Icarus **148**, 324–339 (2000)
doi:10.1006/icar.2000.6509, available online at http://www.idealibrary.com on IDEAL®

The Distribution and Physical State of H_2O on Charon

M. W. Buie and W. M. Grundy

Lowell Observatory, 1400 W. Mars Hill Road, Flagstaff, Arizona 86001
E-mail: buie@lowell.edu

Received November 23, 1999; revised June 30, 2000

We present new 1.4- to 2.5-μm geometric albedo spectra of Charon taken with HST/NICMOS in 1998. These new data provide global coverage of the surface with four spectra at evenly spaced longitudes. The surface of Charon is seen to be globally dominated by H_2O ice. The data indicate that the ice is in the crystalline phase at a temperature consistent with its heliocentric distance. The spectrum of Charon has only weak variations with longitude. There is an indication of a slightly stronger H_2O ice absorption on the leading hemisphere. No variations in the spectrum are seen in response to differing solar phase angles. From model fits there is no indication of any of the volatile species that are seen on Pluto, i.e., CO, CH_4, or N_2. There is spectroscopic evidence for a contaminant with an absorption coefficient which increases with wavelength past $\sim$2 μm. This contaminant is unidentified but is similar to what is seen on other icy satellites in the outer Solar System. We also present a standard spectrophotometric model for Charon that can be used to subtract Charon light from ground-based spectra of the combined Pluto–Charon system.

Key Words: Charon; satellites; spectroscopy.

1. INTRODUCTION

The study of Pluto in the past 2 decades has revealed a complex and compelling planetary body. This progress in understanding was aided in no small way by the mere presence of its only known satellite, Charon. For all the intricacy and difficulty of observing and modeling Pluto, very little attention has been paid to understanding Charon. This situation isn't one of neglect but a consequence of the extreme observing difficulties in measuring properties of a satellite that is always within 1 arcsec of its primary.

All of the major natural satellites in the Solar System have proved to be extremely important in furthering our understanding of the processes of Solar System formation and evolution. In particular, the icy satellites in the outer Solar System exhibit a wide variety of surfaces, interior structure, and varied cratering (or collision) records. There are three principal processes that contribute to the current assemblage of satellites: (1) capture that leaves the new satellite largely untouched (e.g., Phobos, Deimos, jovian irregular satellites), (2) impact capture that involves a substantial alteration of the satellite during the impact process (e.g., the Moon and Triton), and (3) *in situ* accretion (e.g., Galilean satellites, regular saturnian and uranian satellites).

The formation of Charon is most widely believed to be a result of the impact capture process (e.g., McKinnon 1984, Stren *et al.* 1997). This conjecture makes the study of the formation of Pluto and Charon particularly relevant to the understanding of the Earth–Moon system and might predict important differences between Charon and icy satellites which formed via other mechanisms. However, from the standpoint of size, albedo, density, and surface composition, Charon seems most similar to the principal uranian satellites, which, like the regular satellites of other gas giant planets, are presumed to have accreted *in situ* (e.g., Stevenson *et al.* 1986). This formation mode involves a secondary nebular accretion zone that is modified in temperature, pressure, and thus composition by the gravitational influence of the primary body (e.g., Prinn and Fegley 1981, Lunine and Tittemore 1993). One might expect this process to be dependent on the mass of the primary, producing a trend whereby the smaller the primary, the less its satellites differ compositionally and structurally from objects that form directly in the primary solar nebula. If so, then Charon lies at the extreme end of a continuum extending from jovian through uranian satellites. It remains to be determined whether the similarities between Charon and the uranian satellites are a fortuitous coincidence or if there is some "genetic" similarity that is not yet understood.

In two fundamental ways, Charon is an end-member satellite in the Solar System. It and Pluto are the closest in size of any planet–satellite pair. It is also the satellite of the smallest planet. In this specific context we use the term planet as an indicator of differentiation. It remains to be seen whether Charon itself is differentiated. The presence or absence of a strong water ice spectral signature from the surface may well be an indicator of differentiation for icy satellites. Charon may also be relevant as a tie-in to the population of trans-neptunian objects. Of the objects known to date, Charon is at the largest end of the population and may not be as unusual in the trans-neptunian region as is Pluto. We have yet to discern whether Charon is unique among the trans-neptunian objects based on its surface composition, albedo, or interior structure. It may be that its origin and

0019-1035/00 $35.00

present appearance are caused by the binary formation of the Pluto–Charon system. Many, if not all, of these questions will still remain after this paper but the importance of understanding Charon itself is clear.

Prior to the launch of the Hubble Space Telescope (HST), measurements of Charon were possible only in conjunction with the mutual event season in the late 1980s. Total eclipses or occultations of Charon by Pluto were extremely useful. From these events, the size of Charon was determined (cf., Tholen and Buie 1997) along with the optical albedo and color (Binzel 1988, Reinsch *et al.* 1994) and optical (Fink and DiSanti 1988) and near-IR spectra (Marcialis *et al.* 1987, Buie *et al.* 1987). The optical albedo and spectrum shows a bright and neutrally colored surface, although its albedo must drop off into the UV since the B-V color shows that the B reflectance is dropping off slightly. Some observations were tried from the ground without the aid of the mutual events in an attempt to extend the mutual event results to other longitudes on Charon's surface. These observations were very difficult and did not yield tremendously useful results (Bosh *et al.* 1992, Buie and Shriver 1994).

The Hubble Space Telescope provided a new and powerful tool for making direct measurements of Charon. Even before the optical repair mission, its enhanced spatial resolution allowed the first measurement of the optical lightcurve and global optical colors of Charon (Buie *et al.* 1997). These results indicated a very low-amplitude lightcurve and thus a small amount of global contrast on the surface and a constant B-V color with longitude. Additionally, the HST photometric results provided a measurement of the phase function of Charon which was quite different from that of Pluto. This phase function is extremely important for the analysis of our new data and will be discussed again in Sections 3.8 and 5.1.

The gap between ground-based observatories and HST has been closing in recent years. With greater attention to achieving the best possible seeing, some direct imaging observations now compete with HST. The bulk density of Charon has been measured both from HST and from the ground, showing that Charon has a lower density than Pluto (Null *et al.* 1993, Young *et al.* 1994, Null and Owen 1996). Even newer techniques based on adaptive optics on the largest telescopes are also being turned toward imaging, photometric, and spectroscopic investigations.

Regardless of the promise of adaptive optics, HST has two primary advantages that we exploited in the current work. Coupled with the high spatial resolution, the point-spread function (PSF) on HST is extremely stable when compared with ground-based observations. Also, HST is above all the components in the earth's atmosphere that blot out interesting spectral regions. With the addition of the NICMOS instrument to HST, we had an opportunity to make separate spectral observations of Pluto and Charon at wavelengths where both objects are spectrally active. The diffraction limit of HST is sufficient to separate the light of Charon from Pluto but not without non-trivial overlap. However, owing to the stable PSF on HST it is possible to remove the cross-Contamination and obtain individual measurements on both.

In addition to the benefit of the new measurements to the study of Charon, we also hoped to make observations of sufficient quality that they could be used to enhance future ground-based observations. Obtaining time on HST or a ground-based facility capable of splitting Pluto and Charon will be extremely difficult for some time to come. Most ground-based observations will still be forced to measure both Pluto and Charon together. The separate photometric observations of Buie *et al.* (1997) already provide a means for removing Charon from optical photometry. These new NICMOS observations provide the same ability to remove Charon from near-IR spectra of Pluto. This information will improve future efforts of long-term monitoring, measurements of seasonal effects, and support of the planned mission to the Pluto–Charon system. We also hope that our observations will help establish Charon as one of two targets that must be observed by a spacecraft mission rather than as an add-on to a pure Pluto-based mission.

In this paper we describe our efforts to collect and interpret 1.4- to 2.5-μm spectra of Charon with NICMOS on HST. There were considerable challenges in the successful processing of the data but we have succeeded in obtaining new, high-quality albedo spectra that cover all longitudes of Charon. We will discuss the observational strategy in Section 2, the details of our data processing system in Section 3, and direct (Section 4) and indirect (Section 5) results from our new spectra. Substantial details are provided on our data reduction techniques since other HST/NICMOS observations will benefit from what we have learned.

2. OBSERVATIONS

The NICMOS instrument on HST provides the capability for dispersed light observations using a grism. The instrument design was optimized for obtaining observations of multiple objects in a single image and thus does not use an entrance slit. Therefore, all objects in the field of view are seen simultaneously and appear in the image in zeroth, first, and second order. The instrument is optimized for throughput in the first order and thus our reductions are concerned with this order.

All grism observations are done with the NIC3 camera for which three grisms are available: G091, G141, and G206 with central wavelengths of 0.91, 1.41, and 2.06 μm. It was noticed sometime after NICMOS was installed on HST that the NIC3 camera could not be brought into focus simultaneously with the NIC1, NIC2, and the other HST focal plane instruments. As a result, two "campaigns" were scheduled where for an entire week the HST secondary mirror position was adjusted to compensate for the lack of focus range on the NIC3 camera. These campaigns were scheduled for January and June 1998. During the campaigns, NIC3 observations were diffraction limited. Outside the campaigns, observations with NIC3 suffered from broadened PSFs. However, as the cryogens in the instrument were gradually depleted, the focus of NIC3 improved over time. By the time of the second campaign in June 1998, the out-of-Campaign focus was nearly identical to that achieved during the campaign.

Observations of the Pluto–Charon system required special attention to the orientation of Pluto and Charon on the detector. Optimally, the spacecraft must be rotated to place Charon in a direction orthogonal to the spectral dispersion. This angle varies continuously due to the orbital motion of Charon around Pluto. To achieve the best Pluto–Charon separation it was necessary to require a rather specific roll angle (on the plane-of-the-sky). However, there are significant restrictions on permissible roll angles that depend on the target location and the object's solar elongation. Our observing plan required a roll angle that would ensure that the Pluto–Charon separation would be ≥90% of the possible separation. This roll angle constraint combined with the focus campaign schedule was nearly a null set so we were forced to accept observations out of campaign. Predictions for the camera behavior indicated that the PSF would not compromise our observations nearly as badly as accepting a non-optimal roll angle.

Our observations covered four distinct sub-Earth longitudes as summarized in Table I. All observations for a single longitude took place within a single orbit of HST (henceforth referred to as a visit). To maximize the signal-to-noise ratio of the observations, we chose to concentrate our efforts on the most diagnostic wavelength region provided by a single grism, i.e., the G206 grism. We chose to make our observations at four evenly spaced longitudes that would permit sufficient separation at each time. The table shows the mid-time of each visit along with the solar phase angle (g), sub-Earth longitudes on Pluto and Charon (in degrees), sub-solar latitude on Pluto and Charon, and the Charon–Pluto offset in pixels. Note that the x direction is closely aligned with the dispersion direction and the y direction with the spatial direction. Any offset in the x direction is wasted separation. The FWHM of the PSF was just under two pixels for all observations. The out-of-campaign observations in March 1998 (visits 1 and 3) had a somewhat broader FWHM (~10% wider) than that provided for in-campaign observations but did not adversely affect the data. Visit 2 was taken during a NIC3 focus campaign but the FWHM from visit 4, just out of campaign, is identical.

The 0.2 arcsec/pixel scale of the NIC3 camera does not quite resolve the PSF at 2 μm. We were concerned that the undersampled PSF might present some difficulty during data reductions. As a result, we devised a dithering strategy that could provide ample sub-sampling of the grism image. Our strategy combined three separate motions, dithering of a few pixels in the spectral direction, a larger dither in the spatial direction, and a combined fractional pixel dither pattern. The pattern we used is shown in Fig. 1. The fractional portion of the dither pattern provided an evenly spaced, 2-dimensional grid over a single detector as shown in the unit pixel map. The spectral dither (x) was just large enough that any repeatable pixel pattern would be offset by slightly more than a FWHM. The spatial dither (y) was large enough so that the objects moved by many times the FWHM amount to ensure that an accurate sky subtraction could be achieved. Our exposure times (192 s per frame) were long enough that we could only obtain 11 of the 12 dither positions during a single visit. All observations were obtained under fine-lock guiding providing pointing and tracking precision much better than the pixel scale.

TABLE I
Summary of Observations

Visit number	UT date	g (°)	Sub-Earth Lon.		Lat. (°)	Ch–Pl offset	
			Pluto	Charon		Δx	Δy
1	1998/03/17.8	1.80	50	230	21.2	1.66	3.37
2	1998/06/07.3	0.53	128	318	21.7	−0.98	3.08
3	1998/03/27.5	1.68	223	43	21.3	−1.46	−3.17
4	1998/05/28.8	0.42	311	131	21.6	0.04	−3.54

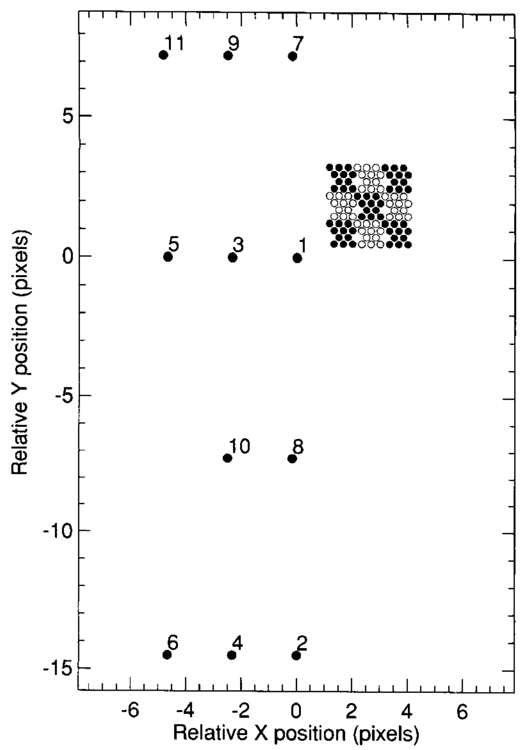

FIG. 1. Dithering pattern for the Pluto and Charon observations. The points labeled with numbers show the total absolute dither spacing for all 11 dithered images relative to the position of the first image. The smaller grid of points shows the fractional pixel dithering achieved by our pattern. Each cluster of points (either open or filled) shows a unit pixel with the fractional position of the dither. For our pattern, 12 positions would have fully filled the unit pixel. The missing position could not be collected in the time available during a single orbit.

3. REDUCTION OF NICMOS GRISM DATA

The extraction of spectra from the grism images presented significant data processing challenges. Numerous systematic errors were encountered at all stages of the processing. Throughout all steps, we relied on the requirement that the 11 spectra at a given longitude should be consistent within the uncertainties of the extracted spectra. Large differences were always seen to arise from some type of systematic error. We developed our own IDL-based extraction software for the extraction of spectra from the grism images. The following sub-sections detail the various steps and algorithms contained within our software.

3.1. Background Subtraction

All NICMOS images suffer from some degree of accumulated dark current and background signal. The dark current is similar to levels seen in detectors from ground-based instruments while the background signal is much smaller and less variable that seen from the ground. However, we still use techniques common for ground-based infrared observations. Most commonly, IR images are pair-wise subtracted with the object of interest offset to distinct regions of the detector. The difference between the two images cancels most of the sky signal leaving a positive and negative image of all non-sky, non-dark flux sources in the image. Alternatively, one can average multiple dither images to create a higher signal-to-noise ratio image of pure sky. For our Pluto–Charon observations we used four image subsets where the object images were not duplicated, e.g., dither positions 1, 2, 7, and 8 shown in Fig. 1. Combining these subsets yielded a sky image which was then subtracted from each individual image in that subset. The data taken on the comparison star were not as well dithered, so we had to use simple image pairs. All of the remaining steps deal with processing these sky-subtracted images.

3.2. PSF

Our observations were background noise limited where the noise arises from a combination of sky signal, dark current, and readout noise. In this noise regime, considerable improvements in the signal-to-noise ratio of the extracted spectra are achieved with an optimal extraction technique described by Horne (1986). The key to the optimal extraction is having an accurate point-spread function. The NIC3 camera PSF is not well represented by any simple analytic function. Also, the width of the PSF was a function of time as the instrument slowly warmed. Therefore, we constructed discretely sampled numerical profiles from the grism images of the objects themselves.

Constructing a numerical PSF is not as straightforward as taking a single cut across the spectrum. Instead, the PSF must be constructed from all the profiles at all wavelengths to ensure a sufficiently high signal-to-noise ratio in the profile. Contrary to our expectations, the PSF did not appear to vary with wavelength in any of the grisms. Therefore we treated all wavelengths equally when constructing the PSF. Each column (wavelength) thus provides one such sampling of the PSF relative to the spectrum centerline. Once normalized by the total object flux at that wavelength, the individual profiles can be stacked. The discrete points in all the stacked profiles must be smoothed to reveal the actual profile. We constructed a smooth curve by using a locally weighted scatterplot smoothing routine (LOWESS) as described by Cleveland (1979). The LOWESS routine uses a local weighting function around each point in the smoothed curve. We use a triangular weighting function that is unity at the output point and tapers to 0.1 at the edge. The half-width of the weighting function was set to 0.5 pixels. We fit a third-order polynomial using the data values and their associated weighting function. The output smoothed point is the value of this polynomial at the location of the output point. This process is repeated for every output point desired in the smoothed profile. We chose to create the numerical profile at 0.1 pixel grid spacing. In this process there are numerous bad pixels in the phased profile so we actually find the smoothed PSF from applying the LOWESS filter twice. Between the two passes any points that deviate unusually from the smoothed curve are omitted. This procedure produces an exceptionally robust numerical representation of the PSF in each image. Some example PSFs are shown in Fig. 2. The profile

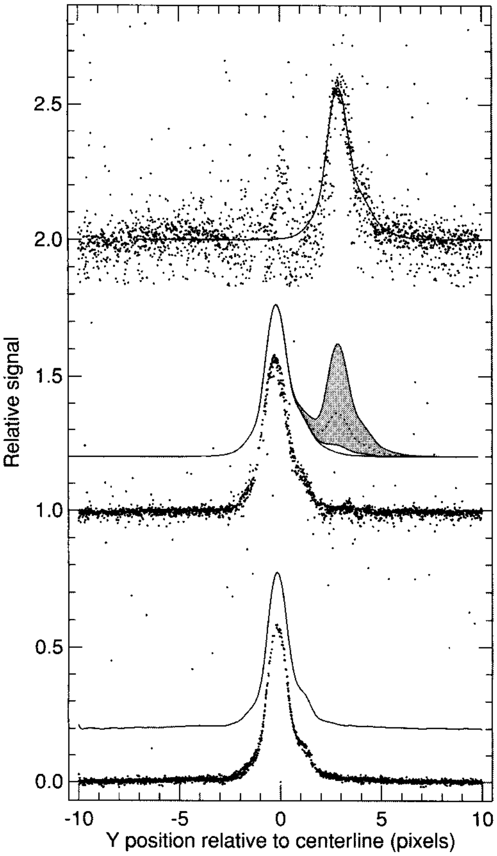

FIG. 2. Spatial profile from sample NICMOS grism images. In all cases, the dots represent intensity samples of the PSF from a single pixel (after suitable scaling). The bottom plot shows a phased spatial profile of calibration star P330E taken during a NIC3 focus campaign. The solid curve shows the smooth numerical profile obtained from the discrete profile. The middle plot shows a phased spatial profile of Pluto after the effects of Charon have been removed. The shaded region indicates the variable amount of flux that Charon contributes in the vicinity of Pluto. The amount contributed depends on the wavelength. The top curve shows the phased spatial profile of Charon after Pluto has been subtracted. The overlain curve is the Pluto curve shifted to match Charon's position. All profiles are shown such that the area of each profile integrates to unity.

at the bottom for P330E shows a simple example of a phased spatial profile and its smoothed representation for a single point source. The more complicated case for extracting a PSF from a two-source image will be discussed later in Section 3.7.

The extraction of the profile depends both on the location of the spectrum and on the spectrum itself. Fortunately, the extraction process converges well. A simple estimate of the centerline and the spectrum can be used to get a first approximation of the PSF. This PSF can then be used to better determine the centerline and spectrum. This process is iterated until the PSF, spectrum, and centerline cease to change by meaningful amounts.

3.3. Flat Fielding

The NIC3 camera has substantial low-frequency sensitivity variations across the image. This pattern is also quite dependent on the wavelength of illumination. However, the construction of a flat field for a slit-less spectroscopic image is non-trivial. To flat field the grism images we first constructed a virtual flat field based on all the narrowband imaging flat fields. The NICMOS filter set has quite a few filters over the wavelength range of the G206 grism but the coverage is quite limited for shorter wavelengths. For each pixel in the array, we fit a polynomial to the set of all flat field values as a function of wavelength. We investigated polynomials from order 1 through 5. Linear and quadratic fits were clearly unable to follow the variations. Benefits are noticeable up to order 4 and thus we chose fourth-order polynomials to construct the virtual flat. From the set of polynomials we could then compute a sensitivity for any pixel at any wavelength.

The final step of flat fielding comes from computing the wavelength of illumination at each pixel. Each pixel's polynomial was evaluated for its illuminating wavelength. The final set of numbers was then divided into the image to apply the flat field. The wavelength calculation assumes that a column is precisely monochromatic. This is an excellent approximation for the object itself. However, it is not particularly good for any residual sky signal since it comes from the entire grism bandpass and is not monochromatic. However, the image being processed comes from a pair-wise subtraction or subtracting the sky flux directly and the remaining sky signal is extremely small (typically 0.25 DN). Any errors caused in flat fielding the sky will be very small and will only distort the sky flux. Whatever sky remains will be taken out by other means in later processing steps.

3.4. Spectrum Centerline

Normally, finding the centerline of a spectral image is a simple process. However, we found some rather strong instrumental effects in the apparent centerline that were caused by an interaction between the detector array and the object PSF. The centerline of the spectrum is slightly inclined relative to the rows of the detector. This tilt causes the spectrum to drift slowly across two to four rows from one end to the other. The extracted centerline of the spectrum (whether done by finding the brightest pixel, taking a simple 1-d centroid, or computing an optimally extracted position) shows significant deviations from a simple straight line. On average, the centerline is linear but there are oscillating deviations relative to this average line.

As noted earlier, the PSF is slightly under-sampled in the G206 grism data. This under-sampling is responsible for the deviations in the centerline in two different ways. An under-sampled image becomes much more sensitive to the spatial distribution of quantum efficiency within each pixel. This effect is known as the pixel response function (PRF). The PRF has been measured in detail on a few optical CCDs (e.g., Kavaldjiev and Ninkov 1998) and reveals a very complex sub-pixel structure caused by the embedded detector electronics. There has been no analogous measurement on the arrays present in NICMOS so we are forced to attempt to deduce the PRF from our data. This effect should be a constant of the detector and thus applies in the same way regardless of the focus or PSF history. This effect is present in our data but is *not* the dominant cause for the centerline oscillations. We will return to a discussion of the PRF correction in Section 3.5.

A second effect is due entirely to the under-sampling of the image and causes an apparent shift of the centerline that is not real. Unlike the PRF, this effect does *not* have any impact on the extracted flux except for its confusing effect on determining the spectral centerline. To understand this effect, consider a spectrum formed from a δ-function PSF. In this case, all the light from the object at a single wavelength would fall completely within a single pixel. As the spectrum centerline slowly crosses a row, the position will not appear to change until the centerline crosses into the next row. Of course, the true centerline in this case is a smoothly varying function of column but it will appear from the image that the centerline looks more like a stair step. The degree to which the stair step is seen will depend on the size of the PSF and thus the effect is strongest at shorter wavelengths (e.g., grism G091).

To obtain an accurate centerline we had to use the following procedure. (1) Fit a first-order polynomial to the object position at each column. (2) Subtract the polynomial fit from the positions creating the first position residual. (3) Compute the fractional row position of the polynomial at each column. (4) Fit a two-term Fourier series to the first position residual as a function of the fractional row position. (5) Subtract the Fourier series fit from the original object positions creating the second position residual. (6) Fit a first-order polynomial to the second position residual as a function of column. This final fitted line is now taken to be the actual centerline of the spectrum for later extractions. Note that this centerline is mildly dependent on the PSF and the spectrum so that the centerline position needs to be iterated with the PSF and spectrum extraction.

3.5. Pixel Response Function

The pixel response function is a term used to refer to the variation in sensitivity across a single pixel. Most array detectors have some degree of variation across a pixel but usually this can be ignored since the average sensitivity will suffice for most data. However, in the case of under-sampled images, this variation cannot be ignored. The precise form of the sensitivity variation in the NICMOS detectors is not known. We have modeled the variation as if the sensitivity of the edge of a pixel is lower than the sensitivity of the center of the pixel and the variation is linear. Thus our correction can be thought of as a flux correction for a triangular function that will rectify the flux in a pixel by the ratio of areas of a uniform pixel divided by the area under the triangular function.

This simple ratio of areas is not sufficient to make a useful correction. Taken by itself, this correction would just multiply all pixels by the same value. For under-sampled data with our idealized function, image areas where the absolute value of the derivative of the intensity is large will have a different correction from flat image areas. So far, this discussion has assumed a 1-dimensional form for the correction. In reality the correction should be 2-dimensional. However, we note that the image derivatives are much stronger in the spatial direction than they are in the spectral direction. Therefore, we based our correction on just the image derivatives in the spatial direction.

Our approach requires the determination of a single number which is the edge sensitivity of a pixel. This number must be inferred from actual grism observations. Given a value, we extract a set of spectra of the same (non-varying) object that are taken at different dither positions. For this work we used the calibration observations of P330E which is a relatively bright solar-type star (Colina and Bohlin 1997). These observations consist of two sets of four observations taken during the first and second focus campaigns. With a perfect extraction, all eight spectra should be the same to within the limits of the intrinsic noise of the observations. We varied the edge sensitivity and found that a value of 0.413 minimized the scatter in the ensemble of star spectra. This correction should be a constant of the detector and should be equally valid regardless of the focus setting for the instrument. Thus this correction should be applied identically to all data regardless of which grism or what the instrument focus was set to.

This discussion may sound somewhat ad hoc. It should be noted that the PRF correction is small, generally amounting to no more than a 5 to 10% adjustment. In the end it is our requirement that all extracted spectra be the same that gives us confidence that the extraction methods are reasonable. We note that this correction appears to be unnecessary for the G206 data but is important for both G091 and G141. Also, for data that are well dithered, the final shape of the spectrum will not be changed by applying the PRF correction. Instead, the noise level in the final averaged spectrum will be reduced. The eight positions available for the P330E actually only give four distinct y dither positions. This set is barely sufficient for a proper reduction while the dithering from our Pluto/Charon observations is much more appropriate. The lesson we've learned from these data is that the best spectra (particularly calibration data) will come from well-dithered data, with at least four to eight positions. More is clearly better.

3.6. Single Source Extraction

Extracting a spectrum from a single source in the grism image is straightforward. A good centerline plus a good PSF will directly yield the desired spectrum using the optimal extraction summation. Getting a good centerline and PSF requires a certain degree of iteration with the extracted spectrum. The only subtlety involved in the spectral extraction is in the case of weak sources where a single frame does not yield a sufficiently good PSF profile. Is these cases, the PSF can be constructed from an average of all the profiles in the given set of dithered spectral images. At all stages of the iterative spectral extraction, the number of iterations or method of extraction is controlled by the internal consistency of the spectra from a single set. This extraction method was used for the solar-type comparison star, P330E, observed as part of the NICMOS calibration observations.

3.7. Double Source Extraction

This mode of spectral extraction was created to handle the case of two slightly overlapping point-source objects where the relative positions of the two sources are assumed to be known. This was precisely the case presented by the Pluto–Charon system. The position of Charon relative to Pluto on the chip was computed from information in the *Astronomical Almanac* (1998, Section F) for the apparent orbit of Charon. The motion of Charon relative to Pluto was slow compared to the time span of observation during a single visit. Thus the information from the *Almanac* was interpolated to the mid-time of the orbit visibility window. The orientation of the sky projected onto the NIC3 camera was taken from the ORIENTAT keyword provided with the pipeline data from STScI. Any errors that may exist on these geometric quantities were much smaller than the pixel scale of the detector and can thus be ignored.

The double source spectral extraction begins from a centerline and PSF derived from a single source extraction of the two sources as one. This step yields an odd looking PSF which is caused by the presence of two sources with differing spectral signatures. Nonetheless, the centerline thus provided is quite close to the true centerline and provides a useful reference guide.

The double source extraction proceeds as follows: (1) make an initial guess at a single source profile. In this step we fold the initial PSF about its peak and take the lowest signal from the folded PSF. This will create a symmetric PSF using the side of the PSF not affected by the second source. (2) At each wavelength, fit the current PSF to the data for two sources. For this step we resample the PSF to match the sampling grid at each wavelength. Then using the known positions of Pluto and Charon we perform a least-squares fit of the PSFs to the data where the areas under each PSF are the only free parameter. These areas are the points in the output spectrum that we desire and formal uncertainties from this fit are retained. This step assumes that the PSFs are identical for both objects. While it is true that Pluto's angular size is twice that of Charon, at these wavelengths the PSFs do not differ appreciably; both can be treated as point sources. (3) Given the new spectrum and current PSF, generate a new PSF. To do this, we subtract a scaled PSF located from Charon's position, leaving behind only the PSF from Pluto. Once Charon is subtracted, the phased spatial profile is created in the same manner as for the single point source extraction. (4) Iterate steps 2 and 3 until the final spectral extraction is reached.

In practice, the extraction process is even more complicated than indicated in the above description. The removal of Charon from Pluto's PSF is never quite perfect. We included numerous ad hoc constraints to try and improve the determination of a master PSF but there were small-level numerical instabilities in the iterative extraction. The Charon flux extracted is *extremely* sensitive to the precise shape of the PSF. Even a small error in the PSF of Pluto near Charon will lead to a rather substantial (usually systematic) change in the overall flux of the spectrum. These problems were not quite as pronounced for the in-campaign data but even they were affected at a non-trivial level.

In the end, we had to combine profiles from pairs of visits to create a master PSF for the spectral extraction. Fortunately, the geometric constraints on scheduling the observations also gave us the PSF data needed. The observations in visits 1 and 3 (and likewise for 2 and 4) had Charon on opposite sides of Pluto. These visit pairs were also taken within a relatively short time of each other. The time differences are small enough that the PSF did not change between observations and could thus be combined. So, for these pairings of visits, we constructed a master PSF by using just the side of the PSF opposite Charon where there are no numerical instabilities in the PSF determination. This new PSF was then used for the final spectral extraction. An example of the separated profiles of Pluto and Charon is shown in the upper two plots in Fig. 2. The solid curve shown with the Pluto profile is the master PSF for visits 2 and 4. This case shows the worst separation of the four visits (No. 2) and is a typical case from the 11 frames in the visit. The shaded area on the middle profile shows the variable degree of Charon signal relative to Pluto. This variation in relative signal is caused by fundamentally different spectral features on the two objects. The Charon/Pluto flux ratio varies between 0.05 and 0.72 across the G206 grism image. If both objects had the same spectral albedo this flux ratio would be more like 0.26, which is the ratio of their projected areas.

3.8. *Final Averaged Spectra*

All of the spectra obtained during a single visit were combined to make a single spectrum. In the case of the two source extractions (Pluto and Charon), the individual spectra were combined in a weighted average using the uncertainties from the extraction. For single source extractions, the data were combined with equal weight and the scatter of the averaged points determined the final uncertainty.

Because the spectra are dithered in the spectral direction as well as in the spatial direction, the wavelength scales and sampling bins are not identical between all spectra. To combine spectra we defined a fixed regular grid of wavelength bins with a dispersion of 0.01152 μm/pixel, which is identical to the dispersion of the grism. The individual spectrum values that fall in each bin are then used to derive an average flux for that bin. No attempt was made to interpolate the data before combining. We felt that the nature of the under-sampled image prevented a reliable interpolation. A slight increase in effective resolution may be possible but this would require a much denser dithering pattern in the spectral direction. The penalty for our approach is that our spectra will show slightly inflated uncertainties in regions of steep spectral slopes.

The average spectra thus created are in units of counts per second but still include any grism efficiency functions and other throughput terms. A ratio of any two spectra will quite effectively remove all of these effects and leave a spectral flux ratio. To derive the geometric albedo spectrum of Pluto and Charon, we applied the necessary geometric corrections to the raw Pluto and Charon fluxes, divided by the raw flux of P330E, multiplied by the absolute flux of P330E, and divided by the absolute flux of the Sun. The absolute flux of P330E was provided by Colina and Bohlin (1997) and is estimated to be good to 3–5%. Their flux spectrum was at a higher resolution than our spectra and on an irregular grid. To get a flux spectrum for our data we interpolated the original flux spectrum to a uniform grid. Then, the data were sinc-filtered to reduce the spectral resolution such that our dispersion would critically sample the spectrum. We applied the same procedure to the solar flux of Colina *et al.* (1996).

The geometric corrections to the Pluto and Charon fluxes included the normal geocentric and heliocentric distance corrections. The pairs of visits were done at nearly the extreme values of phase angle possible and the phase angle effects were quite significant. We used the separate phase coefficients of Buie *et al.* (1997) to correct the fluxes to zero degrees phase angle using a strictly linear correction (as a function of magnitude). Once corrected for distance and phase angle the spectra were converted to geometric albedo using the radii adopted in Buie *et al.* (1992) ($R_{\mathrm{pl}} = 1150 \pm 7$ km; $R_{\mathrm{ch}} = 593 \pm 10$ km).

4. SPECTRAL RESULTS

The final globally averaged geometric albedo spectrum of Charon is shown in Fig. 3. Shown for reference is the geometric albedo of Charon as measured in Johnson V (Buie *et al.* 1997). Given the differences in absolute calibration procedures, the V albedo is consistent with the albedo near 1.8 μm even though V appears to be slightly lower. The uncertainties plotted for each data point include all known noise sources with the exception of the uncertainty on the absolute flux of P330E. The general increase in the size of the error bars at the ends of the spectrum is a consequence of the diminished instrumental throughput defined by the order-sorting cutoff filter used with the G206 grating.

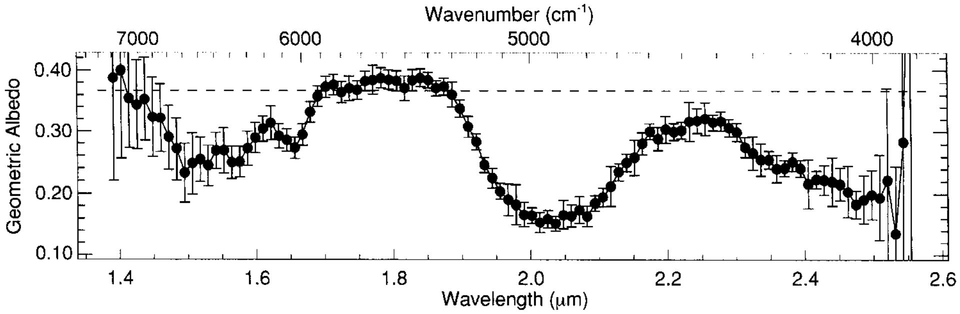

FIG. 3. Global average geometric albedo spectrum of Charon. The dashed line is the *V* geometric albedo of Charon.

This new spectrum clearly shows the unmistakable signature of water ice as had been seen during the mutual event season (Marcialis *et al.* 1987, Buie *et al.* 1987), with better spectral sampling than the older mutual event data. The 1.5- and 2.0-μm bands are nicely confirmed as well as the increased absorption approaching the 3.0-μm H_2O fundamental. A detailed analysis of the overall spectrum is presented in Section 5. One new result is the detection of a minor H_2O ice band at 1.65 μm. This weaker band is quite diagnostic since its strength is strongly correlated with the temperature and phase of the ice.

One advantage of our new data is that they cover four evenly-spaced longitudes over the surface of Charon. In contrast, the mutual event data could only constrain the sub-Pluto hemisphere. Figure 4 shows the spectra obtained at each of the four visits. These four spectra are considerably more similar than they are different. The iterative nature of the spectral extraction forces spectra within a single visit to be self-similar but spectral similarities between visits are completely unconstrained. We now see that the common assumption of the global presence of water ice on Charon is vindicated.

While the spectra are remarkably similar, there are a few subtle differences to be found. Figure 5 shows the continuum flux of Charon as a function of the east longitude. The continuum flux shown is the average of the Charon spectrum from 1.74 to 1.86 μm. If albedo variations are due to variable quantities of opaque continuum absorbers mixed into the surface ice, one might expect the 1.8-μm lightcurve to track the *V*-band lightcurve. However, based on the error bars, there appear to be formally significant differences between the 1.8-μm and visible lightcurves, perhaps providing evidence for a variably distributed, non-gray spectral component on Charon's surface. However, the differences shown are quite small, roughly 3%, and we do not feel confident that they are real. Instead, we believe that mismatches between the 1.8-μm continuum and the visible lightcurve are more likely to be indicative of the level of our systematic errors in the absolute flux determination from one visit to the next. Taking all four spectra together, the internal systematic errors appear to be somewhat random. Therefore, we chose to perform a simple average of the four geometric albedo spectra to get the grand average shown in Fig. 3. An alternate approach would be to correct the geometric albedo spectra to make them consistent with the lightcurve. Since the lightcurve itself is not determined well enough, we preferred a simple average.

Regardless of possible systematic errors influencing *absolute* levels in Fig. 5, the *relative* levels within a single spectrum should remain unaffected. Figure 6 shows the depth of the 2-μm water ice band as a function of longitude. The continuum point is the same as that plotted in Fig. 5. The albedo in the band is the average from 1.98 to 2.10 μm. The plot shows the difference in these two levels such that increasing band depth is toward the bottom of the plot. For reference, the longitudes of the leading and trailing hemispheres are shown as well. While this result is only marginally significant, it would appear that the 2-μm H_2O band is slightly stronger on the leading hemisphere. Such a leading–trailing asymmetry is common on other icy

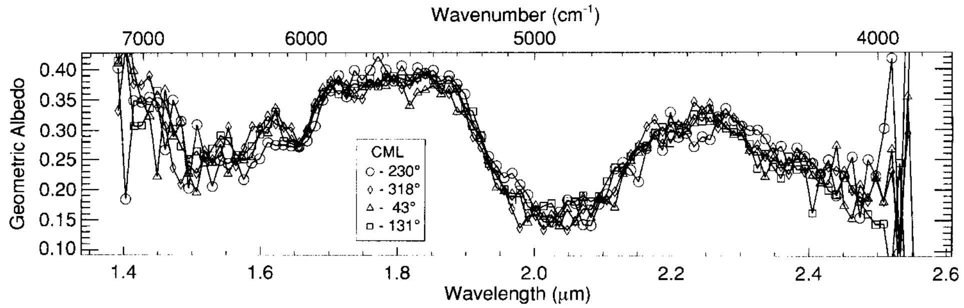

FIG. 4. Spectra of Charon at four longitudes. The legend shows the central meridian (east) longitude on Charon for each spectrum.

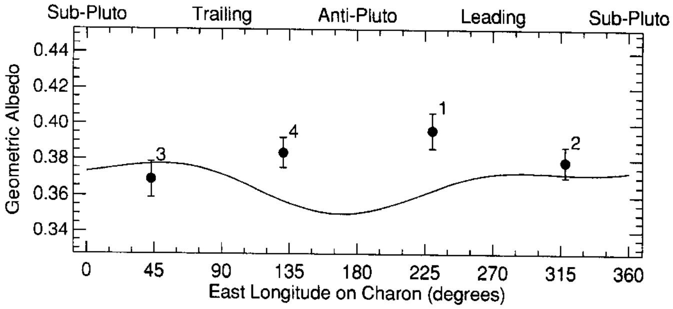

FIG. 5. Comparison of 1.8-μm continuum and V lightcurve. The points represent the average of the 1.8-μm continuum region in the NICMOS spectra. The solid curve is the V lightcurve from Buie *et al.* (1997).

satellites (cf., Cruikshank *et al.* 1998a, and references therein) and indicates a variation in cratering rates caused by the orbital motion of the satellite about its primary. Charon's apparent asymmetry is quite small, which is perhaps no surprise given its relatively slow orbital velocity and the substantially smaller mass of the primary. Regardless, this may be an indication that the surface of Charon is old enough to be well cratered.

One of the key steps in our calibration of the spectra involved using the visible photometry phase behavior to correct the spectral fluxes of our new infrared data. Photometric phase behavior is known to be a function of surface texture and perhaps composition. The correlation of phase behavior with albedo is less understood and one might expect the phase behavior to vary with wavelength. As a check, Fig. 7 shows the water ice band depth (points from Fig. 6) plotted as a function of solar phase angle. This band depth covers a factor of 2 in albedo and is much larger than the difference between the optical and infrared continuum. This figure clearly shows there is no significant spectral dependence of the phase coefficient. The lack of any correlation gives us confidence that our application of the optical phase coefficient is justified.

5. SPECTRAL MODELING

To further understand our spectral data, we must construct multiple scattering radiative transfer models to simulate the interaction between photons and Charon's particulate surface materials as functions of wavelength and viewing geometry. For this purpose, we use the Hapke model (e.g., Hapke 1993) as implemented by Grundy (1995). This model accounts for multiple scattering in multicomponent granular materials, as well as macroscopic surface roughness and the shadow-hiding opposition effect. Reflectance of an unresolved object is calculated by breaking the surface of the object into discrete facets and then integrating the bidirectional reflectance contributions of all facets which are simultaneously visible and illuminated. These model spectra are matched to the globally averaged spectrum of Charon, since the observed longitudinal variations are so slight.

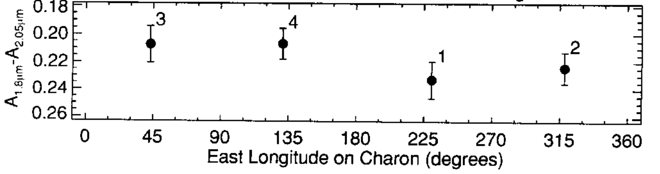

FIG. 6. Water ice absorption band depth versus longitude. The points show the geometric albedo difference between the 2-μm water band and the 1.8-μm continuum as a function of the sub-Earth longitude. For reference the longitudes of the leading and trailing hemispheres are marked.

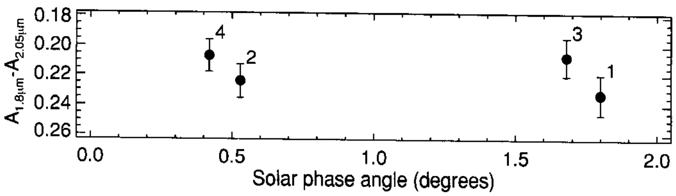

FIG. 7. Albedo dependence of the phase coefficient. The data points from Fig. 6 are plotted against the solar phase angle at the time of observation.

5.1. *History*

Previous efforts to model Charon's spectrum have been based on a low spectral resolution, under-sampled, near-infrared spectrum of Charon's sub-Pluto hemisphere obtained by Buie *et al.* (1987) by subtracting the Pluto-only flux measured during a mutual event from the combined flux measured prior to the event. Roush (1994) was able to match the near-infrared data using Hapke models of H_2O ice, either pure or intimately mixed with grains of other ice species. Because those models ignored continuum absorption, they predicted high albedos inconsistent with values of Charon's geometric albedo measured at visible wavelengths (e.g., Tholen *et al.* 1987, Fink and DiSanti 1988, Reinsch *et al.* 1994, Buie *et al.* 1997). These albedo measurements, especially the CCD spectrum of Fink and DiSanti (1988), show Charon to be spectrally featureless from the visible to 1 μm, with an albedo similar to that of the uranian satellite Ariel. A slight blue slope in the Fink and DiSanti (1988) spectrum is probably spurious, resulting from their use of a solar analog star redder than the sun (e.g., Grundy and Fink 1996). Nevertheless, the Fink and DiSanti (1988) data and other albedo measurements provide extremely valuable constraints on Charon's composition which any proposed Charon model must satisfy.

Grundy (1995) and Grundy and Fink (1996) considered both Fink and DiSanti (1988) and Buie *et al.* (1987) data sets simultaneously, obtaining good fits by including spectrally neutral continuum absorbers intimately mixed with H_2O ice (any species, such as carbon black, is effectively neutral if at all wavelengths its grains have single scattering albedos considerably lower those of the H_2O ice grains). A typical Grundy (1995) model consisted of 50-μm grains, 88% H_2O ice and 12% a continuum absorber with absorption coefficient 100 cm^{-1}. Roush *et al.* (1996) undertook an effort to absolutely calibrate the Buie *et al.* (1987) data set and, after including continuum absorption to match the calibrated data, reached similar conclusions. Roush *et al.* (1996) went on to explore the possibility of other absorbers contributing to Charon's spectrum, but because of the low spectral resolution of the Buie *et al.* (1987) data, had limited success in putting meaningful constraints on Charon's surface composition.

A significant defect of these early modeling efforts was that they relied on optical constants compiled by Warren (1984) for H_2O ice near 270 K. Plausible temperatures for H_2O ice on the surface of Charon are much colder, in the ~40–70 K range. It is now possible to remedy this shortcoming by using recent H_2O ice absorption coefficients from Grundy and Schmitt (1998) and Schmitt *et al.* (1998). The new laboratory data show significant temperature-dependent spectral effects over the entire temperature range from 20 to 270 K. Low-temperature crystalline ice data yield better fits to the Buie *et al.* (1987) data (e.g., Douté *et al.* 1999) as well as to our new data and can even be used to estimate the temperature of Charon's H_2O ice, as will be discussed in Section 5.5.

Another problem with the early models is that their variation with phase angle, g, was somewhat arbitrary. The authors simply fixed unconstrained Hapke parameters at values thought to be reasonable. Both used isotropic single-particle phase functions, $p(g)$, and zero surface roughness, $\bar{\theta}$. Roush (1994) fixed the compaction parameter, h, at 0.05 (the lunar regolith value) and assumed that the opposition surge parameter, B_0, was equal to the reciprocal of w, the single scattering albedo. Grundy (1995) assumed $h = 0.26$ (appropriate for a step-function vertical density profile with a single grain size and 50% porosity at depth) and devised a different ad hoc method of calculating B_0 from particle properties. More recently, based on visible wavelength HST photometry of Charon at phase angles between 0.5 and 2°, Buie *et al.* (1997) derived a linear phase coefficient, $\beta = 0.087 \pm 0.008$ mag deg^{-1}. This value is considerably larger than Pluto's phase coefficient 0.029 ± 0.001 mag deg^{-1} and differs substantially from the phase behavior of the earlier models.

Consistent phase behavior needs to be included in future models of Charon's surface. The Buie *et al.* (1997) result can be used to somewhat constrain Hapke parameters appropriate for Charon's surface. For example, values of $h < 0.02$ are consistent with Charon's phase function between 0.5 and 2°, but h cannot be as large as the lunar value, 0.05. After an exploration of the Hapke parameter space consistent with the Buie *et al.* (1997) β value, we adopted the following: $h = 0.01$, $B_0 = 0.6$, and $\bar{\theta} = 20°$. For $p(g)$, we use a two-parameter double Henyey–Greenstein function with asymmetry 0.615 and forward weight 80%. This configuration fits the Buie *et al.* (1997) photometry and is comparable to parameters reported for other icy satellites, but is by no means unique. Other parameter values exist which give equally good fits to the Buie *et al.* (1997) data, but very different behavior at other phase angles. Extreme caution should be exercised in extrapolating our model beyond $0.5° < g < 2.0°$. In order to avoid such an extrapolation, we henceforth correct all data to $g = 1°$ and simulate observations at that phase angle. We are thus not modeling Charon's geometric albedo (a quantity defined only for zero phase angle), but instead its $g = 1°$ reflection spectrum relative to a normally illuminated and viewed Lambertian disk of the same radius at the same distance from Earth and the Sun. This $g = 1°$ albedo is hereafter indicated by A_1.

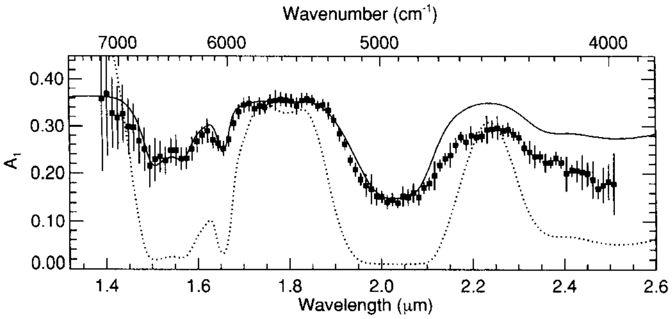

FIG. 8. Models of Charon's reflection spectrum at 1° phase angle (A_1, see text), consisting of water ice and a spectrally neutral continuum absorber, with proportions adjusted to match Charon's continuum albedo around 1.75 μm. The solid curve is a model with 80-μm H_2O grains, which fits the data well at wavelengths shortward of ~2 μm, but becomes too reflective at longer wavelengths. Increasing the grain size increases absorption at these longer wavelengths, but also increases absorption at shorter wavelengths, as shown by the dotted curve, a model with 1-mm grains. Even with multiple grain sizes and multiple terrains, with only H_2O and neutral continuum absorbers, it is impossible to match Charon's water bands at the same time as producing significantly lower albedos at 2.2 than at 1.8 μm.

5.2. H_2O and Continuum Absorbers

Previous Charon models (e.g., Grundy 1995, Roush *et al.* 1996, Douté *et al.* 1999) simulated ice combined with a spectrally neutral continuum absorber. For an initial effort to fit our new Charon data, we tried the same approach. Figure 8 shows two such models, both intimate mixtures of ice and a dark, neutral absorber, with mixing proportions adjusted to match Charon's continuum albedo around 1.75 μm. While these models exhibit spectral features similar to those of Charon, they fail to match all wavelengths simultaneously. Low-temperature crystalline H_2O ice and spectrally neutral continuum absorbers cannot duplicate Charon's albedo contrast between 1.8 and 2.2 μm at the same time as they fit the water ice bands at 1.6 and 2.0 μm, because the absorption coefficients of H_2O ice are quite similar at 1.8 and 2.2 μm (Grundy and Schmitt 1998). Only at extremely large grain sizes does H_2O ice begin to have a lower albedo at 2.2 than 1.8 μm, but large grains are ruled out by the shapes of the H_2O absorption bands at 1.6 and 2.0 μm. Evidently Charon's continuum absorber must not be as spectrally neutral in the infrared as it is in the visible. Beyond about 2 μm, continuum absorption must increase.

This spectral characteristic is not unusual among icy outer solar system objects. Many jovian, saturnian, and uranian satellites are similar in having lower albedos at 2.2 μm than at 1.8 μm (Grundy *et al.* 1999). Curiously, the effect is less pronounced for uranian satellites than it is for Charon and saturnian satellites, so

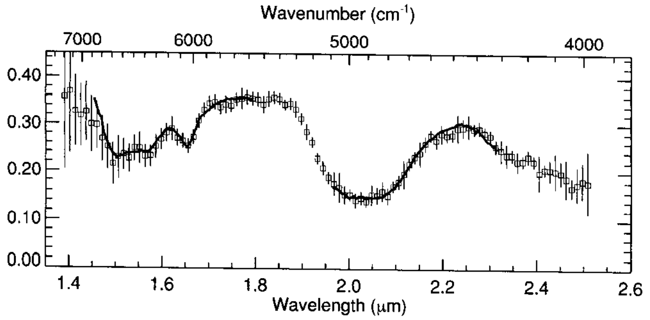

FIG. 9. Charon reflection spectrum (open squares) compared with a spectrum of Tethys (solid line, from Grundy *et al.* 1999). The Tethys spectrum was scaled and a constant offset added to match Charon at 1.8 and 2.0 μm. The errors on the Tethys spectrum are comparable to the width of the line.

there is no clear trend with temperature, heliocentric distance, or mass of the primary planet. Figure 9 shows a scaled spectrum of the trailing hemisphere of Tethys (Grundy *et al.* 1999) superimposed on our Charon average. The two spectra are remarkably similar at these wavelengths, suggesting that similar material may be responsible for depressing albedos at long wavelengths on both satellites. In order to match Charon's spectrum, it is necessary to include a proxy for this unknown continuum absorber, something that is blue at near-infrared wavelengths, is abundant in the outer solar system, and has little if any spectral activity at visible wavelengths, where spectra of icy satellites uniformly lack strong spectral features.

To better determine the spectral characteristics of Charon's unknown absorber, we calculated hypothetical spectral properties consistent with the Charon data, as done by Calvin and Clark (1991) and McCord *et al.* (1999), but for a larger variety of model assumptions (including patchwork, intimate, molecular, and stratified mixtures, various grain sizes and distributions of grain sizes, a wide range of relative abundances, etc.). The process is essentially that of extrapolating from the spectral behavior of water ice (under some set of model assumptions), through the observed spectrum of Charon, to find a family of optical properties which could serve as non-ice end-members. Two examples are shown at the bottom of Fig. 10, with error bars derived by fitting Charon's spectrum $\pm 1\sigma$. Also shown is the McCord *et al.* (1999) hypothetical absorber derived from Galileo NIMS spectra of the jovian satellite Europa. Of course, spectral details of these hypothetical absorbers depend quite strongly on model assumptions, but they do have certain features in common, which offer clues to the nature of the material(s) responsible for the non-neutral continuum absorption on Charon and other icy satellites. For many different models, the derived absorber exhibits abruptly increasing absorption from about 1.9 μm longward, without the steep rise around 2.1 μm seen in water ice spectra. The models also call for increasing absorption beyond 2.3 μm and a red slope in the region of Johnson *B* and *V* filters. Some models show a broad absorption band extending from about 1.9 to 2.3 μm or perhaps a pair of broad bands, one near 2 μm and one near 2.2 μm.

Figure 10 shows spectra of several materials sharing the spectral characteristic of increasing absorption longward of 1.9 μm. These include sulfuric acid ice, two hydrated sulfates, and two phyllosilicates. Sulfuric acid could derive from precursor sulfur compounds in ice on exposure to the space radiation environment (Carlson *et al.* 1999). Two sulfate hydrates, bloedite ($Na_2Mg(SO_4)_2 \cdot 4H_2O$) and epsomite ($MgSO_4 \cdot 7H_2O$), are shown. Hydrated salts like these occur as evaporites on Earth and have been proposed as candidates for the non-ice component of Europa (McCord *et al.* 1999). Similar materials could have erupted onto Charon's surface as briny lavas, although different hydration states likely prevail under such conditions. The phyllosilicate family is represented here by montmorillonite ($Al_{1.67}Mg_{0.33}Si_4O_{10}(OH)_2 \cdot Na_{0.33} \cdot 4H_2O$) and kaolinite ($Al_4Si_4O_{10}(OH)_8$). Phyllosilicates generally result from aqueous alteration of other silicate minerals and could have been formed during a melting episode early in Charon's history.

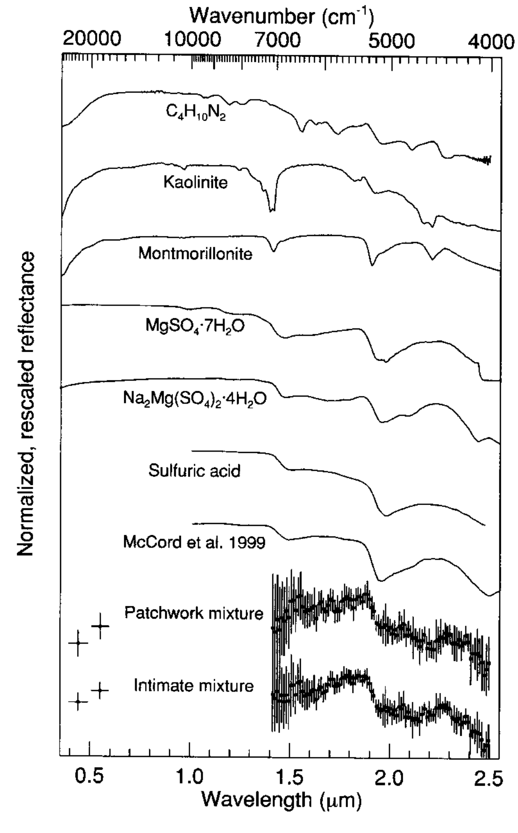

FIG. 10. Comparison between rescaled and offset reflectance spectra of synthetic Charon continuum absorbers and various materials. The synthetic spectra reflect two distinct scenarios. First, we assumed an intimate mixture of 100-μm grains, 50% H_2O ice and 50% the hypothetical continuum absorber. Second, we assumed a patchwork mixture with the surface area divided equally between H_2O ice and the continuum absorber. Also shown are normalized reflectance spectra of several materials with similar spectral characteristics, including an hypothetical non-ice component derived from Galileo NIMS spectra of Europa (McCord *et al.* 1999), sulfuric acid ice (Carlson *et al.* 1999), two hydrated salt minerals, two phyllosilicates (Clark *et al.* 1993), and piperazine (Cruikshank *et al.* 1998b), a nitrogen-bearing organic species ($C_4H_{10}N_2$).

The spectral similarity of the above materials is not a coincidence. They all contain OH and/or H_2O in some form. Vibrational states of isolated OH and H_2O are perturbed (energy shifts, degeneracy breakdowns, etc.) by the local environment when they are incorporated into a solid phase. As a consequence, the positions, widths, and shapes of the OH and H_2O vibrational absorption bands vary subtly from one solid to another. It is these absorptions that are responsible for the bands near 1.5 and 2 μm in all of these materials. We do not claim that one of the materials shown in Fig. 10 is the correct one, but it seems quite possible that a component of the continuum absorber on Charon may simply be H_2O, with absorption bands distorted by close association with other elements or ions. The precise details of the bandshapes depend on the composition, temperature, thermal history, and equilibria established between radiation, erosional, and photochemical processes. More laboratory work is called for to elucidate remote sensing possibilities using the distortions of H_2O bands in low-temperature materials.

We must also consider the possibility of other absorbers being responsible for the increased absorption in Charon's spectrum beyond 1.9 μm. Organic species are abundant on Pluto, and Charon presumably formed from similar material. The *tholin* produced under conditions analogous to Titan's atmosphere (Khare *et al.* 1984) is not a good match because of its strong visible absorption, but piperazine ($C_4H_{10}N_2$) does have increasing absorption from 1.9 μm, just as is required by the Charon data. Precisely what, if any, organic species are likely to be stable (or be produced) under conditions which prevail on Charon's surface remain unknown, and we currently lack the laboratory data which might lead us toward specific compositions.

Table II and Fig. 11 show model configurations which produce reasonable matches to the Charon data, using a phyllosilicate, a hydrated sulfate, and sulfuric acid for the additional absorber.

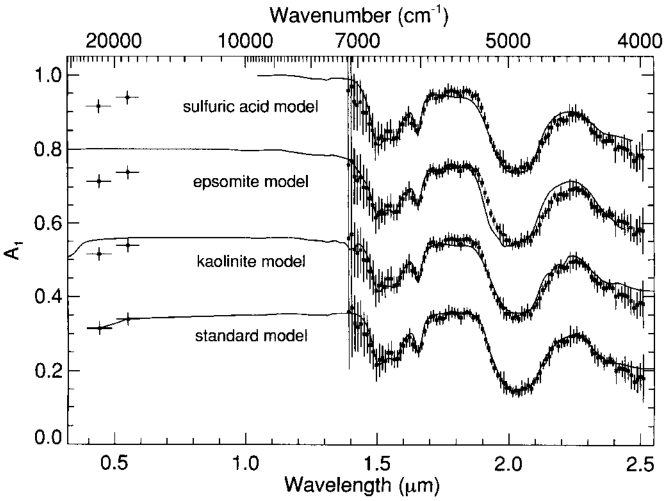

FIG. 11. The models from Table II, using various continuum absorbers, at 1° phase angle, superimposed on the Charon data. The "standard model" is described in the text. The ordinate is labeled for the Charon spectrum and standard model; other models are shifted upward by successive steps of 0.2.

TABLE II

Models of Charon

Model absorber composition	Absorber abundance (% vol)	Absorber grain size (μm)	H_2O abundance (% vol)	H_2O grain size (μm)	"Carbon" abundance (% vol)
40% sulfuric acid	97.2	140	60	—	2.8
50% epsomite	97.3	140	50	—	2.7
Kaolinite	20	130	76	130	4.0
Standard model	40	70	60	100	—

Note. In order to convert reflectance spectra of sulfuric acid (Carlson *et al.* 1999) and epsomite and kaolinite (Clark *et al.* 1993) into optical constants, we assumed grain sizes of 100 μm, 1 mm, and 20 μm and refractive indices of 1.44, 1.45, and 1.56, respectively. We also assumed bidirectional geometry with normal illumination and a detector offset 30° from the zenith (R. Clark 1998, personal communication). These assumptions introduce uncertainties in the scaling, but not the general form of the derived absorption coefficients. Hapke theory's restriction to the regime $k \ll 1$ led us to simulate "carbon" with 10-μm grains having an absorption coefficient of 1000 cm^{-1} (as in Grundy 1995). This approach does not introduce spurious colors, because both real and simulated carbon grains have extremely low single scattering albedos over our entire wavelength range. The sulfuric acid and epsomite models assume grains composed of the absorber plus H_2O, intimately mixed with "carbon" grains. The kaolinite model assumes intimate mixing of three different grain types, and the standard model is an intimate mixture of H_2O and the continuum absorber described in Table III.

Approximate fits can be found using each of these materials, although none is without its problems. For example, the sulfuric acid model has a slope from 1.75 to 1.85 μm which we do not see in Charon's spectrum. The epsomite ($MgSO_4 \cdot 7H_2O$) model has too much absorption at 1.9 μm and not enough absorption at 2.25 μm. The kaolinite ($Al_4Si_4O_{10}(OH)_8$) model matches the 2.0- to 2.6-μm range somewhat better, but predicts too much absorption at 1.8 μm and an additional OH absorption at 1.4 μm which is not apparent in recent data presented by Brown and Calvin (1999). All three models fail to match the observed albedos of Charon in Johnson *B* and *V* filters, but adding a minute quantity of some reddish material such as a tholin or an iron-bearing silicate could rectify that discrepancy easily enough. Clearly, we cannot distinguish between these and other contending absorbers based on the currently available data.

5.3. Standard Model

It would be useful to have a "standard model" of Charon to provide an absolutely calibrated simulation of Charon's spectrum for all observable phase angles (effectively the range $0.5° < g < 2.0°$). This model would permit more accurate removal of Charon's contribution from combined Pluto + Charon observations at arbitrary spectral resolutions and could be useful for comparison with spectra of other icy objects. We propose the following standard model, using the Hapke parameters listed earlier ($\bar{\theta} = 20°, h = 0.01, B_0 = 0.6$, etc.). The model is an intimate mixture of 70-μm-diameter grains, with 60% of the grains being low-temperature crystalline H_2O ice and 40% having the

TABLE III
Standard Continuum Absorber for Charon

Wavelength (μm)	Absorption (cm^{-1})	Imaginary index k
0.44	220	7.7×10^{-4}
0.55	157	6.9×10^{-4}
1.90	105	1.6×10^{-3}
1.93	150	2.3×10^{-3}
2.08	150	2.5×10^{-3}
2.15	280	4.8×10^{-3}
2.27	220	4.0×10^{-3}
2.35	300	5.6×10^{-3}
2.6	800	1.7×10^{-2}

absorption coefficients listed in Table III, with linear interpolation between points and refractive index 1.3. Rather than work with one of the real materials listed earlier, which would have implications for the spectrum of Charon outside of our wavelength limits, we prefer this simple numerical description of the continuum absorber. Our standard model is shown at the bottom of Fig. 11.

While this proposed model provides an excellent match to our average Charon spectrum, it makes no effort to duplicate Charon's modest lightcurve (Fig. 5) or leading–trailing asymmetry (Fig. 6). The amplitude of Charon's lightcurve is ±4% (Buie *et al.* 1997). The largest leading–trailing asymmetry admissible by our data is about ±10%, in terms of the continuum-band center albedo difference, or ±17% in the middle of the 2-μm band, if all of the variation occurs at that wavelength. Consider the errors introduced by neglecting these effects when removing the Charon flux from blended observations of Pluto and Charon. Charon represents about one-fifth of the projected area of the system and, at most wavelengths, its albedo is lower than that of Pluto, so its flux contribution is generally less than a fifth of the total. A 4% lightcurve error in the model Charon flux thus produces less than a 1% error in the derived Pluto flux. At 2 μm, a 17% H_2O ice band albedo error in Charon's flux translates to a about a 1% error in Pluto's albedo, because Pluto is so much brighter than Charon at that wavelength. Even at the worst wavelength, 2.3 μm, where Charon's contribution is about 72% that of Pluto, a 4% Charon error translates to about a 3% error in the derived Pluto flux. If greater precision were required, one could easily modulate our standard model by the lightcurve from Buie *et al.* (1997), but with absolute calibration better than 3% being extremely difficult to achieve, our proposed standard model should be sufficiently accurate for most Pluto-monitoring applications.

5.4. A 2.215-μm Band?

One area where Charon's spectrum appears to deviate slightly from the spectrum of Tethys in Fig. 9 (and also from our standard model) is at 2.215 μm (4515 cm^{-1}), where the Charon spectrum shows two adjacent data points about 1σ low. This dip can also be seen in the hypothetical continuum spectrum from the patchwork mixing model of Fig. 10. Brown and Calvin (1999) have suggested that Charon exhibits an absorption band of ammonia hydrate at this wavelength. While it has not yet been spectroscopically confirmed on any icy solar system object, NH_3 and its hydrates have long been thought of as cosmochemically plausible outer solar system materials (e.g., Croft *et al.* 1988, Kargel 1992).

Schmitt *et al.* (1998) published near-infrared absorption coefficients for an ammonia hydrate at 65 K (of unspecified stoichiometry, perhaps $NH_3 \cdot 2H_2O$), showing an absorption band at about 2.215 μm, as well as a stronger band at 1.99 μm (5015 cm^{-1}). No trace of the latter absorption band is seen in our Charon data, but if the ammonia phase were intimately mixed with grains of pure H_2O ice, the strong absorption of the water ice at that wavelength could suppress this other band to some extent. It is also not known how the spectra of other ammonia hydrate phases differ.

Using the Schmitt *et al.* (1998) lab data, we added ammonia hydrate to our standard Charon model to see what kind of limits could be placed. Our maximum ammonia models are shown in Fig. 12. One is a checkerboard mixture, in which 15% of the area has ammonia hydrate replacing H_2O in the standard model. The other is an intimate mixture in which 8% ammonia hydrate replaces some of the H_2O ice. The two models are indistinguishable at the scale of this figure. While suggestive, our data do not justify an identification of ammonia hydrate, as they are also consistent at the 1σ level with models which include no ammonia at all.

An identification of ammonia in the form of ammonia hydrate would be of considerable interest. Although NH_3 was probably incorporated into icy satellites, it has never been detected in spectra of other icy outer solar system surfaces, despite its very strong spectroscopic features. This scarcity may result

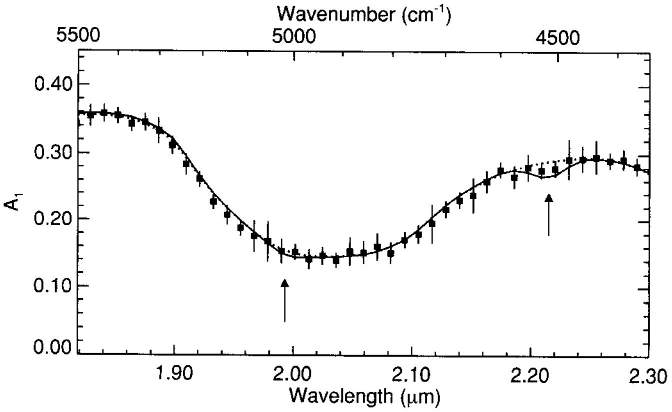

FIG. 12. Charon data compared with the standard model (dotted curve) and modified versions (solid curves) which include ammonia hydrate grains, either intimately mixed with or segregated in patches separate from the ordinary H_2O ice (the two configurations are indistinguishable at this scale). All models are resampled at the same resolution as the Charon data. Arrows indicate wavelengths of ammonia hydrate absorption bands.

from the extreme chemical activity of NH_3, which readily reacts with other species, including CO_2, sulfates, and formaldehyde (Kargel 1992). Ammonia is also more volatile than water and may escape from the uppermost surface layers of molten H_2O lavas erupted into a vacuum. Its stability against the UV and solar wind radiation fluxes impinging on Charon's surface is not known. If higher quality data eventually do establish the presence of NH_3 on Charon's surface, the surprising implication would be that Charon is fundamentally different from all other icy satellites.

It should also be noted that aluminum-bearing phyllosilicates, such as kaolinite and montmorillonite, exhibit absorption bands at 2.215 μm. While abundant occurrence of these sorts of minerals would be unexpected on Charon's surface, they are able to match the spectral data about as well as ammonia hydrate does. More plausible phyllosilicates such as serpentines and smectites do not exhibit absorption bands near 2.215 μm.

5.5. H_2O Ice Temperature and Phase

Grundy *et al.* (1999) introduced a method for determining H_2O ice temperatures from the positions and shapes of the near-infrared H_2O absorption bands, including a strongly temperature-dependent band at 1.65 μm. Using high signal-to-noise ground-based observations of icy satellites, they were able to determine absolute ice temperatures with precisions as high as ±5 K and relative temperatures to within ±2 K. For an unresolved spherical body, the temperature obtained is a disk-average ice temperature, with different parts of the visible disk contributing in proportion to their contribution to the near-infrared spectrum. High albedo regions contribute proportionately more reflected light, so their cooler temperatures tend to dominate the temperatures obtained via the Grundy *et al.* (1999) method. Observations of thermal emission are quite different, in that they are most sensitive to darker, warmer regions.

Our Charon data have lower spectral resolution and considerably less signal precision than the Grundy *et al.* (1999) data, but we tried their method with our Charon spectrum, deriving a disk-average ice temperature for Charon of 60 ± 20 K. This uncertainty is too large to provide much useful constraint on the thermal behavior of Charon's water ice, as it encompasses the entire range of plausible disk-average ice temperatures, as shown by Fig. 13. Assuming that Charon's phase integral is 0.6 (consistent with uranian satellites), the slow-rotator standard thermal model predicts a disk-average temperature of 60 K, and the more plausible fast-rotator model predicts 50 K. Any geographic segregation between Charon's dark continuum absorbers and bright H_2O ice would result in lower ice temperatures.

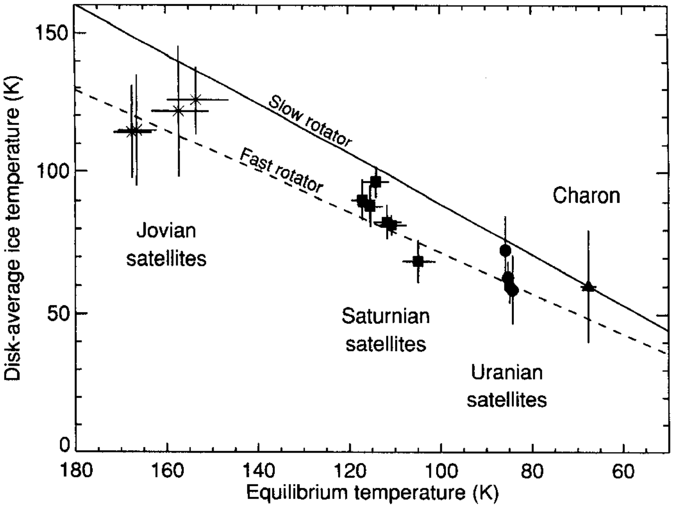

FIG. 13. Charon H_2O ice temperature compared with other satellite ice temperatures reported by Grundy *et al.* (1999, adapted from Fig. 7). Each satellite's ice temperature is plotted against its equilibrium temperature, $[(1 - A_b)S_{\odot}\sigma^{-1}r^{-2}]^{0.25}$, where A_b is the estimated bolometric bond albedo (hence the horizontal error bars), σ is the Stefan–Boltzmann constant, r is the heliocentric distance, and $S_{\odot}$ is the solar constant. The solid line is the locus of points consistent with disk-average temperatures calculated for the temperature distribution of the slow-rotator standard thermal model. The dotted line is for the fast-rotator thermal model, viewed equator-on (as viewed from other latitudes, this line shifts toward the slow-rotator curve, until they coincide for pole-on orientation).

Unlike low-temperature crystalline H_2O ice, amorphous ice lacks an absorption band at 1.65 μm. That Charon's spectrum shows a strong 1.65-μm band proves that Charon's surface ice is predominantly crystalline. However, the action of the space radiation environment (ultraviolet photons, solar wind particles, cosmic rays, etc.) on amorphous or crystalline ice at Charon's surface temperatures is not well known. Radiation over some energy range could trigger crystallization while more energetic radiation is expected to disrupt the H_2O crystal structure, at least on the uppermost surface. Over time, amorphous ice could accumulate if it is produced more quickly than it can recrystallize at Charon's low surface temperatures (e.g., Heide 1984, Kouchi and Koruda 1990, Strazzulla *et al.* 1991, Moore and Hudson 1992). An admixture of amorphous and crystalline ice would spectroscopically resemble crystalline ice at a warmer temperature. If the actual ice temperature on Charon were as cold as 40 K, Charon's spectrum would be consistent with a fine layer of amorphous ice, as much as 10 μm thick, coating the uppermost H_2O ice grains. Future observations coupling thermal and near-infrared spectroscopy may be able to resolve the ambiguity between ice temperature and phase composition.

6. CONCLUSIONS

We have obtained new near-infrared spectra of Pluto's satellite Charon at four well-spaced longitudes. The new data represent a substantial improvement in spectral resolution and longitudinal coverage, relative to previously published spectra, and they advance our knowledge of this small world considerably. The data show that Charon's H_2O ice is globally distributed, although there does appear to be a very slight asymmetry between leading and trailing hemispheres. Charon's H_2O ice is predominantly crystalline in phase and is spectroscopically consistent with ice at temperatures expected for Charon's surface, with no amorphous ice at all. An as-yet-unidentified continuum absorber reduces Charon's albedo at wavelengths longward of about $\sim 2\ \mu\text{m}$, with similar spectral characteristics to what is seen on other icy satellites. No phase-angle-dependent spectral effects were detected. Our data are consistent with a small quantity of ammonia hydrate producing absorption at $2.215\ \mu\text{m}$, but our signal precision does not permit an identification. No spectral evidence whatsoever is seen for the volatile ices abundant on Pluto, CO, N_2, and CH_4.

The picture that emerges is of a satellite similar to comparably sized satellites of Saturn and Uranus. By analogy with these other satellites, a future spacecraft fly-by can expect to return images showing a rich mix of tectonic, cryovolcanic, and impact geology. Charon's end-member status as the icy satellite at the greatest mean distance from the sun, with perhaps the coldest present-day surface temperatures, the most violent formation conditions, and orbiting the smallest primary planet in the solar system, makes it especially interesting in the context of other icy satellites. However, Charon's extreme circumstances have apparently not resulted in a particularly unusual surface composition. The question of whether Charon's geology is unique among icy satellites will have to await a spacecraft mission.

ACKNOWLEDGMENTS

We thank Alex Storrs, who provided valuable input into the planning and execution of our observations on HST. This paper is based on observations with the NASA/ESA Hubble Space Telescope obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA Contract NAS5-26555. Support for Buie was provided by NASA through Grant GO-07818.01-96A from STScI. Grundy was supported by Hubble Fellowship HF-01091.01-97A awarded by STScI.

REFERENCES

Astronomical Almanac for the Year 1998. U.S. Government Printing Office, Washington, DC.

Binzel, R. P. 1988. Hemispherical color differences on Pluto and Charon. *Science* **241**, 1070–1072.

Bosh, A. S., L. A. Young, J. L. Elliot, H. B. Hammel, and R. L. Baron 1992. Photometric variability of Charon at 2.2 microns. *Icarus* **95**, 319–324.

Brown, M. E., and W. M. Calvin 1999. Evidence for crystalline water and ammonia ices on Pluto's satellite Charon. *Science* **287**, 107–109.

Buie, M. W., and S. K. Shriver 1994. The distribution of water frost on Charon. *Icarus* **108**, 225–233.

Buie, M. W., D. P. Cruikshank, L. A. Lebofsky, and E. F. Tedesco 1987. Water frost on Charon. *Nature* **329**, 522–523.

Buie, M. W., D. J. Tholen, and K. Horne 1992. Albedo maps of Pluto and Charon: Initial mutual event results. *Icarus* **97**, 211–227.

Buie, M. W., D. J. Tholen, and L. H. Wasserman 1997. Separate lightcurves of Pluto and Charon. *Icarus* **125**, 233–244.

Calvin, W. M., and R. N. Clark 1991. Modeling the reflectance spectrum of Callisto 0.25–4.1 μm. *Icarus* **89**, 305–317.

Carlson, R. W., R. E. Johnson, and M. S. Anderson 1999. Sulfuric acid on Europa and the radiolytic sulfur cycle. *Science* **286**, 97–99.

Clark, R. N., G. A. Swayze, A. J. Gallagher, T. V. V. King, and W. M. Calvin 1993. The U.S.G.S. digital spectral library: Version 1: 0.2 to 3.0 μm. U.S. Geological Survey Open File Report 93-592.

Cleveland, W. S. 1979. Robust locally weighted regression and smoothing scatterplots. *J. Am. Stat. Assoc.* **74**, 829–836.

Colina, L., and R. Bohlin 1997. Absolute flux distributions of solar analogs from the UV to the near-IR. *Astron. J.* **113**, 1138–1144.

Colina, L., R. C. Bohlin, and F. Castelli 1996. The 0.12–2.5 μm absolute flux distribution of the Sun for comparison with solar analog stars. *Astron. J.* **112**, 307–315.

Croft, S. K., J. I. Lunine, and J. Kargel 1988. Equation of state of ammonia–water liquid: Derivation and planetological applications. *Icarus* **73**, 279–293.

Cruikshank, D. P., R. H. Brown, W. M. Calvin, T. L. Roush, and M. J. Bartholomew 1998a. Ices on the satellites of Jupiter, Saturn, and Uranus. In *Solar System Ices* (B. Schmitt, C. de Bergh, and M. Festou, Eds.), pp. 579–606. Astrophys. Space Science Library, Kluwer, Boston.

Cruikshank, D. P., T. L. Roush, M. J. Bartholemew, T. R. Geballe, Y. J. Pendleton, S. M. White, J. F. Bell III, J. K. Davies, T. C. Owen, C. de Bergh, D. J. Tholen, M. P. Bernstein, R. H. Brown, K. A. Tryka, and C. M. Dalle Ore 1998b. The composition of centaur 5145 Pholus. *Icarus* **135**, 389–407.

Douté, S., B. Schmitt, E. Quirico, T. C. Owen, D. P. Cruikshank, C. de Bergh, T. R. Geballe, and T. L. Roush 1999. Evidence for methane segregation at the surface of Pluto. *Icarus* **142**, 421–444.

Fink, U., and M. A. DiSanti 1988. The separate spectra of Pluto and its satellite Charon. *Astron. J.* **95**, 229–236.

Grundy, W. M. 1995. *Methane and Nitrogen Ices on Pluto and Triton: A Combined Laboratory and Telescope Investigation*. Ph.D. thesis. Univ. of Arizona, Tucson.

Grundy, W. M., and U. Fink 1996. Synoptic CCD spectrophotometry of Pluto over the past 15 years. *Icarus* **124**, 329–343.

Grundy, W. M., and B. Schmitt 1998. The temperature-dependent near-infrared absorption spectrum of hexagonal H_2O ice. *J. Geophys. Res.* **103**, 25809–25822.

Grundy, W. M., M. W. Buie, J. A. Stansberry, J. R. Spencer, and B. Schmitt 1999. Near-infrared spectra of icy outer Solar System surfaces: Remote determination of H_2O ice temperatures. *Icarus* **142**, 536–549.

Hapke, B. 1993. *Combined Theory of Reflectance and Emittance Spectroscopy*. Cambridge Univ. Press, New York.

Heide, H. G. 1984. Observation on ice layers. *Ultramicroscopy* **14**, 271–278.

Horne, K. 1986. An optimal extraction algorithm for CCD spectroscopy. *Publ. Astron. Soc. Pacific* **98**, 609–617.

Kargel, J. S. 1992. Ammonia-water volcanism on icy satellites: Phase relations at 1 atmosphere. *Icarus* **100**, 556–574.

Kavaldjiev, D., and Z. Ninkov 1998. Measurements of the sensitivity within CCD Pixels. *Opt. Eng.* **37**, 948–954.

Khare, B. N., C. Sagan, E. T. Arakawa, F. Suits, T. A. Callcott, and M. W. Williams 1984. Optical constants of organic tholins produced in a simulated titanian atmosphere: From soft X-ray to microwave frequencies. *Icarus* **60**, 127–137.

Kouchi, A., and T. Koruda 1990. Amorphization of cubic ice by ultraviolet irradiation. *Nature* **344**, 134–135.

Lunine, J. I., and W. C. Tittemore 1993. Origins of outer-planet satellites. In *Protostars and Planets III* (E. H. Levy and J. I. Lunine, Eds.), pp. 1149–1176. Univ. of Arizona Press, Tucson.

Marcialis, R. L., G. H. Rieke, and L. A. Lebofsky 1987. The surface composition of Charon: Tentative identification of water ice. *Science* **237**, 1349–1351.

McCord, T. B., G. B. Hansen, D. L. Matson, T. V. Johnson, J. K. Crowley, F. P. Fanale, R. W. Carlson, W. D. Smythe, P. D. Martin, C. A. Hibbits, J. C. Granahan, and A. Ocampo 1999. Hydrated salt minerals on Europa's surface from the Galileo near-infrared mapping spectrometer (NIMS) investigation. *J. Geophys. Res.* **104**, 11827–11851.

McKinnon, W. B. 1984. On the origin of Triton and Pluto. *Nature* **311**, 355–358.

Moore, M. H., and R. L. Hudson 1992. Far-infrared spectral studies of phase changes in water ice induced by proton irradiation. *Astrophys. J.* **401**, 353–360.

Null, G. W., and W. M. Owen Jr. 1996. Charon/Pluto mass ratio obtained with *HST* CCD observations. *Astron. J.* **111**, 1368–1381.

Null, G. W., W. M. Owen Jr., and S. P. Synnott 1993. Masses and densities of Pluto and Charon. *Astron. J.* **105**, 2319–2335.

Prinn, R. G., and B. Fegley Jr. 1981. Kinetic inhibition of CO and N_2 reduction in circumplanetary nebulae: Implications for satellite composition. *Astrophys. J.* **249**, 308–317.

Reinsch, K., V. Burwitz, and M. C. Festou 1994. Albedo maps of Pluto and improved physical parameters of the Pluto–Charon system. *Icarus* **108**, 209–218.

Roush, T. L. 1994. Charon: More than water ice? *Icarus* **108**, 243–254.

Roush, T. L., D. P. Cruikshank, J. B. Pollack, E. F. Young, and M. J. Bartholomew 1996. Near-infrared spectral geometric albedos of Charon and Pluto: Constraints on Charon's surface composition. *Icarus* **119**, 214–218.

Schmitt, B., E. Quirico, F. Trotta, and W. M. Grundy 1998. Optical properties of ices from UV to infrared. In *Solar System Ices* (B. Schmitt, C. de Bergh, and M. Festou, Eds.), pp. 199–240. Astrophys. Space Science Library, Kluwer, Boston.

Stern, A. S., W. B. McKinnon, and J. I. Lunine 1997. On the origin of Pluto, Charon, and the Pluto–Charon binary. In *Pluto and Charon* (D. J. Tholen and S. A. Stern, Eds.), pp. 605–663. Univ. of Arizona Press, Tucson.

Stevenson, D. J., A. W. Harris, and J. I. Lunine 1986. Origins of satellites. In *Satellites* (J. A. Burns and M. S. Matthews, Eds.), pp. 39–88. Univ. of Arizona Press, Tucson.

Strazzulla, G., G. Leto, G. A. Baratta, and F. Spinella 1991. Ion irradiation experiments relevant to cometary physics. *J. Geophys. Res.* **96**, 17547–17552.

Tholen, D. J., and M. W. Buie 1997. Bulk properties of Pluto and Charon. In *Pluto and Charon* (D. J. Tholen and S. A. Stern, Eds.), pp. 269–293. Univ. of Arizona Press, Tucson.

Tholen, D. J., M. W. Buie, R. P. Binzel, and M. L. Freuh 1987. Improved orbital and physical parameters for the Pluto–Charon system. *Science* **237**, 512–514.

Warren, S. G. 1984. Optical constants of ice from the ultraviolet to the microwave. *Appl. Opt.* **23**, 1206–1225.

Young, L. A., C. B. Olkin, J. L. Elliot, D. J. Tholen, and M. W. Buie 1994. The Charon–Pluto mass ratio from MKO astrometry. *Icarus* **108**, 186–199.