark 5 systems and show that we the results are the same with both systems.
7. Recorder initial performance evaluation
   a. Setup test using the current version of the XCube recorder and the current RDBE frontend. The test should be done so that we can record data at (2, 4, or 8 TBD) recording speed with appropriate software. The data on disks should be verified to have the right number of packets and correct timestamps.
   b. The test could be extended to be run in parallel with a Mark 5C recorder and the packet streams compared against each other.
   c. Software needed: script to verify packets, and script to compare packets. For test (B) we need to setup a switch between RDBE and the Mark 5 and XCube recorder.