MONTSERRAT VOLCANO OBSERVATORY

GOVERNMENT OF MONTSERRAT

Rebuilding MVO seismic monitoring, January - February 2004

Glenn Thompson

MVO Open File Report 04/04

MONTSERRAT VOLCANO OBSERVATORY

GOVERNMENT OF MONTSERRAT

Rebuilding MVO seismic monitoring, January - February 2004

Glenn Thompson

MVO Open File Report 04/04

Montserrat Volcano Observatory
PO Box 292
Montserrat
West Indies

Tel: 1 664 491 5647
Fax: 1 664 491 2423
Email: mvomail@mvo.ms

* Affiliation inside cover

Affiliation of author

Glenn Thompson
Montserrat Volcano Observatory
Flemings
Montserrat
West Indies

TABLE OF CONTENTS

Executive Summary

In January 2004 I returned to MVO after a 6 month break to find the seismic monitoring in a heavily compromised state. About half the seismic stations were working, the worst network health since the devastating pyroclastic flows in 1997. Since November 2003 at least all acquisition and event triggering systems had been exhibiting problems, with broken paths and drive mounts.

The alarm systems which public safety relied on were non-operational. This was because Cable and Wireless had stopped supporting pagers, and because the wrong cellphone numbers were programmed into the alarm systems. Had those problems not existed, the alarms would have been impacted anyway by the problems with data acquisition.

The diagnostic alarm system designed to alert the Seismologist to problems with any operational software has not been maintained, and the seismic monitoring webpage to which is used to echo these status data was deleted on 27 June 2003, within days of me leaving MVO. A vast assortment of Matlab software and associated data was also deleted.

Computer hardware has also been removed, including one Linux workstation and two Windows PCs. Previously these had run in parallel with a twin for data acquisition, alarm and event analysis purposes. Their removal has made MVO much more vulnerable.

Checks had been abandoned after I left and consequently PC clocks had drifted with correction, RSAM data were been incorrectly time stamped, and waveform data files from EarthWorm were corrupted. There had been some hardware failures and no spares had been acquired.

Finally, none of the computers automatically reboot, logon and restart critical software anymore. This was an important part of the design of the seismic monitoring through the 2001 – early 2003 period that minimised downtime.

Clearly there has been a substantial degradation in the seismic monitoring at MVO in the last several months. Recognition of these problems has undoubtedly been made difficult due to the turnover of seismologists and other staff in recent months and the consequent loss of corporate knowledge, coupled with abnormally low seismicity in the last several weeks.

This report documents the status of the seismic monitoring as I found it, some of the steps undertaken to rectify problems, and recommendations for the resident Seismologist to take forward with the Software Engineers. Finally, some of the reasons that contributed to this decline in seismic monitoring are explored.

1 INTRODUCTION

The seismic monitoring programme at MVO had to rebuilt following Y2K problems, and designed in a more robust manner due to the fact that with a reduction in staffing levels, MVO no longer manned the Operations Room around the clock, leading to much greater reliance on automated systems.

To achieve this increased level of robustness several safety nets were incorporated into the design of the seismic monitoring systems, given the importance of MVO seismic monitoring for public safety:

1 Operational volcano alarms for safeguarding the population of Montserrat and aviation.

2 A Space Shuttle approach – all PCs have a twin running in parallel, or sitting on the shelf. Failure of any one PC does not impact the mission.

3 All PCs should autoreboot, autologon and restart software. So if there is a brown-out, all systems reboot.

4 All possible points of failure in the system are monitored by a diagnostic alarms (including the diagnostic alarm system itself!). So if anything fails, the Seismologist is notified within minutes by pager/cellphone. All status messages are replicated on a web page which is monitored at least twice daily.

5 All PCs run on a UPS (and there is a backup generator).

6 Critical spares were acquired for all seismic network equipment, ILI cards, ADC cards, DAT drives and other periperhals.

The worrying fact is that between early 2003 and now, most of these safety nets have gone (5 is the exception). Resources (mainly computer hardware) intended for effective operation of the seismic monitoring and early warning systems are being used for running the database. Strategies that were developed to safeguard the seismic monitoring have been abandoned. Other important software has been deleted or disabled and replaced by less effective software, impacting alarm response, data analysis and data archival. The seismic networks are in their worst ever state, and there have also been hardware failures, including the loss of two computers, and changes to the computer network.

2 STATUS OF THE SEISMIC MONITORING SYSTEMS

On arrival, Resident Seismologist Art Jolly asked me to troubleshoot problems with corrupt Earthworm data files he had noticed in recent weeks. I look into this and discovered a wide range of other problems, most of which had been masked by the almost total lack of seismicity. I documented my discoveries each day in my notebook during my 5 weeks at MVO, and in Appendix A I reproduce my notes from the first week in the form of a technical diary.

By 13th January 2004 I recognised I needed to present an overview of these problems for MVO Director Peter Dunkley and Project Manager Bill McCourt, so that they could tell me how they wished my to proceed under these unusual circumstances; as Cover Seismologist, I was cautious about making sure I got the proper authority to proceed. The reasons for this cautious approach might be clear after reviewing section 5. The email is shown here, and the problems as I documented them below in subsections:

From: Thompson, Glenn
Sent: 13 January 2004 10:08
To: 'peter@mvo.ms '
Cc: McCourt, William J; Thompson, Glenn
Subject: Status of seismic monitoring

Hi Pete,

Although I had hoped to concentrate on science on this trip, I spent my first week discovering and troubleshooting a number of problems with the seismic systems. I know the details of the seismic systems are always hard for others to understand, so I'll spare you the details - the attached document summarises the main problems.

I anticipated that problems like these were likely to occur when Joel tookover the seismic systems, without my consent.

HOW DO YOU WISH ME TO PROCEED? Possibilities:

1. I could just document the problems, and leave them for others to solve, and get on with other observatory work.

2. I could fix the superficial problems, and document the deeper problems for others (this has been my approach so far). The systems would inevitably still be very fragile, but I could turn predictable problems into procedures for Art to follow.

3. I could make a start on putting the seismic systems back together properly - this would definitely impact the database, but make the seismic systems more robust. But I would probably run out of time to complete this work.

4. I could write some detailed documentation on how the software is supposed to work (but much of it has changed), including the design principles that guided this work from the outset. This is probably something I could only cover in superficial detail on this stint.

Cheers,

Glenn

The response from Bill McCourt was:

Glenn

Things look to be in a mess I'll talk to Pete and he'll coordinate a way forward with/through you. My initial response (after 2 mins scanning the docs) is that you follow option 2 since there isn't time to finish 3 and option 1 is not a realistic professional approach - which I know you're fully aware of.

Thanks & cheers

Bill

2.1 Seismic networks: (health: Marginal)

In 2001, there were 8 operational stations on the digital network, 7 on the analog network. But we've got 4 working stations only (including the Multimo station!), on each network. There are another 2 stations on the digital network that are intermittent. This is the worst state networks have been in since late 1997 (when PFs took out several stations). This is marginal for event detection, alarms and earthquake location.

I have asked Pyiko to advise on which station problems are likely to be long-term (perhaps only solved at the time of the seismic upgrade) and which could be fixed in the short-term (no feedback has been received).

The purchase of an EarthWorm ADC for $1600 is long overdue. This was proposed as long back back as December 2002, and would have allowed both networks to be acquired through EarthWorm (making QNX obsolete), in which case we'd now have 6 independent station locations, which would be rather more comfortable. Given that the seismic network upgrade is unlikely to be complete for more than 1 year, I would still strongly recommend this option. Moreover, the software could be greatly simplified, since currently a lot of software complexity is for the sole purpose of merging datastreams prior to the data analysis stage.

Another step is to merge the Calipso seismic data into the MVO EarthWorm system, as originally intended. This would give us some extra stations to enhance our currently understrength monitoring.

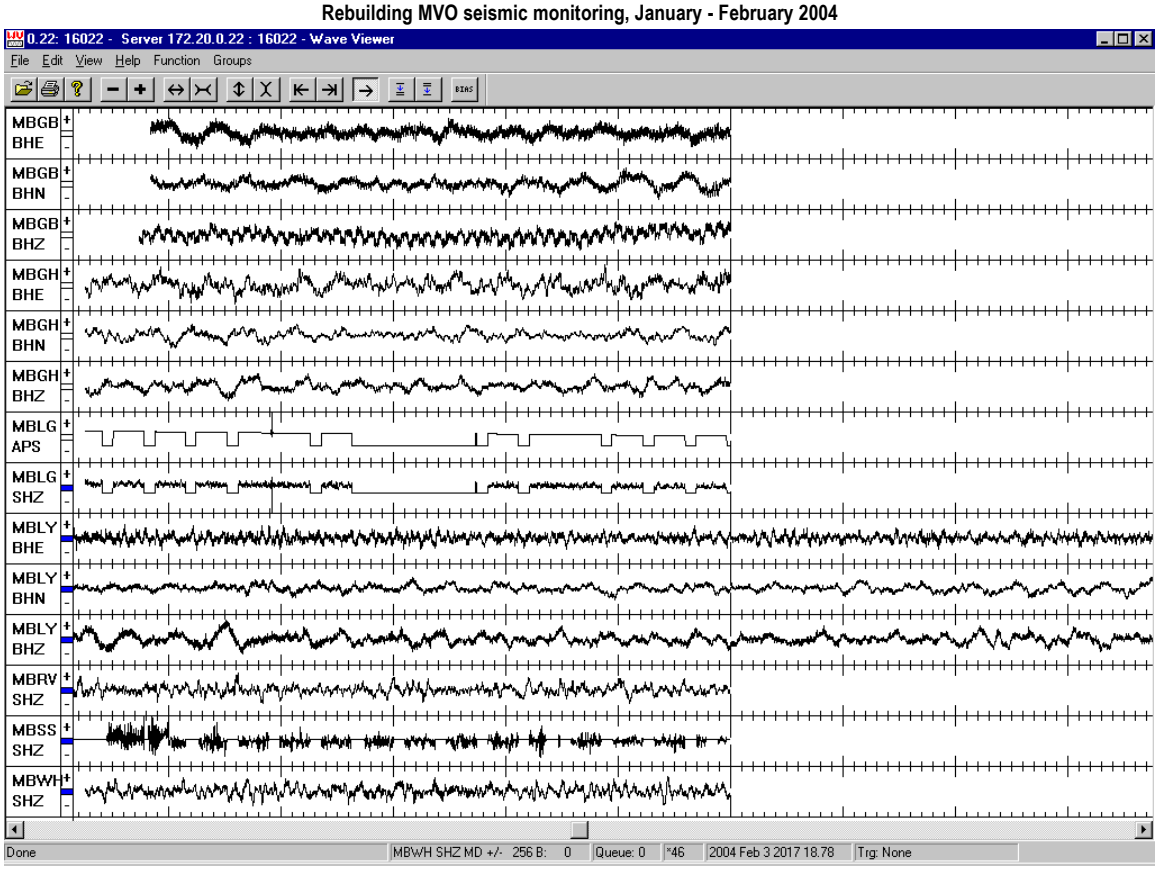

Figure 1: EarthWorm data acquisition, 2004-02-03 20:17 UTC. Window length is about 90 seconds. Main things to note: (1) High frequencies on MBGB-BHZ indicating damaged vertical component – no useful data; (2) square waves and dropouts on MBLG APS and SHZ – no useful data; (3) dropouts on MBSS SHZ – no useful data. MBRY is also dead (and not currently being acquired on EarthWorm at all). So working stations are MBGH, MBRV, MBWH + the Multimo station MBLY. (Note that data from MBLY arrive first – this is because data from MVO network get buffered on SA24 for 1 minute before being made available to EarthWorm).

2.2 Acquisition systems: (Fragile, but now operational)

Analog network: QNX data acquisition problems - hard drive problem, meaning no events recorded since mid-December. Think I've solved this – problem was being caused by mounting these drives under Linux. Got “Stale NFS file handle” error.

Digital network: Earthworm data acquisition problems - blank event files since early November - caused by PC clock errors. Also fixed problems with continuous data acquisition.

The acquisition systems all used to reboot. But they don’t now. So in the event of power brownout, all seismic systems would be dead.

2.3 Alarm systems: (No pager service; more work needed)

The main thing to note about the alarm systems is that there is currently no operational method for getting a *message* to a cellphone (the SMS messenging which Joel got working last summer is no longer working). So while a cellphone might ring, you don't know whether its an event or tremor alarm, or how big it is. And you wouldn't know if it's a volcano alarm or diagnostic alarm. Pagers are no longer supported by Cable & Wireless.

The RSAM volcano alarm system, which we've used since 1998, is obsolete because it cannot send messages to cellphones.

The BSAM (and QSAM) volcano alarm system, started in 2000 as a replacement for the fragile RSAM alarm. The system has become convoluted since last May, and appears to be fragile. I found the numbers were wrong – now fixed.

The BSAM test alarm, is just a false alarm at 7am and 9pm each day to test if the communication method used by the MVO volcano alarm system is working. It is still running, but I had to fix the cellphone numbers.

The diagnostic alarm system, began in 2001, and is supposed to check all dependencies in the seismic acquisition and alarm systems, many times a day. But new dependencies have been introduced to the seismic systems since last May, and these are not being checked. Moreover, this system is still sending alerts by pager (not cellphone), so staff were not aware of problems anyway. Consequently, when the volcano alarm system crashed on Saturday night (the Linux PC had crashed), nobody was aware until Monday morning.

2.4 Analysis systems: (Loss of capability)

Much of our best analysis software has been lost, including the most useful tool, a near-real-time seismic webpage, which automatically updated with all kinds of seismic data including seismograms, spectrograms, hypocentral plots, counts, energy, rsam data, and state of health information. This also included the tools which were used for generating reports on various timescales (weeks, quarters etc). These programs, data and web pages were removed on June 27th 2003. Its interesting to note that at the same time, the cover Seismologist Lars Ottemoller developed Fortran routines to plot event counts and RSAM data. These are convoluted and impractical to use. The GUI based Matlab tools that existed at MVO since late 2000 were not only much easier to use, but far faster and highly flexible. Other capabilities lost include near-real-time estimation of equivalent magnitudes (which were reported in alarms) and determination of rockfall trajectories (the rockfall location system to which 'ampmap' was the GUI front-end).

To go some way to rectifying this situation I have recovered archived versions of the Matlab programs of greatest utility, and rebuild the appropriate datasets. I have installed the simple programs "MBOB" and "RSAMGUI" on the desktop on PROGRAMMER, ANALYST and Graham Ryan's PC, to provide simple rsam/bsam and counts plots. However, I have not attempted to reinstall a much wider array of analysis software that had previously been used for plotting spectrograms, hypocentres, or estimating the equivalent magnitudes or energy of events. Nor have I attempted to reintroduce tools that were previously used for integrated monitoring of

seismic, gas and deformation data, since the resident Seismologist did not feel this would be an effective use of my time.

2.5 Archival systems: (Loss of data)

An archival system which had evolved over many years to very robust and user friendly, was replaced by Jan/Feb 2003 without my consent. Lars Ottemoller who was then covering for me decided to showcase the capabilities of a language called 'Expect' and decided that data archival software was a good application for this language. Unfortunately the program he and Software Engineer Joel Maranhao wrote did no checking whatsoever. It was inferior to the previous application in several important ways:

1. It does not check to see if there is a tape in the DAT drive before archiving data. Consequently, there have been times (as I witnessed last week) where the program was run without a tape in the drive, and there was no error message to the user to indicate that no data had been saved!
2. It does not forward to the end of data on the tape before archiving. Consequently, it is easy to overwrite old data.
3. It does not verify tarfiles as they are written.
4. It does not check for an end-of-tape condition.

Together these glaring omissions with this data archive program mean that some data that has been acquired in the past 12 months, has not been saved to tape. Given the capacity of the hard drives, hopefully all the data is still online.

I had not been keen on using this software, but Joel nevertheless demanded that the seismic team use it, and my authority was never underlined by the MVO Director at the time.

2.6 Loss of computers

In the past several months 3 seismic monitoring computers have gone missing or have been taken over for other purposes.

There used to be a spare SA24 computer on the shelf, ready to plug in should the operational system fail. This was taken by the Software Engineer in March 2003 without anyone's consent, and he used it at his home as his personal computer, and deleted all remnants of SA24 and the operating system it required, Windows 98, off the system. In July 2003, there was a failure of the primary SA24 system and I was contacted by email and asked what should be done to get another SA24 system operational.

Until July 2003 at least, there were two Earthworm systems running in parallel. A problem with SA24 meant that if a station went down, it would begin sending corrupt data for that channel, which could break Earthworm acquisition if that station were being acquired. So I always had one Earthworm set to acquire all stations, and another set to acquire all stations that I was confident about – i.e. those that rarely exhibited problems. The logic for deciding which system to use for waveform data was simple. The primary would produce larger files (because there were more

stations), and these would be used by default. But if a file was missing (each file represented 20 minutes of data) from the primary, the corresponding file from the backup would be used. This parallel approach was instrumental in resolving SA24/Earthworm problems in the months following the installation of these systems in March 2001. Since since July 2003 this backup functionality has been lost, but has now been reintroduced. The question is whether it will be maintained.

Likewise, there were two Linux workstations running in parallel. These were used for the volcano alarm system, the diagnostic alarm system, the equivalent magnitude computation, the rockfall location system, event classification and hypocentral determination. Failure of either PC was not critical, since the other computer still performed all functions. However, in February 2003, the Software Engineer announced that he was taking over one Linux workstation which would be for his personal use only. All software and data were removed or disabled.

In recent months, the diagnostic alarm system on the one remaining Linux workstation utilised for seismic monitoring has not been maintained as network paths have changed. The net result is that the alarm system can become compromised for several days or more (as happened at the weekend) without any notification to the staff.

An important part of the early warning system was the seismic monitoring webpage, which as noted above was deleted in June 2003. In order to utilise this the Seismologist had a desktop computer at his house which was used for checking seismicity in response to alarms and seismic monitoring status at least a couple of times a day. Since July 2003 this computer has apparently been in the bubble as is no longer available for seismic monitoring.

3 RECOMMENDATIONS

Three weeks later, I produced an update of the status of various components of the seismic monitoring, and a list of recommendations for each:

3.1 MVO digital seismic network:

Status of the digital seismic network:

- MBWH, MBGH and MBRV are working.
- MBRY is dead.
- MBLG and MBSS are intermittent.
- MBGB has a damaged Z component. The site has becoming increasingly more noisy. A potential site exists just beneath the peak, and a vault could be built for us if we move on this within the next week.
- MBBY has been down since late 2001.

Recommendations:

- Decide ASAP on whether site near peak MBGB is suitable. If so, ask vault to be constructed. Otherwise, relocate MBGB (lower priority).
- Fix MBLG
- Fix MBRY
- Fix MBSS
- Move ASAP on network upgrade, so can get MBBY back online.

3.2 MVO analog seismic network:

Status of the analog seismic network:

- MJH, MLY, MGH and MWH are working.
- MLG is dead.
- MRY is no longer deployed.

Recommendations:

1. Fix MLG
2. Move ASAP on seismic upgrade, so Short Period network will become obsolete.

3.3 Seismic acquisition systems:

Status:

- Broadband net acquired through SA24/EarthWorm, Short Period net through QNX. Triggering is not as effective, and event files from different networks do not coincide. Hypocentral determination is also not as effective, because fewer stations are available.
- All acquisition systems used to auto reboot. But now, SA24, EarthWorm3 and RSAM no longer reboot. So in the event of a power outage, somebody has to notice and manually reboot those systems.
- SA24 doesn't record data properly for dead / intermittent channels. Downstream this impacts EarthWorm modules that try to write out data in Seisan format – and leads to massive data loss. For this reason, two EarthWorms are run in parallel. EarthWorm3 should only be acquiring stations we know are working well. This means at least some data is captured (data collection software on Seisan PC has been written to use the best data available).

- EarthWorm and RSAM software use PC clock for data recording purposes. With EarthWorm, clock drift causes data dropouts in event files. With RSAM, the timestamping of the data is incorrect.
- Each EarthWorm module has its own configuration file. So when adding/removing channels from EarthWorm, many files have to be modified, and an error in any one can cause acquisition problems. Develop program to write all Earthworm/GlowWorm configuration files based on 1 new parameter file.
- Multimo datastream is currently being acquired through EarthWorm.
- Calipso datastream is not currently being acquired through EarthWorm. Indeed, it isn't being telemetred at all (it gets manually downloaded directly from field stations instead, in much the same manner as MVO GPS data). University of Arkansas still have a lot of software development work to complete to enable this to occur.

Recommendations:

1. Purchase and install EarthWorm ADC. This will enable both seismic networks to be acquired through EarthWorm. This will make QNX obsolete, and lead to more effective event triggering, classification and hypocentral determination. (Moreover it will allow massive simplification of downstream software, making the seismic systems more robust and easier to maintain).
2. RSAM, SA24 and EarthWorm3 should be set up once more so that the acquisition software auto-reboots.
3. Acquisition datastreams on EarthWorm and EarthWorm3 should be checked daily, using the wave_viewer program. Any stations with dropouts should be removed from EarthWorm3.
4. Move ASAP on the seismic network upgrade. SA24 will then be obsolete (replaced by Scream). QNX will also be obsolete. (The same benefits as in (1) above will be obtained, and the software can be simplified in exactly the same manner).
5. PC clocks on EarthWorm, EarthWorm3 and RSAM should be checked daily and corrected if more than 10 seconds out.
6. Encourage and help University of Arkansas to sort out Calipso telemetry problems and get datastreams into EarthWorm.
7. (Low priority) Develop program to write all Earthworm/GlowWorm configuration files based on 1 new parameter file.
8. NEVER MOUNT QNX DRIVES REMOTELY.

3.4 Seismic alarm systems:

Status:

- The RSAM alarm system is running, but since Cable & Wireless pager service is down, and since RSAM rarely manages to call phones successfully, the system is ineffective.
- The AEF volcano alarm system on Seisan is running. But SMS messaging is not working. Patch is investigating. The software has also become convoluted and difficult to maintain.
- Attempts have been made to use the GlowWorm alarm system, but again SMS messaging is not working. This approach is particularly worth pursuing, since this alarm system does not depend on the Linux PC running – by removing this dependency, the system is more robust.
- Reimplement the AEF volcano alarm system to use network MAG data instead. This would make the thresholds easier to tune for the Seismologist, and reduce the complexity of the alarm algorithm.
- Reimplement the implement an adaptive network MAG threshold. This will result in more meaningful alarms in future.
- The main diagnostic alarm system is running. However, failure points have been introduced to the seismic systems since last year, and are not being monitored.
- The diagnostic alarm which checks if the Seisan PC is running, is flawed, because the program won't run if the Seisan PC is not running! This is because /live/check_seisan is on a remote mount from the Seisan PC itself!

Recommendations:

1. Need to enable SMS messaging.
2. Further explore the GlowWorm alarm system.
3. Rationalise the AEF volcano alarm software – 1 program (+ modules) with 1 parameter file.
4. Rewrite AEF volcano alarm algorithm to use network MAG instead. Put magnitude thresholds and phone lists into a single parameter file. Do nothing at this stage about adaptive thresholds.
5. Identify weak points in the seismic systems, and add them to the main diagnostic alarm systems, so they are monitored.
6. Install check_seisan on Seisan4 directly, and run it from there.

3.5 Backups and archiving:

Status:

- Backups of raw seismic data (in Seisan format) are currently made to DDS3 DAT tapes with the "backup" program. System backups are made using the same program. However, the backup program does not verify tar files as it writes (as previous software used to) or report end of tape condition. Consequently, the backup procedure is unreliable.
- Tapes are stored in the fireproof data safe. However, this safe no longer closes, and so in the event of fire at MVO, tape backups will be destroyed along with original data on hard drives.
- Although DDS3 and DDS4 tape drives are supposed to be backwards compatible with DDS1 (on which most old seismic data still reside), only 1 of 4 tape drives actually reads these old tapes.
- There is an automated daily backup. However, programs and datasets have been moved, and other things changed, so this no longer backs up relevant files.
- There has been a project since 2000, for research purposes, to get all seismic data copied off tapes and into online Seisan databases. There are still many gaps, and "A" copies of tapes may be in MVO or BGS MH.

Recommendations:

1. Until the fireproof safe door is fixed, system and data backups should be stored at the Seismologist's home.
2. The backup program should be fixed so it verifies tar files as its writes, and stops and reports clear error messages if there is a problem.
3. I intend to take one of the DDS4 DAT drives back to BGS and let the Unix guru there investigate incompatibility issues.
4. Decide what files need to be backed up each night, and add to the daily backup program.
5. (Low priority) On-line Seisan databases should be completed.

3.6 Other software considerations:

Status:

- MVO software on Seisan PC is organized under /development. Each program is supposed to have its own directory, makefile and test program. Programs use common chunks of code stored in Perl modules. There should only be one copy of each piece of code.
- Chunks of code seem to be replicated in more than one module, e.g. global variables and paths from Settings.pm are replicated elsewhere.
- Not all makefiles work.

- Test programs are not up to date.
- Get errors reported about the database when using make with seismic programs. **This is alarming. Problems with database should not be interfering with (higher priority) seismic systems.**
- Various drive mounts have been added to the Seisan PC since last July, including many Windows drives. The QNX drive mounts had to be removed because they were causing the alarm computer and QNX data acquisition to crash. But Windows drives have also caused problems in the past. So decision was made a long time ago to transfer all data by ftp/Kermit, and have no unnecessary drive mounts between the seismic computers, or between a seismic computer and any other computer. However, no documentation has been left explaining why these drive mounts have been added. So removing them could cause problems with the database or alternative website.
- Seisan4 is supposed to be the twin of Seisan. Indeed Seisan4 was bought as a more powerful replacement for Seisan, the idea being that Seisan would sit on the shelf, and be brought in if Seisan4 ever failed. To date, Seisan4 never has been set up as a twin, and if Seisan fatally crashed, it could take weeks to recover the seismic monitoring.
- Hard drives have become cluttered with random files and directories.

Recommendations:

1. Clean up and test Perl modules.
2. Clean up and test “make”, “make test” and “make install” for Perl programs.
3. Data transfer should all be by ftp/Kermit.
4. Reasons for mounting Windows drives under Linux should be documented.
5. Drive mounts unnecessary for seismic acquisition, alarms and archival should be removed from /etc/fstab.
6. Seisan4 should be set up to replicate Seisan, with the same software and drive mounts. The two machines should not be linked in any way, and ideally Seisan should be put on the shelf, and used as an emergency backup only.
7. Someone should systematically go through hard drives, particularly those on Seisan, and remove obsolete files and data. This should be repeated perhaps every month. Someone needs to check that no seismic hard drive fills up, otherwise systems can crash.

4 MATLAB SOFTWARE

4.1 Matlab software

Between 2001-2003, a wide variety of seismic software was written in Matlab, and was run automatically to update a seismic monitoring webpage. Much of this functionality was deleted in on 27th May 2003 for some unknown reason just weeks after I had left MVO, and it is not in my terms of reference to replace all this software or the webpage on this short visit. However, two forms of data are particularly important to look at: RSAM data and detected earthquake counts. Procedures for plotting these with Seisan are rather convoluted, so I have re-established my Matlab programs for plotting these data. The Matlab programs are stored at //dome/seismology/MATLAB/mfiles.

4.2 Installing Matlab

Matlab release 13 Product CD should be in the data safe in the Ops Room. Install it on your computer. The installation directory should default to:

C:\MATLAB6p5

For my Matlab programs to work, the following drives should be mounted (all Windows computers used by Team Seismic should have these settings already):

P:/ //dome/seismology

J:/ //rsam/win_rsam

X:/ //seisan/data3

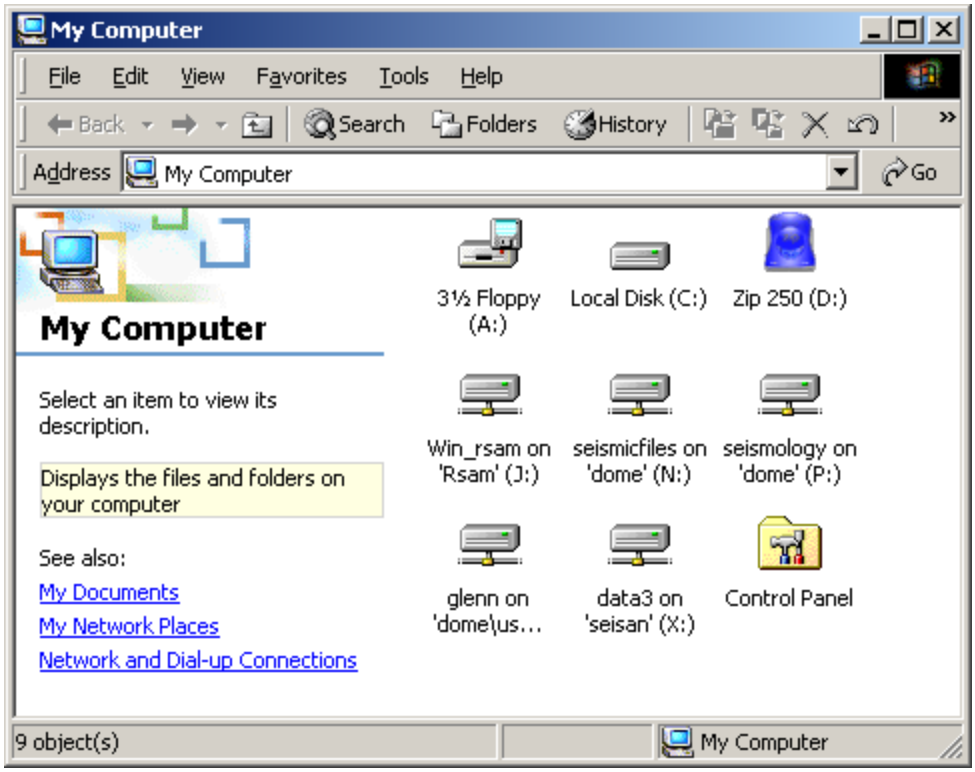

My Computer, showing drive mounts needed for successful running of MVO seismic software written in Matlab.

To set these drive mounts, use Tools -> Map Network Drive under My Computer. Then copy:

P:/MATLAB/mfiles/Shortcuts to your Desktop, and

P:/MATLAB/mfiles/startup_gt.m to C:/MATLAB6p5/work/startup.m

4.3 Running Matlab

You should now have three Matlab shortcuts on your Desktop.

MATLAB6p5	- use this just to start Matlab
MBOB	- this starts Matlab and then runs mbob.m
RSAMGUI	- this starts Matlab and then runs rsamgui.m

Icons on the Desktop in the Operations Room. RSAM GUI is a more sophisticated interface. For a rapid assessment of recent seismicity, MBOB is easier.

Whenever you start Matlab, it will look for a file called startup.m under C:\MATLAB6p5\work. This file is used for setting global variables, adding paths, and changing the working directory. After this has run, the current directory should be P:/MATLAB/mfiles.

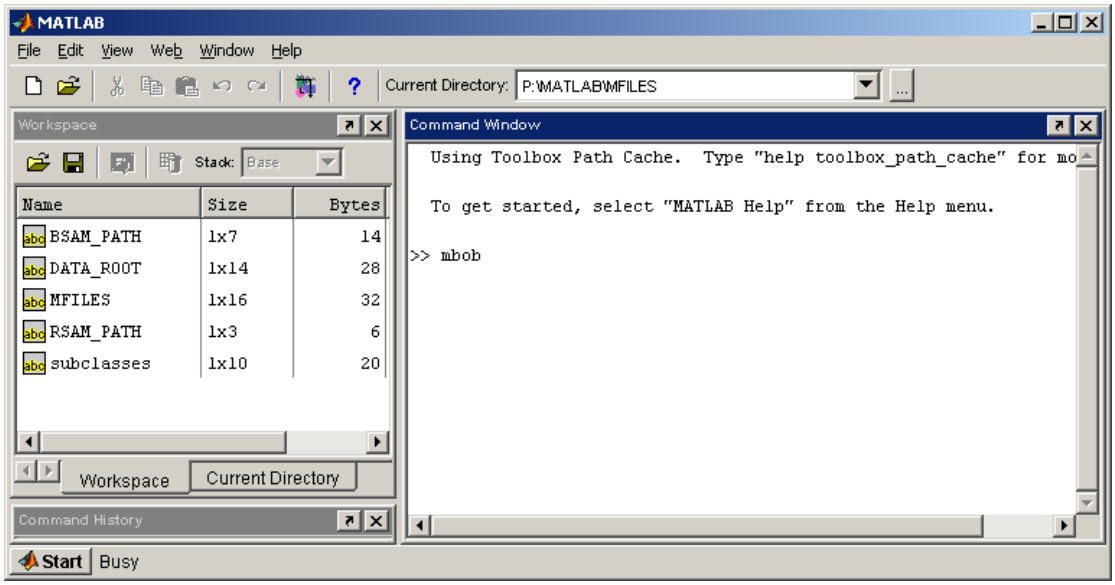

MBOB is particularly useful for real-time monitoring. With it you can plot COUNTS or RSAM data.

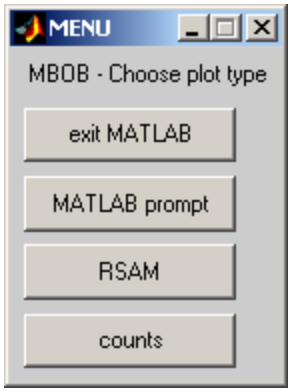

The MBOB main menu. Use this to choose RSAM plots or counts (detected events) plots.

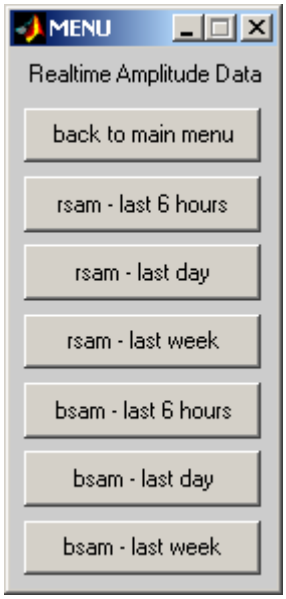

The RSAM/BSAM menu. RSAM data correspond to the analog network, BSAM to the digital network. Its handy to have these quick buttons to look at recent data over different time periods, particularly while monitoring a dome collapse and trying to establish if seismicity is increasing or decreasing.

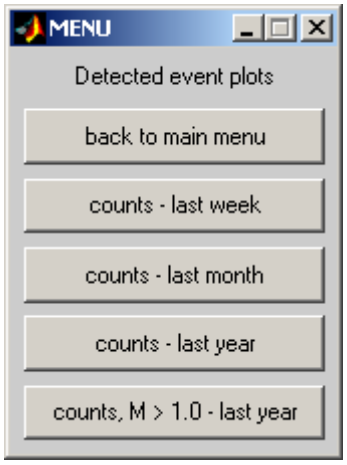

The counts menu. Its useful to look at counts data over different periods to establish if seismicity is increasing or decreasing.

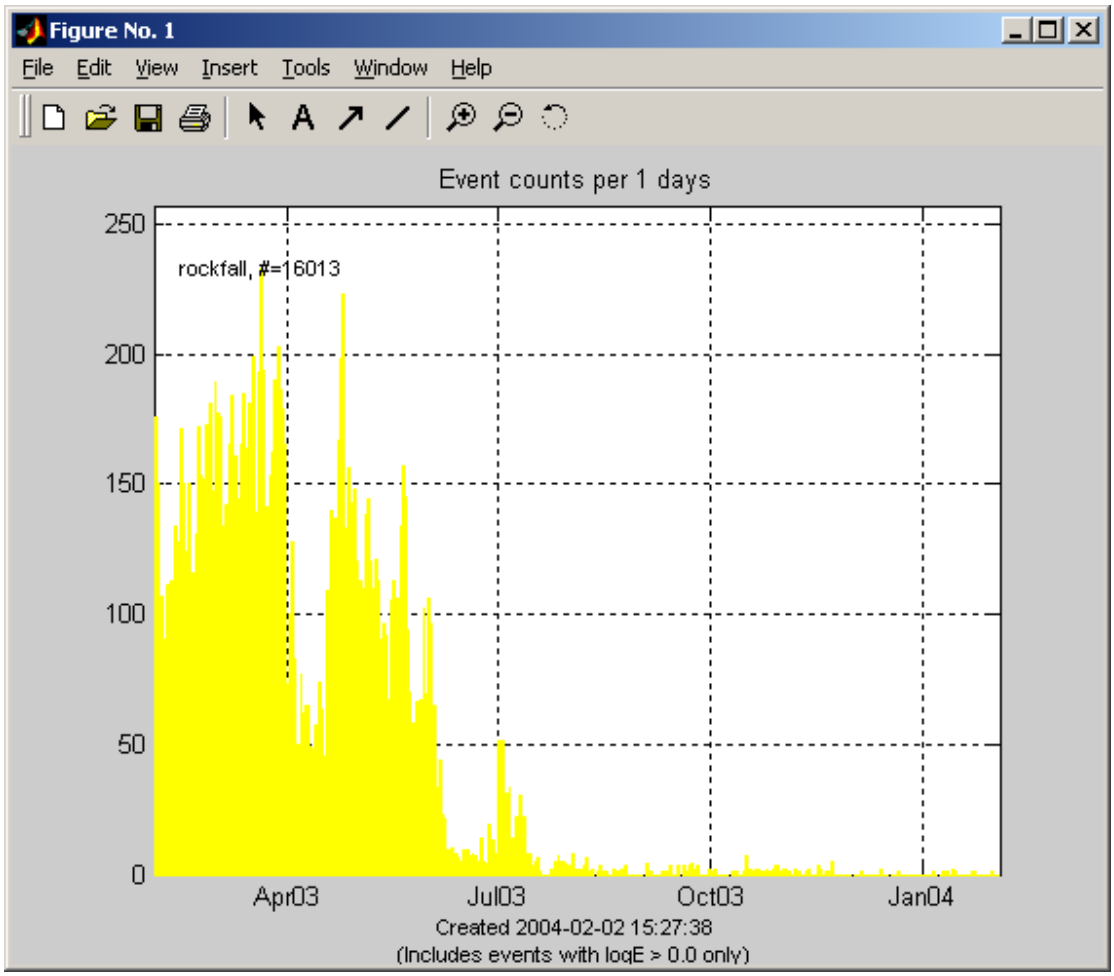

Example of a counts plot. It shows the sharp decline in the number of detected rockfall signals following the 12 July 2003 dome collapse. Rockfall seismicity appears to peak in February and March 2003, resurge in late April 2003, and then decline until early July.

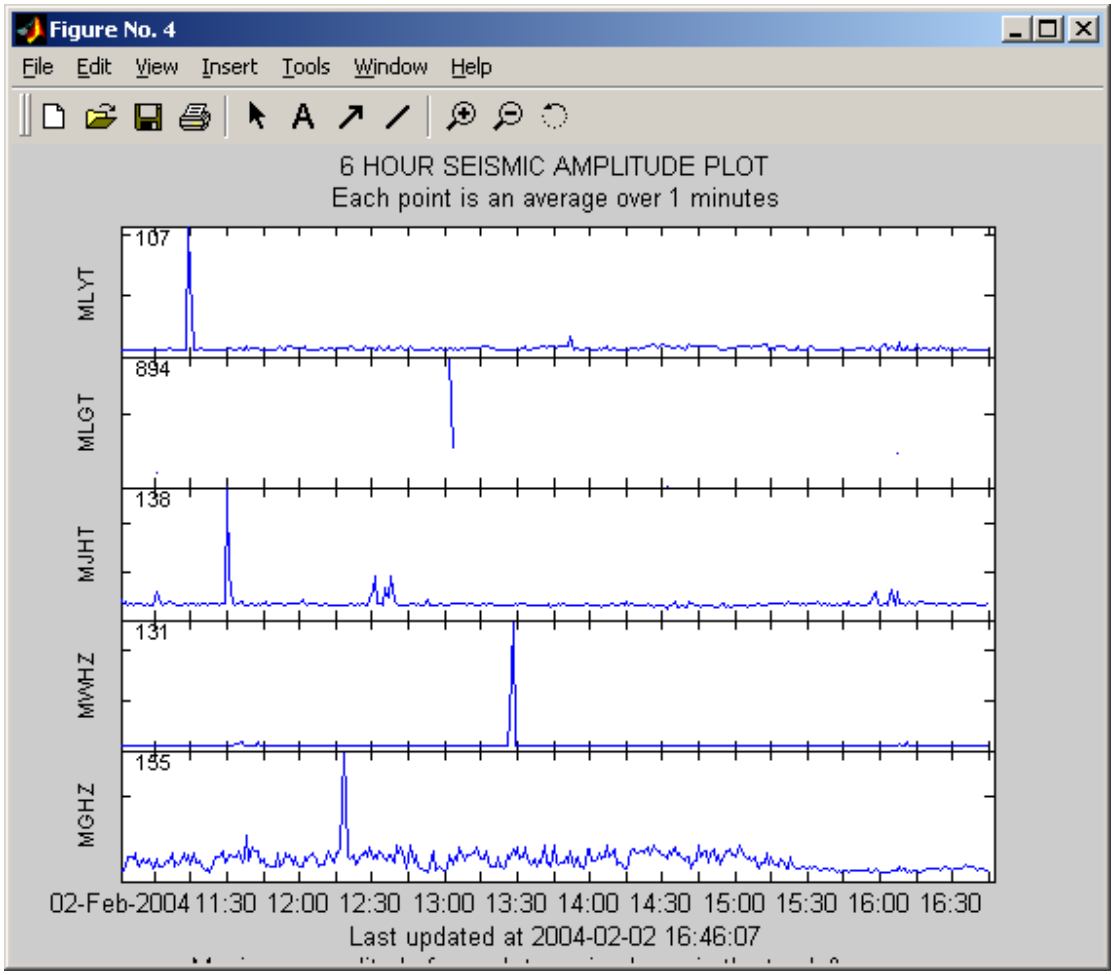

Example of an RSAM plot. There is essentially no seismicity here and the only spikes shown are due to telemetry noise. In general RSAM provides an excellent way of determining seismicity levels and network health.

5 SPECIFIC RECOMMENDATIONS FOR THE MVO SEISMOLOGIST AND SOFTWARE ENGINEERS

Shortly prior to leaving, MVO Director Peter Dunkley asked me to write a list of specific recommendations for Seismologist Art Jolly, and software engineers David Silcott and Joel Maranhao (who was returning for 6 weeks after I left). These were presented on my last day, February 7th 2004:

5.1 General design considerations:

1. All acquisition and alarm systems should auto-reboot.
2. Alarm systems should have as few dependencies as possible. Practically this means they should run on the EarthWorm PCs.
3. Database should be disassociated from seismic acquisition and alarm systems. That way, problems with the database can't interfere with MVO real-time monitoring.
4. Diagnostic alarm systems should check all points of weakness within the seismic acquisition and alarm systems.
5. Data transfer to analysis PCs should be manual, so any problems can be detected.
6. All software should be simple, robust and easy to maintain. Elaborate software in specialist languages can be difficult to maintain once its author has left. This has been the case with some of the Expect, ASP and MySQL software in recent months.
7. There should be duplicates of all seismic acquisition and analysis PCs. These should either be run in parallel, or with one as a backup on the shelf.
8. A minimum level of spares of 25% should be maintained for all seismic network and computer hardware. As soon as the level dips below 25%, new spares should be ordered. Otherwise, long-term problems are inevitable.

5.2 Work that could be taken forward by Patch:

- Continue to investigate SMS messenging. Now that Cable & Wireless have discontinued the pager service, we have to rely on cellphones. SMS messenging is the way to get alarm messages to cellphones. But help may be needed from Dave and Cable & Wireless. Document.
- Reconfigure the Seisan-4 Linux workstation so it is an operational spare for the Seisan Linux workstation (this was a machine Joel had taken over for his personal use months earlier). This means there will have to be a /live directory on both, and both will need to mount //dome/seismicfiles. But they should not be mounted to each other. The Seisan software will need to be installed on Seisan4. Cronjobs will also be the same. Come up with some procedures to do a daily ftp copy of /development, /live and any other relevant

files from each machine to the other. The idea should be that if one machine goes down fatally, the other can be used, and the software will all be up to date. But only one machine should routinely collect data, and be used for processing.

- Investigate the GlowWorm alarm system. Joel has worked on this before, so will be able to help. Document.
- Set up the acquisition PCs (RSAM, SA24, EarthWorm3) so that they autologon and the main software autoreboots in the event of power failure. (Involves changes to Windows Registry).

5.3 Work that could be taken forward by Joel:

- Clean up and test the Perl modules. Eliminate duplicate code. Test key programs. He dropped this work back in January/February 2003 leaving a substantial part of the seismic software in a limbo state.
- Rewrite the alarm system to use magnitude data. This was a proposal which dates back to early 2003, and Joel decided it would be good to create a database to hold the magnitude data. This led to Joel's database, but the alarm system has not been rewritten yet. Doing so would greatly simplify the alarm system, and in particular, make it much easier for the Seismologist to tune. (A copy of the original proposal is attached). Document.
- Use the above as an opportunity to rationalize the whole alarm system. Use a single parameter file containing thresholds for event and tremor alarms, and list of cellphone numbers (or email aliases) to call. Document.
- Explore possibility of setting up the alarm system to run on the EarthWorm and EarthWorm3 PC. This is the ultimate goal, since it means there will be fewer dependencies – data do not need to be transferred to the Linux PC and processed there. The Linux PC will then not be mission critical in terms of real-time seismic monitoring. The EarthWorm PCs will be, but they already are. Document.
- Document what all the mounts on Seisan (and Seisan4) are for. Remove any that are not mission critical. Above all, do not remount the QNX drives, as these have caused extensive problems with data acquisition and alarm systems. And do not mount Seisan4 on Seisan, or vice versa.
- Look into the diagnostic alarm system(s). Think about what other weaknesses in the seismic systems should be added, and add them if possible.
- Document procedures for system and data backups to DAT tape.
- Fix backup.exp so that is verifies tar files as it writes to tape. Currently it does not even alert the user to an end of tape condition, and merrily pretends to be recording data, when it isn't. Data loss can result.

5.4 Work that could be taken forward by Art:

- Get and install the EarthWorm ADC. I asked for an updated quote on 29 Jan. I have phoned since, but not gotten through. Anyway, this is the key to simplifying the seismic systems, and making them more stable, since there is currently a massive overhead to combine data from both networks. I can simplify the software on my next visit if you request it.
- Get the seismic upgrade ASAP, and make arrangements for support to get it installed as necessary.
- Decide where to redeploy MBGB (Peter is asking for a replacement CMG-40T). Pyiko and I identified one possible site. It needs to be moved because of Cable & Wireless plans to run a large generator at the present site.
- Work with Pyiko to get MBLG, MBRY, MBSS, MLGT and MBBY back up (in that order).
- Until the fireproof safe is fixed, store data and system backup tapes at your house. Otherwise, if the MVO burns, we'll lose all the data and software. Also check how often Dave is making DLT backups of //dome/seismicfiles, and know where they are (at MVO is no good if it burns).
- Check the spares/equipment situation. Make sure we've got enough spare field equipment (to get through the move), helicorder spares (and make sure all 4 helicorders operational) etc. Get a new Seismologist PC, to replace that one that blew up in June 2003.

6 DISCUSSION

6.1 The conflict between the MVO database and the seismic data acquisition and alarm systems

The MVO database grew out of the software specification for tremor alarms I presented in February 2003 to the then Software Engineer. He felt a MySQL database would be helpful in implementing this specification. He also quickly saw how it could be used to implement other things – such as an electronic logbook, with attached photos, and so it grew. I was enthusiastic about and very supportive of this development.

The problem with the MVO database is the manner in which it had been implemented. It has been installed on the same computer used for the alarm system, which is mission critical. Shortcuts taken with the database, e.g. mounting of multiple remote hard drives (something we learned not to do with the seismic systems years ago because this could make software hang whenever there were problems with the computer network), mean that the alarm system is much more vulnerable than before. Indeed, the alarm system crashed in April 2003 for this very reason.

Moreover, the seismic code has been infused with database code – there is no clear separation between the two, so problems with the database itself now impact seismic monitoring.

6.2 How did this happen?

Between August 2000 and May 2003 my authority for the seismic monitoring programme was repeatedly eroded. This was particularly so in early 2001, when the Electronics Engineer demanded that he take authority for all Operations Room equipment and operating systems used for the seismic monitoring, and the MVO Director capitulated. Technicians spent less and less time at MVO and were under no pressure from the Director to provide vital technical support to the monitoring programmes. Physically threatening and abusive behaviour was often directed at BGS staff, particularly the Seismologist, and went unchallenged. These were far from ideal conditions in which to achieve an urgently needed complete upgrade of failing seismic systems. But with great perserverance, assistance from Software Engineer David Silcott, and vital contributions from Brian Baptie, Simon Flower and others in BGS Seismology, these upgrades were achieved, along with the move to the new observatory, and the development of innovative analysis tools amongst other accomplishments.

Having developed this excellent foundation, being at the new observatory and having a 5 year contract to look forward to, MULTIMO and CALIPSO coming on-stream the future looked really bright. There seemed to be an excellent opportunity to publish work, collaborate widely and do research. I left for vacation in January 2003 looking forward to spending several more years in Montserrat.

Unfortunately, old problems resurface and more seriously than ever before. On my return from vacation in February 2003 a new Software Engineer made a hostile takeover of the seismic systems. He demanded I no longer have any say in the future maintenance or development of the seismic monitoring systems, and being highly abusive and threatening, and unwilling to cooperate or discuss these demands. Over the next 10 weeks his behaviour became more extreme and included spontaneous outbursts in the Operations Room threatening me and using profane language. He also sent an email to all MVO staff accusing me of tampering with his software, when all I had done was fixed a drive mount which was preventing alarms being sent out. Data acquisition, alarm, analysis and archival systems were all compromised as a result of him removing, deleting or modifiying critical seismic hardware, software and data without my consent. I felt public safety had been put at unnecessary risk, given the concerns at that time about a major collapse of the dome, with pyroclastic flows to the north-west, and the potential for volcanic ash above 10000 feet.

The MVO Director throughout this time was embroiled in legal issues, troubled by the volcanic activity and rather eager to escape the island as soon as possible, and do so without further difficulty. Unfortunately that was not possible given what was going on with the seismic monitoring. For several weeks I appealed for him to intervene and hold a three-way meeting, and he felt he didn't have time. By late April 2003 things had reached a level where I felt I would have been neglecting my duties as MVO Seismologist not to inform the Montserrat Project Manager back in the UK of the serious undermining of the seismic monitoring. I feared for my job, but my job was untenable in the face of having no authority to do it, and facing open abuse in the workplace without any hope of intervention by the Director.

Ultimately I did not get the support I needed, and this brought to a close my involvement with the Montserrat project. I stressed at the time how important it would be for me to overlap with my

successor to train him/her up in the complex seismic monitoring systems at MVO, and the need for me to write a report about those systems, but neither request was given consideration. Over subsequent months there was backfilling by highly competent Seismologists, but each was largely unfamiliar with the progress made in recent years, and the erosion of my authority and undermining of the seismic monitoring.

The state of the seismic monitoring in January 2003 compared to that in January 2004 serves to underline this. Figure 1 shows the status of the seismic monitoring systems as of January 2003 (though they had changed little over the previous year), following the move to the new observatory. All computers were on UPS and would automatically reboot critical software following a brown-out. There were backups for all computers, and a diagnostic alarm system that monitored all potential points of failure, updated status information on the remote monitoring webpage, and paged the Seismologist if any problems were detected. A live backup to linux1 (called 'Seisan' elsewher in this report) was linux2 (called 'Seisan4' elsewhere in this report) and its main job was to check if the diagnostic alarm system was running on linux1, and page the Seismologist if a problem was found: so even the diagnostic alarm system had a backup. This provided a robust backbone which should have served MVO superbly for many years.

Near-real-time seismic data was also displayed on the remote monitoring webpage to expediate alarm response. These data included spectrograms, sonograms, helicorder plots, RSAM plots on various timescales, hypocentral plots and summaries of event counts and energy and equivalent magnitudes by subclass on different timescales, as well as graphs of alarms sent versus time. A wide array of Matlab tools existed for data analysis, and there was an extremely robust data archival system, which was made idiot-proof by a wide range of error checking.

Figure 2 shows the status of the seismic monitoring systems by January 2004. The linux1 PC was now unstable thanks to the poor implementation of the MVO database. The linux2, QNX2 and SA242 PCs had all been lost to the seismic monitoring. None of the operational computers automatically rebooted anymore. The detected events and continuous waveform databases were compromised by the loss of all online data around in July or August 2003, and the unstable nature of the seismic acquisition systems. The rockfall location system had been lost. The diagnostic alarm system was no longer maintained, volcano alarms no longer work. The remote monitoring webpage and the extensive Matlab analysis tools had been blatantly deleted. Judging by the last data displayed, these deletions took place on the 27th June 2003, while I was officially at least still the MVO Seismologist (I left on 14th May 2003, but remained the Resident Seismologist officially until 1st July 2003).

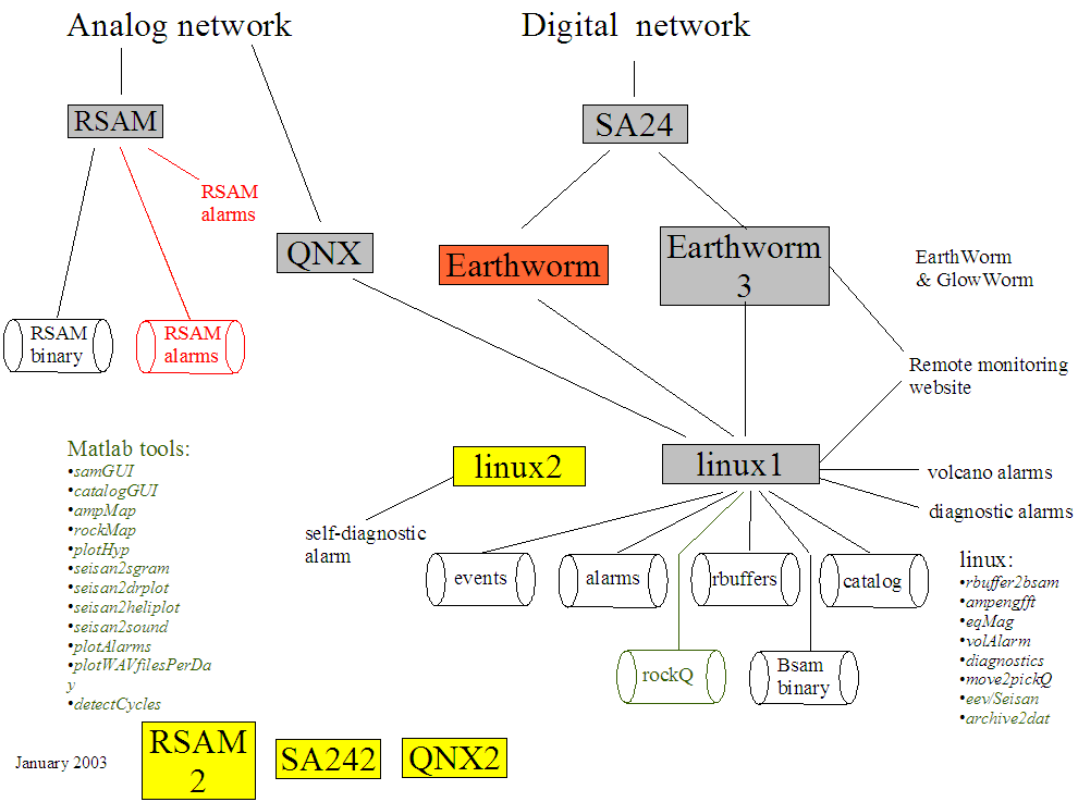

Figure 1: A logical overview of the seismic monitoring systems circa January 2003, following the move to Flemings. The boxes represent computers. The cylinders represent databases or datasets. Black denotes an operational real-time system. Green denotes an operational manually-driven system. Orange denotes a system which was showing instability. Red denotes a system which was compromised. Yellow boxes denote operational or on-the-shelf backups.

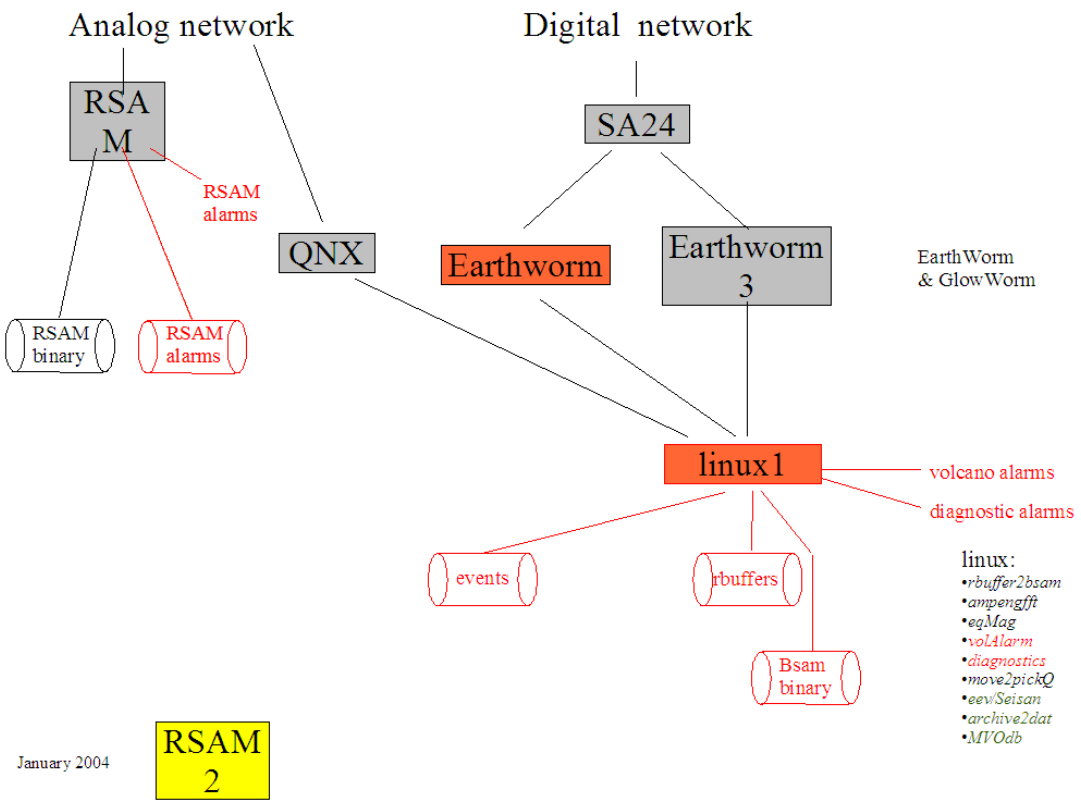

Figure 2: A logical overview of the seismic monitoring in January 2004. The symbology is the same as in Figure 1.

APPENDIX A - DIARY

The following is a diary of problems detected with the seismic monitoring I identified in the period 5-12 January 2004.

05/01/2004 (My first full day at MVO since I left for vacation on 14 May 2003)

Basic Ops Room check:

- MGH and MLY are on the drums.
- ILI is working fine. MBRY is dead. MBLG and MBSS intermittent.
- SA24 is working fine.
- QNX/Seislog is running. MGH, MLY, MJH and MWH are working, though data intermittent. MLG is dead. MRY is no longer deployed.
- Fire-proof data-safe no longer closes. This means its pointless storing tapes at MVO, since if MVO burns, so will the tapes. Until safe is fixed, a copy of the seismic data should be stored at the Seismologist's house.
- EarthWorm PC clock is 4 minutes fast. EarthWorm3 PC clock is 2 minutes fast. Hence event data are corrupted. Instructed Art to modify the latency parameter in EarthWorm module, and time corrected. Need to re-establish procedure for daily checking EarthWorm PC clocks. This used to work with an internet time server.
- Stations being acquired on EarthWorm PC: MBRY, MBSS, MBGH, MBWH, MBGB, MBLG, MBRV, MBLY. But no data being saved because MBRY is dead, and hence corrupt data in EarthWorm wave-server (a problem in SA24 or sa242ew module).
- Stations being acquired on EarthWorm3 PC: MBSS, MBWH, MBGB, MBLG, MBRV, MBLY. But not a lot of data being saved because MBLG and MBSS are intermittent (again, SA24 or sa242ew is to blame – and this is the reason for running two EarthWorm's in parallel, one with just the stations we know are good).
- RSAM PC clock is 3 minutes slow, so data have been incorrectly timestamped probably for many months. Time corrected. Need to re-establish procedure for daily checking RSAM PC clock.
- Seismic monitoring webpage (most useful real-time seismic analysis software) no longer exists. Matlab cronjobs, on EarthWorm3, were disabled on 27 June 2003, when Lars was Seismologist. Why?
- Heli_ewII and Sgram modules are back running again on EarthWorm3. These modules have been unstable in the past and need to be monitored closely.

- Seismologist PC no longer exists. Died in June 2003? Monitor, mouse and keyboard still in Ops Room. Whereabouts of parts unknown. No replacement sought.
- Multimo PC and monitor gone from Ops Room. Presumably in server room.

Drive mounts & Hard drive space (tables not reproduced here, see book):

- Lots of drive mounts on Seisan. This looks bad. Windows drive mounts on this PC in the past have caused whole system to crash, meaning no alarms. Better policy is to use ftp/Kermit to transfer data as needed. The reason for this many mounts is not documented, so not going to change anything. Just need to monitor it. Hard drive space on /mnt/qnx/events is 0GB – not good!
- Drive mounts on Seisan4 PC are mostly the same as on Seisan PC.
- Drive mounts on EarthWorm and EarthWorm3 quite different. Different Seisan drives mounted on each. This is not good. Acquisition PCs should have as few drive mounts as possible, and just use ftp/Kermit for data transfer. We have experienced Linux drive mounts failing in the past, and causing problems with the EarthWorm acquisition software.
- RSAM PC OK.

On Seisan, cronjobs are same as I remember them:

- Getqnxdata – every minute
- Register.all.csh – 3 times an hour
- Seisan2matlab_wrapper WEBUPDATE – once an hour
- Seisan2matlab_wrapper – every Friday (for weekly report)
- Seisan2matlab_wrapper UPDATEALL – every night
- Volcano_alarm – every minute
- Autotests – every minute
- Test_alarm – every day at 7am & 9pm
- Daily_backup – once a day
- Clean old log files – once a day
- Delete old 1 minute rbuffers – once a day
- Bzip2 DSNC_ and MVOE_ - once a day
- Count_wavfiles_per_day – once a day

On Seisan4, there is only 1 cronjob running, called check_seisan. This is essentially an alarm system that checks if the Seisan PC are running (and therefore the diagnostic and volcano alarm systems are running). But the program is currently set up to run from /live, which is a mount from

the Seisan PC. So if the Seisan PC goes down, the program will not run! THIS MUST BE CHANGED !!

On EarthWorm PC and EarthWorm3 PC there are two cronjobs running:

- Gwstatus – every minute
- Transfer_seisdata – every minute
- The Matlab cronjobs on EarthWorm3 that updated content on the seismic monitoring webpage have been disabled (the webpage is no longer accessible either).

Multimo project:

- VNC icon on EarthWorm PC connects to Multimo server (172.20.0.7). Multimo PC clock is 2 minutes fast. Time corrected.

Seismic networks (notes from Art):

Multimo:

- Brian Baptie replaced Multimo CMG-3T (broken) with MVO CMG-40T in vault. When was this? June 2003?
- Station failed on July 12 (ash?).
- Art got seismics going again in October. Scream is solid.
- Still no pressure sensor data.
- Multimo data downloaded every night from MVO to Udine.

MVO:

- Everything failed after July 12.
- All back except MBLG, MBRY by mid-September.
- Silver Hills repeater problem for MBRY.
- MBLG pressure sensor – should have DC offset but doesn't. Rubber ducky waveforms look reasonable. (There were signals with July 12 which presumably Art wants to calibrate).
- MLGT – battery problem?

Acquisition problem:

- No Sfiles for 2004 (Venus pointed this out). Found problem is that location of Seisan database was changed, and entire directory structure for 2004 onwards had not been replicated. Modified /development/RegisterWAV/register.new.pl to create directories as necessary.
- IP Addresses (information from ping command) 172.20.0 + .27 SA24, .22 EARTHWORM, .24 EARTHWORM3, .25 SEISAN, .26 SEISAN4, .29 RSAM, .7 MULTIMO, .80 CALIPSO, .73 PROGRAMMER, .46 PHOTOPC

06/01/2004

Seismic upgrade:

- Discussion with Art about acceptable latency on drums with Guralp system. I suggested digital audio signal, in combination with Scream display, should suffice. It's the audio signal that really alerts Ops Room staff, and it should be trivial for Guralp to extend Scream to put out an FM audio signal (its trivial in Matlab).
- Alternative may be to keep MLYT going with scanner. But think that should only be a stop-gap measure. Better to be able to get audio signal from ANY station at will.

Backup program:

- Venus demonstrated backup of December 2003 rbuffers. End of tape had been reached, and no data were being saved, but backup program carried on regardless and gave no warning. THIS IS VERY POOR AND NEEDS TO BE CHANGED. Previous data archival software verified tar files as it wrote, and checked for end of tape conditions.

Acquisition problem:

- No QNX events collected since 16 December 2003. Took a long time to troubleshoot. Found that /home/events on QNX randomly disappeared, or came up as a bad directory (or "Stale NFS file handle"). Recreated several times. However, once I removed /mnt/qnx/events from /etc/fstab on Seisan, the problem disappeared. So somehow by mounting QNX drive under Linux, QNX filesystem gets corrupted, and events don't get recorded.

07/01/2004

- Still no EarthWorm events (turned out to be EarthWorm acquisition problems).

- System clock on EarthWorm3 reset incorrectly, presumably by network time server. Could interfere with data acquisition.
- Rbuffers not accessible through eev, but okay with mulplt. Problem with Sfiles.
- RSAM plots on Alternative website are not being updated (past new year 2003). Why?
- Counts plots on web (via Lars' volcstat program) have stopped. Why?
- Had chat with Art. Lots of problems with seismic systems discovered so far. Not supposed to be here to fix software, but I can do it. What does he want? We agreed best compromise would be for me to fix problems that were skin deep, and make recommendations about deeper issues for the software engineers. I then cleared this with Bill and Pete.

Spares:

Showed Art location of:

- ILI cables
- SA24 card
- RSAM card
- DAS1202 (QNX) card
- PC-SEIS card
- GPS clock

Cronjobs revisited:

- Testing seisan2matlab_wrapper.pl: Problem: /data3/MATLAB_DATA/ has been deleted. Later found this was removed in July 2003 a few days after I left. This would have caused all Matlab software to crash. Directory, and datasets recreated. Fixed.
- Test_alarm doesn't seem to work: wrong numbers were in Expect script. Cable & Wireless service is intermittent. These should be in a configuration file.
- Autotests is sending alarms about: (i) No EW events, (ii) Backups mount down, (iii) Alarm system not running. Checks need to be updated to reflect current weaknesses in seismic systems.
- Count_wavfiles_per_day: Also not working because /data3/MATLAB_DATA removed in July. Fixed.
- Daily_backup: not working – directory missing. What is being backed-up doesn't make sense.

Hardware:

- No dial-in PC (given to Graham)
- 1 x 500 GB drive missing.
- 2 x 250 GB drives missing (one later found, but without case).
- No Seismologist PC.
- Spare GPS clock need fixing (doesn't have correct ports for QNX & SA24, but those systems should be obsolete soon).
- Need drum time signals back.
- None of my old Matlab software is running. Data streams all stopped 27 June 2003, suggested the software was all deleted on that date.
- Went through and re-established drive mounts, load_rsam, bob, and bobsgui.
- Strategy (not on this trip) should be to create a Matlab application for plotting all MVO datastreams. Much easier to maintain than a database.
- Found no BSAM data for 2004. Found problem with data acquisition for 1-minute rbuffers.

08/01/2004 Field trip

09/01/2004

- Autoreboot test – which acquisition systems will autoreboot in case of power failure? All acquisition PCs used to do this OK. How about now?
- EarthWorm – YES
- EarthWorm3 – NO (need manual logon, then manual start GW)
- RSAM – NO (have to logon as seisan)
- SA24 – NO (have to logon as seisan, and then manual start SA24)
- EarthWorm and EarthWorm3 – shortcuts to scripts to easy reboot GW software, and wave-viewer recreated. Had been deleted from Desktop since May, leading to complex procedures.
- Pager test – Pete supplied me with 4 old pagers. Numbers unknown. I tried all numbers we previously used. Found number to 2 (494 9790, 494 9514). But neither work reliably in observatory. Later found that C&W regard pager service as obsolete.
- EarthWorm acquisition problems: Modified lists of acquired stations to:
  - ➔ EarthWorm PC: MBSS, MBLG, MBGH, MBWH, MBGB, MBRV, MBLY (all stations)
  - ➔ EarthWorm3 PC: MBGH, MBWH, MBGB, MBRV, MBLY (good stations)
- Testing RSAM alarm system (Many hours). Works intermittently with cellphone.

- Seisan PC kept hanging today. Wouldn't reboot. Problems with QNX drive mounts were to blame.
- Also problems with transfer_data.pl not connecting to Seisan. ALARM SYSTEM NEEDS TO BE ON EARTHWORM!!

12/01/2004

- Seisan down for 36 hours. No diagnostic alarms sent. Appallingly fragile system.
- Many more reboots of Seisan and Seisan4. Problems with QNX drive mounts to blame. Removed from /etc/fstab on both systems.
- SEISAN4 STILL NOT SET UP IN PARALLEL WITH SEISAN
- Switched MGH drum to MWH, since MGH dominated by truck noise during day.

Daily troubleshooting work went on like this, with numerous problems with data acquisition, volcano alarm and diagnostic alarm systems, until 19 January 2004, when I got things stable. All these problems had been present for weeks/months, but had not been recognised.