The Research Council of Norway (NFR)

NORSAR

NORSAR Scientific Report No. 2-93/94

Semiannual Technical Summary

1 October 1993 – 31 March 1994

Kjeller, May 1994

APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED

7.2 Monitoring ESAL time-lag

Background

The Intelligent Monitoring System (IMS) is a software system of "a new generation of systems for automated and interactive analysis of data from a network of seismic stations to detect and locate seismic events" (Bache et al, 1991). A version of IMS that accepts data from an arbitrary number of arrays and single 3-component stations was installed at NORSAR in October 1991. The system currently processes data from the 6 arrays ARCESS, NORESS, FINESS, GERESS, Spitsbergen and Apatity.

One of the objectives of the IMS development was to improve the quality of the automatic event solutions and thus reduce the need for interactive, manual analysis. In addition to being accurate, the automatic solutions from IMS have to be quickly available to be useful in monitoring underground nuclear explosions. Ideally, the system should be operated in close to real-time to provide up-to-date explosion (and earthquake) alarms.

When IMS failed to define the Chinese test explosion on 5 October 1993 until 8 hours after the event occurred, an initiative was taken to start recording the time-lag in the ESAL (Expert System for Association and Location, the knowledge-based system which locates events in IMS) processing, as well as identifying the cause for each of these lags. It was expected that such an exercise would provide considerable insight into the complexity of the task of keeping a system like IMS running at close to real-time. In the following, the results from the first six months of monitoring time-lags (October 1993 - March 1994) are reported. The first three weeks have also been described in detail earlier (Baadshaug, 1993).

IMS architecture and design

Fig. 7.2.1 gives a broad outline of the IMS components relevant to this discussion.

The waveform data from the seismic stations are transferred to circular diskloops on acquisition computers. Signal processing programs extract detection parameters and store them in the SigPro database. Up to this point, there are separate programs for each station and all stations are processed in parallel. From here, data from all stations are processed together. Once an hour, detections from the SigPro database are copied to a second database used by ESAL. If one of the stations for some reason falls behind, the detection copying waits for that station, i.e., only detections with arrival time less than or equal to the last detection from the delayed station are copied. (See example in Tables 7.2.1 and 7.2.2) The progress of ESAL is vulnerable to all preceding processing steps. A delay at one station will propagate down the processing pipeline and delay event location.

This processing scheme is used because the IMS results are mainly used for research and bulletin production. The emphasis has been on completeness (wait for all detections from all stations) rather than speed (skip a station if it is too far behind). These are often conflicting criteria and in a continuous monitoring situation, the focus may have to be shifted to speed to provide early event notices.

As can be seen from Fig. 7.2.1, there are several hardware and software components in the system and therefore lots of error sources. In the current implementation, there is no redundancy or duplication of equipment. This may also have to change in a continuous monitoring situation with tight requirements on system reliability.

Collecting the time-lag data and checking ESAL progress

Every ten minutes, the UNIX system time (based on a GPS satellite clock) is written to a file together with the detection time last processed by ESAL (found in a database table - timestamp). The difference between these two times is the ESAL time-lag, i.e., the number of seconds ESAL is behind real-time.

In Fig. 7.2.2, the time-lag is shown in six plots, with one month of data in each.

The maximum time-lag was 204241 seconds, close to 57 hours, on 14 December 1993, when the communication to Apatity was down because of a satellite dish positioning problem.

The minimum time-lag was 868 seconds or 14 minutes 28 seconds. This should not be interpreted as a delay one can expect to see during normal operation with the current setup, but rather as a result of a combination of favorable circumstances.

The median time-lag was 4770 seconds or approximately 1 hour and 20 minutes.

Table 7.2.3 shows examples of common reasons for delay.

In addition to the time-lag data, full reports of system progress are collected at less frequent intervals. These reports include the status of all processes which ESAL depends on: The last data time recorded for each station, the last detection processed, etc. Normally, it is possible to determine from these reports what caused ESAL to fall behind. This is usually done through a backward-trace where each processing step is checked in reverse succession:

1) Discover that ESAL processing is not up-to-date

2) Check if the latest detections have been copied from the SigPro- to the ESAL database.

3) Check if some signal processing job is delayed.

4) Check if the last recorded raw data are current.

It is not always possible to determine why ESAL falls behind. There may be no errors reported in log-files and the delay may have happened between two instances of the full system report.

An example showing both the backward-tracing of processing steps and some unexplained delays is seen in Figure 7.2.3. This example from April 1994 (after the reporting period) illustrates the mentioned points. The weekly ESAL time-lag plot had several peaks, meaning that location processing had fallen behind real-time. To find the reasons for the delays, the delays of the signal processing for each station were plotted. It was found that all delays were caused by the Spitsbergen SigPro. This shows how a single station can hold up all processing. The reason for the delays on April 5/6 and on April 10 were found to be large numbers of detections, overloading the computers where the signal processing software runs. The peaks on April 4 and April 8, however, remain unexplained. The one on April 8 is most likely caused by some network problem, as the SigPro programs for two other stations stopped as well. On April 4, there was no clue to the reason for the delay. The detection processor slowed down for no obvious reason and then caught up again by itself.

Problem classes and possible remedies

The reasons for delays in the ESAL-processing can largely be divided into four groups, each with their own possible remedies:

1) **Hardware**. Caused by errors and maintenance on computers, disks and communication equipment.

As mentioned above, there is no redundancy in the current IMS implementation. One way to minimize hardware downtime, would be to duplicate equipment. Fault-tolerant computer systems are commercially available with different degrees of redundancy. Parallel execution of the same program on several CPUs, disk mirroring and alternative network routes are measures that could improve the IMS uptime and reduce the time-lag.

2) **Software**. Stops due to program bugs or upgrades.

At the moment, programs in the IMS system are started manually. When a program aborts on an error, it will stay down until manually restarted. The downtime could be significantly reduced by implementing Manager programs that will detect stops and do an automatic restart. This may not, however, be desirable. Some error situations should cause a program to halt until an operator has cleared the underlying cause. A Manager program may also introduce another error source into the system as happened at the aforementioned Chinese test explosion on 5 October 1993. At that time, the detection processors for each station were checked automatically (by a program in the UNIX crontab file) on an hourly basis. This check failed, i.e., a running detector was seen as stopped, and a second detector was started. This overloaded the computer, causing the detection processing to fall hours behind in addition to creating duplicate phase detections.

3) **Data**. Missing or delayed waveforms. Most of these delays can usually be traced back to malfunctioning hardware, and belong under 1) but some are results of field work, power outages or weather conditions (e.g snow on satellite dish) and can be treated as a separate category.

Smarter or more flexible processing algorithms would handle the situations not caused by faulty hardware. At the moment, no event location or association is performed until data are available from all stations. A processing scheme where a station is left out of the processing if it is more than a predetermined interval behind, would allow events to be formed at a closer to real-time rate. Such a scheme could be adopted, and would be in accordance with the GSETT-3 IDC Alpha/Beta station processing where stations that do not have data available in real-time could contribute in the next refinement of the event solution.

4) **Seismicity**. Large number of detections, often because of temperature changes resulting in ice-cracking or thawing.

Since each phase detection initiates several processing steps (quality assurance, beaming, onset estimation, broadband FK-analysis, polarization analysis and numerous database accesses), bursts of detections will slow down the signal processing. The number of detections vary both with the time of day (working hours, cultural noise) and with the season (ice-cracking, see example below). From the start, the signal processing programs for several stations used the same SUN SparcStation 1 computer. During periods of normal seismicity, this worked well, but peak periods made the machines struggle. This problem was greatly reduced when a SUN SparcStation IPC was allocated to each SigPro process, but can probably be eliminated altogether if a better distribution of processing load is introduced. This would call for some kind of Manager program, similar to the one outlined under 1), that would detect processing bottlenecks and move the struggling programs to higher capacity computers.

At ARCESS seasonal variations related to ice-cracking/-thawing have been observed (Fyen and Hansvold, 1992) which can make the number of detections rise to 1200 - 1500 a day compared to a few hundred on a normal day, see Fig. 7.2.4a). When the ARCESS signal processing ran on a SUN SparcStation IPC, the processing load of these bursts would make the processing fall as much as 19 hours behind (See Fig. 7.2.4b). After a SparcStation 10 was introduced in the system and used for ARCESS processing, this problem has disappeared. This machine will not help when other arrays have detection bursts, unless their signal processing are moved there manually.

There seem to be two possible ways to handle increased number of detections: Upgrades to faster processing hardware or introduction of automatic distribution of the processing load.

Relevance to continuous monitoring.

We have only studied time-lags in the ESAL processing. This gives us a measure of the delay in the automatic definition of events. In a continuous monitoring situation, other delays further down the processing pipeline may be equally important. Retrieval of raw waveforms from the diskloop machines, beamforming and event plotting are tasks which have to be up-to-date to allow early event review by a human analyst.

In a monitoring situation, there will probably also be stricter limits on the delays allowed. The current deadline for including an event in the GSETT-3 IDC Alpha Event List (AEL), is about 50 minutes after recording (IDC staff, 1994). With the inherent delays in the IMS processing implementation at NORSAR (Baadshaug, 1993), lots of events will not be defined until two hours after recording, even in normal, error-free processing, and would miss this deadline.

Stead (1994) explains how the time limit policy has been enforced at the GSETT-3 IDC:

"Originally, the system functioned with a drop-dead time. This was enforced by requiring esal to run on pre-determined time intervals at pre-determined times (using cron). For example, the AEL was run every 20 minutes for a time window 20 minutes long ending at the GMT when it was run. With a 30-minute lookback, this means that the oldest data esal could look at would be detections associated with an event that had detections within 50 minutes of real time. No new events could be formed unless all the detection were within 50 minutes of realtime (at the start of the esal run). Esal normally runs in less than 10 minutes, so the events were all but guaranteed to be in the database before the hour."

In January 1994, the IDC adopted a new model where time segments of possibly non-uniform length are generated for ESAL automatically. There may also be more changes. The drop-dead times (no data will ever be included in a bulletin if it arrives late) still in use may give way to targets (IDC only tries to get the bulletin in final form by the deadline, but may be delayed, and may include late-arriving data).

When deciding which model to use, issues discussed above should be relevant. If the waveform data arrives on time but the signal processing computer is overloaded and falls behind because of detection bursts, should the detections be included? They probably should, but a decision has to be made regarding what to do with this and other types of late-arriving data.

U. Baadshaug
B. Ferstad

References

Baadshaug, U. (1993): ESAL time-lag statistics, in *NORSAR Semiannual Tech. Summ. 1 Apr 1993 - 30 Sep 1993*, Scientific Rep. No.1-93/94, Kjeller, Norway.

Bache, T.C., S.R. Bratt, J.W. Given, T.D. Schroeder, H.J. Swanger, J. Wang (1991): The Intelligent Monitoring System Version 2, *NMRD Quarterly Technical Report #7,* SAIC-91/1137, San Diego, CA.

Fyen, J. and K. Hansvold (1992): Correlation between temperature and number of detections, in *NORSAR Semiannual Tech. Summ. 1 Oct 1991 - 31 Mar 1992*, Scientific Rep. No. 2-91/92, Kjeller, Norway.

IDC Staff (1994):, *IDC Performance Report, March 5 - 18, 1994*, March 23, 1994

Stead, R (1994): *Personal communication*

Station	Detection time
NORESS	1994-050:11.58.17
GERESS	1994-050:11.58.00
ARCESS	1994-050:11.55.49
Spitsbergen	1994-050:11.52.00
Apatity	1994-050:08.29.25
FINESS	**1994-049:22.58.30**

Table 7.2.1. Times of the last processed detection for each station. (Taken from the **sigpro_time** table in the sigpro database). FINESS has fallen behind real time.

Process Class	Process Name	Time
ARS	ANALYST	1994-043:00.00.00
DetMag	ANALYST	1994-043:00.00.00
EvPlot	ANALYST	1994-043:00.00.00
bull	ANALYST	1994-043:00.00.00
bull	NRSN	1994-049:00.00.00
Beamer	APA0	1994-049:21.58.30
GetData	APA0	1994-049:21.58.30
Beamer	ARA0	1994-049:21.58.30
GetData	ARA0	1994-049:21.58.30
Beamer	FIA0	1994-049:21.58.30
GetData	FIA0	1994-049:21.58.30
Beamer	GEC2	1994-049:21.58.30
GetData	GEC2	1994-049:21.58.30
Beamer	NRA0	1994-049:21.58.30
GetData	NRA0	1994-049:21.58.30
DetMag	NRSN	1994-049:21.58.30
EvPlot	NRSN	1994-049:21.58.30
Beamer	SPA0	1994-049:21.58.30
GetData	SPA0	1994-049:21.58.30
GET_DET	NRSN	**1994-049:22.58.30**
ESAL	NRSN	1994-049:22.58.30

Table 7.2.2. Times last processed by IMS programs, taken from the **timestamp** table in the IMS database. 'ARS ANALYST' shows how far manual analysis has advanced. The 'EvPlot' entries give the time of the last analyzed ('ANALYST') and automatic ('NRSN') event plotted. 'GetData <station>' gives the time of the newest data retrieved to the IMS from the diskloop. 'GET_DET NRSN' is the GetArrivals-process (see Fig. 7.2.1) that fetches detections from the EP-SigPro database to the IMS database. It never proceeds beyond the oldest time in Table 7.2.1.

Station/Program	Error	Reason
ESAL	Machine down	Changed disk on tjalfe
ESAL	Slow processing	Unknown reason, network problem?
ESAL	Stopped	Bad Apatity detection
GetArrivals	Failed (twice)	
IMS	Database down	Parity error on vile
Sigpro	Database down	Changed disk on ve
apatity	Delayed data	
apatity	No data	Satellite dish positioning problem
apa_ep		Number of detections, gaps in the data
arcess	No data	Snow on antenna?
arcess_ep		Number of detections
arcess_ep	Stopped	Changed disk on rein
finess	No data	
finess_ep		Diskloop read error
finess_ep	Stopped	Bad data
geress		Communication down
geress_ep	Beamforming failed	Sonic boom
noress_ep		Diskloop read error
spitsbergen	Delayed data	
spitsbergen	No data	Broken power supply
spi_dp	Delayed	Unknown reason
spi_dp & apa_ep	Delayed	
spi_ep	Stopped	Bad data?

Table 7.2.3 Examples of reasons for stops or delays in ESAL processing

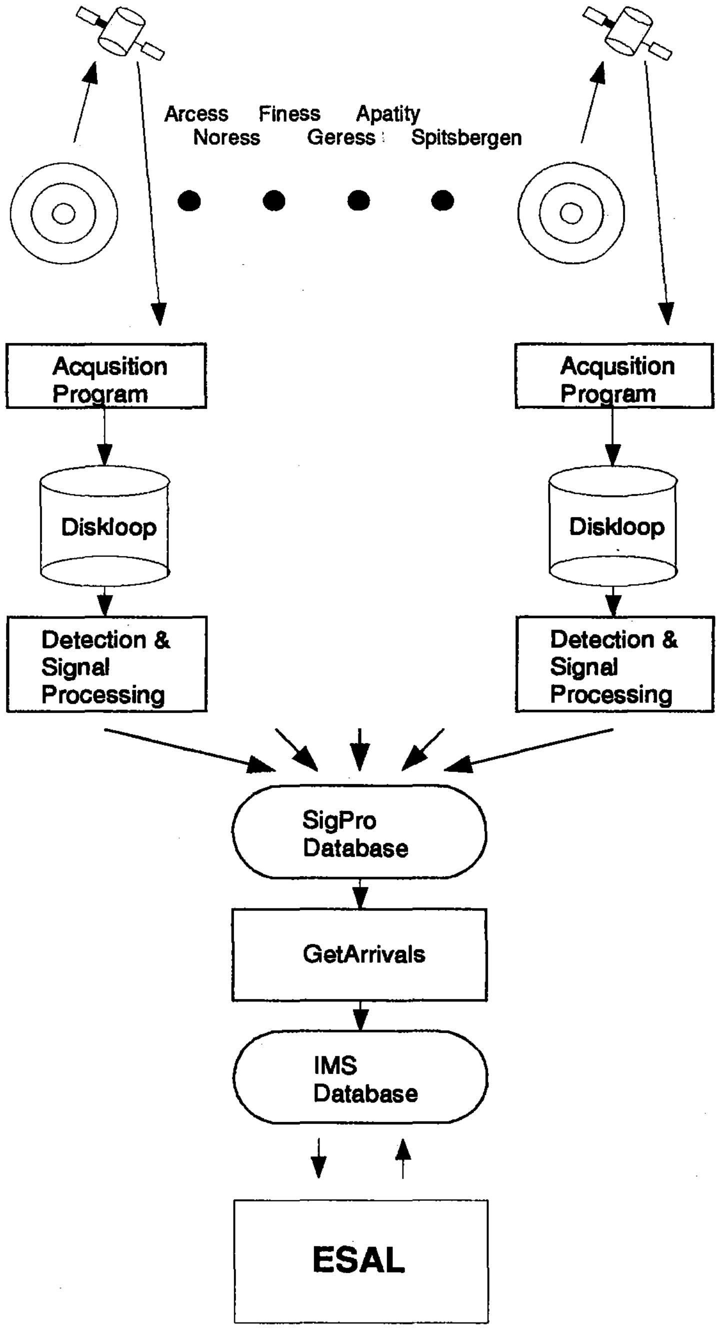

Fig. 7.2.1. Major IMS hard- and software components from seismic station to event association and location. All stations are processed separately, but in parallel, until the detections are entered into the SigPro database. From there, the detections from all stations are processed together.

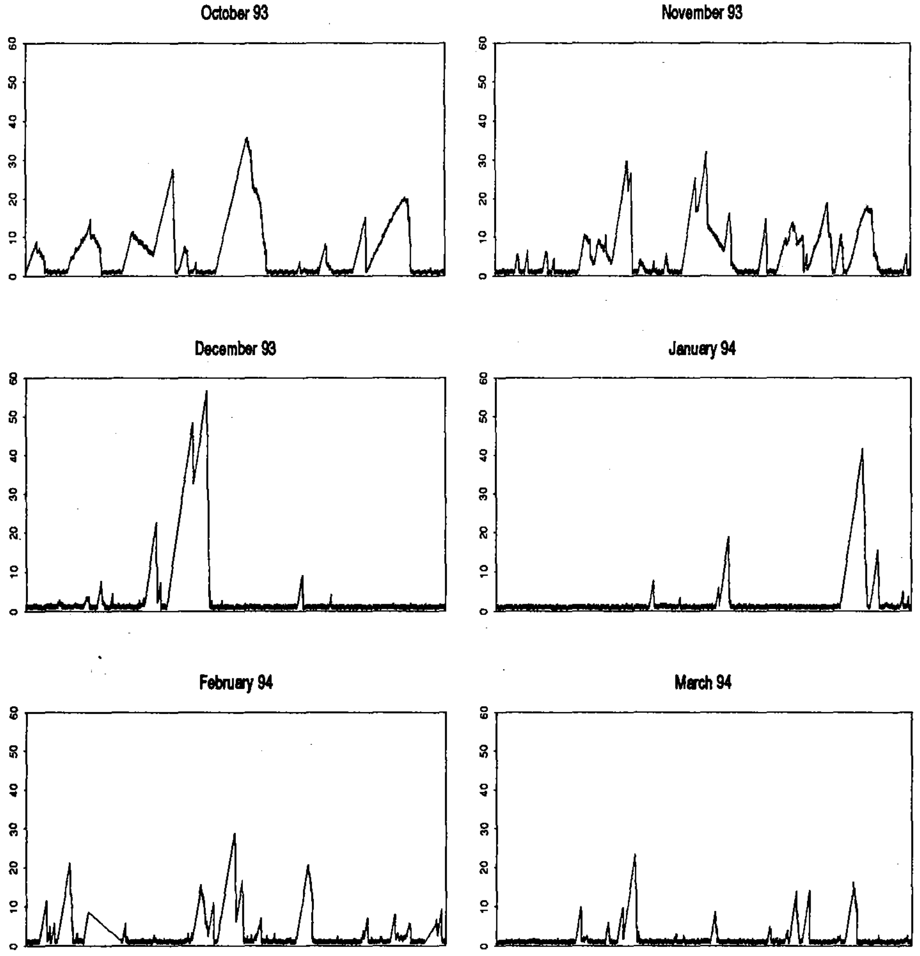

Fig. 7.2.2. ESAL time lag, October 1993 - March 1994. Each plot contains one month of time-lag data and shows the difference in hours between ESAL processing time and real time.

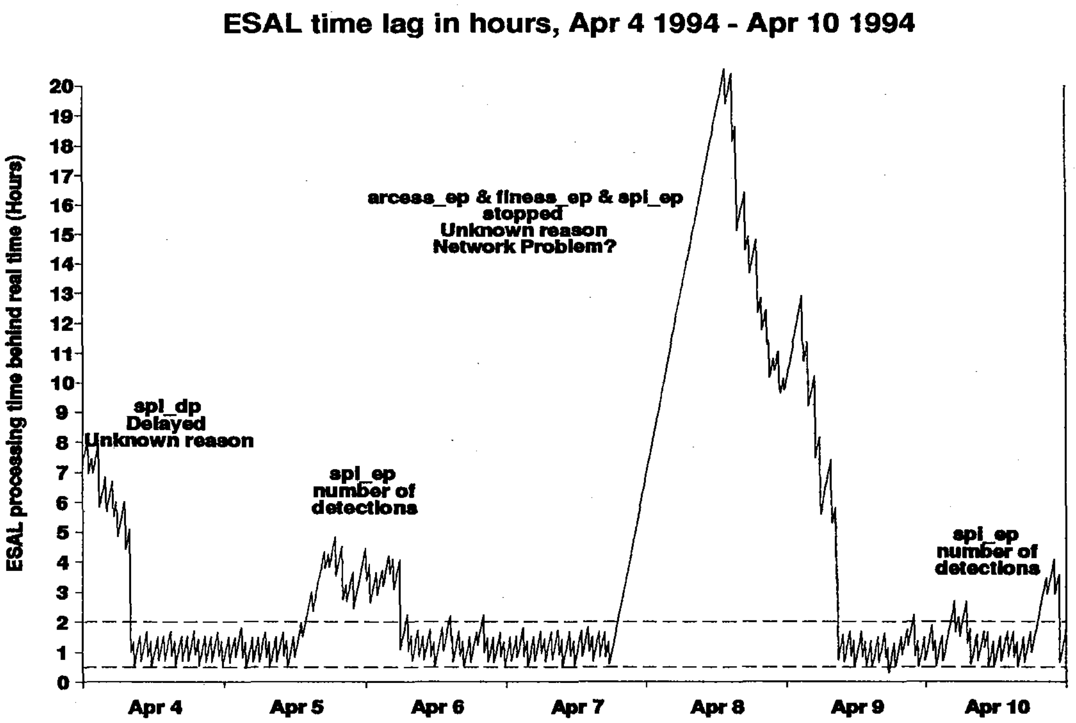

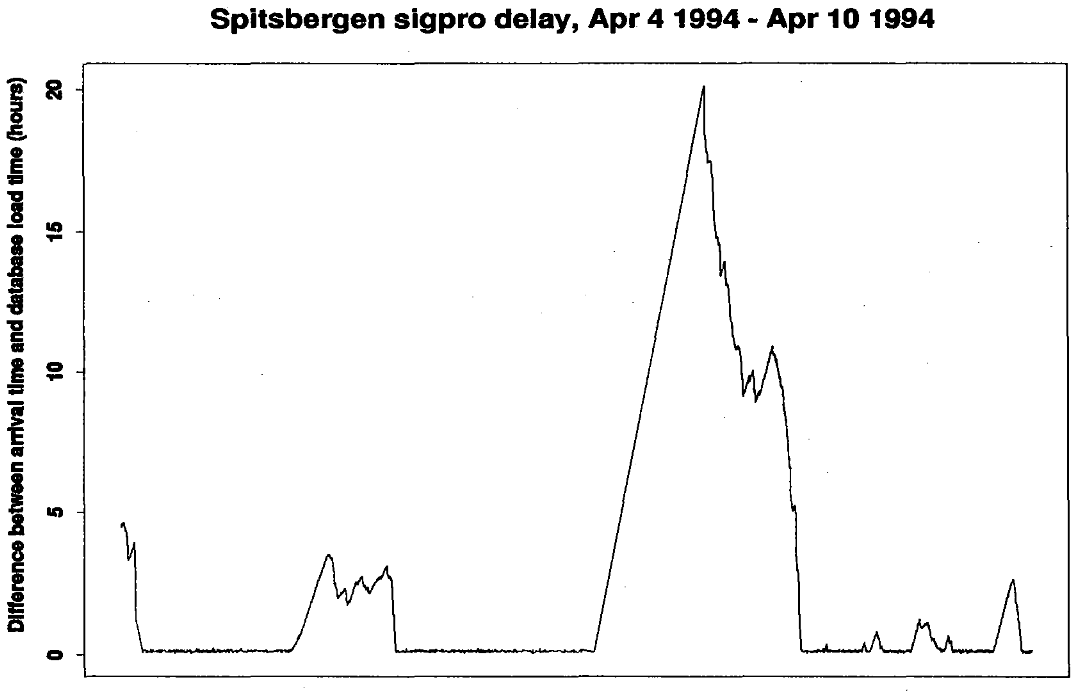

Fig. 7.2.3. ESAL time lag, April 4 - April 10, 1994 (top) and Spitsbergen signal processing delay for the same period (bottom).

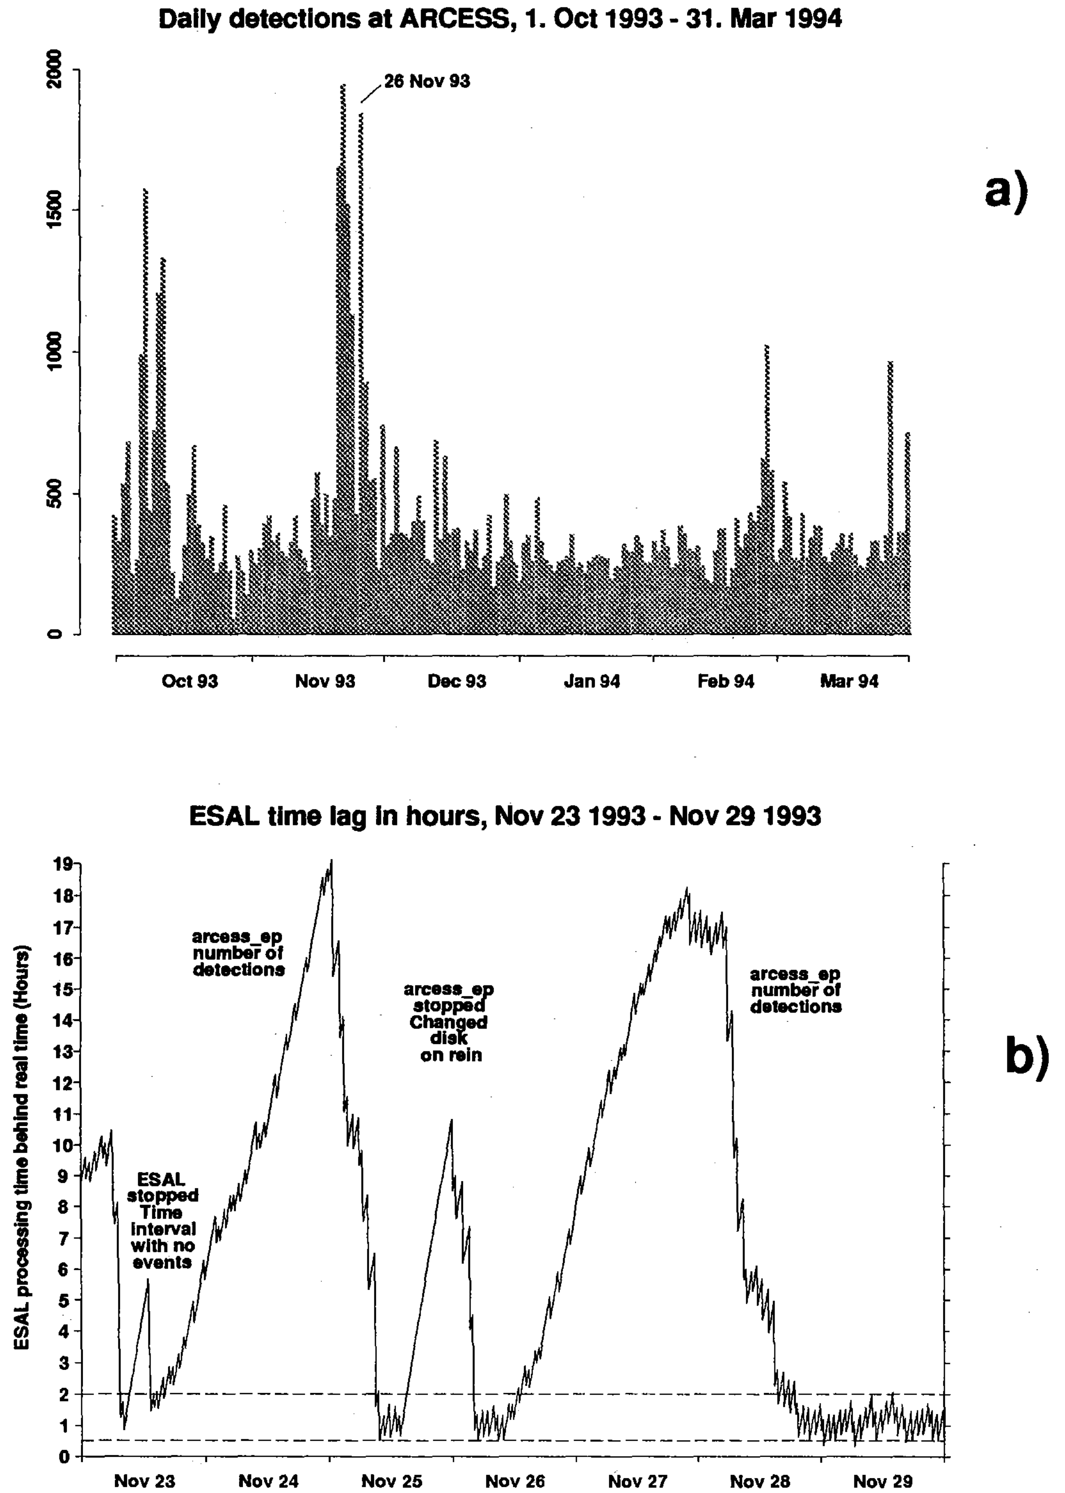

Fig. 7.2.4. a) The number of detections per day at ARCESS for the reporting period and **b)** ESAL time lag for the period with many detections in late November 1993.