SLICER: SCANNING LIDAR IMAGER OF CANOPIES BY ECHO RECOVERY INSTRUMENT AND DATA PRODUCT DESCRIPTION

DRAFT

Additions to complete this draft will include:
1) currently missing values for instrumentation specifications
2) figures illustrating many of the concepts discussed in the text
3) reference list

David J. Harding, J. Bryan Blair, David L. Rabine, Kathy Still
NASA's Goddard Space Flight Center
harding@denali.gsfc.nasa.gov
June 2, 1998
v. 1.3

PURPOSE

This document describes the Scanning Lidar Imager of Canopies by Echo Recovery (SLICER), an airborne laser altimeter and surface lidar instrument that has acquired topographic and canopy structure data sets in support of NASA Earth Science programs. Supported programs include the Topography and Surface Change Program, the Terrestrial Ecology Program, and the Boreal Ecosystems Atmosphere Study (BOREAS). SLICER was developed in, and operated by, the Laboratory for Terrestrial Physics at NASA's Goddard Space Flight Center. Funding for the development of SLICER and acquisition and processing of data products was provided by the Goddard Director's Discretionary Fund, and NASA's Solid Earth and Natural Hazards Program, Terrestrial Ecology Program, and Mission to Planet Earth Program Office. This document describes the operation of the SLICER instrument, the characteristics of the acquired data, and the structure of SLICER data sets produced and distributed by Goddard's Laser Altimeter Processing Facility (LAPF). An archive showing data location maps and profiles of processed SLICER data can be viewed via the World Wide Web at http://denali.gsfc.nasa.gov/lapf. Full SLICER data sets are available by distribution on CD-ROM.

INSTRUMENT DESCRIPTION

The principles of surface lidar, as implemented in SLICER, are rather simple and are based on precise timing of the round-trip travel time of a short-duration laser pulse. The ranging component of the instrument consists of a laser transmitter, scan mechanism, receiver telescope and detector, threshold-detection and digitization timing electronics, and an instrument control and data collection system. The ranging instrumentation is augmented by an Inertial Navigation System for precise determination of laser beam pointing, GPS receivers for differential, kinematic determination of aircraft position, and videography and/or still photography for documentation of the ground track. SLICER evolved from a profiling system described by Blair et al. (1994). The profiling instrument itself evolved from the Airborne Topographic Laser Altimeter System (ATLAS), described by (ref??), via the integration of surface backscatter digitization with threshold-detection laser ranging. SLICER was an evolutionary instrument intended as an engineering test-bed to evaluate digitization and scanning surface lidar technologies. As such, certain instrument features are non-optimal, limiting ranging accuracy and precluding calibration of digitized return amplitudes. Experience derived from SLICER was incorporated by Blair et al. (1996) in a completely redesigned, next-generation scanning surface lidar system referred to as the Lidar Vegetation Imaging Sensor (LVIS) which has superseded SLICER. LVIS is primarily used for algorithm development and validation activities in support of the Vegetation Canopy Lidar (VCL), the first mission selected by NASA's Earth System Science Pathfinder Program (ESSP). VCL is scheduled for launch into Earth-orbit in 2000.

The sharp pulse laser transmitter (Coyle and Blair, 1995; Coyle et al., 1995) utilized in SLICER is a cavity-dumped, Q-switched, diode pumped ND:YAG oscillator-amplifier. The laser outputs a 800 mJ, 1.064 mm pulse with a 4 nsec full-width at half-max pulse width at rates up to 80 Hz. The cavity-dump architecture of this transmitter was specifically designed to achieve a fast leading-edge rise time, on the order 1.5 nsec, and high peak power in order to maximize ranging accuracy via threshold-crossing detection and vertical resolution to multiple targets within a single laser footprint. The trailing-edge energy of this cavity-dump pulse is designed to fall off asymptotically, yielding a Raleigh distribution of pulse energy in time. However, with non-optimal alignment of laser transmitter elements, the asymptotic fall-off of laser energy was at times corrupted, yielding a pronounced step-like "back porch" on the pulse trailing edge. The best characterization of pulse shape is achieved by examining the digitized backscatter pulses reflected from flat surfaces such as calm water bodies. Such data is routinely acquired during pointing angle roll and pitch calibrations (described in detail below). Spatially the laser beam is circular in cross-section with a Guassian distribution of energy that decreases radially outward from the center. The divergence of the laser beam is approximately 2 mrad at the angle where the pulse energy has fallen off to 1/e2 that of the beam center. Lenses can be inserted into the transmit beam path that provide divergence amounts of approximately 3 and 9 mrad. Precise measurement of far-field laser beam divergence is difficult. For the SLICER laser transmitter divergence was measured in the laboratory via several methods yielding inconsistent results; the reported divergence values are accurate only at the 10% level.

The laser transmitter includes a photodiode detector which samples a small fraction of the output pulse energy at a point within the transmit beam. The signal from the photodiode detector is accumulated in an integrator and recorded as a single byte value for the output pulse energy. However, the resulting measure is a highly noisy, uncalibrated measure of total output energy. Shot-to-shot variations in the recorded output energy show variations on the order 50% percent, whereas laboratory calibration of the laser transmit energy indicates actual shot-to-shot variations at the 10% level or less. In addition, long term drifts are evident in the recorded output energy during a data acquisition which may be due to variations in laser transmitter efficiency as a function of temperature, but might also be due to thermally-induced mechanical flexure of the photodiode detector mount or other unknown sources. In some instances the long-term drift of recorded laser output energy is abruptly offset in amplitude. The amplitude offset appears constant and apply to large, continuous segments of laser shots (100's to 1000's of shots). The abrupt offsets are probably not actual changes in transmit energy . The cause of the offset is not known but is likely due to a gain or bias shift in the start pulse energy detection electronics. In rare instances, individual laser shots have markedly lower recorded output energy values than the majority of the other shots. These single shots do appear to indeed have significantly lower transmit energies, possibly due to Q-switch mistiming in the transmitter.

The horizontally transmitted laser pulses are reflected downward and are scanned in a direction perpendicular to the aircraft flight direction by means of a small, galvanometer-mounted mirror. The ?? cm diameter mirror is rotated to fixed positions with an angular accuracy of ?? mrad and a settling time of ?? nsec by computer control of the galvanometer. The laser transmitter is triggered to fire a single laser pulse by computer control after the mirror has been rotated to a specified, fixed position. The absolute time at which the laser is commanded to fire is recorded with ?? resolution by utilizing a timer card with a GHz oscillator that is synched to a 1 pulse per second time-tag from an on-board GPS receiver. The resulting pattern of laser footprints on the ground depends on their size and along-track and cross-track spacing. The footprint size depends on the laser divergence and aircraft altitude above ground. In normal operations 5 cross-track footprints are produced with 2 mrad divergence from an above-ground altitude of 5000 m, yielding nominal footprint diameters at the $1/e^2$ energy point of 10 m. The aircraft is typically flown at a level altitude above sea level, so any relief variations of the underlying ground cause variations in footprint diameter. Cross-track footprint spacing is determined by the angular separation between successive transmitted pulses and aircraft altitude above ground. The angular separation is determined by the programmed galvanometer angular positions and aircraft roll. The galvanometer is typically controlled to yield footprints which are contiguous cross-track during level flight (e.g., 2 mrad angular steps, or 10 m spacing for the above example) . However, the combination of low-frequency roll excursions of several degrees during data acquisition and high-frequency galvanometer scanning yields non-uniform cross-track spacing. Along-track spacing is determined by laser pulse repetition rate, number of cross-track footprints, and aircraft ground speed and pitch attitude. Changes in pitch attitude are usually small and thus have an insignificant effect on along-track spacing. For typical ground speeds of 120 m/sec and 5 cross-track footprints, a laser pulse repetition rate of 80 Hz is used yielding an along-track spacing of 7.5 m, which is slightly over-sampled for the above example. Wind speed and direction at the aircraft altitude can have a large effect on ground speed and the resulting along-track spacing.

Laser illumination backscattered from reflecting surfaces within each footprint is collected in a receiving telescope which has a diameter of ?? cm. The transmitted laser pulses are aligned so as to be within the 10 mrad field of view (FOV) of the telescope. Typically, 5 contiguous laser footprints are fit cross-track within the telescope FOV when using a 2 mrad beam divergence. Three and 1 cross-track footprints are typically used with the 3 and 9 mrad beam divergences, respectively. The receiving telescope's optical efficiency falls off towards the edges of the FOV by approximately 15%, so the backscatter energy collected for the outer scan footprints is slightly lower than for the central footprints. Furthermore, any misalignment between the transmitted scan pattern and the receiver FOV causes a further reduction in receive energy for footprints falling partially outside the receiver FOV. The laser scan pattern is aligned to the receiver FOV in-flight, typically at the start of a flight mission, by manually adjusting mirror mounts so as to maximize the return energy of each footprint. Returns off of a water surface are normally used as it provides a uniform albedo, flat reference target. Laser to receiver alignment drift of unknown cause can occur during a flight, causing varying, unmonitored return signal strength, especially of the two outer footprints in the 5 footprint configuration as they can fall partially to wholly outside the receiver FOV.

The backscattered laser light and any background reflected solar illumination collected by the receiver telescope is focused through a 2 nm wide bandpass filter, centered at 1064 nm, onto a 1 mm diameter silicon avalanche photodiode detector (Si:APD) with a sampling rate of ??? Mhz. The Si:APD detector converts input optical energy into an output voltage with a timing resolution of ????. At a constant operating temperature, the optical energy to voltage conversion is linear as long as the detector is not saturated (i.e. exceeding the maximum optical energy that can be measured). However, the sensitivity of the Si:APD detector is temperature sensitive and the detector temperature was not monitored nor controlled during flight. Thus, there are unaccounted for variations in detector sensitivity due to detector temperature, which is itself a function of ambient temperature inside the aircraft cabin and duration of detector operation. The sensitivity variability could be as large as a factor of 10 for the full range of operating conditions encountered. Detector saturation is prevented by reducing the transmit pulse energy. This is accomplished by manually inserting neutral density absorbing filters into the transmit laser pulse path prior to the galvanometer scan assembly. An appropriate filter is selected in-flight by the instrument operator (as a function of flight altitude, atmospheric transmissivity, target reflectivity, and pulse broadening due to within-footprint surface height variability) in order to reduce the maximum received backscatter laser energy below the saturation level of the detector. The filter selected by the operator is not recorded by the altimeter data system, nor are the filter absorption coefficients well calibrated. Filters are sometimes changed during acquisition of a single data segment in order to compensate for changing atmospheric or surface conditions.

In some cases the selected filter is insufficient to prevent detector saturation. When saturation occurs, the detector output voltage remains at a high, nearly-constant level (but with detector electronic noise still superimposed). Due to the limited bandwidth of the detector this saturated output voltage usually continues later in time than the end of the actual received optical energy reflected from the target. Thus, when the detector leaves the saturated state, there is no optical input and the output voltage abruptly drops to the background noise level, yielding characteristically truncated ends to saturated returns. The abrupt voltage drop also often leads to detector overshoot (output voltages extending below the background noise level) and ringing (output voltages oscillating about the background noise) after the end of the saturated return. It is thought that for saturated returns the integrated voltage output by the detector is approximately equivalent to the integrated voltage that would have been output if the return had not been saturated (i.e., the resulting decrease in peak output voltage is compensated for by the extended duration of the output).

The output voltage from the Si:APD detector is split between a time interval unit (TIU) and an analog-to-digital ??? Mhz digitizer. The TIU precisely measures the time interval between the transmitted laser pulse and the received backscattered pulse based on a threshold detection scheme applied to the detector output voltage. The TIU contains a high frequency internal oscillator that yields 156.25 picosecond timing resolution, equivalent to ?? cm resolution in round-trip ranging. A start-pulse fiber optic pickoff transfers a small portion of the transmit laser pulse energy into the receiver telescope where it is focused onto the Si:APD detector. The TIU starts when the leading edge of the transmit pulse pick-off rises above a start-pulse threshold value which is established by the instrument operator. The TIU stops when the output voltage of the detector rises above a stop-pulse detection threshold which is also established by the instrument operator. The detector output voltage can exceed the stop-pulse detection threshold due to a valid return from a reflecting surface or due to instrument electronic noise and/or background solar illumination. A range gate is used to define a ranging time over which the stop-pulse detection threshold will be applied. Signal greater than the detection threshold occurring before or after the range gate does not stop the TIU. The range gate is established by the operator to bracket the distance to the anticipated target (e.g. land surface) given the known flight altitude above the ground. In some cases, no detector output voltage above threshold occurs within the range gate, yielding a no range result. The stop-pulse detection threshold and range gate are established so as to maximize the frequency of ranging to valid returns and minimize the number of noise and no range results. Valid returns are received from the first encountered surface which backscatters sufficient transmitted laser energy to exceed the stop-pulse detection threshold. This can be an optically dense cloud layer, a vegetation canopy top, bare ground in unvegetated areas, or an ice or water surface.

Because SLICER employs a threshold detection method for ranging, the range data is affected by an error source known as range walk. Low amplitude returns typically have slower rise times on the leading edge of the return than do high amplitude returns. Thus, for targets at equal distance from the laser transmitter, the detection threshold is crossed slightly later in time for low amplitude returns as compared to high amplitude returns. As a result, the target with a low amplitude return will appear to be at a slightly greater ranging distance and thus lower in elevation. A range walk calibration curve is measured during each SLICER flight by transmitting laser pulses of varying intensity through a fiber optic delay cable of fixed length. However, SLICER range data in the geolocated distribution data sets is not corrected for range walk. The resulting uncorrected ranging error is typically at the 10's of cm level but can be at the meter level for the greatest excursions of return pulse amplitude. There are many sources that can cause variation in return pulse amplitude, including transmit pulse energy, atmospheric transmissivity, 1064 nm reflectivity of the target, the radially variable efficiency of the receiver telescope, misalignment between the laser scan pattern and the telescope FOV, and the temperature sensitivity of the Si:APD detector. In addition, the relevant amplitude governing the return pulse leading-edge rise time is the return signal strength from the first encountered surface within the laser footprint. Discrete, continuos surfaces such as a flat ice sheet result in abrupt returns with fast leading edge rise-times as compared to diffuse, vertically-distributed surfaces such as an open vegetation canopy which result in broadened returns with slow leading edge rise-times. The SLICER range data has not been corrected for range walk due to the complexity of these factors affecting leading edge rise time.

In addition to return-pulse range-walk, the SLICER TIU ranging can be affected by range-walk on the threshold detection of the start pulse if the transmit pulse energy exhibits large amplitude variations. This should not normally be the case, but as noted above in the description of transmit pulse energy monitoring, in some rare instances transmit pulse energy does appear to be lowered significantly. In these cases, range-walk will cause the TIU start to occur late in time and the resulting range measurement will be slightly shorten than normal. The derived elevation will be slightly higher (10's of cm) than would be determined with a transmit pulse of normal output energy. Because the transmit pulse energy is not accurately monitored in SLICER, the range data has not been corrected for start-pulse range-walk.

Upon detection of a signal above the stop-pulse threshold, the TIU measure of the start to stop time interval is recorded, which corresponds to the round-trip travel time of the transmitted laser pulse for detection of valid returns. A value of zero is recorded for no range results. In addition to TIU ranging to the first detected return, SLICER implements a surface lidar function by digitizing the time history of the energy backscattered from the surface. Upon detection of the stop pulse, a constant number of digitization bins are read from the digitizer, yielding a waveform for that laser shot that records return amplitude as a function of time. The digitizer samples the detector output voltage at a rate of ??? sec per bin, with a resulting vertical sampling resolution of 0.1112 m per digitizer bin (one-way travel time). The digitizer converts the detector analog voltage to an 8-bit digital number. The scale factor between voltage and digital number (i.e., the gain) is established by manually selecting an attenuation setting on the digitizer. The digitizer attenuation, which is normally left unaltered for a SLICER mission, is set so that maximum voltage output by the detector (i.e. saturation) falls just below the 8-bit digital limit (i.e., 255). If the digitizer gain is inadvertently set too high, the detector output voltage can exceed the 8-bit limit and the recorded digital signal will "flat-line" at a constant value of 255. Thus, this flat-line digitizer "saturation" can be differentiated from detector saturation which is characterized by a high-amplitude signal that varies slightly due to superimposed instrument noise. The digitized waveform is stored as a byte array, normally at the full digitizer bin resolution. In some cases digitizer bins are averaged together before storing in the waveform array, yielding a vertical waveform resolution reduced by an integer averaging factor. The bins read from the digitizer correspond in time to a range of bins referenced to the TIU stop time. A small number of bins correspond to the detector output voltage recorded prior to TIU detection of the stop-pulse threshold crossing (i.e. background noise before the return). The time at which the TIU stop-pulse detection threshold is exceeded theoretically corresponds to a constant digitizer bin position which is assigned in software. However, due to inexact timing synchronization between the TIU oscillator and the digitizer oscillator, the TIU stop-pulse event can jitter between several adjacent digitizer bins from shot to shot. The TIU stop-pulse bin position assigned in software is constant for a single data file but can vary significantly between files depending on operational parameters. The bin position corresponding to the TIU stop-pulse time is not recorded in the data system, and is normally found for a data file by visually assessing multiple, over-plotted waveform records to determine at which bin the signal level first increases above the background noise. The appropriate bin position can only be determined with this method to within a few waveform bins due to the above-mentioned bin jitter from shot to shot, and the variable character of the return pulse leading edge as a function of surface conditions.

Digitizer bins after the bin corresponding to the TIU stop-pulse time record the detector voltage continuing after the threshold detection is exceeded. The return signal is broadened in time, and thus occupies more than one digitizer bin, for a number of reasons. First, the transmit pulse energy is not instantaneous but is distributed in time with a pulse width of 4 nsec (full-width at half the peak amplitude), corresponding to ?? digitizer bins. Second, the Si:APD detector operating at a bandwidth of ?? Mhz has insufficient temporal resolution to measure a 4 nsec pulse, thus blurring the return pulse in time. The combination of transmit pulse width and broadening due to detector bandwidth yields an impulse response which is the theoretical return signal, as measured by the detector output voltage, for a reflection from a perfectly smooth and flat surface. The amplitude and width of the impulse response increases with increasing optical energy received. The return signal can also be broadened by the distribution in height of reflecting surfaces within the laser footprint. Unvegetated surfaces (bare ground, ice, water) with a rough and/or sloped surface cause a broadened return pulse due to greater travel times to those parts of the surface in the footprint which are at greater distance from the laser transmitter. For vegetation canopies, the vertical distribution of canopy elements (leaves, needles, twigs, branches) leads to reflections from surfaces at multiple heights within the laser footprint and a resulting broadening in time of the return signal. The resulting return signal captured in the waveform is a convolved measure of transmit pulse width, detector bandwidth, vertical distribution of vegetation, and the roughness and slope of the surface. The strength of the return signal as a function of time also depends on the reflectance at 1064 nanometers of the vertically-distributed backscatter surfaces. Thus the waveform amplitude also depends on the convolution of the Gaussian spatial distribution of laser energy across the footprint with the reflectivity and horizontal and vertical distribution of backscatter surfaces. The number of digitizer bins recorded per laser shot, which is constant for a data file, is selected to provide a digitizer height range that is larger than the anticipated greatest broadening of the return pulse. Thus, the waveform will normally include the end of the return signal and additional bins which record instrument noise following the last return. In some cases, a greater than anticipated pulse broadening will cause the end of the return signal to extend beyond the range of recorded digitizer bins. For example, unexpectedly tall trees might be taller than the height range recorded by the digitizer, causing the return from the ground beneath the trees to occur after digitization has ended (i.e. the ground return will be "missed").

In some cases the recorded waveforms can be extremely noisy. The source, or sources, of the noise is uncertain. One source may be rf interference from pilot radio communications. Often the noise is present as alternating high and low amplitudes between even and odd digitizer bins. Digitization is actually accomplished by two digitizer circuits each operating at ?? Mhz which are then combined, possibly accounting for the observed alternation. Usually no digitizer averaging is done when recording a waveform bin, so a waveform bin corresponds to a single digitizer bin and this alternating, dual-digitizer noise is observed in the waveform. Averaging of adjacent bins in the waveform data can be applied by the data user to greatly reduce this noise. In some cases, two or more digitizer bins are averaged together in the data system to form the value recorded in a waveform bin. In the case of digitizer averaging by two (or an even number), the dual-digitizer noise will not be evident in the waveform. Digitizer averaging by an odd number will reduce, but not eliminate, the dual-digitizer noise.

For no range shots where no stop-pulse above the TIU detection threshold is found, no signal is transmitted to the digitizer and the bin amplitudes remain stored in the digitizer from the previous laser shot. For the no range shot, the data system reads these previous values when recording the waveform. The waveform for no range shots should be ignored. No range shots are identified by TIU range values of 0.

Due to the numerous instrument factors described above that effect the digitized record of backscatter return, the SLICER waveform amplitudes are uncalibrated and should not be used as an absolute measure of return signal strength. Some of the factors effecting return signal strength are constant. In particular, the outward fall-off in receiver telescope sensitivity is constant and the resulting lower return strength of the outer beams as compared to the inner beams could be established by a comparison of returns from a surface of uniform reflectivity, such as water. However, many of the relevant instrument factors are unmonitored and vary slowly over time, including transmit pulse energy, the temperature-dependent detector sensitivity, and alignment drift between the scanned transmit pulses and the receiver telescope FOV. Over short segments of data (seconds to minutes) these slowly varying effects probably due not significantly change instrument performance and thus successive shots, along a single beam position, can be compared in a relative sense. Other unmonitored factors abruptly change digitized signal strength, including filters used to attenuate the transmit beam strength and the digitizer attenuation setting. Changes in these factors should appear as instantaneous, and usually pronounced, changes in receive energy from shot to shot. However, the attenuation differences between some of the filters is quite small and the resulting difference in received signal may not be easily recognized. Because of these numerous effects on digitizer amplitude, the most reliable use SLICER waveforms is to treat each independently and use it only as a relative measure of the vertical distribution of return signal strength for that footprint. Furthermore, for cases of detector saturation, preventing the correct measurement of peak return amplitudes, the waveform record should not be used even in this relative sense.

The location of the laser "bounce point" corresponding to the first detected reflection is determined by combining the laser altimeter TIU ranging data with knowledge of the laser pointing angle and the absolute position of the aircraft derived from a GPS trajectory, interpolated to the time of the laser firing. The TIU timing data is converted to travel distance based on the speed of light through an atmosphere of standard pressure and temperature. The TIU timing data is corrected for a system delay, due to optical and electronic signal delays which vary between installations of the instrument. The system delay is found by iteratively solving for the location of laser reflections from surfaces at a known distance. Typically data acquired in a flight across a site with an established geodetic benchmark is used, yielding an end-to-end system delay measurement. In some cases, ranging is done on the ground to a fixed target by firing the laser over a horizontal path using a turning mirror placed beneath the aircraft. The distance to the fixed target is independently determined with a steel tape. The ground calibrations of system delay have proven to be less reliable than the benchmark overflights, possibly due to the short ranging distances used.

Laser pointing angle knowledge is obtained from an onboard Inertial Navigation System (INS) which reports roll, pitch, and yaw parameters computed from gyroscope-determined accelerations. The attitude parameters are read in-flight from the INS and recorded by the altimeter data system for each laser shot. The trajectory is determined by post-flight processing of 2 Hz differential kinematic GPS data employing a fixed Ashtech Z-12 receiver at a base station and a roving Ashtech Z-12 receiver on the plane. The base station receiver utilizes a ground plane GPS antennae mounted on a tripod over a benchmark and the roving receiver utilizes an avionics GPS antennae mounted on top of the aircraft fuselage. Trajectories are produced using proprietary Ashtech software called PNAV which is capable of centimeter-level carrier-phase processing with dual-frequency, full-wavelength data. PNAV is a Kalman filter-based program which processes GPS observables sequentially. It does "on-the-fly" resolution of carrier-phase integer ambiguities (cycle slips). Forward and backward kinematic solutions were computed and combined in order to optimize the resulting accuracy. However PNAV is intended for short (< 10 km) separation between the receivers; errors increase for the long baselines typically used on SLICER flights. Comparison of several long-baseline PNAV solutions to trajectory solutions produced for identical data by Bill Krabill at Goddard using software optimized for long-baseline flights over

Greenland show differences at the meter level. Thus long baseline PNAV trajectories probably have only meter-level absolute accuracy. Reprocessing of the trajectories with more sophisticated software, like that developed by Bill Krabill, could yield absolute accuracies at the sub-decimeter level. However, for individual flight segments the relative accuracy of the PNAV trajectories should nominally be at the sub-decimeter level. There are events, however, that can cause large errors even in a relative sense for brief periods such as a change in satellites observed or cycle slips. The trajectory file gives quality metrics for each epoch reported by the PNAV solution which can be used to assess the reliability of the trajectory. Roll and pitch maneuvers throughout the duration of a flight are usually kept to be below 10 degrees in order to minimize loss of lock on GPS satellite transmissions due to blockage of the signals by the airframe.

Post-flight processing software implemented in the Interactive Data Language (IDL) uses rotation matrices to compute the elevation, latitude, and longitude of the laser "bounce point" in WGS-84 coordinates from laser TIU range, INS pointing, and GPS trajectory inputs. The IDL geolocation routines were developed by Bryan Blair, based on adaptation of routines described by Vaugn et al. (199?). The spatial offset between the GPS avionics antennae mounted on the aircraft fuselage and the laser transmitter/receiver is measured and accounted for in the matrix rotations. Time delays between firing of the laser pulse and the recorded roll, pitch, and yaw parameters are also accounted for. The time delays for each of the attitude parameters consists of two components, an INS instrument reporting delay and a synch-up delay between the constant INS data rate and the software-controlled laser fire rate. The INS instrument reporting delay is constant for each parameter and is determined during post-processing by analyzing roll and pitch calibration maneuvers over water surfaces. The reporting delay parameters are varied in order to yield a "flat" water surface (exclusive of wave structure). The synch-up delays change between each data file and are determined by an iterative shifting and interpolation of the recorded attitude parameters with respect to the laser fire times. Roll and pitch alignment biases between the INS and the laser transmit direction are also determined by roll and pitch calibration maneuvers over water surfaces and accounted for in the geolocation matrix rotations. Roll and pitch biases are iteratively varied during post-processing in order to produce optimally "flat" water surfaces. Ideally the alignment biases should remain constant for the duration of an installation. However, roll and pitch calibration maneuvers are typically done on each flight day, often several times, and alignment biases appropriate for that flight day are used in the geolocation processing. Small alignment bias variations are observed from day to day, and even during a single flight, but the cause and rate of change is unknown. These alignment bias variations are likely related to the unexplained drift between the transmit pulse and receiver FOV.

The accuracy of the geolocation depends on the accuracy of the ranging data, pointing knowledge, platform trajectory, and determination of the system angle and time biases. For the high altitudes typical of SLICER acquisitions (nominally 5 km above ground) pointing errors are the largest source of geolocation errors. Based on the specified accuracy of the INS and the magnitude of the unmonitored alignment drifts, the pointing knowledge accuracy should be on the order 2 milliradians, yielding position accuracies equivalent to the laser footprint diameter in the nominal 5 beam mode. Geolocation accuracies of this order have been verified by comparing SLICER elevation profiles to corresponding image features (e.g., edges of building, forest clear cut boundaries) in several independently georeferenced high-resolution images and to topography in independently georeferenced digital elevation models (DEMs). Note that to compare SLICER elevations, referenced to the WGS-84 ellipsoid, to topographic data (maps or DEMs) which typically use a vertical datum referenced to mean sea level, an ellipsoid to geoid correction needs to be applied. Because the ellipsoid to geoid separation varies spatially, the correction needs to be applied to each laser footprint as a function of its latitude and longitude. The National Geodetic Survey distributes software called GEOID which computes the ellipsoid to geoid separation within the conterminous United States as a function of latitude and longitude. In a few cases, comparison of SLICER profiles to georeferenced images show unexplained horizontal offsets along the data track at the scale of 100's of meters. These are likely due to intermittent errors in the SLICER data system time-tagging, possibly due to miscounted second pulses from the GPS 1 pps time-tag.

In addition to using the on-board GPS receiver as a time source and for post-flight, kinematic trajectories, the real-time GPS positioning information is used to navigate the aircraft on predetermined flight paths. Data lines across specific targets and at specific orientations are preprogrammed into a flight navigator using waypoints. The flight navigator, developed by Wayne Wright at Goddard's Wallops Flight Facility, displays the waypoints and the desired and current flight path to the pilot. The pilot adjusts heading to follow the desired flight path. Due to the standard Anti-spoofing (AS) and Selective-availability (SA) degradation of the military GPS signal, the accuracy of the real time position information supplied to the flight navigator varies with time with a maximum error on the order of 100 m. The ability of the pilot to follow the desired flight path usually exceeds this imposed GPS error source, so the pilot in fact tracks a displayed flight path that deviates in time by a variable amount from the preprogrammed path. Real-time reception of broadcast RTCM corrections to account for AS and SA-induced position errors was not implemented on SLICER flights. In addition to 100 m scale flight path position errors, the track of the laser footprints across the ground deviates from the nadir ground track due to aircraft roll. The SLICER instrument was not mounted on a gyro-stabilized platform, so aircraft roll causes meandering of the ground track. Roll is commonly at the level of several degrees on a data collection line, causing up to 100s of meters of ground track deviation from the nadir track at the typical SLICER flight altitude of 5 km above the ground. The ground track deviation due to uncompensated roll is large as compared to the SLICER swath width (typically 50 m at 5 km altitude), making it difficult to ensure data acquisition across small (10 to 100 m) targets.

DATA PRODUCT DESCRIPTION

All CDs created by the LAPF are produced in a SUN Workstation UNIX Environment, but are readable in UNIX, DOS, Windows or Macintosh environments. However, text file end of line indication and binary data file byte order follow the UNIX conventions . File names are restricted to be 8 characters or less, and extensions are restricted to be 3 characters or less, to be compatible with the DOS operating system. In some cases, padding characters extending shorter file names to 8 characters appear when reading the CDs in Macintosh or PC environments, but not in UNIX.

CONTENTS OF .GEO CDs:

.geo CDs contain the raw altimetry data and post-processed geolocation results for each laser footprint. The .geo CD contents include:

XXXX.TXT	where XXXX is a alphanumeric mission code which may include the Julian calendar day of the year for the flight; this is a text file with flight log notes for each data segment acquired
PRDATA	directory containing processed data files
PROGS	directory containing IDL procedures and scripts used to process the data (IDL = Interactive Data Language)

PRDATA directory contents:

Each data segment is named as YYMMDD## where YYMMDD refers to year, month, and day of the flight and ## refers to a sequential numbering of the data segment acquired on that day. A data segment can be a topography flight line, an in-flight instrument calibration, or an on-the-ground instrument calibration.

For each flight line, there are the following files indicated by suffixes:

.geo	binary data as documented in detail below
.gfl	gif map of the flight line location
.gif	gif plots of location map, roll & pitch on the line, and profiles of elevation and total pulse width
.ps	ps plots of location map, roll & pitch on the line, and profiles of elevation and total pulse width
.psb	same as .ps but elevation profile is plotted as a connected line to highlight segments of non-surface returns (i.e. clouds or 0 range returns which yield a elevation at the plane altitude - 0 range returns occur when no laser backscatter signal above the detection threshold was received)
.tff	text file of the roll, pitch, and yaw time offsets (derived during processing)

For flight lines done as over-water roll and pitch calibrations there are two additional files derived during processing:

.bia	text file of roll and pitch biases
.rpc	postscript plot of bias determination

The PRDATA directory also contains a text file labeled YYMMDD.trj which is the trajectory file for the flight. The first line of the trajectory file indicates the number of records that follow. Each record consists of the following data elements per 2 Hz (twice a second) epoch:

time as seconds past GMT midnight
decimal latitude in WGS-84 reference frame (0 to ±90°)
decimal longitude in WGS-84 reference frame (0 to 360° East)
elevation in meters above the WGS-84 ellipsoid
number of services (GPS satellites) observed
pdop (position dilution of precision)
rms one sigma position error in meters

flag for solution acceptability

This is an example trajectory data record:
63521.00 18.43989221 66.00056597 -37.0566 4 44.5 3.27 1

The PNAV solution is reliable only with a minimum of 5 satellites observed and a PDOP of less than 4. The RMS position error is only valid in an absolute sense when certain conditions are met (which is probably not the case for long baseline flights), but it can give you a relative sense of where the solution is of poorer accuracy.

Note that the GMT times in the aircraft trajectory file are less than the GPS times in the .geo files by a small, constant number of seconds due to the difference between GMT and GPS time in effect at the time of the flight. GMT progressively falls behind GPS time as leap seconds are added to GMT but not to GPS. During the period of SLICER data acquisitions, the time difference GPS-GMT was as follows after the listed date:

1990 DEC 31	+7 sec
1992 JUN 30	+8 sec
1993 JUN 30	+9 sec
1994 JUN 30	+10 sec
1995 DEC 31	+11 sec

STRUCTURE OF .GEO FILES:

Each binary .geo files starts with a header consisting of:

parameter name in idl created by @rdsldgeo	data type	description
A.DATETAKEN	8 CHAR. STRING	date of SLICER flight
A.DATEPROCESSED	8 CHAR. STRING	date .geo file produced
A.AVERAGE	LONG	digitizer bins averaged in waveform
A.WAVEPTS	LONG	# of bins in waveform
A.BEAMS	LONG	# of cross-track footprints
A.GALV_CENTER	LONG	galvanometer position for center beam
A.NUMSHOTS	LONG	# of laser shots that follow header

After the header, the following binary data elements are stored for each laser shot:

C.GPSTIME	DOUBLE	seconds past GPS midnight for laser firing
C.TIU	LONG	round-trip travel time in milliseconds
C.VERNIER	LONG	high-precision determination of return time
C.STARTEN	LONG	uncalibrated measure of transmit energy
C.ROLL	FLOAT	laser optical bench roll in degrees
C.PITCH	FLOAT	laser optical bench pitch in degrees

C.YAW	FLOAT	laser optical bench yaw in degrees
C.GALV_STEP	LONG	cross-track beam position (1=left, 3 or 5=right, 0 = profile)
C.GALV_SET	LONG	galvanometer position
C.LAT	DOUBLE	aircraft latitude (deg) wrt WGS84 ellipsoid
C.LON	DOUBLE	aircraft longitude (deg) wrt WGS84 ellipsoid
C.ALT	FLOAT	aircraft altitude (m) wrt WGS84 ellipsoid
C.FLAT	DOUBLE	footprint latitude (deg) wrt WGS84 ellipsoid
C.FLON	DOUBLE	footprint longitude (deg) wrt WGS84 ellipsoid
C.ELEVATION	FLOAT	first return elevation (m) wrt WGS84 ellipsoid
C.WIDTH	FLOAT	distance (m) from first return to end of return
C.WAVEFORM	BYTE ARRAY	uncalibrated return echo waveform as byte array

Note that there is no parameter for laser shot number. In IDL, the index to the C data array is used to indicate the laser shot number (starting at 0 and increasing to A.NUMSHOTS -1).

C.ELEVATION is the elevation of the highest surface detected within the laser footprint, as determined by the round-trip travel time recorded by the altimeter time-interval-unit (TIU) combined with the pointing of the laser beam and the absolute position of the aircraft platform. The highest surface would normally be the ground in unvegetated areas and the canopy top in vegetated areas. However, anomalous elevations can result from several sources:
1) TIU triggers on noise yielding elevation above or below the ground
2) processing math errors (usually the first few shots before proper synching of the data system) yielding incredibly high elevations (e.g. 7.28e+08).
3) no range data yielding an elevation nearly equal to the plane's altitude (but offset by the system delay); this can be especially common for outer beams (1 and/or 5) if the scan pattern is misaligned with respect to the receiver FOV, causing low return signal strength for the beams, because the same TIU detection threshold level is used for all the beams
4) clouds, yielding an elevation somewhere between the plane's altitude and the ground
In all these cases the footprint latitude and longitude will be in error.

As described in the instrument section, a range walk correction has not been applied to the TIU range data so the absolute accuracy of the C.ELEVATION data varies from shot-to-shot at the 10's of cm level. Also, drift errors in the determination of pointing attitude can cause absolute elevation errors at the 10's of cm to meter level, but the drift should occur over long times scales (minutes to hours) so short segments of data should not be significantly affected by pointing drift in a relative sense. GPS determination of the aircraft trajectory, as computed by PNAV, will have unidentifiable absolute position errors over long time scales (minutes to hours) at the meter level but short-term relative accuracy usually at the sub 10 cm level. However, short-duration (seconds to minutes) absolute errors in the trajectory can also occur at the meter level but these periods are flagged by the trajectory PDOP and rms quality metrics. Instantaneous trajectory errors (from one epoch to the next) can also occur due to, for example, changes in satellites viewed or cycle slips. These instantaneous trajectory errors are often associated with instantaneous, large excursions in the PDOP or rms values. Such instantaneous trajectory errors can be detected by identifying abrupt steps in a plot of shot-to-shot difference in aircraft altitude. Potential errors in the determination of pointing biases and delays can effect the C.ELEVATION data at the 10's of cm level over time scales equivalent to the rate of change of roll and pitch. An error in determination of the system delay will impose an essentially constant bias in the C.ELEVATION of equal magnitude to the system delay error.

C.WIDTH, the distance from the first detected return to the "last" return, is derived by the following steps:
1) smooth C.WAVEFORM with a running average to reduce background noise using a filter width appropriate for the noise characteristics of the data file
2) visually determine a threshold level that is just above the background noise level in the smoothed waveforms, usually by over-plotting a large number of waveforms
3) for each laser shot determine the last crossing of the threshold level in C.WAVEFORM (corresponding to the final signal above the threshold)
4) determine the number of digitizer bins between the bin corresponding to the TIU stop-pulse time and the last return bin
5) convert number of bins to distance in meters based on digitizer bin resolution and the number of digitizer bins averaged together per waveform bin
6) decrease this computed distance by the impulse response pulse width (the pulse width from an idealized, flat surface) thus correcting for pulse broadening effects due to the transmit pulse width and the bandwidth of the detector; the impulse response pulse width is assumed to be the minimum valid pulse width for that data file

If the pulse reflection is from an unvegetated surface, C.WIDTH is a measure of the pulse broadening caused by varying return amplitude and surface slope and roughness. If the "last" return is a reflection from the ground surface beneath a vegetation canopy, then C.WIDTH is a measure of vegetation height and the ground elevation is equal to C.ELEVATION - C.WIDTH. This simple derivation assumes that the laser pulse is pointed at nadir. If the laser pulse was transmitted at an off-nadir angle, C.ELEVATION - C.WIDTH is an increased "slant range" distance which must be reduced to account for the off-nadir pointing. Usually, data is acquired in a near nadir orientation and this effect is small. The "last" return should be the ground in cases where there were sufficient holes through the canopy at nadir extending all the way to the ground. However, if the canopy is nearly or totally closed, the ground return may be too small to detect with this threshold approach or it may be entirely absent. In those case, the "last" return would be from a canopy layer and thus be above the ground. In some instances, the "last" return is instrument noise following the real ground return, and the "last" return is thus below the true ground surface.

The determination of C.WIDTH is based on selection of a smoothing filter width and a threshold level above noise that is based on the typical signal-to-noise properties of the data file as a whole. However, signal-to-noise properties change from pulse-to-pulse, due to many factors as described in the instrument description section, and thus this C.WIDTH determination is not optimized for individual pulses, leading to a fairly large proportion of shots with incorrectly identified "last" returns (either selecting a higher canopy return or noise following the ground return). An effective way to assess this is to overplot profiles of C.ELEVATION (i.e. the canopy top) and C.ELEVATION-C.WIDTH (i.e. the "ground") versus laser shot number. From context along the profile, one can see where the "ground" profile is pulled up into the canopy or extends below the ground, missing the true ground return. Note that for scanning data a simple profile plot of all the laser shots does not represent a linear trace along the ground. Rather, in map-view the profile follows a saw-tooth pattern across the scan pattern from beam 1 to beam 5 and then stepping forward to the start of the next scan. The resulting elevation profiles can have saw-tooth patterns if the laser beam scan was across a sloped surface. To avoid this effect, individual beams can be selected and plotted using the C.GALV_STEP parameter. It is frequently observed that C.WIDTH for the outer beams (1 and/or 5) is most often in error due to the lower signal-to-noise of these beam positions.

Detector saturation can induce errors in the determination of C.WIDTH, either due to the broadening of the saturated pulse and/or detector ringing causing signal above the background noise after the end of the pulse. Both effects cause detection of the "last" return after the actual end of the backscatter signal and an overestimated C.WIDTH. Transmit pulse quality can cause variable error in the determination of C.WIDTH. Although the amplitude of the trailing edge of the transmit laser pulse is designed to fall off asymptotically, misalignments within the laser transmitter can cause a pronounced step-like "back porch" on the pulse trailing edge. The magnitude of this back porch varies from shot-to-shot depending on return amplitude thus affecting the position where the last return threshold intersects the end of the return signal. Variations in C.WIDTH, and thus in ground elevation, due to this back porch effect are at the meter-level. The quality of the transmit pulse shape is essentially constant for all data acquired during a SLICER installation. Examination of impulse response waveforms from flat surfaces such as smooth water show the quality of the transmit pulse and the significance of this back porch effect.

IDL PROCEDURES FOR READING .GEO DATA:

The binary .geo files can be read into IDL with rdsldgeo.pro (run as @rdsldgeo), with the header data stored in the structure named "a", and the data stored in a structure named "c":

```
.run slcrd
print,'Select Data file!'
datafile=pickfile(path='')
print,'Opening ',datafile,' for input'
openr,datanum,datafile,/GET_LUN
readu,datanum,a
c=replicate({slicerdata},a.numshots < 200000)
print,'Reading ',a.numshots < 200000,' records'
readu,datanum,c
```

rdsldgeo calls the structure defintion contained in slcrd.pro:

```
a={slicerhdr,datetaken:'11/22/99',dateprocessed:'33/44/99',average:0L,wavepts:0L,beams:0L,gal
v_center:0L,numshots:0L}
b={slicerdata,gpstime:0.0D,tiu:0L,vernier:0L,starten:0L,roll:0.0,pitch:0.0,yaw:0.0,galv_step:0L,
galv_set:0L,lat:0.0D,lon:0.0D,alt:0.0,flat:0.0D,flon:0.0D,elevation:0.0,width:0.0,waveform:bytar
r(600)}
end
```

The text gps trajectory files can be read into IDL with rdgpsr+.pro (run as @rdgpsr+), with the data records stored in a structure named "t":

```
;Read GPS trajectory file
.run gps+
print,'Select GPS file!'
gpsfile=pickfile(path='')
print,'Opening ',gpsfile,' for input'
openr,gpsnum,gpsfile,/GET_LUN
wc=0L
readf,gpsnum,wc
print,'Reading ',wc,' records'
t=replicate({gpshdrplus},wc)
readf,gpsnum,t
t.lon = 360.0 - t.lon
```

rdgpsr+.pro calls the structure defintion contained in gps+.pro:

```
gps={gpshdrplus,gmt:0.0D,lat:0.0D,lon:0.0D,alt:0.0D,servs:0,pdop:0.0,rms:0.0,flag:0}
end
```

STRUCTURE OF .DAT FILES:

The determination of C.WIDTH provided in the .geo files is noisy, both in terms of correct identification of the ground return (causing potentially 10's of meters error) and proper determination of the end-of-signal when the ground return is properly identified (causing meter level errors). More sophisticated software has been developed which uses the noise characteristics of individual waveforms to establish a last-return detection threshold. This software also extracts the component of the waveform inferred to be the reflection from the ground beneath canopies and determines the beginning, peak, and end of the ground return. Errors in identification of the ground return and determination of its position are greatly reduced. Invalid laser shots (no range data, clouds and noise above an elevation threshold set above the highest expected ground, noise below an elevation threshold set below the lowest expected ground, and shots with anomalously low or high start pulse energy) are excluded. The software also derives the inclination and azimuth of the transmit pulse from the roll, pitch and yaw data, simplifying correction for waveform slant range distances. The diameter of each laser footprint, based on the laser divergence and the ranging distance, is also provided. These data products, referred to as .dat files, are not routinely produced and are not included on the .geo CD distributions.

The .dat file structure was created to be compatible with the BORIS data archive established for the BOREAS boreal forest project, and thus all parameters are stored as scaled long integers or bytes. The .dat files contain all parameters relevant to the end-user, but exclude .geo engineering parameters not directly relevant to end use applications.

SLICER file structure for BOREAS

file naming convention: YYMMDDLL.DAT
(e.g., 96072904.DAT is flight line 4 acquired on July 29, 1996)
binary files created on a SUN workstation

Parameter	Stored As	Bytes	Scale Factor	Scale Factor Precision	Output As	Bytes
header record:						
tiu_bin	long	4	none	-	long	4
dig2wf_average	long	4	none	-	long	4
wvfm_bins	long	4	none	-	long	4
numshots	long	4	none	-	long	4
Header Record Total:		16				16
data records:						
shotnum	long	4	none	-	long	4
beam	long	4	none	-	long	4
starten	long	4	none	-	long	4
gpstime	long	4	1.00E+04	double flt	double flt	8
diameter	long	4	1.00E+06	float	float	4
azimuth	long	4	1.00E+06	float	float	4
inclination	long	4	1.00E+06	float	float	4
latitude	long	4	1.00E+06	double flt	double flt	8
longitude	long	4	1.00E+06	double flt	double flt	8
elevation	long	4	1.00E+06	float	float	4
grndstart	long	4	1.00E+06	float	float	4
grndpeak	long	4	1.00E+06	float	float	4
grndend	long	4	1.00E+06	float	float	4
waveform	byte array	600	none	-	byte array	600
Data Record Total:		652				664

Parameter	Description
header record:	
tiu_bin	waveform bin corresponding to TIU stop-trigger event
dig2wf_average	number of digitizer bins averaged together to form a waveform bin
wvfm_bins	number of bins per waveform
numshots	number of data records in the file
data record:	
shotnum	laser shot number from the .geo source file
beam	cross-track beam position: 1 (left) to 3 or 5 (right) in scan mode; 0 in profile mode

starten	energy of transmitted laser pulse (uncalibrated units)
gpstime	time of laser pulse transmission (seconds past GPS midnight)
diameter	$1/e^2$ diameter of footprint = range x 0.02 mrad (approx.) laser divergence (m)
azimuth	azimuth of the laser pointing vector (degrees from north, 0 to 360 clockwise)
inclination	inclination of the laser pointing vector (degrees from horizontal; 90 = vertical)
latitude	latitude of the footprint at highest detected surface (decimal degrees N in WGS-84)
longitude	longitude of the footprint at highest detected surface (dec degrees E in WGS-84)
elevation	elevation of the highest detected surface (m w.r.t. WGS-84 ellipsoid)
grndstart	distance along laser vector from elevation to inferred start of ground return (m)
grndpeak	distance along laser vector from elevation to peak of ground return (m)
grndend	distance along laser vector from elevation to end of ground return (m)
waveform	raw return amplitude (uncalibrated units) per waveform bin

Like the determination of C.WIDTH in the .geo structure, the end of the last return (grndend) in the .dat files is determined by a threshold method. However, the threshold is derived on a per shot basis by establishing the mean and maximum background noise occurring in the last 10% of each waveform record. A potential error is introduced by using the end of the return to define the noise level if a greatly broadened signal, such as from very tall trees, extends all the way to the end of the waveform record. However, for the BOREAS region with its short vegetation, this effect does not occur. The start of the waveform record, prior to the bin corresponding to the beginning of the return signal, could also be used to establish the noise level. However, in some SLICER acquisitions the signal start bin is very near the beginning of the waveform, leaving an insufficient number of bins to adequately characterize the noise. The end-of-signal threshold level for a waveform is set to some scaling factor times the greatest noise value in the final 10% of the record. The scaling factor is adjusted depending on desired detection sensitivity (a higher scaling factor causes less false triggers on noise but more misses of small ground returns, whereas a smaller scaling factor causes more false triggers on noise but fewer misses of small ground returns). The appropriate scaling factor is determined based on visual examination of ground profile results for a selected set of test cases. If the waveform bins were derived from the digitizer bins without averaging, three adjacent waveform bins are summed together prior to establishing the shot noise level in order to increase the waveform signal-to-noise. The effective resolution of the ground distance measurements (grndstart, grndpeak, and grndend) is thus triples to 0.3336 cm. After threshold determination of grndend, the peak of the ground return is determined to be the location where the first change from negative slope to positive slope occurs prior to grndend (i.e., where the derivative of the waveform crosses 0). The start of the ground return (grndstart) can not be uniquely determined from the waveform record because vegetation directly above the ground can cause return signal that is not distinctly separated from the beginning of the ground return. Identification of the start of the ground return is accomplished by using the half-width of the return pulse from grndpeak to grndend. This half-width for the back half of the pulse is multiplied by a constant scaling factor, to account for the non-symmetric character of the Raleigh-shaped transmit pulse, in order to derive the half-width of the front half of the pulse. The scaling factor is determined from impulse response reflections from smooth water surfaces. With this technique, returns from vegetation directly above the ground should be separated from the ground return. This half-width determination of the start of the ground return is noisy, at the sub-meter level, due to the sensitivity of the grndend threshold determination. The grndend threshold technique is sensitive to the asympoptic fall-off in energy of the transmit-pulse trailing edge, and to the potential presence of a variable amplitude backporch during periods of non-optimal beam quality. A peak fitting approach to selection of the last return likely would achieve a more accurate identification of the ground signal.

In some instances of very low return amplitudes throughout the waveform the algorithm described above finds no signal above the grndend detection threshold. In those instances, the grnd start, grndpeak, and grndend positions are assigned to the last three bins at the end of the waveform as a flag to indicate failure to detect any waveform signal above threshold. _ In other instances of low ground return amplitudes, the threshold method misses the ground return but identifies a higher canopy return as the ground return. In some rare cases, the grndend position is properly identified but the grndpeak position misses the maximum ground return and instead identifies a maximum return from a higher canopy layer. The grndstart position is then placed higher up into the canopy by the half-width method. A final error is the identification of waveform noise occurring after, and thus below, the ground surface as the ground return. These various anomolous ground return detections are usually readily identified by examing a profile plot of the canopy top, grndstart, grndpeak, and grndend elevations.

The above described determination of the ground return position fails for saturated ground returns. In saturated cases, the threshold determination of grndend often corresponds to a small detector ringing peak after the return. Software checks are included to minimize detection of these false peaks. Because the saturated peak does not have a well defined inflection point, grndpeak is simply set to the mid-point of the saturation. Thus, the grndpeak to grndend half-width does not have physical significance. In these cases, grndstart is defined to be the point prior to grndpeak at which the slope of the waveform record decreases below a constant, but arbitrary, magnitude. The cutoff slope magnitude is selected so that the resulting grndstart positions look qualitatively correct. This determination of grndend, grndpeak, and grndstart is less accurate and more noisy when the ground return is saturated. Also, the grndstart technique for saturated returns will not effectively separate returns from vegetation directly above the ground. However, in most cases saturation of the ground return occurs where the surface is unvegetated and the intensity of the transmit pulse reaching the ground is not diminished by any overlying vegetation.

CDs containing .dat data include an IDL browser for reading and viewing the data. See the README.TXT file on the CD for a description of this data browser.

BOREAS SPECIFIC INFORMATION:

The GPS base station benchmark used was a survey mark selected by the GPS technician (Earl Frederick, NASA GSFC EG&G contractor) near the Prince Albert Airport runway. The following published position for the survey mark, in WGS-84 coordinates, was used to compute the trajectories: N 53 deg 12 min 59.160480 sec, W 105 deg 39 min 37.141840 sec, elevation 402.6360 meters. The GPS tripod and antennae was reinstalled over the benchmark for each flight and the height of the antennae above the benchmark was measured on each flight day. The time offset between GPS and GMT was 11 seconds during the mission. The GPS trajectories for each flight day show a large PDOP and rms excursion near 17 hrs GMT, due to a non-optimal satellite geometry. Differencing of the aircraft altitude in the GPS trajectory shows a small (less than 1 m) or no instantaneous error associated with the onset of these events, but short-term (seconds to minutes) trajectory errors at the meter-level could be associated with these events due to the PNAV Kalman filtering. The aircraft used was a C130-Q turboprop stationed at NASA Goddard's Wallops Flight Facility in Wallops Island, Virginia. The Inertial Navigation System used for pointing determination was a Litton LTN-92 laser ring gyro operating at 64 Hz and mounted directly to the laser altimeter optical bench. The system delay was determined to be 15.4 m based on altimetry data acquired across the GPS basestation site.

Waveform digitization parameters:

no averaging of digitizer bins (i.e., 1 waveform bin = 1 digitizer bin)
vertical resolution = 0.1112 m per bin
600 bins x 0.1112 m = 66.72 m height range
bin 28 corresponds to TIU detection of return from highest detected surface
(i.e., bin 28 = elevation)
bins 0 to 27 record background noise prior to first return
bins from end of the return to 599 record background noise after return signal
filter size for smoothing waveform prior to .geo width thresholding = 11 bins
threshold level to determine .geo pulse width = 23