The NTT Service Observing Programme: Period 60 Summary and Lessons Learned

D. SILVA (dsilva@eso.org), Data Management and Operations Division

This is the second in a regular series of articles about VLT Data Flow Operations (DFO). In this article, the NTT service observing programme, a VLT DFO prototype, is discussed.

1. Introduction

Between February 1997 and March 1998, ESO ran a service observing programme at the NTT. The primary goal of this programme was to develop and test prototype operations tools and concepts for the planned VLT service observing programme. The NTT programme has been halted for now and service observing is not scheduled for Periods 61 or 62. However, the Observing Programmes Committee and the Users Committee have recommended that NTT service observing resume during Period 63, provided ESO can devote resources to this programme without hindering the development or implementation of VLT science operations. ESO is also studying what is necessary to provide limited service observing support for use of the Wide Field Imager at the 2.2-m starting in Period 63.

Before continuing service observing at the NTT or implementing it at other ESO facilities, it is appropriate to ask: what have we learned? In this article, the concepts and history of the NTT service observing programme are described. Then some lessons learned and how they will affect VLT planning are discussed.

2. Why Service Observing?

When assessing the success of a prototype programme, it is important to keep in mind the original high level goals. Service observing (as previously discussed by Silva & Quinn, 1997) has three main goals:

2.1 Maximise science efficiency

Service observing attempts to maximise science efficiency in two ways. First, at any given moment, the highest OPC ranked programme with Principal Investigator (PI) specified observing conditions requirements which match the current observing conditions is given highest priority. In this way, the most scientifically meritorious programmes are executed first and under the PI defined optimal observing conditions. Remember, however, that no matter how highly a programme is ranked by the OPC, if a PI requests rare conditions, such as exceptionally good seeing, their programme may never be executed if those conditions do not arise.

A second goal is to acquire scientifically useful datasets, as defined by the programme Principal Investigator (PI). For example, in a multi-cluster sample, it may be more useful to have a complete dataset for one cluster than randomly incomplete datasets for all clusters. Over the course of some time interval, these goals should assure that the highest ranked proposals are completed first as long as they do not require rare observing conditions.

A third goal is to facilitate the scheduling of a broader range of science programmes, such as synoptic and Target of Opportunity programmes as well as programmes that require co-ordination between several different facilities. Such programmes are difficult to support in a standard operations model.

2.2 Maximise operational efficiency

Service observing tries to maximise operational efficiency in two ways. First, whenever possible, scientifically appropriate, and consistent with OPC recommendations, observations that require common calibration data are executed as a group so that the calibration data can be shared between several different programmes. Second, experienced staff observers are used to execute observations. Since they are more familiar with the facility, these observers should be more efficient relative to visiting astronomers who only use the facility once or twice a year. Furthermore, most visiting astronomers, no matter how experienced, are less efficient on their first night. For the two- or three-day observing runs typical at heavily used facilities like the NTT, this first night inefficiency can have a significant impact on the overall productivity of the observing run. Staff astronomers, given their familiarity with the facility, can usually minimise this problem.

2.3 Maximise ability to re-use data

An important goal for the VLT and other ESO facilities like the NTT is to maximise the ability of future researchers to re-use data to address different science problems. It is easy to capture all acquired data and store them, but these data will be worthless without suitable calibration data and proper records. Anecdotal evidence as well as a review of the NTT data archive suggests that many visiting astronomers do not perform a calibration and operations plan rigorous enough to produce uniform datasets for future use. During service observing, however, the observatory has total control and can assure that proper calibration data are acquired and sufficient records kept. This process is facilitated by the implementation of calibration plans for each instrument by the scientist responsible for that instrument.

3. Why the NTT?

As part of the NTT Big Bang process, ESO installed a VLT Data Flow System (DFS) prototype. The DFS is a set of software tools and procedures to manage VLT science programmes from the initial proposal to future archival research (Quinn, 1997). These tools were designed with the needs of both service observing and visiting astronomer programmes in mind.

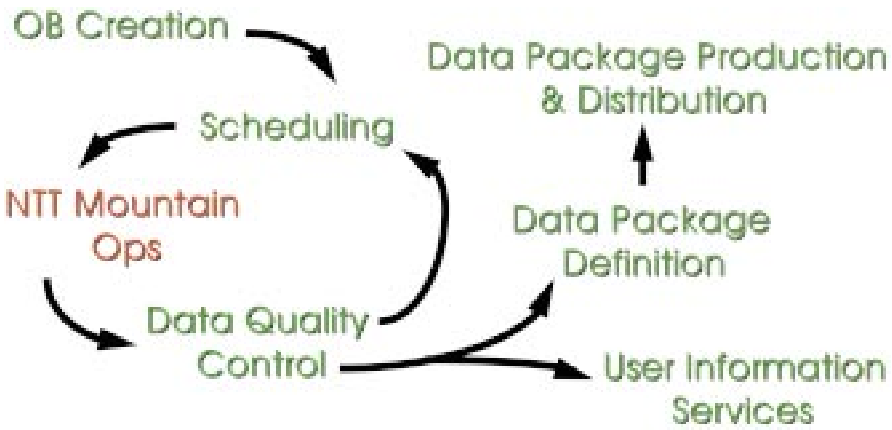

Figure 1: Period 60 NTT Service Observing task flow.

But do these tools work? Do they facilitate or hinder service observing? Do they make general observatory operations more efficient or less efficient? These are issues that ESO wanted to address and resolve before VLT operations began.

Furthermore, VLT science operations are not simply a matter of implementing the right software and hardware. The proper procedures must also be developed and documented. Some of the most important procedures involve providing adequate user support, which involves not only answering user queries and providing adequate general documentation, but also providing the right level of real-time information about observatory operations. Again, ESO saw that it was better to develop these procedures as early as possible, so that VLT science operations would be as efficient as possible right from the start.

4. Periods 58 and 59: Learning the Hard Way

NTT service observing began at the NTT during February 1997. During the second half of Period 58 and the first half of Period 59 (February – June 1997), service observing was executed by the NTT Team on a "shared risk" basis as Big Bang activity and weather permitted. Despite heroic efforts by the NTT Team (led at the time by Jason Spyromilio), operations were plagued by a number of technical problems, as well as a particularly strong El Niño weather pattern, reducing the amount of useful data acquired to a minimum. To make matters worse, users were required to come to Garching to prepare Observation Blocks (OBs)[1], assisted by the DMD User Support Group (led by the author), using prototype software. Several frustrated users who invested many hours in preparing OBs never received any data. Due to poor planning and bad records (e.g. faulty FITS headers), data quality control and distribution by the USG fell behind. Although this period had been advertised as "shared risk" and "best effort", most users were unhappy. Nevertheless, some interesting science programmes were successfully supported, such as the SUSI Deep Field project (D'Odorico, 1997) and follow-up spectroscopy of Galactic bulge microlensing events (Lennon et al., 1997).

The second half of Period 59 (July – September 1997) (often called NTT2 or Period 59.2) was supposed to be a return to full NTT operations. On the mountain, technical operations proceeded much more smoothly, thanks again to the hard work of the NTT Team (led then by Gautier Mathys). Unfortunately, observing conditions continued to be poor, especially the seeing, due to the strong El Niño effect. In Garching, although the OB creation and scheduling process went much better, USG data quality control and distribution activity continued to be very inefficient. Much of the Period 59.2 data was not distributed to the proper end users until December 1997, i.e. 3–6 months after the data were acquired. Simply put, the USG did not have sufficient personnel during Period 59.2 to accomplish its assigned tasks in a timely manner.

These were painful months for everyone, users and ESO staff members alike.

[1]Observation Blocks (OBs) describe how simple datasets are acquired and record the status of those datasets. OBs are simple – they consist of just one telescope pointing and the acquisition of a single dataset. They are the smallest items that can be created, scheduled, and executed by the Data Flow System.

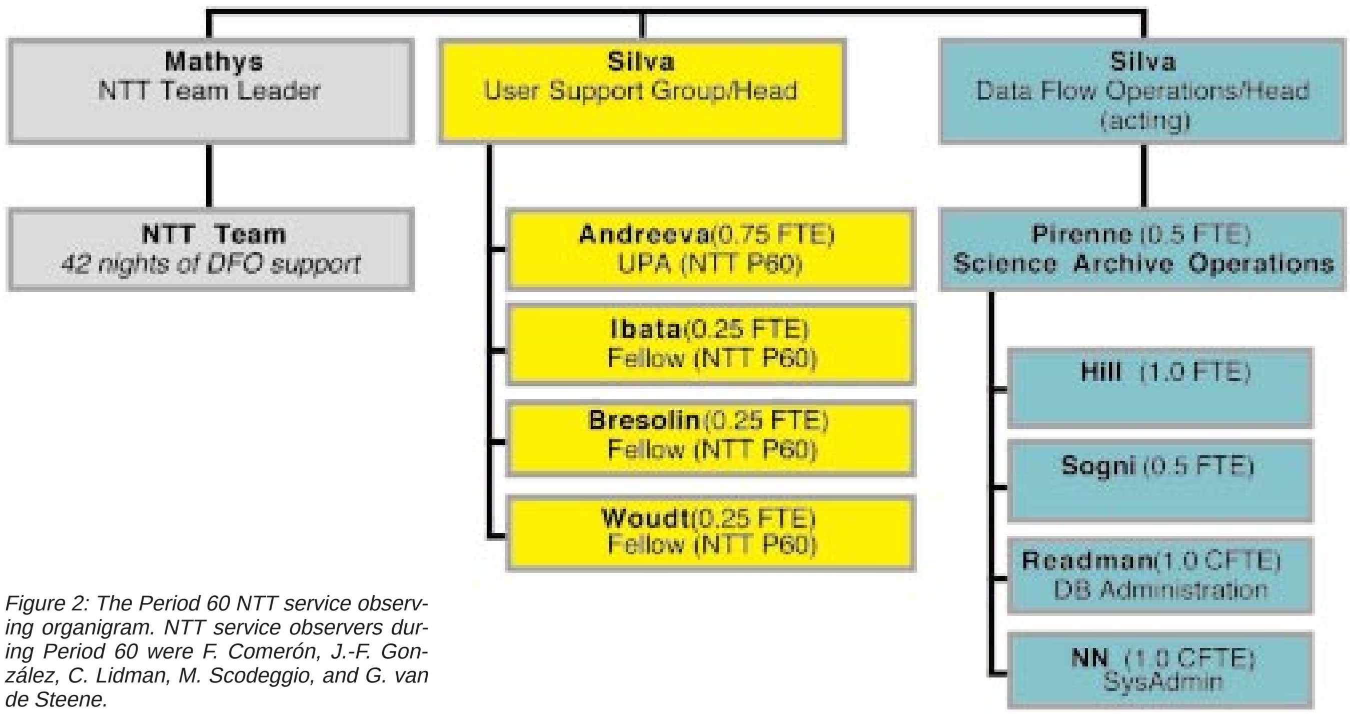

Figure 2: The Period 60 NTT service observing organigram. NTT service observers during Period 60 were F. Comerón, J.-F. González, C. Lidman, M. Scodeggio, and G. van de Steene.

5. Period 60: A New Beginning

In many ways, Period 60 represented a new beginning for the NTT service observing programme. With improved software tools, adequate staffing, and improved weather conditions, ESO was able to provide much better service to our users.

5.1 The Tasks

Figure 1 illustrates the NTT service observing tasks during Period 60.

Observation Blocks (OBs) were created by PIs assigned NTT service observing time using the Phase 2 Proposal Preparation (P2PP) tool. Most Period 60 users were able to create their OBs at their home institutions without significant assistance from the USG. Only one user had to come to Garching and the USG did create OBs for two programmes. After creation, OBs were submitted directly across the Internet to ESO by the P2PP tool for scheduling.

Scheduling was a two-part process. For any given service observing run, the USG created a Medium-Term Schedule (MTS) which contained a list of all executable OBs for the given Local Sidereal Time range, lunar phase, and available instrumentation. The MTS also contained information about relative OPC priorities, user specified observation condition requirements, and required instrument configurations. Typically, the MTS was created less than 36 hours before the start of a service observing run. The second phase of scheduling, creating the Short-Term Schedule (STS), was done by the NTT Team at the telescope. In real-time, the NTT service observer would select an OB from the MTS pool with the highest OPC priority which could be executed under the current observing conditions. For most of Period 60, this was done manually, although during early February 1998, a prototype STS creation software tool was tested.

NTT mountain operations was the responsibility of the NTT Team. Tasks included science OB selection and execution, calibration plan execution, real-time data quality checking, and a variety of record-keeping tasks.

Once acquired, data were transferred back to Garching, usually immediately via FTP transfer. The next day, all data were reviewed and classified manually using a standardised data quality control process. Frames were graded:

- A – all user specifications met
- B – all user specifications met within 25%
- C – all user specifications violated by 25% or more

If an OB (or some part of an OB) failed to pass data quality control, attempts were made to re-execute it later.

Once a logical dataset was collected for any given programme (e.g. all observations in a given filter, all observations of a given cluster), a CD-ROM was prepared containing all the science data plus the relevant calibration data. Users were sent all data taken for their programme, no matter what the quality control grade. In practice, Period 60 data were distributed three times: in December 1997 when SUSI was decommissioned (all SUSI PIs were issued whatever data were available for their programmes), in February 1998, when a number of EMMI programmes were completed, and in April 1998 (i.e. the end of Period 60), at which point all remaining data were distributed. In short, within one month of the completion of Period 60, all data had been distributed, a vast improvement over the Period 58 and 59 problems.

Finally, user information was regularly updated on our Web site, which is still available at http://www.eso.org/dmd/usg. Information posted included nightly activity summaries and updates on the progress of individual programmes.

5.2 The Team

Although many improved tools and procedures were used during Period 60, the most significant improvement was allocating the right number of people. The Period 60 team is shown in Figure 2. Keeping in mind that the DMD Science Archive Operations team was supporting NTT and HST/ECF archive activities as well as VLT archive operations start-up work, the Period 60 NTT service observing programme personnel cost was 2.0–2.5 Full Time Equivalents (FTEs), including all USG, NTT Team, and Science Archive Operations activities.

5.3 The Results

During Period 60, we tried to quantify as many operations processes as possible so we can analyse where we had been successful and where we needed to improve. Here, three examples of the information available are discussed.

In Figure 3, the distribution of OB grades based on the post-observation quality control check is shown. Most Grade B OBs were observed under slightly worse seeing than specified, typically because the seeing deteriorated during an exposure or because these were the best OBs available during a period of marginal seeing. Grade C OBs represent more serious failures: poor tracking near the zenith or sudden change in sky transparency are typical causes. Fortunately, we were able to repeat most Grade C OBs. The generally high success rate was due to the fact that OBs which spanned a large range of observing conditions were almost always available. Rarely did we have problems with not having enough OBs for a given night.

Our records also allow us to accurately account for the time we used on the NTT. Figure 4 shows the fractional breakdown.

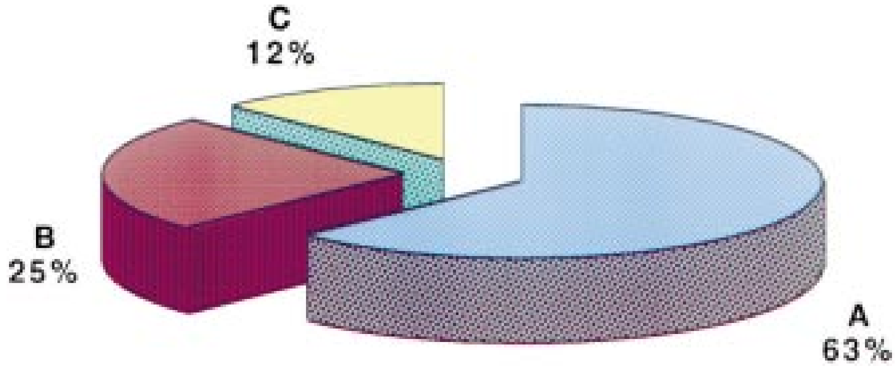

Figure 3: Period 60 OB Grade Distribution. Total number of science OBs executed: 495. See text for grade definition.

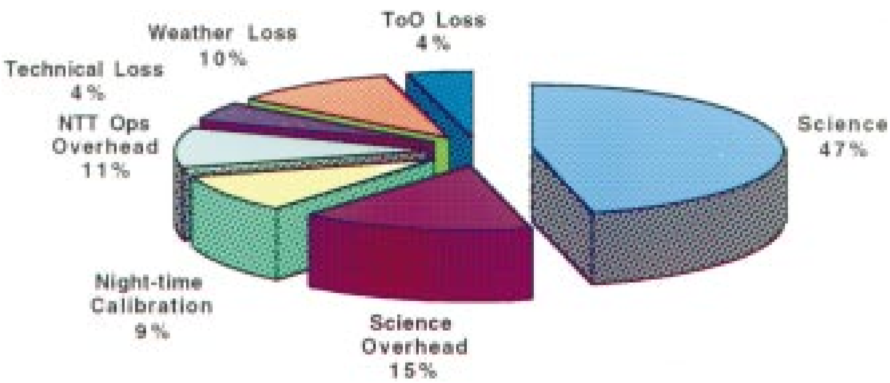

Figure 4: Period 60 Service Observing Programme NTT Time Usage. Science time = integration time only. Science Overhead = slews, CCD read-down, instrument set-up. NTT Ops Overhead = focus and active optics updates. ToO Loss = time used on service observing nights for Target of Opportunity observations (including all overheads).

ProgID	Inst	User Requirements Setup	Moon	See	Sky
High OPC Priority					
59.B-0352	EMMI	REMD	D	1.0	C
60.B-0711	EMMI	RILD/S	D	5.0	TC
60.C-0212	EMMI	RILD/S	G	1.4	P
60.E-0818	EMMI	RILD/I	any	1.4	C
60.A-0129	SUSI		D	0.8	C
60.B-0575	SUSI		any	0.8	P
60.B-0576	SUSI		D	0.8	P
60.C-0757	SUSI		D	0.8	C
Medium OPC Priority					
60.A-0794	EMMI	RILD/I	D	0.8	P
60.B-0740	EMMI	RILD/S	D	0.8	C
60.B-0843	EMMI	BIMG/I	D	5.0	C
60.D-0261	EMMI	RILD/S	G	5.0	P
60.A-0748	SUSI		D	0.6	C
Low OPC Priority					
60.A-0369	EMMI	BLMD	G	1.4	TC
60.A-0609	EMMI	REMD/S	D	5.0	TC
60.A-0700	EMMI	BIMG	D	0.8	P
60.B-0319	EMMI	Echelle	any	5.0	TC
60.D-0496	EMMI	REMD	D	0.8	C
60.D-0682	EMMI	BLMD	G	0.8	C
60.A-0713	SUSI		D	0.8	P
60.D-0326	SUSI		D	5.0	C/P
60.D-0331	SUSI		D	1.4	C
60.D-0682	SUSI		G	0.8	C

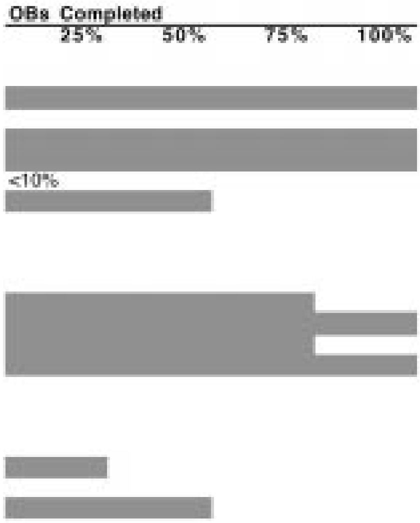

Figure 5: Period 60 Observing Programme Final Status Summary. For Moon (for "moon phase"), D = dark, G = gray. Under See (for "seeing"), these are upper limits on acceptable PSF FWHM in arcseconds. Under Sky (for "sky transparency"), P = photometric, C = clear, and TC = thin cirrus.

Several interesting things are revealed. First, the weather down time was only 10% because during service observing runs there is almost always an OPC approved programme available to take advantage of worse than median seeing or reduced sky transparency. Second, the 15% science overhead, is mostly EMMI CCD readout time, will be more than cut in half when EMMI is upgraded to FIERA controllers. Finally, the NTT operations overheads during Period 60 were mostly caused by inefficient telescope focus and active optics update procedures. However, these procedures have recently been made more efficient – hopefully, Period 61 users have already benefited from this! Finally, it is remarkable (at least to the author) that the NTT is only suffering 4% technical downtime so soon after the Big Bang.

A common question asked is "how many programmes did you complete?" This issue is addressed by Figure 5. Again, this figure has several interesting features. First, no SUSI programme was completed because: (1) many SUSI programmes required excellent seeing; (2) the actual seeing was seldom sub-arcsecond during the first half of Period 60; and (3) SUSI was decommissioned during the first week of December, halting SUSI operations. Second, programmes with less stringent user requirements were more successful. Finally, convincing the OPC to assign a high grade is very helpful! Given this latter point, we did not request OBs from all programmes in the low OPC priority group unless there was a high probability that we could execute some part of their programme. Our projections were not always accurate (i.e. some PIs submitted OBs which we never executed) but we spared some users unnecessary work.[2]

6. Lessons Learned

NTT service observing has been halted until at least Period 63 and ESO is now preparing for the VLT service observing programme. What have we learned from the NTT experience?

6.1 Queue Management

It is important that service observing queues span a sensible range of observing conditions. For example, if on average only 10% of the available time is expected to have seeing less than 0.5 arcseconds, 30% of the available time should not be allocated to programmes that require such good seeing. Or, if good bright-time instrumentation is unavailable, do not assign many bright-time programmes to the service observing queue. Similarly, it may be more desirable to have a much larger relative fraction of dark-time than bright-time assigned to service observing if suitable bright-time programmes are unavailable. Be prepared for abnormal conditions by over-subscribing service observing – not by numbers of targets or hours, but by observing conditions. For example, since the NTT usually delivers good image quality, the ESO OPC tends to allocate a large fraction of time to programmes that require sub-arcsecond seeing. Unfortunately, it was not anticipated that 1997 would be a strong El Niño year and the OPC continued to allocate time to good seeing programmes. As a result, the 1997 NTT service queues contained too many sub-arcsecond seeing programmes and not enough programmes that could tolerate worse than arcsecond seeing.

6.2 User Information Management

Users with programmes in the service observing queue typically want to know the following things: why is their programme not being executed right now; if not now, when is their programme going to be executed; what OBs have been executed to date, and when are they going to receive their data. At the end of a schedule period, if their programme was not initiated or completed, they want to know why. If they had been at the telescope, they would know the answers to these questions (e.g. it was cloudy, there was an instrument failure). In service observing, the answers can be more complicated (e.g. the conditions were never right, your programme did not have a high enough OPC scientific ranking) and sometimes the final answers are not available until the entire scheduling period is completed.

The NTT experience demonstrates that service observing user anxiety can be greatly reduced by publishing the OPC recommended scientifically ranked queue at the beginning of the period and nightly summaries of service observing activity during the period, including programmes serviced, observing conditions, usable time vs. lost time, and updated summaries of individual programme progress. In combination, this information allows individual users to understand for themselves why scheduling decisions have been made about their programme. The basic information desired is actually quite simple and is easily published via the Web (see http://www.eso.org/dmd/usg/). In addition, the VLT Data Flow System will provide tools for users to retrieve information about individual OBs as their programme progresses.

6.3 OB Management

Philosophically, Observation Blocks are supposed to be indivisible objects. Users are instructed to keep OBs as simple as possible – no more than one instrument configuration per OB and a total execution time of no more than one (1) hour, preferably as short as possible. Nevertheless, it is very tempting for users to make OBs more complicated. This may seem more natural to users because during traditional observing runs users are familiar with the need to interrupt and re-arrange their programme in real-time to match observing conditions. Users expect that service observers will do the same. However, service observers are not restricted to executing any particular

[2] Regretfully, no data were acquired for 60.B-0711, despite having a high grade and less restrictive requirements, because the PI did not submit any OBs. But, of course, other programmes benefited from this decision.

programme during the night. They have the freedom to pick observations to meet conditions from a variety of programmes. To make this process efficient, service observers want OBs to be as simple and as short as possible to maximise scheduling flexibility.

Two example OBs illustrate this point: consider OBs which request N imaging integrations through M filters of the same target or N 3600-second spectroscopic integrations of the same target. During service observing operations, such complex OBs are difficult to schedule efficiently. Furthermore, the schedule problem is exacerbated when part of a complex OB produces bad data and the observation must be re-scheduled. Since the DFS is only supposed to re-schedule complete OBs, not parts of OBs, re-scheduling complex OBs reduces the overall productivity of the queue.

Such inefficiencies can be avoided in two ways. First, users must be educated to make OBs which are easy to schedule, so they must be told how OBs are scheduled and given OB construction guidelines. In the examples above, the simplest and easiest to manage choices would have been $M \times N$ imaging OBs and N spectroscopic OBs. Second, the tools for OB construction must be able to make OBs efficiently so the user is willing to make many simple OBs as opposed to a few complex OBs.

6.4 Data Management

One of the most important experiences from the NTT is the need to review and distribute service observing data as quickly as possible. Timely data review, preferably the next day after a service night, is critical for uncovering data problems (e.g. deficient calibration data, instrument performance anomalies) early enough that many subsequent nights of data are not corrupted. Timely review maximises schedule flexibility by providing information about which OBs need to be re-scheduled and which programmes need to be serviced (or can stop being serviced because enough OBs have been completed). Finally, timely review also allows more rapid data distribution. One VLT goal is to have 90% of the data distributed within one month of the observation execution, with the biggest anticipated bottleneck being the transfer of data from Paranal to Garching on digital media. Improved tools to make data quality control activities more efficient, more informative, and more uniform, are planned for UT1.

Efficient data management also relies on correct and complete FITS headers. During the first half of 1997, the NTT headers had a number of problems which exacerbated our data distribution problems. It is obviously critical to stabilise FITS headers early and then rigorously maintain them.

6.5 Operations Management

A requirement for the success of VLT science operations is the need to co-ordinate science operations across transcontinental distances. The NTT service observing programme demonstrated that the DFS will be able to handle the technical co-ordination issues. During 1997, it was possible: for OBs made by users at their home institutions to be submitted electronically to Garching and then automatically forwarded to La Silla; for Medium-Term Schedules to be generated in Garching and transferred electronically to La Silla; for operational problems to be resolved via e-mail; and for acquired data to be automatically transferred from La Silla back to Garching for ingestion into the science archive.

However, the NTT programme also illustrated that not all relevant information is or can be encoded in the OBs, logs, and FITS files transferred between Chile and Germany, that real-time decisions must sometimes be made when the time zone difference and/or workshift inversion make real-time communication between the two groups difficult to co-ordinate; that it is usually easier for the local team to anticipate and solve local problems than the remote team; and that deviations from standard procedures or the original plan must be globally communicated and explained.

Applying these lessons to VLT science operations implies a high degree of personal operational awareness and responsibility from the members of the operations teams to assure success, especially by the on-duty service observer. One of the biggest challenges of service observing is assuring that the service observer has enough information to make good real-time scheduling decisions. Service observers cannot just blindly follow a schedule made remotely – they must be trained and empowered to make the right real-time decisions to use the VLT in a scientifically effective and operationally efficient manner.

Clearly stated, globally communicated lines of responsibility and authority are also important for a distributed operations model. Early operations efficiency at the NTT was reduced because the tasks and responsibilities of the various sub-teams were not sufficiently defined. By Period 60, this issue had been resolved but it is unacceptable for this to happen again at the VLT. During 1998, ESO will be working on defining carefully the VLT science operations roles and responsibilities, as well as the VLT operations command structure.

Finally, task load must be matched to available resources. ESO did not adequately support NTT service observing during Period 59 and the result was deficient user services, particularly in the area of data distribution. We believe the VLT science operations task and available resources are properly matched.

6.6 Operations Tracking

Every oversight committee in the world interested in telescope operations wants to have detailed statistics about service observing programmes, especially in these early days as this technique is developed at ground-based observatories. It is not that they are skeptical, they are just being cautious. Common requests are: total integration time vs. total available time, programme completion status vs. TAC/OPC ranking, programme completion status vs. actual observing conditions, and relative fractions of operational overheads, calibration observations, technical downtime, weather-related downtime, and science observations. Other combinations are possible.

It is far better to anticipate these reporting requirements early and build up the statistical databases progressively during operations than to try to recover this information from observing logs, e-mail, OBs, etc. after the fact. Although not an explicit design requirement, as a by-product of operations, the Data Flow System generates all the relevant raw information. Tools for automatically generating the anticipated reports are planned.

7. Conclusions

We conclude with the obvious question: did we achieve our original high-level goals at the NTT as stated at the beginning of this article? In the area of science efficiency, as Figure 5 illustrates, we did concentrate on more highly ranked programmes and we were successful at completing programmes we executed. We were also operationally efficient (see Figures 3 and 4) – we were able to minimise weather down-time, most of the time we only had to execute an OB once to acquire data acceptable to the user, and we did share calibration data between programmes (although only for imaging programmes). Finally, because we did have a standard calibration plan and we did execute it faithfully, the data delivered to the archive should be quantitatively useful to archival researchers once the data become public.

However, other questions are important as well. Was service observing more or less efficient than standard visiting astronomer observing? Are the PIs who received data satisfied with their data and has it made them more productive? Does service observing increase or decrease the science impact of the NTT? ESO will try to answer the first question by analysing records from visiting astronomer runs. The second question will be addressed later this year via a survey of PIs assigned service observing time. The final question will only be answered over time as we see what and how many papers derived from service observing programmes are published in the scientific literature.

In the end, ESO made mistakes during the Period 58–60 NTT service observing programme, but we also learned many valuable lessons. VLT science operations will benefit significantly from this initial prototype programme.

In the next article of this series, the OB creation process will be discussed.

8. Acknowledgements

The NTT service observing programme was a team effort. During this programme, the NTT Team was lead by Jason Spyromilio and Gautier Mathys, both of whom contributed enormous amounts of their time and energy into this programme. Significant service observing support and operations feedback was provided by the following NTT Team members: Fernando Comerón, Jean-François Gonzalez, Chris Lidman, Pierre Martin, Marco Scodeggio, and Griet van de Steene. In Garching, NTT Data Flow Operations support came from the following DMD User Support Group members: Carlo Boarotto, A. Maurizio Chavan, Gino Giannone, Steve Roche; Petia Andreeva, Fabio Bresolin, Rodrigo Ibata, and Patrick Woudt. Data Flow Operations support also came from the DMD Science Archive Operations Team: Benoit Pirenne, Susan Hill, and Fabio Sogni. As always, Bruno Leibundgut struggled valiantly to keep us on the True Path. Special thanks to Stephanie Cote, and especially, Albert Zijlstra who contributed so much during the grim, early days.

Extra special thanks to all the astronomers who were awarded NTT Service Mode time during P58–P60. Hopefully, all your hard work has paid off in useful NTT data now and better VLT science operations in the future.

References

D'Odorico, S. 1997, *The Messenger*, **90**, 1.
Lennon, D.J., Mao, S., Reetz, J., Gehren, T., Yan, L., and Renzini, A. 1997, *The Messenger*, **90**, 30.
Quinn, P. 1997, *The Messenger*, **84**, 30.
Silva, D. and Quinn, P. 1997, *The Messenger*, **90**, 12.

SOFI Infrared Images of the 'NTT Deep Field'

Deep infrared J (1.25 μm) and Ks (2.16 μm) band images of a 5 × 5 arcmin field centred on 12^h 05^m 26^s; –07 43 27 (J2000) obtained during the commissioning of SOFI (Moorwood, Cuby and Lidman, 1998, *The Messenger*, 91, 9) at the NTT in March 1998 will be made available via the Web (under Science Activities on ESO's Homepage) in early June. The Ks image is shown here in Figure 1. This field contains the smaller region observed with SUSI (D'Odorico, 1997, *The Messenger*, 90, 1) for which visible images are already available on the Web. The infrared images have been constructed from jittered observations totalling 4.3 hours in J and 10.4 hours in Ks and have an average point source FWHM of about 0.75 arcsec. Limiting magnitudes (3σ within a 1.5 arcsec diameter aperture) are J = 24.66 and Ks = 22.87. Full details of the observations and data reduction will be put on the Web together with instructions for retrieving the images.

A. MOORWOOD

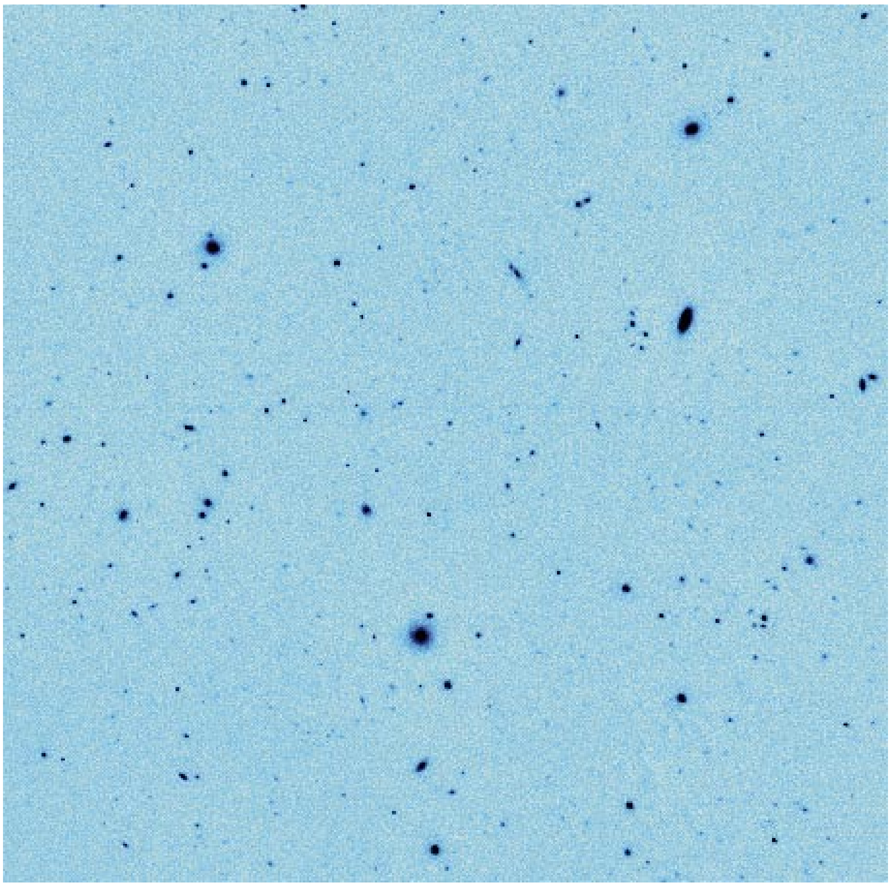

Figure.1: Ks (2.16 μm) image of the NTT Deep Field. The field is ~ 5 × 5 arcmin, seeing is ~ 0.75 arcsec and the 3σ limiting magnitude in a 1.5 arcsec diameter aperture is 22.87 (data reduced by P. Saracco).