

THE 1998 NOVEMBER 14 OCCULTATION OF GSC 0622-00345 BY SATURN'S ATMOSPHERE

JOSEPH HARRINGTON¹

326 Space Sciences Building, Center for Radiophysics and Space Research and
Cornell University, Ithaca, NY 14853-6801

AND

RICHARD G. FRENCH¹

Astronomy Department, Wellesley College and
Wellesley, MA 02481

(Received 2004 October 4)

Submitted to ApJ. DRAFT of October 5, 2004 11:46.

ABSTRACT

On 1998 November 14, Saturn and its rings occulted the star GSC 0622-00345. We observed atmospheric immersion with NSFcam at the National Aeronautics and Space Administration's Infrared Telescope Facility on Mauna Kea, Hawaii. Immersion occurred at 55.5° S planetocentric latitude, the first occultation reported for Saturn south of the equatorial region. A $2.3\text{-}\mu\text{m}$ filter suppressed reflected sunlight. We present a lightcurve whose signal-to-noise ratio per scale height is 267, the thermal profile obtained by numerical inversion, a wavelet analysis, and an empirical significance test for waves in occultation inversions. The thermal profile is valid between 1 and 60 μbar . The isothermal fit to the lightcurve has a temperature of 143 ± 3 K, consistent with the inversion. We show that even our low level of noise can produce temperature swings of over 7 K in inversions, and that spurious periodic features in the "reliable" regions of wavelet transforms can exceed 0.3 K. The peak wavelet amplitude in our reliable region is 0.62 K, with no significant evidence for vertical propagation to or through that peak, so we do not claim a wave detection. The vertical temperature gradient is >0.2 K km^{-1} on the stable side of adiabatic. It does not show the alternating-rounded-spiked appearance seen in some temperature gradient profiles and attributed to breaking buoyancy waves, consistent with our wave non-detection. Atmospheric emission and ring data were not successfully obtained. We present improved techniques for aperture positioning, removal of scintillation effects, and timing.

Subject headings: atmospheric effects, atmospheric structure, methods: data analysis, occultations, planets and satellites: Saturn, waves

1. INTRODUCTION

On 1998 November 14, Saturn and its rings occulted GSC 0622-00345, as predicted by Bosh & McDonald (1992). We obtained a lightcurve for atmospheric immersion, based on infrared imaging observations at the NASA Infrared Telescope Facility (IRTF) on Mauna Kea, HI. The high signal-to-noise ratio (S/N) allowed us to determine the vertical temperature profile of Saturn's stratosphere at 55° S latitude, a region not sampled by previous stellar occultation observations (see Fig. 1 and Table 1).

Earth-based occultations remain an attractive method for measuring the thermal profile in the 1–100- μbar region of a planetary atmosphere, and there are now many such profiles for Saturn. Most were recorded during the 28 Sgr occultation of 3 July 1989, which sampled the equatorial region from 6.6° N– 15.2° S latitude (Hubbard et al. 1997). There is a single profile for the north polar region (Cooray et al. 1998, 82.5° – 85° N), and a northern low-latitude profile from the same event (French et al. 1999, 19.16° N). Saturn's central flash probes much deeper, around 2.5 mbar; Nicholson et al. (1995) obtained IR images of the flash during the 28 Sgr event, from which they inferred the zonal wind profile of the sampled latitudes along Saturn's limb.

Occultation observations from spacecraft near a giant planet probe the troposphere from the cloud deck (~ 1 bar) to the mbar level at radio and infrared wavelengths, and probe the upper stratosphere and thermosphere (<1 μbar) in the ultraviolet. Earth-based visual and infrared occultations measure the thermal structure of the intervening mesosphere and stratosphere regions that are not well sampled by spacecraft experiments. The Pioneer and Voyager observations sensed only the equatorial region (Kliore et al. 1980; Lindal et al. 1985; Smith et al. 1983), but Cassini mission plans call for occultation observations at a variety of latitudes.

We obtained our low-noise lightcurve by applying several new analysis techniques, which we describe in the next section. Subsequent sections cover the timing solution, thermal inversion, wavelet analysis, noise tests for the wavelet analysis, and our conclusions.

2. OBSERVATIONS AND LIGHTCURVE

The NSFcam InSb array camera (Shure et al. 1994) recorded the event in "movieburst" mode. This mode writes a time sequence of images to computer memory, saving the set to disk after the observations finish. We recorded three sub-images derived from the full 256×256 -pixel readout. These boxes contained the star (28×32 pixels), Rhea (28×32 pixels), and a piece of blank sky (16×16 pixels), placed as shown in Fig. 1. The larger boxes were originally 32×32 pixels, but a camera software error overwrote four edge columns. No critical data were lost and the error is now fixed. Table 2 gives additional observation and lightcurve parameters.

The filter selects a methane band. Saturn's atmosphere

Electronic address: jh@oobleck.astro.cornell.edu
Electronic address: rfrench@wellesley.edu

¹ Visiting Astronomer at the Infrared Telescope Facility, which is operated by the University of Hawaii under Cooperative Agreement no. NCC 5-538 with the National Aeronautics and Space Administration, Office of Space Science, Planetary Astronomy Program.

TABLE 1
 OCCULTATION PARAMETERS

Description	Value	Units	Comment
Star	GSC 0622-00345		
RA (J2000)	1h49m54. ^s 36		
DEC (J2000)	8°23'12."7		
K magnitude	8.4		2MASS
Type	G2		est. from 2MASS magnitudes
Projected diameter at Saturn	~0.3 km		
Half-light latitude	55.5 ° S		planetocentric
Fresnel scale	1.2 km	$\sqrt{\lambda d/2}$	
Sky-plane velocity	17.940 km sec ⁻¹		
Perpendicular velocity, v_{\perp}	12.802 km sec ⁻¹		vs. oblate limb
Ring opening angle	-14.6772 °		
Pole position angle (J2000)	1.5134 °		
Saturn's geocentric distance, d	1.25168 × 10 ⁹ km		

 TABLE 2
 OBSERVATION AND LIGHTCURVE PARAMETERS

Description	Value	Units	Comment
Telescope	IRTF		
Instrument	NSFcam		
Image scale	0.30 "	pix ⁻¹	
Wavelength, λ	2.28 μ m		"Spencer 2.3" filter
Bandpass, FWHM	0.17 μ m		
Time between frame starts	0.213 s		
Exposure time	0.128 s		
Start of first frame	11:05:02.615 UTC		
Start of upper baseline	11:16:44.855 UTC		
Start of lower baseline	11:26:19.415 UTC		no evidence of star in frames
Start of last frame	11:56:06.722 UTC		
Number of frames	14400		
Number of frames per baseline	300		
Photometry (upper baseline interval):			
Star full flux	5439 DN		above lower baseline
Rhea full flux	5216 DN		above lower baseline
Sky box	12.55 DN	pix ⁻¹	
Frame-to-frame variability (upper baseline interval):			
Star	2.2 %		RMS, scintillation uncorrected
Star	1.5 %		RMS, scintillation corrected
Sky box	3.8 %		RMS
S/N per frame	66		
S/N per scale height	267		

strongly absorbs sunlight in the band and emits minimally there, so contrast between the star and Saturn is very high. The rings are bright at this wavelength, but their separation from the immersion latitude made it practical to subtract their scattered light (see below).

At the half-light time (see below), seeing was 0.9" FWHM and the star was at 1.3 airmasses and setting. The night was clear, but Rhea's constant flux varied by up to 8% in successive frames (2.3% RMS) due to scintillation. Rhea was nearly as bright as the unocculted star (see Table 2) and had a point-spread function (PSF) that very closely matched that of the star in each frame. We thus used Rhea as a standard for flux and position, and attempted to use it as a PSF standard. In contrast to the scintillation, sky emission was very steady and low, based on the separately-recorded sky box.

To obtain sub-pixel accuracy in image shifts, photometry aperture placement, and scattered-light subtraction, we increased the spatial resolution by a factor of 10 using nearest-neighbor sampling (scaling by a factor of 5 gave essentially identical results). We performed all operations on the ex-

panded images, except as noted below. The expanded pixels allowed a reasonable approximation to a circular photometry aperture that included partial pixels.

After debiasing, flattening, and interpolating three hot pixels that were not near the star or Rhea, we centroided Rhea in all images. The relative position of the star in two well-separated sets of 100 pre-occultation frames determined the Rhea-star offset and the offset's drift with time. This allowed accurate placement of the Saturn and ring background template (described below) and photometry apertures, even when the star was fully occulted. A third set of stellar centroids confirmed that the predicted stellar positions were accurate to within 0.045" (RMS), or 0.15 unscaled pixels. To predict the position of Saturn (for template placement, below), we added the Rhea-Saturn relative ephemeris motion to Rhea's derived position. All ephemeris calculations come from the Jet Propulsion Laboratory's Horizons On-Line Ephemeris System website², which is based on the DE-0406LE-0406 and

² <http://ssd.jpl.nasa.gov/horizons.html>




FIG. 1.— Occultation geometry, subarray placement, and photometry apertures at the half-light time. The light line marks the path of GSC 0622-00345 relative to Saturn, with crosses every 20 minutes. The heavier portion gives the span of the dataset. The star, Rhea, and sky subarrays were 28×32 , 28×32 , and 16×16 pixels, respectively, and tracked the star with the relative placement shown. Tethys’s motion relative to Saturn is evident by the position mismatch between the half-light data and the scattered-light template. We computed the template from post-occultation frames and subtracted it from the data to produce the star frame shown here.

SAT136 ephemerides. Our derived Rhea-star drift rate closely matched the ephemeris prediction.

The small amount of scattered light from Saturn and the rings complicated stellar photometry by requiring the subtraction of a scattered-light template. We made a template by mask-averaging (see below) a set of post-immersion images, each divided by Rhea’s flux in that frame and shifted to align the predicted position of Saturn in all frames. All shifts were in units of integer scaled pixels, so no interpolation occurred.

To perform the shifted mask average calculation, we made two 3D “stacks,” one of the images and one of repeated copies of the bad pixel mask. The mask was an image sized like the data images, where each pixel had a value of one except those corresponding to bad data pixels. The latter had values of zero. We added sufficient zero padding so that no data would shift outside the boundaries of the stack in the following procedure. We multiplied the image stack by the mask stack to set bad pixels to zero, shifted the images to align them, and applied the same shifts to the masks. We collapsed each stack into a 2D image by summing along the image index dimension. At each pixel location, we divided the image sum by the mask sum. Each image shifted differently, so the bad pixels shifted out of alignment with one another. The result has no bad pixels and was calculated without interpolation. The mask sum gave the number of good pixels contributing to each pixel in the average.

To apply the template, we shifted it according to Saturn’s predicted position (see above) in each occultation frame, multiplied it by Rhea’s flux in that frame, down-sampled to the original spatial resolution, subtracted from the calibrated stellar frame, up-sampled the result to the scaled resolution using nearest-neighbor sampling, and performed aperture photometry using the predicted stellar position.

The template-subtracted stellar frames were quite flat except in the corner that included some ring light (see Fig. 1). Pixels in this region were in the sky annulus but not in the photometry aperture, and were efficiently rejected by taking the median of all good pixels in the sky annulus. Rhea photometry also used the median, for consistency.

TABLE 3
ISOTHERMAL FIT RESULTS

Description	Value	Units
Half-light time	$11:18:46.88 \pm 0.18$	s UTC
Full flux	1.0145 ± 0.0026	
Background	-0.0139 ± 0.0012	
Scale height, H	50.3 ± 1.1	km
Temperature, T	143 ± 3	K

The stellar and Rhea photometry apertures both had diameters of $3.6''$, four times the seeing FWHM. The abutting sky annuli had diameters of $7.2''$ (see Fig. 1). The photometry for the unocculted star and for Rhea closely followed each other, indicating that the main noise contributor was scintillation. We thus divided the stellar flux by Rhea’s flux in each frame (just as we adjusted the template’s flux). The lightcurve’s upper and lower baselines are averages of sections of the data before and after immersion (see Table 2). The lower baseline and the template come from the same frames.

The lightcurve’s per-frame S/N is 66 for the unocculted star. The French et al. (1978) atmospheric occultation noise parameter $\epsilon_\phi = \epsilon \sqrt{v_\perp / H} = 0.00375$, where ϵ is the standard deviation of the unocculted stellar flux in one-second bins divided by its mean and Tables 1 and 3 give the other parameters. The reciprocal of this value is the S/N per scale height. Our value of 267 exceeds even that of the extremely bright 28 Sgr occultation. For example, the IRTF 28 Sgr immersion S/N was 192, using the fractional flux standard deviation per frame of 1.7% and exposure time of 0.25 s reported by Harrington et al. (1993) and $H/v_\perp = 2.66$ reported by Hubbard et al. (1997). The raw S/N for these occultations is limited by scintillation rather than by photon statistics, eliminating 28 Sgr’s eight-magnitude brightness advantage over the event reported here. Our slightly slower event, improvements in infrared array technology, and the analysis methods described herein account for the rest.

A histogram of our upper baseline residuals closely follows a Gaussian distribution to $\pm 3\sigma$, but at $> +3\sigma$ it is slightly above the Gaussian. The standard deviation in the difference between adjacent-frame fluxes is $> 93\%$ of that expected for uncorrelated errors, and reaches $\sim 99\%$ over an interval of 1 s. The noise spectrum is thus still slightly red (see below). In the occultation itself, one would expect correlated structure on the scale of the projected stellar diameter or the Fresnel diffraction scale (see Table 1), but both of these are below the 2.7-km sky-plane resolution of the lightcurve.

Table 3 presents the results of an isothermal lightcurve model fit. We computed the nominal occultation track from the ephemeris, and then offset by an additional 500 km E and 500 km N, relative to Saturn’s center, so that the absolute radius of the half-light level matched the oblate half-light surface determined by Hubbard et al. (1997) from 28 Sgr occultation observations. This shift is within the astrometric accuracy of the star relative to Saturn.

Since the PSFs of Rhea and the star were nearly the same, we attempted to perform optimal stellar photometry, analogous to the optimal spectral extraction of Horne (1986) and others. We divided each background-subtracted stellar image by a normalized PSF template. Each pixel thus became an estimate of the total flux, with an error that increased (rapidly) with distance from the stellar centroid. We also estimated the




FIG. 2.— Lightcurve for occultation of GSC 0622-00345 by Saturn, observed at the NASA IRTF on 1998 November 14 UTC. The lightcurve has stable upper and lower baselines (dashed lines) and quite low photometric and terrestrial atmospheric scintillation noise. This makes it possible to determine the thermal profile of the atmosphere by numerical inversion. The lightcurve is characterized by many sharp, narrow spikes, due to small-scale refractive focusing by density fluctuations in Saturn’s atmosphere. In some cases, these spikes have an amplitude of several times the ambient intensity in the lightcurve. We fit an isothermal model (smooth solid line) to the observations, primarily to establish inversion parameters (see Table 3).

variance per pixel, and computed an error-weighted (i.e., optimal) average of all the good pixels in the frame.

In principle, this process can reduce the noise by up to 70%, but it requires an accurate PSF model and good centroids. The highly-varying and non-Gaussian PSFs of these short-exposure images precluded any smoothing or analytic PSF model fit, so the PSF templates were just the background-subtracted, normalized images of Rhea, shifted to the centroided position of the star in its images and down-sampled to the original resolution. For both Rhea and the star, we used the background level from aperture photometry.

Regrettably, while this procedure yielded a lower baseline that was three times less noisy than in the standardly-derived lightcurve, the upper baseline was about four times noisier. Close examination of the PSFs showed that noise and the fractional-pixel centroiding errors were sufficient to compromise this method. However, a much brighter PSF standard might have substantially improved our photometry. Obtaining such a standard may be difficult in occultation experiments with bright stars, but we note that large satellites could still play this role in occultations of fainter stars, potentially increasing the number of useful occultation events.

We did not acquire useful egress data. Atmospheric emission occurred in the rings (see Fig. 1), so, following Harrington et al. (1993), we employed a 3.34- μm filter (measured warm) in which both planet and rings are dark due to overlapping bands of water ice and methane. The NSFcam movieburst mode cannot display data during acquisition, so it was not possible to see until afterwards that the frames were empty. Possible reasons include a rise in the telluric water column (which would absorb all the light at this wavelength) or a filter wheel problem. The declination recorded in the header was close to the star’s. It is possible that a pointing or tracking error occurred, in right ascension. The telescope reached its horizon limit a few minutes before the end of the observation, so the recorded right ascension differed from the star’s by $\sim 6\text{m}$. This difference could hide an earlier error. Movieburst data are not shown in real time so that the camera computer can record frames as fast as possible. However, two processes using shared memory on a dual-CPU camera computer would accomplish this goal while simultaneously displaying the data.




FIG. 3.— A portion of the timing fit derived from the LED observations. The heavy line is the LED lightcurve. The upper light line is the model. The offset between them is a drift in the background that was not included in the model and that is positive in this part of the timing lightcurve. The lower light line presents the residuals.

Accurate and precise frame timing is key to constructing a meaningful lightcurve from a sequence of occultation frames. The NSFcam instrument computer synchronizes the start time of the first movieburst frame with a Global Positioning System (GPS) timing receiver. It then records the images in sequence. By observing a flashing light-emitting diode (LED) that is connected to the GPS and located in the instrument, one can estimate the mean frame rate and the delay between the computer’s notion of the start time and the actual first collection of photons. We made such timing measurements before and after the occultation observations, using exactly the same camera settings as those for the occultation. Results are included in Table 2.

We performed simple box photometry on the frames in the timing sequences and fit a model of the LED intensity *vs.* time (see Fig. 3). The model was constructed by making two lists: frame event times and LED event times. The frame times have parameters of delay after nominal start UTC, exposure time per frame, and dead time per frame. The LED turns on and off on the second and half second, UTC, respectively. One sorts the concatenated lists according to time and interprets the sequence in terms of LED on time per exposure. The model’s five free parameters are the frame event parameters and the upper and lower baseline levels. We fit the model to the LED data by minimizing the RMS residuals. This procedure very accurately determines the frame rate, which converts frame times to distance along the occultation’s track through Saturn’s atmosphere. The per-image timing uncertainty (“jitter”) is difficult to estimate, as fluctuations in the LED brightness, read noise, etc. all masquerade as jitter. This parameter is important in studies of detailed lightcurve structure, such as the shape and location of spikes, but these are not the focus of our present investigation. A better timing system is now available on the SpeX instrument at the IRTF.

4. TEMPERATURE PROFILE

We applied a numerical inversion to obtain the atmosphere’s temperature (T), pressure (P), and number density as a function of height above the half-light level (z). Table 4 and Fig. 4 give our parameters and results. We applied an Abel transform to the normalized lightcurve under the usual as-

TABLE 4
INVERSION PARAMETERS AND RESULTS

Description	Value	Units	Comment
Gravity at half-light	11.05	m s^{-2}	
Refractivity at STP	1.3×10^{-4}		
Mean molecular mass	2.135	AMU	
Local radius of curvature	64,307.9	km	along line of sight
T/H	2.840	K km^{-1}	
Adiabatic lapse rate	-0.901	K km^{-1}	

sumptions that the atmosphere is radially symmetric and that ray crossing is not substantial (French et al. 1978). Cooray & Elliot (2003) point out that in grazing occultations (unlike ours) these assumptions may be violated. The derived vertical refractivity profile is proportional to the density profile. The T and P profiles are then found by integrating the hydrostatic equation and applying the ideal gas law (French et al. 1978).

French et al. (1978) showed that, for high-quality, Earth-based stellar occultations, the valid region of the temperature profile ranges from about 0.5 to $-3.5 H$ above the half-light level. Initially, the inversion process is strongly affected by noise in the upper baseline of the lightcurve. Here, the actual refractive effects of the tenuous upper atmosphere are completely overwhelmed by photometric noise in the data. The noise introduces ~ 100 K swings in the thermal profile above the valid region, and those unphysical swings bias the curve at the top of the valid region. As deeper levels are probed, the uncertainties associated with this initial condition of the inversion are less significant. Below the region of validity, unphysical inversion fluctuations and trends arise because the decreasing stellar flux increases the fractional noise level (see, e.g., Raynaud et al. 2004).

To minimize the effects of the uncertain initial condition on the resulting temperature profiles, we replaced the noisy upper baseline and the initial part of the occultation itself by the isothermal model presented in Table 3 and Fig. 2. This is a variant on the strategy developed by French et al. (1978) of replacing the upper part of the lightcurve by the best isothermal fit to that restricted part of the data. Effectively, we are assuming that the upper atmosphere above the inversion’s reliable region does not have large-scale temperature fluctuations, and is comparable in mean temperature to the valid region. A comparison of 28 Sgr and Voyager UV stellar occultation observations bear this out for Saturn (Hubbard et al. 1997). The stability of the hybrid-lightcurve inversion depends on the length of this isothermal “cap.” If the cap extends only from the upper baseline to the 99% level (in units of normalized stellar flux), then the noise in the subsequent upper part of the observed lightcurve will still produce spurious temperature variations at the onset of the inversion. As the cap length increases, the inversion stabilizes, eventually contaminating the inversion’s valid region. The optimal cap is large enough to give a stable inversion but ends far above the valid region.

The left panel of Fig. 4 shows the thermal structure derived from a suite of lightcurve inversions with seven different isothermal caps ending at 99%–75% of the full stellar signal. The order of the curves is not systematic with cap length, and instead reflects the sensitivity of the inversion to features in the data (both due to noise and actual atmospheric structure) that immediately follow the isothermal cap. The degree of uncertainty in the inversion is shown by the spread between the




FIG. 4.— Family of thermal profiles derived by numerical inversion from lightcurves with isothermal caps ending at 99%, 97.5%, 95%, 90%, 85%, 80%, and 75% of full flux. The data are plotted vs. pressure (left axis), with the approximate corresponding altitude scale given on the right axis. The left panel gives the temperature. The middle panel shows the same data after removal of a linear fit to each profile. This shows that atmospheric structure varies by ± 4 K. The right panel gives the vertical temperature gradient vs. altitude, derived from the left panel’s data. The atmosphere is statically stable over this altitude range, since the temperature gradient is well removed from the adiabatic lapse rate (dashed line). This is consistent with the lack of alternating rounded and spiked structures, which are seen in many other profiles when the vertical temperature gradient approaches adiabatic. They have been attributed to breaking inertia-gravity waves. See Inversion section of text for discussion.

profiles. We have used the 95% cap in all subsequent analyses since it lies in the middle of the group and since the cap length is minimal. The average temperature is slightly warmer than results from the equatorial 28 Sgr occultation, which stayed below ~ 140 K in the valid region (see Table VII of Hubbard et al. 1997).

The temperature variations are more clearly seen in the middle panel of Fig. 4, which shows the deviations of each of the profiles from a linear fit to that profile. The right panel shows the vertical temperature gradient for each of the inversions. The vertical dashed line corresponds to the adiabatic lapse rate. The atmosphere is locally stable against convection from 1–60 μbar . Some occultation temperature gradient profiles (e.g., Fig. 10 of Raynaud et al. 2003, Fig. 7 of Raynaud et al. 2004, and others cited therein) have oscillations with a rounded shape on the low $\Delta T/\Delta z$ side and a narrow, spiked shape on the high side. Our profile has no evidence of this pattern. The pattern has been attributed to gravity wave breaking as the profile’s gradient approaches the adiabatic lapse rate. Saturn’s adiabatic lapse rate is separated from the negative-side extrema in our profile’s gradient by 0.2 K km^{-1} everywhere, and generally by larger amounts, apparently sufficient to prevent wave breaking, if the extrema are indeed propagating waves.

5. WAVELET ANALYSIS

The derived temperature profiles all show fluctuations with amplitudes as large as 4 K. To determine the wave nature of these structures in more detail, we performed a wavelet transform of the temperature profile. This technique allows one to study the variation of the power spectrum with atmospheric depth. Small-scale, quasi-periodic structures in atmospheric profiles are often interpreted as inertia-gravity waves (French & Gierasch 1974; Young et al. 1997; Cooray et al. 1998; Raynaud et al. 2003, 2004). The waves would change their wavelengths if the wave speed changed and they would change amplitude if pumped or damped. Such behavior would be evident

in the wavelet transform as tilted or curved ridges, whereas uniform waves show up as vertical ridges of constant amplitude. Torrence & Compo (1998) provide a quantitative and accessible wavelet tutorial with software (see Acknowledgements).

Following Raynaud et al. (2003, 2004), we used the Morlet wavelet with parameter value 6. The left panel of Fig. 5 gives the inversion data, while the right panel shows the wavelet amplitudes. We have removed the unit-energy normalization of Torrence & Compo (1998, their Eq. 8) so that amplitudes are in Kelvins. This adjustment and the orthogonal basis set of the transform allow us to recover accurately the amplitudes of synthetic sinusoidal signals inserted into the input data. The cross-hatched region is the so-called “cone of influence” (COI) of the edges of the data. This is the region in which the wavelets are significantly off the edge of the data, with the effect of averaging in zeros from outside the data. Structure within the COI is unreliable, so we ignore it.

As noted above, inversions suffer from spurious wave-like structures on the length scale of their valid regions. To study waves, one generally removes a constant or slowly-varying basic state. This is not necessary with wavelets. As Raynaud et al. (2003, 2004) point out, one can zero the short scales (e.g., less than H) and inverse transform to derive the basic state. However, after subtracting that from the inversion, the orthogonal basis set ensures that the result is identical to zeroing (or simply ignoring) the large scales. We have thus chosen to ignore scales larger than ~ 32 km, both for this reason and because only a small portion of the data are outside the COI at these scales (see Fig. 5).

The profile is dominated by structures with wavelengths longer than 32 km. Few-km, irregular wiggles are superposed, but 5–30 km features are absent. The wavelet transform quantifies this: wavelet amplitudes drop substantially outside the COI. The contour lines encircle regions of 95% confidence with respect to the global wavelet power spectrum (i.e., the

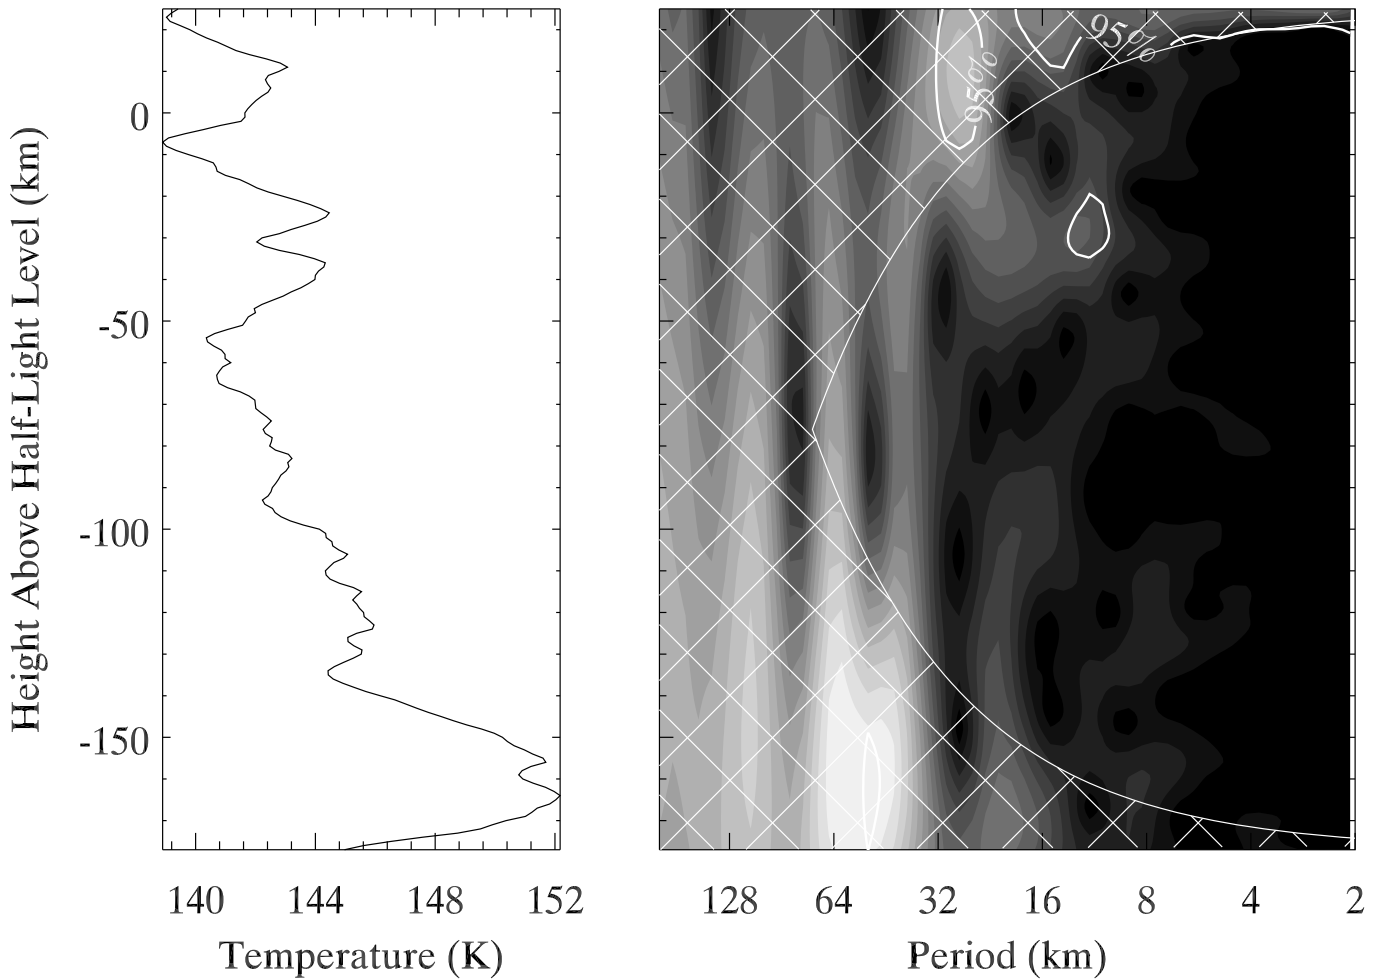


FIG. 5.— Wavelet analysis of the thermal profile. Left: Thermal profile from the 14 Nov 1998 occultation by Saturn. Right: Corresponding wavelet amplitude (not power) spectrum, with same vertical axis. The data are not normalized. The cross-hatched region is the unreliable “cone of influence” (COI) of edge effects that alter the amplitudes. The solid contour is the 95% confidence level based on the global wavelet power spectrum (Torrence & Compo 1998). The large and variable degree of correlation in occultation inversion profiles renders such statistical significance tests inadequate, however. Empirical tests with real and Gaussian noise produced features at least as strong as the 0.62 K peak enclosed by the single contour outside of the COI. See Wavelet section of text for discussion.

wavelet power spectrum averaged over height). Only one of these areas is outside the COI; it has an amplitude of 0.62 K, a period of 32 km, and is located ~ 30 km below the half-light level. However, we will show in the next section that such features can be generated by noise in the lightcurve, indicating that the large and variable degree of correlation in occultation inversions takes such data out of the realm of the statistical wavelet significance test of Torrence & Compo (1998).

6. NOISE TESTS

It is not a simple matter to quantify in detail the effect of lightcurve noise on the numerical inversions. The inferred structure at a given atmospheric level is contaminated by errors in the derived refractivity of all overlying levels. French et al. (1978) showed that the correlation length scale of the inversion process extends to well over a scale height above any given pressure level, as dictated by the width of the kernel in the integral equation for the Abel transform. Thus, even uncorrelated white noise in the lightcurve results in correlated errors in the derived thermal profile.

The problem is even more complex if the input lightcurve noise is correlated, and almost *any* baseline drift on time scales longer than the exposure time will introduce signifi-

cant correlation. Causes of such drifts include scintillation in the Earth’s atmosphere and pointing drifts. The latter can be problematic in spacecraft with few-spatial-channel instruments.

Although previous investigators have explored the consequences of white noise to the *mean* temperature determined by inversion (e.g., French et al. 1978; Elliot et al. 2003), there has been no systematic study of the effects of noise on wave analyses, and little consideration of correlated noise. To assess the significance level of waves in our derived profiles, we added both Gaussian and real noise (taken from the unocculted portion of our lightcurve) to an isothermal model lightcurve. We created 25 realizations of Gaussian noise with the same standard deviation as our upper baseline. To eliminate the baseline uncertainty issue, we added this noise only below the 80% light level. For the real noise tests, we took a section of our (slightly red) upper baseline, subtracted the mean, removed a low-order polynomial, and repeated the section several times. The polynomial ensures that the sections meet without a discontinuity. We shifted the the result by three different amounts and added it to the isothermal lightcurve starting at each of the seven caps levels, to create

21 sets.

Both types of noise induced large oscillatory structures into the profiles. Inversion ranges were typically 2–3 K for the Gaussian noise (a single outlier had a 6.25 K range). Within a given shift of the real noise, the profiles and their ranges were mostly similar. Ranges for the three sets averaged 7.4, 5, and 4.5 K. Fortunately, these large-amplitude oscillations generally have periods of at least a scale height, which puts them in the COI in a wavelet transform (similar to the transform in Fig. 5). At shorter periods, the transforms had much lower amplitudes. Maxima for the Gaussian set averaged ~ 0.13 K, and for the three real-noise sets were 0.19, 0.30, and 0.34 K, respectively.

Our real noise had a dramatically larger effect on the inversions than Gaussian noise. However, in both cases about half the tests contained regions outside of the COI which passed the Torrence & Compo (1998) 95% global wavelet spectrum significance test. These regions were larger for the real noise tests, often substantially larger than the one in Fig. 5. Based on our tests, we conservatively estimate a 0.3 K uncertainty level due to the real noise in our dataset. The 0.62 K feature is thus only a 2σ detection, and it is highly localized, meaning that there is no evidence for continuous vertical propagation. We would consider believing waves larger than 1 K in amplitude, or somewhat less in the case of continuous vertical propagation through the valid region.

We conclude that real noise in the data, even for a high-S/N event such as this one, can result in false-positive wave detections simply due to the correlated nature of scintillation. One must apply caution when making wave interpretations of fluctuations in ground-based occultation inversions, particularly for longer spatial wavelengths. We note that numerous reported wave features in the literature have amplitudes smaller than 1 K, often in noisier datasets (e.g., Raynaud et al. 2003, 2004). In such cases, modeling additional wave parameters, such as the effects of damping on amplitude and phase (Raynaud et al. 2004), may help to make the wave identifications more robust.

7. CONCLUSIONS

We have presented a lightcurve, an atmospheric thermal profile, and a wavelet analysis based on IRTF observations

of the 1998 November 14 occultation of GSC 0622-00345 by Saturn. Techniques used to improve the photometry include using Rhea as a scintillation and pointing standard, subsampling the data by a factor of 10 for accurate photometry aperture placement, and removing a template of scattered light. The lightcurve has a per-frame S/N of 66 and a per-scale-height S/N of 267, placing it among the best ground-based atmospheric lightcurves.

The thermal profile is only slightly warmer than those presented for the equatorial region. Table VII of Hubbard et al. (1997) shows temperatures ranging from 134–140 K, with two inflections in this range. This information is useful for modeling the seasonal and latitudinal variations of Saturn's stratosphere. The stratospheric region sounded by the occultation is statically stable: the vertical temperature gradient is removed by more than 0.2 K km^{-1} from the adiabatic lapse rate. Our thermal gradient profile does not show the alternating-rounded-spiked appearance of some profiles that closely approach the adiabatic lapse rate, including one for Saturn (Cooray et al. 1998). This shape has been interpreted as evidence of gravity wave breaking (Raynaud et al. 2003, 2004).

Our noise analysis shows that even with a statistical 95% confidence test it is possible to get false positives when presenting inversion analyses of atmospheric lightcurves. It is thus important to present empirical noise tests using several realizations of both Gaussian and observed upper baseline noise. Consistent with the symmetric thermal gradient profile, our wavelet analysis detected no significant wave activity.

We thank P. Nicholson for his pre-occultation calculations and K. Matcheva for wavelet discussions. We thank the NASA Astrophysics Data System, JPL Solar System Dynamics group, and the free and open-source software communities for software and services. Free wavelet software (<http://paos.colorado.edu/research/wavelets/>) was provided by C. Torrence and G. Compo. This investigation was supported by Wellesley College under NASA Contract #961169 and by the NASA PGG program.

REFERENCES

- Bosh, A. S. & McDonald, S. W. 1992, *AJ*, 103, 983
 Cooray, A. R. & Elliot, J. L. 2003, *ApJ*, 587, L121
 Cooray, A. R., Elliot, J. L., Bosh, A. S., Young, L. A., & Shure, M. A. 1998, *Icarus*, 132, 298
 Elliot, J. L., Person, M. J., & Qu, S. 2003, *AJ*, 126, 1041
 French, R. G., Elliot, J. L., & Gierasch, P. J. 1978, *Icarus*, 33, 186
 French, R. G. & Gierasch, P. J. 1974, *Journal of Atmospheric Sciences*, 31, 1707
 French, R. G., Nicholson, P. D., & Matthews, K. 1999, in Proc. 5th workshop on observations of mutual events PHEMU97 held in Catania, Italy on March 4–6, 1997, ed. J. Arlot & C. Blanco, Institut de Mécanique Céleste, Bureau des Longitudes, Paris, 163–168
 Harrington, J., Cooke, M. L., Forrest, W. J., Pipher, J. L., Dunham, E. W., & Elliot, J. L. 1993, *Icarus*, 103, 235
 Horne, K. 1986, *PASP*, 98, 609
 Hubbard, W. B., Porco, C. C., Hunten, D. M., Rieke, G. H., Rieke, M. J., McCarthy, D. W., Haemmerle, V., Haller, J., McLeod, B., Lebofsky, L. A., Marcialis, R., Holberg, J. B., Landau, R., Carrasco, L., Elias, J., Buie, M. W., Dunham, E. W., Persson, S. E., Boroson, T., West, S., French, R. G., Harrington, J., Elliot, J. L., Forrest, W. J., Pipher, J. L., Stover, R. J., Brahic, A., & Grenier, I. 1997, *Icarus*, 130, 404
 Kliore, A. J., Lindal, G. F., Patel, I. R., Sweetnam, D. N., Hotz, H. B., & McDonough, T. 1980, *Science*, 207, 446
 Lindal, G. F., Sweetnam, D. N., & Eshleman, V. R. 1985, *AJ*, 90, 1136
 Nicholson, P. D., McGhee, C. A., & French, R. G. 1995, *Icarus*, 113, 57
 Raynaud, E., Drossart, P., Matcheva, K., Sicardy, B., Hubbard, W. B., Roques, F., Widemann, T. H., Gladstone, G. R., Waite, J. H., Nadeau, D., Bastien, P., Doyon, R., Hill, R., Rieke, M. J., & Marley, M. 2003, *Icarus*, 162, 344
 Raynaud, E., Matcheva, K., Drossart, P., Roques, F., & Sicardy, B. 2004, *Icarus*, 168, 324
 Shure, M., Toomey, D. W., Rayner, J., Onaka, P., Denault, A., Stahlberger, W., Watanabe, D., Criez, K., Robertson, L., & Cook, D. 1994, *Experimental Astronomy*, 3, 239
 Smith, G. R., Shemansky, D. E., Holberg, J. B., Broadfoot, A. L., Sandel, B. R., & McConnell, J. C. 1983, *J. Geophys. Res.*, 88, 8667
 Torrence, C. & Compo, G. P. 1998, *Bull. Amer. Meteorol. Soc.*, 79, 61
 Young, L. A., Yelle, R. V., Young, R., Seiff, A., & Kirk, D. B. 1997, *Science*, 276, 108