SLR2000 Automated System Control Software

Jan McGarry
NASA Goddard Space Flight Center

Jack Cheek
Tony Mallama
Nick Ton
Raytheon STX Corporation

Brion Conklin
Tony Mann
Mattie Sadeghighassami
Mike Perry
AlliedSignal Technical Services Corporation

Randy Ricklefs
The University of Texas McDonald Observatory

Presented at the Eleventh International Workshop on Laser Ranging in Deggendorf, Germany, September 21-25, 1998.

SLR2000 will be a totally automated Satellite Laser Ranging System, capable of making the necessary decisions to protect its health, optimize tracking / ranging performance, reschedule satellites based on cloud conditions, and keep its home facility apprised of its status. To do all of this SLR2000 will rely on new developments in both hardware and software. This paper gives an overview of the software and describes some of the algorithms developed in the areas of signal detection, scheduling, sky clarity, and others. Problems and potential solutions in the areas of backscatter and point-ahead are also discussed. Detailed discussions of the various software and hardware subsystems of SLR2000 are presented in other related papers in this Proceedings.

Background

SLR2000 is the next generation of NASA's Satellite Laser Ranging (SLR) Systems. The goal is to build a totally automated system whose ranging abilities (quantity and quality) are comparable to existing SLR systems. Specifically SLR2000 requirements are:

- unmanned, eyesafe operation,
- 24 hour satellite laser tracking (weather permitting),
- ~1 cm (RMS) single shot tracking,
- minimum of 100 ranges per normal point,
- mean time between failures of > 4 months,
- automated communication with a central facility,

- ability to operate for days without communication, and
- free of optical, electrical, and chemical hazards.

To produce such a system requires the use of eyesafe laser beams instead of the current 100 milliJoule per pulse green Nd:YAG lasers that NASA systems use. To make use of existing detector technology it was decided to continue using a laser operating in the visible spectrum, but reduce the power to eyesafe levels (~100 microJoules). In order to make this system work, a high laser fire rate is required (> 1 kilohertz) to ensure that enough signal returns are received to generate an accurate normal point. A narrow laser divergence also has to be maintained to concentrate the low laser energy on the satellite. This narrow divergence (+- 4 arcseconds) implies the need for very accurate pointing which requires both accurate predicts and the use of a quadrant detector to correct for mount errors. A Signal Detection Algorithm has to be used to distinguish the signal in this very low signal to noise environment. Lastly, all of the functions that an operator normally performs will have to be done by the software, including the following:

- detecting the signal,
- acquiring the satellite,
- maintaining satellite track,
- determining system health and safety,
- calibrating and assessing the system performance,
- scheduling, and
- communicating with a central facility (routine and emergency).

How the SLR2000 system has approached the solutions for all of these requirements will be addressed in the sections below.

Hardware Overview

An overview of the ranging and tracking subsystems will be given here as a background for an understanding of the software. A more complete description of the SLR2000 system is given in John Degnan's paper "SLR2000 Project: Engineering Overview and Status" in this Proceeding.

The ranging system determines the times of events relative to the start of the fire interval (500 microseconds for the 2kHz fire rate). Figure 1 shows the timing for a single fire. Up to four returns can be captured for each fire. Each event is tagged to indicate whether it's a fire or a return, and for returns, which quadrant of the quadrant detector the event came from. Due to the high laser fire rate, there will be multiple fires in flight. A circular buffer in the software (Figure 2) correctly associates returns with fires, and calculates the roundtrip time of flight. Actual clock time (for time stamping the range data) is kept within the computer; GPS time is used to update the computer time approximately once per minute.

▸ All measurements are made relative to the start of 2kHz interval.

▸ Actual measurements, dt_F and dt_R, are only pieces of roundtrip range measurement.

▸ Software must calculate number of 2kHz intervals (M) between fire and return, and then use M, along with dt_F and dt_R, to compute roundtrip range.

Number of intervals from fire to return: M = (predicted_range + dt_F)/500usec

Measured roundtrip range = M*500usec + dt_R - dt_F

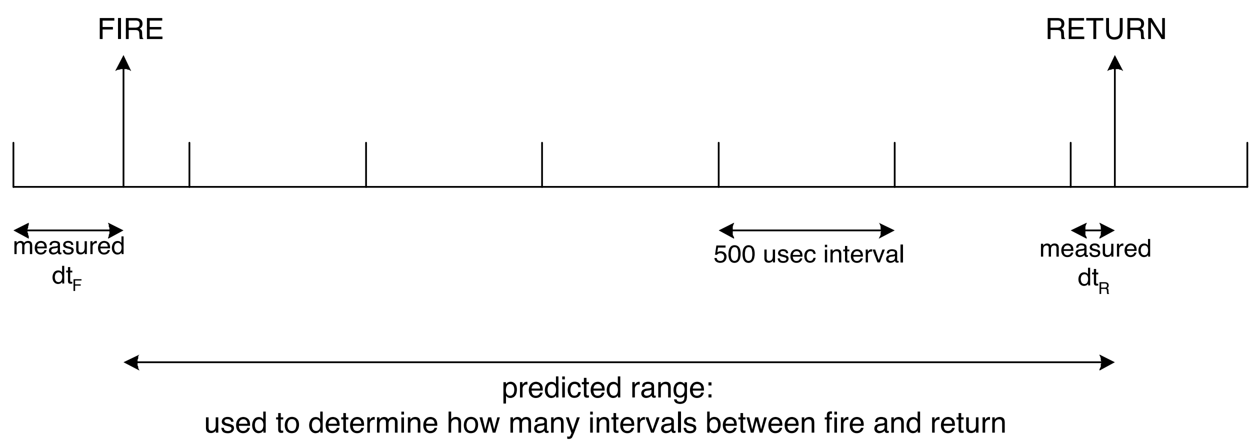

Figure 1: Range Timing

The laser is passively Q-switched and the software has no direct control over when the laser fires; it can only control the fire rate, also known as pulse repetition frequency (PRF). To prevent collisions of the incoming return with an outgoing fire, the software will have to constantly monitor the expected return from each fire and adjust the PRF as needed to move the fire out of the way of any incoming return. The interval that must be kept clear is ~10 microseconds before the fire to ~100 microseconds after. This period is referred to as the blanking time and is a function of the laser jitter, the transmit/receive switch requirements, and the atmospheric backscatter from our laser.   Backscatter from our laser off of the atmosphere (to roughly 70km altitude) will create a much higher noise rate than our Signal Detection Algorithm can handle. For a related discussion on backscatter and avoiding fire/return collisions, see the paper by Paul Titterton, “System/Usage Impact of Operating the SLR2000 at 2kHz,” in these Proceedings.

To keep the cost of the system low, the requirements on the telescope aperture and the pointing accuracy have been somewhat relaxed. The telescope primary optical diameter will be between 30 and 40 cm, which limits the transmitted energy to 70 - 130 microJoules due to eyesafety constraints. The mount repeatability has been specified as $\leq$ 5 arcseconds; this will be an error that cannot be modeled out. To offset this pointing limitation we will ensure that (1) acquisition searches a large enough angular area, and (2) once we detect the signal, we will use the quadrant detector information to form an angular correction (see Figure 3).

```
// Ranging Circular Buffer: integers are in psec - reals are in microsec
// All data, except where specified, is associated with THIS interval's fire

struct RngElem{
long   I2K;         // # of 2KHz interval past the buffer start time
long   dtF;         // delta fire time (for this interval's fire) in psec
double Rp;          // predicted range for this fire (usec)
double Rpdot;       // predicted range rate for this fire (usec)
long   MR;          // number of intervals to this fire's expected return
long   dtE;         // expected time into THIS interval for return from previous fire
(psec)
long   MF;          // # of interval back to fire for return(s) in this interval
long   dtR[4];      // measured range delta(s) from start of interval for returns
                    //    from this fire (psec)
double R[4];        // measured roundtrip range for this fire (usec)
char   Q[4];        // tag indicating quadrant for each return (1,2,3,4, or 0)
char   N[4];        // flag indicating signal(=1), noise(=2), or no data(=0)
double rbias;       // range bias (usec)
long   tbias;       // time bias (usec)
long   rwin;        // range window width (psec)
double dRref;       // refraction correction (usec)
long   tblank;      // blanking time for this interval (psec)
};

struct RngCircBuf{
double RngT0SOY     // Start of circular buffer in seconds of year
struct RngElem[240000];
};
```

Figure 2: Range circular buffer

Due to the narrow beam divergence, the transmit beam will have to be pointed ahead of the satellite's location at the actual time of fire. Similarly the received returns will arrive from the satellite's past location. While the differences in angle between point-ahead and point-behind are small, for SLR2000 they are a significant fraction of the beam divergence and receiver field of view. Thus, the system will have to be able to independently point the outgoing laser pulse ahead of the telescope; the telescope will be pointed behind, to receive the returns from the satellite. The point-ahead angles are dependent on the orbit. Table 1 gives a list for some representative satellites.

Table 1: Difference between point-ahead and point-behind angles

Satellite	**Elevation=20°**	**Elevation=90°**
STARLETTE	6 arcsec	11 arcsec
LAGEOS	7 arcsec	8 arcsec
ETALON	5 arcsec	5 arcsec

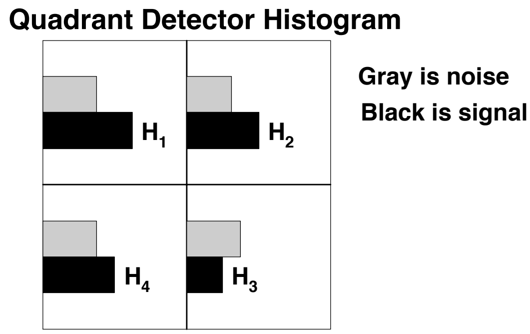

This histogram is used to compute the angular corrections (α,β), which then must be scaled and rotated, due to the Coude path, to get biases in (Az,El).

$$\alpha = \frac{(H_2 + H_3) - (H_1 + H_4)}{Q} \qquad \beta = \frac{(H_1 + H_2) - (H_3 + H_4)}{Q} \qquad where \quad Q = \sum_{i=1}^{4} H_i$$

Figure 3: Quadrant detector angular correction

Detecting the Signal

As in other SLR systems, the actual laser roundtrip return times will be differenced with the expected to produce an Observed Minus Calculated (O-C) data set. In current NASA SLR systems, this data is plotted for the operator to determine signal acquisition. In SLR2000 this determination must be made by the software using the Poisson statistics of the noise and the linear nature of the signal error in O-C space. The noise will appear uniformly distributed across the range window. The signal, however, will appear as a straight line, although it may be a weak line. Figure 4 shows a simulated LAGEOS pass near 20 degrees elevation taken by an SLR2000 system with a 40cm telescope. It takes roughly 85 seconds for the system to search for and find the satellite (due to mount induced pointing errors). In this plot, the timebias is determined and corrected for at approximately 100 seconds into the pass.

For satellites that are regularly tracked, the small timebiases (shown in Table 3) will ensure that the signal lines in O-C space have a small slope. The range window can be divided into bins and the data histogrammed. For most satellites 0.5 nanosecond bin sizes will ensure that the signal lies within one or two bins for a long enough period to collect the "k" returns needed to acquire. "k" is typically 6 to 14. The histogram bin with the most counts and whose count is greater than "k" is determined to contain the signal. The period of time required to collect "k" returns with a high probability of success is called the Frame Time.

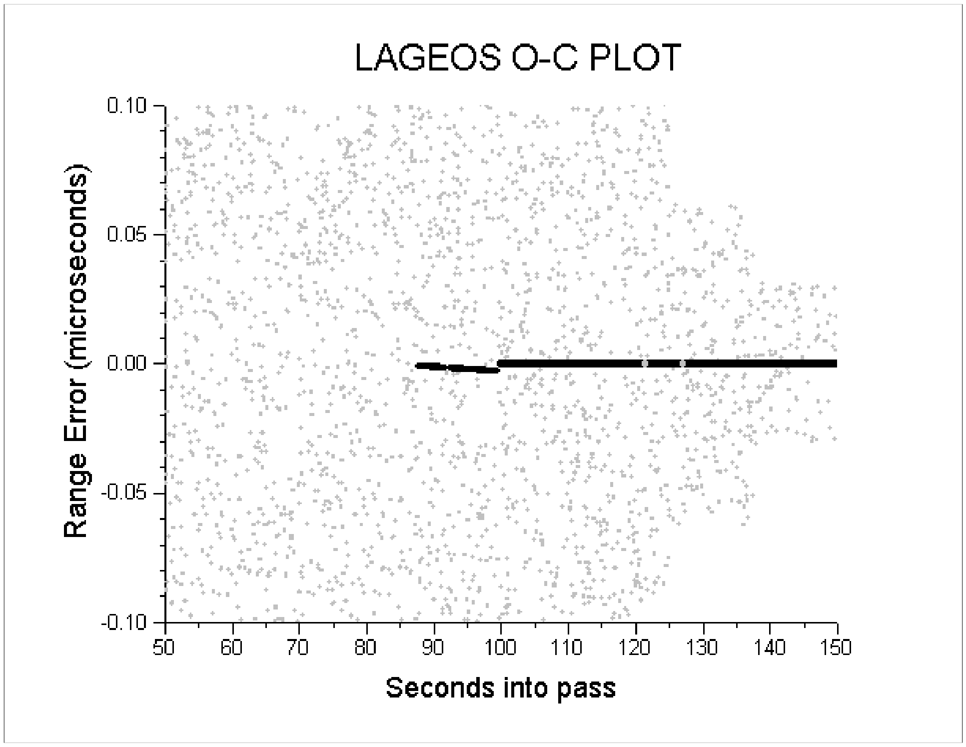

Figure 4: Simulated SLR2000 LAGEOS pass.
Black is signal and gray is noise.

Picking the right Frame Time depends on assumptions about both the signal level and the background noise level. At night this isn't an issue, but during the day, when the background noise is high, there are two options. We could tailor the algorithm to the worst case scenario where the atmospheric transmission is very bad, or we could use the visibility sensor (described in the section on Determining System Health & Safety) to give an indication of the atmospheric clarity. This can be used to deduce the affect of the atmosphere on both background noise and signal strength. The Sky Camera might also be used to determine cirrus cloud cover. Using these existing sensors may allow us to calculate the optimum Frame Time to fit the prevailing conditions -- reducing the chance of tracking on noise, and speeding up signal detection during clear conditions.

Table 2 below gives the expected number of returns per second for various satellites at elevation angles of 20 degrees and 30 degrees for both a 30cm and a 40cm telescope. STARLETTE is only listed for the worst case, since its signal is very strong. ETALON is listed only for 30 degrees, since no acquisition will be attempted for any high earth orbiting satellites below 30 degrees. The results in this table are based on a moderately clear atmosphere. The detector quantum efficiency was taken to be 12% and a system transmission of 34% was used.

Acquisition occurs when the signal detection algorithm correctly finds the signal. False acquisition occurs when the signal detection algorithm mistakes noise for signal. The larger the range window, the more noise is introduced, and the higher the probability of mistaking noise for signal. To reduce the probability of false acquisition, the maximum width of the range window will be limited to 200 nanoseconds.

The noise rate in Table 2 was taken to be 4 x 10^5 photoelectrons (pes) per second for both size telescopes. This is higher than the expected values of 1 x 10^5 for the 30cm telescope and 2 x 10^5 for the 40cm telescope. Even with this larger value for noise, the expected signal level is still much higher than the corresponding noise level as can be seen by comparing columns 2 and 4.

The Frame Time calculated in Table 2 give a 93% chance of getting k=6 signal returns for the given "Expected signal pes/shot" shown in column 2 of the table. The last column gives the probability of mistaking noise for signal by getting 6 or more noise counts in any histogram bin in the range window. In all cases such a probability is small.

Table 2: Expected signal photoelectrons/shot, Frame Time, Noise rate, and Probability of False Alarm for various combinations of satellites, telescopes and elevation angles for a moderately clear atmosphere.

Satellite	**Expected signal (pes per shot) (counts/frame)**	**Frame Time: (seconds) (shots)**	**Noise rate: (pes/bin per shot) (counts/frame)**	**Prob(FalseACQ) in 200nsec range window**
STARLETTE 20° / 30cm	0.230 pes = 10 cts/frame	0.022 sec = 44 shots	0.0002 pes/bin = 0.009 cts/frame	negligible
LAGEOS 20° / 30cm	0.004 pes = 10 cts/frame	1.250 sec = 2500 shots	0.0002 pes/bin = 0.500 cts/frame	0.6%
LAGEOS 30° / 30cm	0.016 pes = 10 cts/frame	0.313 sec = 626 shots	0.0002 pes/bin = 0.125 cts/frame	negligible
LAGEOS 20° / 40cm	0.014 pes = 10 cts/frame	0.357 sec = 714 shots	0.0002 pes/bin = 0.143 cts/frame	negligible
ETALON 30° / 30cm	0.003 pes = 10 cts/frame	1.667 sec = 3334 shots	0.0002 pes/bin = 0.667 cts/frame	2.7%
ETALON 30° / 40cm	0.008 pes = 10 cts/frame	0.625 sec = 1250 shots	0.0002 pes/bin = 0.250 cts/frame	negligible

Lower relative signal strengths (compared to the noise rate) require higher values of "k" to correctly distinguish signal from noise. In the case of an expected signal return of 0.001 pe/shot with the same noise rate as above, a k=14 must be used. The Frame Time then must be 10 seconds in order to achieve a probability of acquisition of 93%. Here the probability of false acquisition is 3%. For lower signal strengths or higher noise rates, this technique eventually fails, and the N of M technique described below must be used to detect the signal.

Very low earth orbiting, sparsely tracked, and newly launched satellites present a tougher problem. Due to the poorer quality of the predicts, the slope of the signal in O-C space may not be small. We are currently working on a technique that shortens the Frame Time, but uses multiple Frames together (called a SuperFrame) to make a signal decision. In this case, N of M consecutive Frames would have to pass the test before the Signal is determined to be present. While the N of M technique does take longer to make a determination of signal, it is better able to reject noise. Also, if the bins in each of the "N" Frames are not required to have the same index, but are rather allowed to form a line, then this technique is better able to handle signal with a larger slope. Further discussion on this topic is outside the scope of this paper.

<u>Acquiring the Satellite</u>

To increase our ability to acquire satellites, improvements were made in the generation of the orbital predictions. The SLR2000 orbital ephemerides are now updated daily, thus keeping the biases small (see Table 3), and the onsite calculation limitations have been avoided by changing the type of predictions used. Instead of using daily Tuned Inter-Range Vectors (TIVs) as in the current SLR systems, which require sophisticated gravity models and complicated numerical integration schemes onsite, the new predicts are untuned vectors spaced at one minute intervals throughout the day. The only calculations now required on site are interpolation and translation. Good predicts not only limit the time required to search for the satellites, but allow for smaller range window widths which reduces the noise and thus the chance of false acquisition.

Table 3: Errors in the daily updated orbits of some commonly tracked satellites

Satellite	**Altitude (km)**	**Timebias(msec)**	**Angle Error(arcsec)**	**Range Error(nsec)**
STARLETTE	806	10	0.40	50
LAGEOS	5912	2	0.16	30
ETALON	19115	2	0.10	5
GPS	20000	100	0.60	200

For most of the regularly tracked satellites, the SLR2000 predictions should be good enough to put the laser beam directly on the satellite, however unmodelable errors in the mount pointing (on order of 5 arcseconds) will prevent this from happening. Because of this, no matter how good the predicts are, the system must as a matter of course be prepared to search for the satellite. The search will be in angle, not in timebias. A range window search will also be used to allow the software to find the satellite in the presence of range biases, without having to open the range window beyond 200 nanoseconds. Because each satellite is different, the software must keep a list of satellite characteristics needed for acquisition, signal detection, and tracking.

Once the signal detection algorithm has detected returns, the satellite is considered acquired. A time bias and a range bias are calculated from the detected signal ranges (if a solution is possible) and the corresponding contribution of these biases is removed from the angular and range biases accumulated during acquisition. All of the biases are then placed in shared memory for access by the tracking and logging tasks.

The Sky Camera (thermal IR camera) reports the temperature of the sky and is used to determine the cloud cover. Cooler temperatures imply clear skies and warmer areas of the sky mean denser cloud cover. The camera takes an entire sky image (approximately 20 degrees above horizon to zenith) once every minute as reflected from the convex mirror shown in Figure 5. There are 120 x 120 pixels in the image giving a temperature for roughly every 1° x 1° section of the sky.

Figure 5: Sky camera and convex mirror

If the majority of the satellite pass is in the clouds, then the pass will be abandoned in favor of another target that is in a clearer part of the sky. If the majority of the pass is clear or somewhat cloudy, the pass will be taken and the sky map will be used to correctly determine a course of action if no signal is detected. In this case (some cloud cover where no returns are detected) the software will wait until a clearer part of the pass is reached before making another step in the acquisition search.

Acquisition ends successfully with a transition to tracking. Acquisition can also end in a rejection of the object being tracked due to: (1) the pass being in a cloudy part of the sky, (2) exceeding the search time limit, or (3) the schedule indicating a change of objects.

<u>Maintaining Satellite Track</u>

Once the satellite has been acquired and the correct biases determined, the software closes the range window down to its minimum size (around 40 nanoseconds). The software continues to run the Signal Detection Algorithm described above, and continues to use this signal information to correct the time and range biases. Once in tracking mode, the software also uses the quadrant detector information to correct the pointing.

Tracking ends when one of the following occurs: (1) the signal is lost for more than the given time limit during a clear sky, and we transition back to acquisition, (2) the majority of the remainder of the pass becomes cloudy, and we abandon the satellite in favor of another, or (3) the schedule indicates that we should change objects.

<u>Determining System Health & Safety</u>

The SLR2000 system is housed in a dome with a closeable slit that tracks the telescope's azimuthal position. The software can monitor all of the relevant weather conditions outside of the dome with the slit shut, and can determine whether or not to open the system. Table 4 gives the external weather information available to the system software.

Table 4: Weather instrumentation

Measurement	**Manufacturer**	**Use**
Temperature, Barometric Pressure, Humidity	Paroscientific	Measurement -- Range & El correction. Safety -- dome closed for high temp, high humidity, or pressure too low.
Wind	Belfort - Young	Safety -- dome closed in high wind.
Visibility	Vaisala	Safety -- dome closed in fog Acquisition -- used in Frame Time calc.
Precipitation	Vaisala	Safety -- dome closed in rain or snow.

The software also can monitor all of the relevant temperature, humidity, voltage, and current information inside the dome and within the electronics. It also has the capability to check the security of the system by monitoring motion sensors and interlocks; security cameras are triggered to take pictures automatically when security events occur. The Health & Safety software monitors all relevant parameters from once per minute to once per hour depending on need. Every minute a decision is made on the health and safety of the system, and a status is determined. The possible statuses and corresponding actions taken by the system are given in Table 5 below. Changes in system status are automatically recorded in the Daily Diary (see the section on Communication with a Central Facility); hazard and emergency conditions also cause an immediate e-mail and phone message be sent to the Central Facility.

Table 5: System status conditions and resulting actions

System Status	**Meaning**	**Resulting Action**
WHITE	Not Ready	Wait -- not all subsystems have reported on power-up
GREEN	OK	Operate normally
YELLOW	Warning	Operate normally but monitor condition
ORANGE	Error	Operate with restrictions on certain subsystems
RED	Hazard	No operation -- turn power off to effected subsystems
BLACK	Emergency	Turn all power off to entire system

<u>Calibrating and Assessing System Performance</u>

As in current SLR systems, ground calibrations to fixed targets will be performed periodically to determine the system delay and the system stability. The data from these calibrations will be placed in a database for use during data processing and system performance assessment. Calibrations will be scheduled at least once every two hours by the scheduling software.

Star calibrations will be used to determine the coefficients of a mount model which describes the nonorthogonalities and misalignments in the mount. Current SLR mount models assume fixed errors in the mount, however, experience indicates that the mount errors change as a function of temperature. It is unknown at this time whether more frequent star calibrations than are done in current systems (no more than once per month) will be needed for the SLR2000 pointing requirements, or whether a dynamic model which is a function of temperature will be required. Automated star calibrations using 50 - 60 stars are currently being performed at NASA's 1.2 meter telescope using a CCD camera, and this software will be ported to the SLR2000 system. The actual number of stars used in the star calibration will be a function of the number of mount modeling coefficients that the SLR2000 system requires.

Single star updates to the mount model will also be performed at the start of every satellite pass when a bright enough star is within easy reach of the mount (<10 degrees). At the start of a new satellite pass the system will check for such a star, and if one is found, the software will drive the mount to that star, center the star in the FOV, and perform a single star Kalman Filter update to the mount model. The software will then focus the system on this star, before moving on to the satellite pass.

The system will continually monitor its performance to determine if any corrective action is needed. Initially this performance assessment will monitor: (1) ground calibration system delay RMS and stability, (2) star calibration RMS, (3) satellite and ground calibration acquisition successes versus attempts (in clear skies), and (4) percentage of returns versus fires. Corrective actions will include overriding the schedule and performing either a ground calibration or a star calibration, and sending a message home to its Central Facility indicating that the system performance appears to be degrading.

Scheduling

The system will be self-scheduling based solely on the predicts, a list of satellite priorities, and the required maintenance. The schedule includes satellites to be tracked, ground calibrations, star calibrations, and any maintenance or diagnostics to be performed. The system can also schedule "down-time" when either the laser cannot be fired or the dome must be shut. The granularity of the schedule is one minute. This schedule is generated daily and contains the object number (Satellite ID if satellite), the date, the starting time (in hours and minutes), the durations of the event (in hours and minutes), the object's priority, the starting azimuth and elevation, and the three-letter name associated with that object. Note that there can be multiple objects for a given time period. The software picks the highest priority object, moving on to the next priority object only when higher priorities have been rejected (due to sky conditions or other reasons). Figure 6 shows a part of one daily schedule.

The realtime code checks the schedule to determine what operation it should be performing. If the sky is clear and the system status allows, the software will attempt to carry out the scheduled task. If the sky is cloudy where a satellite pass is scheduled, the realtime code can reject this object and ask the scheduling routine for the next. The realtime code can also override the schedule to either complete a scheduled ground calibration or star calibration, or to perform an unscheduled one.

OBJ#	DATE	DOY	Start	Duration	Prior.	AZ	EL	NAME
9998	19980717	198	00:00	00:01	98	0	0	cal
4377	19980717	198	00:02	00:12	1	344	17	top
5986	19980717	198	00:02	00:12	6	103	30	la2
3535	19980717	198	00:02	00:12	7	95	28	g35
5986	19980717	198	00:14	00:04	6	72	23	la2
3535	19980717	198	00:14	00:04	7	89	30	g35
9998	19980717	198	00:18	00:01	98	0	0	cal
9990	19980717	198	00:19	00:16	90	0	0	mnt
9998	19980717	198	00:35	00:01	98	0	0	cal
1500	19980717	198	00:36	00:16	5	241	16	aji
1155	19980717	198	00:36	00:16	6	88	21	lag
3535	19980717	198	00:36	00:16	7	77	31	g35
1155	19980717	198	00:52	00:23	6	41	45	lag

Figure 6: Part of the daily schedule

Communication with a Central Facility

The SLR2000 system gets predicts via the Internet daily, and sends out Normal Point and Merit II data hourly. The system log, called the Daily Diary, is e-mailed out daily. This file contains a list of all important events that happened during the day. Items included would be such things as when the dome was opened and closed, when the system's status changed, and why certain satellites could not be tracked. All of the weather information is also sent out in a separate file at the end of each day.

The software maintains a database of important system information which can be queried via the Web. The database information includes weather, ground calibrations, satellite statistics, and star calibration statistics. Any interested person will have access to this information over the Web.

Software Design

Figure 7 shows the software design and data flow. The software is divided into four major categories:

1) Interface and Control Computer (ICC)
2) Pseudo-Operator (POP)
3) Data Analysis (DAN)
4) Remote Access Terminal (RAT)

The ICC performs all of the realtime hardware interface that occurs faster than 1Hz. This includes the ranging electronics, the laser, and the mount. The ICC is also the interface to the CCD camera used for star calibrations.

POP is essentially the replacement for the human operator as well as the interface to the ICC. POP makes all of the realtime decisions, does all of the realtime prediction update, interfaces with the ICC at 2kHz, and records most of the realtime data.

DAN performs the data analysis and interfaces via Internet with the external world. DAN performs all of the functions of the current HP Data Analysis computer on NASA systems. DAN gets the predicts for the system, generates the schedule, provides the final data products, maintains the databases, and controls the Web site. DAN also hosts the RATSNEST which is the local server software for RAT.

The ICC, POP and DAN are the operational parts of the SLR2000 software system. However, in order to debug this system, and to allow maintenance personnel who visit the site to work with the system, some operator interface is required. The RAT software fulfills this requirement. RAT resides on a laptop which is interfaced to the rest of the SLR2000 system via the Internet. RAT allows the user to view what is happening in the system, and to control the system to a certain extent. Most importantly for the software development team, RAT is a required tool for system debugging. See the companion papers in these Proceedings on POP, DAN, and RAT for more detail on each of these systems.

LynxOS was chosen as the Operating System for POP and DAN because of its realtime capabilities, its good development environment, and its proven track record on existing NASA SLR systems. LINUX was chosen as the operating system for RAT. LINUX is UNIX freeware that is enjoying great popularity at this time. It does not, however, have realtime capability, but is a robust operating system with standard UNIX capabilities and features.

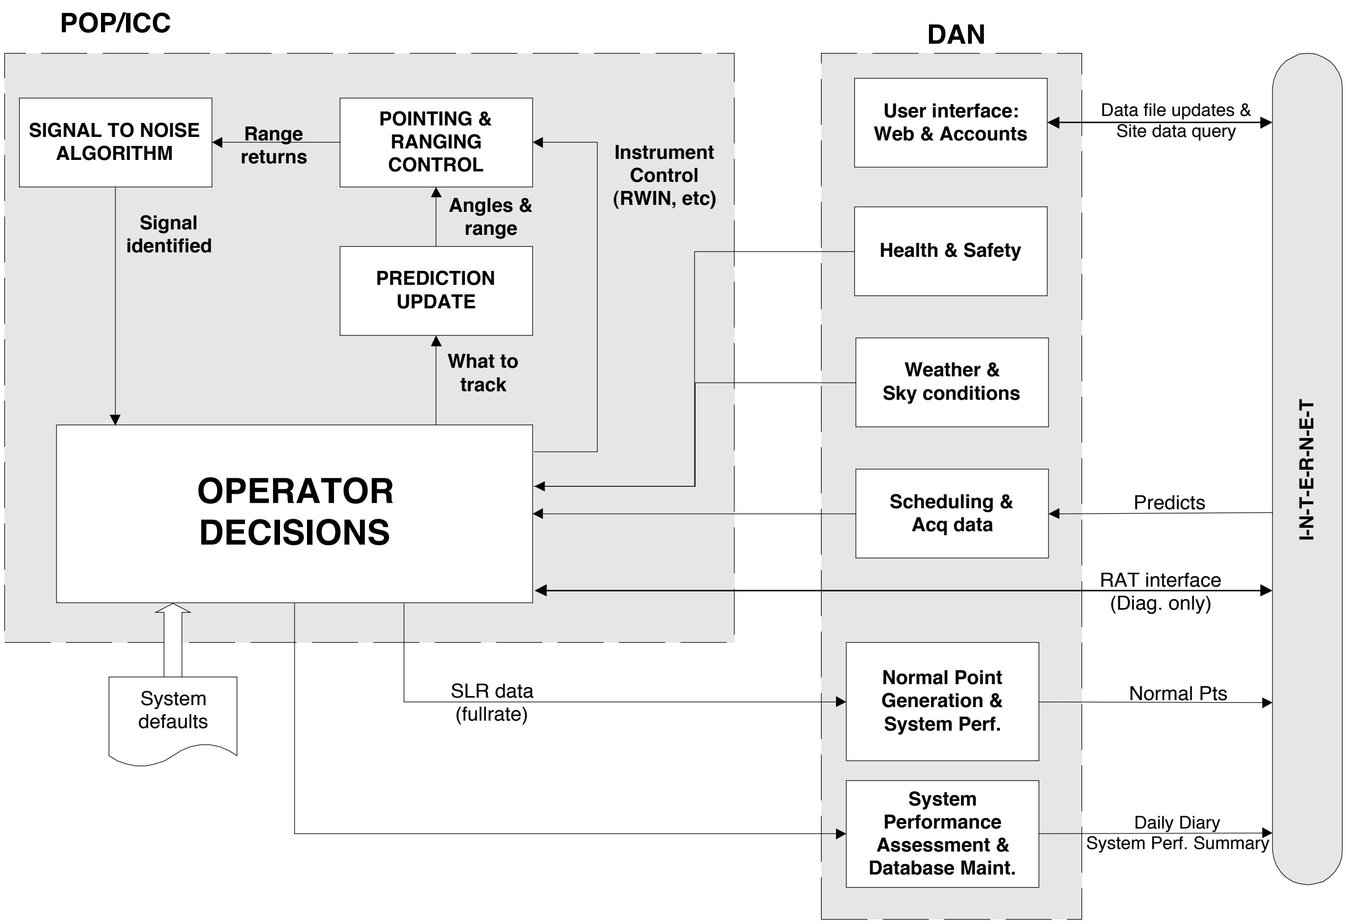

Figure 7: Software block diagram and data flow

Conclusion

SLR2000 is an ambitious project, especially for the software team. We anticipate over 20 man years of software development from the project start to the completion of the first operational system. Much effort has already gone into the software design, and the Software Design Document has been published. The design doesn't end here, but will continually require modification throughout the development process. Coding is currently underway for all four systems (ICC, POP, DAN, and RAT), and the first test involving all of the software systems is scheduled for mid-February of this year. Star and ground calibration testing with the prototype SLR2000 system is expected to take place in the fall of 1999.

Acknowledgements

The authors would like to thank John Bosworth and David Carter of NASA's Laboratory for Terrestrial Physics Space Geodesy and Sensor Calibration Office for their continued support of this project, and John Degnan of NASA's Laboratory for Terrestrial Physics Geoscience Technology Office for his leadership and technical expertise in the development of this system as a whole.

The authors would also like to thank Paul Titterton and Hal Sweeney of The Electro-Optics Organization for their analysis of SLR2000's performance and their development of the Signal Detection Algorithm.