Inverse modeling of GPS multipath for snow depth estimation – Part II: Application and validation

Felipe G. Nievinski and Kristine M. Larson
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Boulder, CO 80309-0429

Corresponding Author:
Felipe G. Nievinski
fgnievinski@gmail.com

ABSTRACT

GPS multipath reflectometry (GPS-MR) is a technique that uses geodetic quality GPS receivers to estimate snow depth. The accuracy and precision of GPS-MR retrievals are evaluated at three different sites: grasslands, alpine, and forested. The assessment yields a correlation of 0.98 and an RMS error of 6-8 cm for observed snow depths of up to 2.5 m. GPS-MR under-estimates *in situ* snow depth by ~10-15% at these three sites, although the validation methods do not measure the same footprint as GPS-MR.

1. INTRODUCTION

A new method to measure snow depth based on GPS multipath present in signal-to-noise ratio (SNR) observations was introduced in [1]. An improved forward/inverse methodology has been formulated in [2]. It capitalizes on knowledge about the antenna response and the physics of surface scattering to aid in retrieving the unknown snow conditions. Although it had been demonstrated with simulations [2], its application to actual measurements collected under real-world conditions poses challenges that must be overcome to prove the concept.

Here we examine the snow depth retrievals from the GPS multipath retrieval algorithm and assess both the precision and accuracy of the method. Multiple metrics are developed to assess the quality of the results. The accuracy of the method is evaluated by comparing with *in situ* data over a multi-year period. The three field sites were chosen to highlight different limitations in the method, both in terms of terrain and forest cover.

We start with a general development applicable to all sites, in which intermediate results are explored in more detail. Then we proceed to show the final snow depth time series at each of the three sites, validating them against independent *in situ* measurements.

2. General development

In this section we examine the satellite coverage, over time and space; the matching of model and measurements, in terms of observation residuals; the quality control procedures used to mitigate anomalous results; and combinations of estimates obtained from different satellites.

2.1 Satellite coverage and track clustering

All GPS-MR retrievals reported here are based on the newer GPS L2C signal. Of the ~ 30 GPS satellites, 8-10 L2C satellites were available between 2009-2012 (8, 9, and 10 satellites at the end of 2009, 2010, and 2011, respectively). Satellite observations are partitioned into ascending and descending portions, yielding approximately twenty unique tracks per day at a site with good sky visibility. GPS orbits are highly repeatable in azimuth, with deviations at the few-degree range over a year, translating into ~ 50-100 cm azimuthal displacement of the reflecting area (corresponding to the first Fresnel zone at 10-15° elevation angle for a 2-m high antenna, see Appendix in [3]). This repeatability permits clustering daily retrievals by azimuth. It also allows the simplification that estimated snow-free reflector heights are fairly consistent from day to day, facilitating the isolation of the varying snow depth during the snow-covered period.

For a given track, its revisit time is also repeatable, amounting to practically one sidereal day. The deficit in time relative to a calendar day results in the track time of the day receding ~ 4 min and 6 sec every day [4]. This slow but steady accumulation eventually makes the time of day to return to its starting value after ~ one year. As all GPS satellites drift approximately at the same rate, the time between successive tracks remains nearly repeatable. Its reciprocal, the sampling rate, has a median equal to approximately one track per hour, with a low value of one track within two hours and a high of one track within 15 min; both extremes occur every day, with low-rate idle periods interspersed with high-rate bursts. The time of the day reduced to a fixed day (e.g., January 1st 2000) could also be used to cluster tracks. Neighboring clusters, that are close in azimuth and/or in reduced time of the day are expected to be more comparable, as they sample similar conditions and are subject to similar errors. We materialize this notion in terms of diagrams of track clusters, having azimuth and reduced time of day as horizontal and vertical axes, respectively; it will be applied to each of the three sites below (Figure 4).

2.2 Observations

Figure 1 shows several representative examples of SNR observations. A typical good fit between measured and modeled is shown in Figure 1(a), corresponding to the beginning of the snow season. Generally the model/measurement fit is good when the scattering medium is homogeneous; it deteriorates as the medium becomes more heterogeneous, particularly with mixtures of soil, snow, and vegetation. Here we discuss genuine physical effects as well as more mundane spurious instrumental issues that degrade the fit but do not necessarily cause a bias in snow depth estimates.

2.2.1 Secondary reflections

Throughout we have assumed the existence of a single specular reflection, which matches large planar surfaces. Finite and/or non-planar surfaces represent a departure from this assumption; the first case transforms a simpler reflection into a more complicated diffraction phenomenon, while the second case might as well introduce secondary reflections, originating from disjoint surface regions. Interference fringes become convoluted with multiple superimposed beats, see Figure 1(b). As long as there is there is a unique dominating reflection, the inversion will have no difficulty fitting it, as the extra reflections will remain approximately zero-mean.

2.2.2 Interferometric power effects

Random deviations of the actual surface with respect to its undulated approximation – called roughness or residual surface height – will affect the interferometric power, P_i. SNR measurements will exhibit a diminishing number of significant interference fringes, compared to the measurement noise level Figure 1(c). This facilitates the model fit but the reflector height parameter may become ill-determined – its estimates will be more uncertain. Changes in snow density also affect the fringe amplitude.

2.2.3 Direct power effects

Snow precipitation attenuates the satellite-to-ground radio link which affects SNR measurements through the direct power term. First, this shifts the SNR measurements up or down (in decibels); second, it tilts the trend tSNR as attenuation is elevation angle dependent; third, fringes in dSNR will change in amplitude because of the decrease in the coherent component of the direct power.

Partial obstructions can affect either or both direct P_d and interferometric powers P_i. In this case, SNR measurements, albeit corrupted, are still recorded. This situation is in contrast to complete blockages as caused by topography. The deposition of snow and the formation of a

winter rime on the antenna are a particularly insidious type of obstruction, as their presence in the near-field of the antenna element can easily distort the gain pattern in a significant manner. In the far-field, trees are another important nuisance, so much so that their absence is held as a strong requirement for the proper functioning of multipath reflectometry [3].

2.2.4 Instrument-related issues

Satellite-specific direct power offsets and also long-term power drifts are to be expected as spacecrafts age and modernized designs are launched. In addition, noise power depends on the state of conservation of receiver cables and on their physical temperature. Less subtle incidents are sudden ~ 3-dB SNR steps, hypothesized to originate in the receiver switching between the L2C data and pilot subcodes, CM and CL [5].

2.3 Quality control (intra-clusters)

Anomalous conditions may result in measurement spikes, jumps, and short-lived rapidly-varying fluctuations. For snow depth sensing purposes, it is necessary and sufficient to either neutralize such measurement outliers through a statistically robust fit or detect unreliable fits and discard the problematic ones that could not otherwise be salvaged.

The key to quality control (QC) is in grouping results into statistically homogeneous units, having measurements collected under comparable conditions. In our case, azimuth-clustered tracks are the natural starting unit. Secondarily, we must account for genuine temporal variations in the tendency of results, i.e. from beginning to peak to the end of the snow season. The detection of anomalous results further requires an estimate of the statistical dispersion to be expected. Considering that the sample is contaminated with outliers, robust estimators – running median instead of the running mean, and median absolute deviation over the standard deviation – are called for, if the first- and second-order statistical moments are to be representative. Given

estimates of the non-stationary tendency and dispersion, a tolerance interval can then be constructed such that it bounds, say, a 99% proportion of the valid results with 95% confidence level. We also desire QC to be judicious, or else too many valid estimates will be lost. Notice that in the present intra-cluster QC we compare an individual estimate to the expected performance of the track cluster to which it belongs; later we complement QC with an inter-cluster comparison of each cluster's own expected performance.

Based on our practical experience, no single statistic detects all the outliers. In the sub-sections that follow we present four particular statistics that we have found to be useful. The corresponding time series of estimated reflector heights are shown in Figure 2 (top panel); these raw reflector heights will be transformed into snow depth in the next section. Even in raw form, the snow season can be clearly distinguished by the dramatic increase in reflector height retrievals (only the negative bias with respect to the *a priori* value is shown). Several snow precipitation events are discernible as sudden rises in reflector height, followed by a slower decay indicating snow settling. At the end of the season, melting occurs until bare ground eventually becomes exposed. Over the remainder of the year, reflector heights varies just a few centimeters at this site, which indicates that our *a priori* value for the height of the antenna above the ground was accurate and that the site is devoid of vegetation with large amounts of water content.

Statistics that involve a sum of squared values are expected to follow a chi-squared distribution. To accommodate this characteristic, we handle such values in logarithm form, restoring normality. For these second-order statistics, a single-tailed tolerance bound was found more appropriate than the double-tailed intervals applied to first-order statistics.

2.3.1 Degree of freedom

The simplest statistic is the degree of freedom, essentially the number of observations per track (modulo a constant number of parameters). The time series shown in Figure 2 (bottom panel) is stationary, so a global low-order polynomial (instead of a running average) suffices for this statistic. It does a good job detecting problems related to data outages.

2.3.2 Goodness of fit

We use the scaled root-mean-square error (RMSE, denoted $\hat{\sigma}_0$ in Part I, sec. 5.1) to test for goodness-of-fit, i.e., how well measurements can be explained adjusting the unknown values for the parameters postulated in eq.(24) of Part I. Figure 2 (second panel from top) shows a time series of $\hat{\sigma}_0$ and its tolerance interval. Tracks near days -25 reject the null hypothesis of statistically equivalence against the long-term non-stationary average thus are deemed outliers. Notice we neutralized the risk posed by outliers in distorting the tolerance intervals.

2.3.3 Reflector height uncertainty

Sometimes the fit is good but the reflector height uncertainty $\hat{\sigma}_{\hat{H}_B}$ is bad as the multipath modulation is gone – there are not many oscillations greater than the measurement noise level, Figure 1(c). Such cases are missed by the previous statistics but are detected by this one. High uncertainty is more common during heavy snowfalls.

2.3.4 Peak elevation angle

The peak elevation angle (defined in Part I, sec. 6.2) behaves much like a random variable, as it is determined by a multitude of factors. We can thus form a tolerance interval to detect outliers. This statistic was found to perform especially well in cases that were particularly challenging for the previous statistics. For example, some fits yielded small residuals and produced unsuspecting reflector height uncertainties, but the result corresponded to a substantially different peak

elevation angle. This was the case not just because the SNR oscillations that are typically found in SNR turned out to be missing on that day. The aggravating factor is the presence of oscillations at atypical elevation angles. Despite being well fit and well determined, they are generated by different reflecting conditions, compared to those found on most other days, see Figure 1(d).

2.4 Combinations (inter-clusters)

In the previous section we dealt with results on a cluster-by-cluster basis. Now we are interested in combining multiple clusters. The main purpose is to average out random noise. Noise mitigation aims at not only coping with measurement errors but also compensating for model deficiencies, to the extent that they are not in common across different clusters.

2.4.1 Vertical datum

Before we combine different clusters, we have to address their long-term differences. Recall that total reflector height $H = H_A - H_B$ is the difference between its *a priori* value H_A and the estimated bias H_B; it is positive downward, reckoned from the horizontal plane containing the receiving antenna. For easier comparison we denote H_{ij} the i-th day and j-th cluster estimate; $\mathcal{W}$ will denote the set of Winter days, indexed by i (see Figure 3). The initial situation is that snow surface heights H_{ij} will be greater downhill and smaller uphill; we amend it as follows. Cluster-wise uncertainty-weighted median of estimated reflector heights H_{ij} (evaluated in the snow-free period $i \notin \mathcal{W}$) is taken as the ground height, valid for all days on a cluster-by-cluster basis:

$$\bar{H}_j = \underset{i \notin \mathcal{W}}{\text{med}}\, H_{ij} \tag{1}$$

Subtracting ground heights $\bar{H}_j$ from their respective snow surface heights H_{ij} results in snow thickness values,

$$T_{ij} = H_{ij} - \bar{H}_j \tag{2}$$

which is a completely physically unambiguous quantity. Snow thickness is more comparable than snow heights across varying-azimuth track clusters. Yet snow tends to fill-in ground depressions, so thickness exhibits variability caused by the underlying ground surface, even when the overlying snow surface is relatively uniform. Further cluster homogeneity can be achieved by accounting for the temporally permanent though spatially non-uniform component of snow thickness. This is achieved defining snow depth as:

$$D_{ij} = T_{ij} - \delta\bar{T}_j \tag{3}$$

where the median thickness deviation $\delta\bar{T}_j = \bar{T}_j - \bar{\bar{T}}$ removes the cluster-wise median thickness $\bar{T}_j = \mathrm{med}_{i\in\mathcal{W}}\, T_{ij}$ (evaluated in the snow-covered period, $i \in \mathcal{W}$) and restores the site-wide median thickness $\bar{\bar{T}} = \mathrm{med}\, \bar{T}_j$. In effect, the vertical datum is established having the topography of the median snow surface but lowering its mean level closest to the ground level (though without conforming to the ground topography). As the deviations $\delta\bar{T}_j$ sum up to zero, average depth retains the same physical meaning as average thickness, albeit showing a smaller inter-cluster variability.

2.4.2 Averaging

The averaging of snow depths collected for different track clusters employs the inversion uncertainties to obtain a preliminary running weighted median D_k, calculated at a given spacing $k = 1,2,\ldots$ (say, daily postings), with overlapping or not. The preliminary post-fit residuals $\delta D_{ij} = D_{ij} - D_k$ then go through their own averaging – necessarily employing a wider averaging window (say, monthly) –, which produces scaling factors for the original uncertainties. The running weighted median is then repeated, producing final averages. The variance factors reflect the fact that some clusters are better than others.

Thus the final GPS estimates of snow depth follow from an averaging of all available tracks, whose individual snow depth values were previously estimated independently. A new average is produced twice daily utilizing the surrounding 1-2 days of data (depending on the data density), i.e., 12-h posting spacing and 24-h moving window width. The averaging interval must be an integer number of days, so as to minimize the possibility of snow depth artifacts caused by variations in the observation geometry, which repeats daily.

In the site results below (Figure 6, Figure 8, and Figure 10), a dark-gray band denotes the 95% confidence interval for the average, which follows from the averaging least-squares uncertainty, scaled by the RMS of inter-cluster residuals, then expanded based on Student's t-distribution for the number of tracks available. The light-gray band denotes the 95% simultaneous prediction interval for a random observation, utilized to detect and reject outliers. The relationship between the two statistics is analogous to the standard error of the mean $\sigma/\sqrt{n}$ and the standard deviation σ in an n-element sample. Individual GPS tracks that passed quality control are shown as gray dots.

3. Site-specific results

GPS-MR snow depth retrieval is now further explored at three stations (Table 1) and over a longer period (up to 3 years). Throughout, we assess the performance of the GPS against independent nearly co-located *in situ* measurements (Table 2). We also compare the GPS estimates to the nearest SNOTEL station [6]. Although not co-located with GPS, SNOTEL data – widely used for operational snow monitoring in the U.S – are important because they provide accurate information on the timing of snowfall events.

Table 1: Site summary; last two columns refer to its separation from nearest SNOTEL station.

Code	Environment	Latitude (degrees)	Longitude (degrees)	Altitude (m)	Horizontal (km)	Vertical (m)
P360	Grassland	44.317852	-111.450677	1857.9	~ 12	-60
RN86	Forested	41.864856	-111.502514	2590.7	0.350	~ 0
NWOT	Alpine	40.055387	-105.590527	3522.5	4	+380

Table 2: *In situ* data quantity; replication indicates the number of values sampled per epoch.

Code	Duration	Interval	# Epochs	Type	Replication
P360	6 mon.	~ 6 h	500	webcam	1
RN86	7 mon.	~ 2-3 week	9	manual	20-150
NWOT	3 yr.	~ 2-3 week	60	manual	1

Table 3: Regression coefficients, with *in situ* snow depth taken as the explanatory or independent variable and GPS snow depth as the response or dependent variable.

Station	Intercept (cm)	Slope (m/m)
P360	5.9 ± 1.5	1−0.04 ± 0.05
RN86	15.4 ± 9.1	1−0.14 ± 0.09
NWOT	2.1 ± 2.8	1−0.11 ± 0.03

3.1 Forested site (RN86)

GPS site RN86 was installed at the T. W. Daniel Experimental Forest in Fall 2011 so we show results for only one water-year (which is the period starting October 1st through September 30th of the following year, encompassing the northern-hemisphere winter; a water-year is designated by the calendar year in which it ends). Topographical slopes range from 2.5° to 6.5° (at the 2-m spatial scale), with average of ~ 5° within 50-m radius around the GPS antenna. RN86 was specifically built to study the impact of trees on GPS snow depth retrievals (Figure 5). Ground crews manually collected *in situ* measurements around the GPS antenna approximately every other week starting in November, 2011. Measurements are made every 1-2 m from the GPS up to

25-30 m. In the second half of the year, the sampling protocol was changed to azimuths of 0° (N), 45° (NE), 135° (SE), 180° (S), 225° (SW), and 315° (NW). With these data it is possible to obtain *in situ* average estimates, with their own uncertainties (based on the number of measurements), which allows a more meaningful comparison.

The diagram of track clusters (Figure 4) indicates a reduced visibility at the current site, compared to other sites. Also the cluster uncertainty is not as favorable, as indicated by the small-radius circles. Notice clusters are concentrated due south, with only two clusters located within ± 90° of north. Therefore, the GPS average snow depth is not necessarily representative of the azimuthally symmetric component of the snow depth. In the presence of an azimuthal asymmetry in the snow distribution around the antenna, the GPS average is expected to be biased towards the environmental conditions prevalent in the southern quadrant. To rule out the possibility of an azimuthal artifact in the comparisons, we have utilized only the *in situ* data collected along the SE/S/SW azimuths.

The comparison shows generally excellent agreement between GPS and *in situ* data (Figure 6). The first four and the last one *in situ* data points were collected with coarser spacing and/or smaller azimuthal coverage, which may be partially responsible for different performance in the first and second halves of the snow season. The correlation between GPS and *in situ* snow depth at RN86 amounts to 0.990, indicating a very strong linear relationship. Inspecting the individual differences at each of 9 visits (Table 4), we find that all are within the corresponding uncertainty, which is somewhat large given the propagation of GPS and *in situ* uncertainties. The GPS uncertainties are generally slightly smaller than the *in situ* ones, with a few exceptions when the number of usable tracks decreases. It is important to bear in mind, though, that there is a trade-off between uncertainty and temporal representativeness, in that wider averaging windows

(sec. 2.4.2) yield smaller uncertainties but also smooth out potentially genuine rapidly-changing snow depth variations.

Carrying out a regression between *in situ* and GPS values, the RMS of snow depth residuals improves from 9.6 cm to 3.4 cm. The regression intercept and slope (with corresponding 95% uncertainties) amount to 15.4 cm ± 9.1 1 cm and 0.858 m/m ± 0.09 m/m, respectively. According to these statistics, the null hypotheses of zero intercept and unity slope are rejected at the 95% confidence level. This implies that at this location GPS snow depth estimates exhibit both additive and multiplicative biases. The latter is proportional to snow depth itself, meaning that, compared to an ideal one-to-one relationship, GPS is found to under-estimate *in situ* snow depth at this site by 14% ± 9%, albeit the uncertainty is somewhat large.

The SNOTEL sensors are exceptionally close to the GPS at this site, ~ 350 m horizontally with negligible vertical separation. Yet the former is located within trees, while the latter is located in the periphery of the forest and senses the reflections scattered from an open field. Therefore only the timing of snowfall events agrees well, not the amount of snow. Although forest density is generally negatively correlated with snow depth [7], exceptions are not uncommon [8], [9], especially in localized clearings exposed to intense solar radiation, where shading of the snow by the trees reduces ablation.

Table 4: Snow depth at the forested site (RN86) for each sampled day. The 95% uncertainty is given after ±. Column "DOY-W" denotes day of water year and "num. obs." is number of independent *in situ* observations (unknown numbers are indicated by *, in which case the *in situ* uncertainty is taken as the median value).

Year	Month	Day	DOY-W	GPS (cm)	*In situ* (cm)	Diff. (cm)	Num. obs.
2011	11	18	49	40 ± 11	27 ± 11	13 ± 15	4
2012	1	5	97	60 ± 18	51 ± 19	9 ± 26	*
2012	2	10	134	107 ± 22	123 ± 11	-16 ± 24	10
2012	2	24	148	138 ± 10	140 ± 12	-2 ± 16	15
2012	3	15	168	129 ± 8	135 ± 10	-6 ± 12	14
2012	3	30	183	129 ± 14	123 ± 10	7 ± 17	74
2012	4	11	195	100 ± 43	94 ± 11	6 ± 45	76

2012	4	25	209	68 ± 10	65 ± 11	4 ± 15	74
2012	5	10	224	20 ± 10	5 ± 19	15 ± 21	*

3.2 Grassland site (P360)

P360 (Figure 7) is one of 1100 GPS stations that make up the EarthScope Plate Boundary Observatory (http://pbo.unavco.org). It was installed with the purpose of studying crustal deformation in the western U.S. The typical setup is a 2-m tall metal tripod drilled into bedrock. At the apex rests a choke-ring antenna (boresight facing zenith), both housed within a radome.

P360 is located in an open field; visibility to the ground is unobstructed. The nearest trees are ~ 200 m away due west; visibility to the sky is also excellent. The ground is mostly flat, with topography deviating no more than 3.5 m above and below the mean horizontal plane within a 100-m radius of the antenna. Topographical slopes range from 0.7° to 1.7° (at the 2-m spatial scale), with average of ~ 1° within 50-m radius around the GPS antenna. At the sub-meter scale the terrain is rugged, with exposed rocks and littered with loose cobbles. Land cover classification is grasslands. There is a watercourse 200-m away due NE-W.

At P360 we show three years of estimated snow depth. Collection started at the time when L2C tracking was enabled in the receiver. Throughout this period there are sonic snow depth measurements available from a SNOTEL station. However, this SNOTEL site is 60-m higher in altitude and at a distance of 12 ± 4 km. In the third year there are up to four-times daily *in situ* validation data from a pole co-located with the GPS antenna; tick marks can be read in from the photographs with a precision of 3 cm. Figure 8 shows results separately for each water-year; in the second year, satellite PRN 25 had been launched, and for the third water-year, satellite PRN 01 was also launched, each new satellite adding potentially four new track clusters.

This is not to say that all are equally good; the diagram in Figure 4 includes circles that are proportional to the weight of each cluster (i.e., larger circles indicate more important clusters).

Figure 8 shows that generally the timing of snowfall events is comparable for the GPS and SNOTEL sensors, although the amount of snow is not. A salient temporal feature that is well captured by both GPS and SNOTEL is the sharp transition between accumulation and settling that happens when precipitation stops and snow depth starts to drop.

For water-year 2012, *in situ* measurements become available, and we find a markedly improved agreement in terms of absolute amounts of snow depth, in contrast to the SNOTEL comparison. Still, this type of *in situ* data is not expected to be exactly comparable with the GPS estimates. The pole data stem from one-time readings at a fixed location with no coverage area, so it lacks statistical replication necessary to quantify the variability of snow depth. Clearly, the pole readings are not estimates of the mean and as such they are not expected to fall within the GPS confidence bounds for the mean. Furthermore, its spatial footprint is orders of magnitude smaller than the GPS, so the GPS prediction bounds also do not apply.

The GPS/*in-situ* bias that remains constant throughout time could be ascribed to an unchanging spatial trend in snow deposition, likely controlled by the underlying ground topography is this open field environment. The periods of improved GPS/*in-situ* agreement could be explained based on the sample randomization offered by temporal variability, even if sampling takes place at a fixed location. During quiet periods, the GPS/*in-situ* discrepancy exhibits serial autocorrelation, which complicates their regression analysis, as the frequent though non-replicated *in situ* measurements are not statistically independent. To overcome this difficulty, we form monthly averages; the resulting regression intercept is statistically significant (5.9 cm ± 1.5 cm) while the regression slope (0.959 m/m ± 0.05 m/m) is not: the deviation from

a one-to-one relationship (-0.04 m/m), albeit indicating a 4% under-estimating on the part of GPS compared to *in situ,* is smaller than its 95% confidence interval (0.05 m/m).

During the snow-free period we find that reflector height does not remain exactly zero. Variations occur mainly when the scattering medium is transitioning, from snow to slush/mud and eventually grass-covered soil. This issue is both a challenge and an opportunity. On the one hand, it poses the risk of being mistaken for snow depth events. In fact, the identification of the site-overall (i.e., non-cluster specific) zero-level or bare-ground reflector height is perhaps the weakest link in the whole GPS processing chain of snow depth retrieval, as it relies on only a few data points – we compute it as the 5th percentile of site-average reflector heights over the snow-free period. On the other hand, such observations attest to the prospects of using GPS reflector heights for monitoring environmental targets other than snow, such as vegetation biomass [10]. If successful, the estimation of non-snow targets would contribute to guaranteeing that snow depth remains non-negative.

3.3 Alpine site (NWOT)

We finally consider the data collected at the Niwot Ridge Long-term Ecological Research site in Colorado. Topographical slopes range from 2° to 7° (at the 2-m spatial scale), with average of ~ 5° within 50-m radius around the GPS antenna. At an altitude of 3,500 m, it is located in a saddle-like mountain-top, in an alpine tundra environment. A 3-meter tall continuously-operating GPS system was established there in 2009 (Figure 9). Poles are staked at 50-m intervals making up a 120-by-400 meter Cartesian grid at which snow depth is measured manually using a snow sampling tube, approximately every two weeks; we use *in situ* data collected at the pole nearest to the GPS antenna (shown in Figure 9). The ground at the present site is not as planar as in the previous two sites, but visibility to the sky is good, with no trees and only minor topographical

obstructions, predominantly due east and west. Indeed, the diagram of track clusters in Figure 4 shows a nearly uniform azimuthal distribution.

The nearest SNOTEL location is more than 4 km away and ~ 380 m lower in altitude. As is typical for SNOTEL stations, this one is found among trees, whereas the GPS is above the tree line. Therefore their comparison comes with caveats. Similar as for the forested site, the comparison against co-located *in situ* data is more favorable. Results for the pole co-located with the GPS antenna is in good agreement for the first two years, less so in the third year (Figure 10).

In the last year the peak snow depth is much smaller than in previous years, and the performance of the GPS deteriorates. This is partially because the amount of snow, subject to redistribution by the wind, is not sufficient so as to fill-in the ground depressions during most of the season. Therefore the multiple GPS track clusters are not as comparable as in previous years, when the inter-cluster variability leveled out as the air/snow interface became more planar than the snow/ground interface. Furthermore, the spatial variability affects the *in situ* data, too, as the single pole co-located with the GPS is not as representative of the area sampled by the GPS. The remaining poles are more than 50 m from the GPS, so their potential contribution is questionable under such low-snow conditions.

The inter-annual variations in snow depth at the NWOT site are more drastic than at P360. Indeed, it exhibits a five-fold difference in peak snow depth, from ~ 0.5 m in 2011-2012 to ~ 2.5 m in 2010-2011. The timing of the end of the season varies by more than a month over this three-year period. The exact beginning of the season is less clear as the snow that accumulates from the initial precipitation events can be totally dissipated if the snowpack is not replenished with more frequent and vigorous snowfalls.

Figure 11 shows a scatterplot of the GPS vs. *in situ* snow depth for the three-year period at NWOT. The correlation is 0.980, which indicates a very strong linear relationship. Carrying out a regression, RMS of residuals improves from 10.7 cm to 7.8 cm. Notice residuals are more dispersed at smaller snow depth values, especially when considered in proportion or relative to the *in situ* values (Figure 11, bottom panel). The intercept 2.1 ± 2.8 cm is not statistically significant vis-à-vis its 95% confidence interval. The regression slope 0.89 ± 0.03 m/m indicates a small though statistically significant deviation from unity (-0.11 m/m), corroborating a similar finding first detected in the forested site, albeit now by a wider margin – four times smaller uncertainty, owing to the three times longer time series. Therefore, at this site, GPS estimates are ~ 10% lower than *in situ* snow depth, although their footprints are not overlapping.

At NWOT we had extra *in situ* measurements, so we checked whether the GPS/*in-situ* snow depth discrepancies depend on snow density, and found a very small correlation: coefficient value of -0.18. We also tried augmenting the regression, using *in situ* density in conjunction with *in situ* depth, thus estimating a total of three regression coefficients: a GPS/*in-situ* additive bias (in meters); a GPS depth vs. *in-situ* depth multiplicative bias (in m/m); and a GPS depth vs. *in-situ* density linear coefficient (in $\mathrm{m}/(\mathrm{kg} \cdot \mathrm{m}^{-3})$). Compared to the two-coefficient regression, the introduction of density left the multiplicative bias unaltered: changes smaller than 1% for coefficient as well as for its uncertainty.

4. Conclusions

We have demonstrated a statistical inverse model for estimating snow depth based on GPS multipath present in SNR observations. The model performance was assessed against independent *in situ* measurements and found to validate the GPS estimates to within the limitations of both GPS and *in situ* measurement errors after the characterization of systematic

errors. The assessment yields a correlation of 0.98 and an RMS error of 6-8 cm for observed snow depths of up to 2.5 m, with the GPS under-estimating *in situ* snow depth by ~ 5-15% (Table 3). This latter finding highlights the necessity to assess effects currently neglected as recommended in [2].

The fit of SNR observations described in Part I provided parameter estimates and their covariance matrix, as well as observation residuals, for each satellite track. In part II we analyzed the resulting parameters and residuals. We examined a few representative fits, illustrating and discussing the origin of a variety of good and bad conditions, such as measurement noise, well- vs. poorly-determined reflector heights, instrument-related issues, etc. Then we discussed a methodology to quality control these estimates based on track clusters; the thousands of tracks retrieved in a year can thus be analyzed in terms of only 10 to 20 units. We introduced a specially designed diagram as a convenient summary of the track clusters available in a site. Such a repeatable sensing configuration allowed us to compare tracks belonging to the same cluster with the purpose of detecting and rejecting anomalous conditions. This principle leads to a number of strategies for quality control of results, which are needed for operational use of GPS snow sensing. Daily site averages are then compiled after different clusters are homogenized. This entailed accounting for genuine azimuthal asymmetry in the distribution of snow around the antenna, and also dealing with issues such as assigning statistical weights to varying quality track clusters.

The suitability of GPS SNR measurements for snow monitoring was found to be heavily influenced by the site conditions; this lends weight to the finding of [3] that clearance to the satellite line-of-sight as well as to the ground are strong requirements for GPS multipath reflectometry (GPS-MR). Therefore the quality of retrievals may vary enormously over different

azimuths at the same site. Furthermore, for a high-quality track cluster, a few percent of the daily retrievals might be discarded for a variety of reasons, such as receiver failures and heavy snowfall. Quality control (QC) is therefore mandatory for operational exploitation of GPS-MR.

A multi-test QC strategy – including goodness-of-fit, reflector height uncertainty, peak elevation angle, and statistical degree of freedom – was found to work best, as no single test detected all outliers. Combining multiple track clusters further improved the precision of GPS retrievals, yielding daily averages that captured remarkably well the temporal dynamics of snow accumulation and ablation, including sudden changes associated with new snow. Statistically robust methods – e.g., median instead of the mean – were adequate in achieving a reasonable level of processing automation and dispense with frequent manual intervention. Continuity of the time series as new satellites were launched every year indicate no obvious satellite-dependent biases; this stability is paramount for future utilization of GPS-MR results in climate studies.

Turning attention to aspects that would require more care in the future, our treatment of the azimuthal asymmetry exhibited by snow depth was admittedly cursory, in the sense that we only tried to minimize its impact on the daily site averages by making clusters more comparable. This treatment worked well when the amount of snow was enough to fill in the ground depressions. Yet when the amount of snow was insufficient to make the air/snow surface more planar than the snow/ground surface (alpine site, 2011-2012), the treatment failed to improve the dispersion around the site average.

Finally, further investigation is needed for the definition of the bare soil reflector height. The challenge is that although we can measure reflector heights precisely, we cannot unambiguously attribute an individual track estimate or even a daily site average to a specific target, i.e., to distinguish between snow vs. vegetation vs. soil moisture changes manifested in

reflector height. We have relied on the temporal dynamics of reflector heights along with reasonable assumptions about the snow behavior and optionally ancillary information (photographs, temperature records, climatic expectations, etc.) to determine the snow-covered period. This strategy worked very well for large amounts of snow, but it becomes less reliable for smaller amounts. As a rule-of-thumb, a 10-cm reflector height change would be a reasonable cutoff value for distinguishing snow, based on the behavior expected from other targets. So this issue is more serious for ephemeral snow sites, but it remains relevant for all sites.

ACKNOWLEDGEMENTS

This research was supported by NSF (EAR 0948957, AGS 0935725), NASA (NNX12AK21G), and a CU interdisciplinary seed grant. Mr. Nievinski has been supported by a Capes/Fulbright Graduate Student Fellowship and a NASA Earth System Science Research Fellowship (NNX11AL50H). Some of this material is based on data, equipment, and engineering services provided by the Plate Boundary Observatory operated by UNAVCO for EarthScope (http://www.earthscope.org) and supported by the National Science Foundation (EAR-0350028 and EAR-0732947). We thank Jobie Carlisle at Utah State University data collecting field data at RN86. Mark Williams at the University of Colorado Boulder provided field and logistical support for NWOT through the NSF-funded Niwot LTER project. The receiver used at NWOT was leant to the project by Trimble Navigation, with engineering and archiving support provided by UNAVO. Jim Normandeau at UNAVCO provided additional oversight for the installation of the validation pole and camera at P360. SNOTEL data shown in this paper were retrieved from http://www.wcc.nrcs.usda.gov/nwcc/. We thank PBO for providing the site photographs and Google Earth for satellite images.

REFERENCES

[1] K. M. Larson, E. D. Gutmann, V. U. Zavorotny, J. J. Braun, M. W. Williams, and F. G. Nievinski, "Can we measure snow depth with GPS receivers?," *Geophys. Res. Lett.*, vol. 36, no. 17, p. L17502, Sep. 2009. doi:10.1029/2009GL039430

[2] F. G. Nievinski and K. M. Larson, "Inverse modeling of GPS multipath for snow depth estimation – Part I: Formulation and simulations," *IEEE Trans. Geosci. Remote Sens.*, 2013.

[3] K. M. Larson and F. G. Nievinski, "GPS snow sensing: results from the EarthScope Plate Boundary Observatory," *GPS Solut.*, vol. 17, no. 1, pp. 41–52, Mar. 2013. doi:10.1007/s10291-012-0259-7

[4] D. C. Agnew and K. M. Larson, "Finding the repeat times of the GPS constellation," *GPS Solut.*, vol. 11, no. 1, pp. 71–76, Aug. 2006. doi:10.1007/s10291-006-0038-4

[5] R. D. Fontana, W. Cheung, P. M. Novak, and T. A. Stansell, "The new L2 civil signal," in *Proc ION GPS*, 2001, pp. 617–631.

[6] M. C. Serreze, M. P. Clark, R. L. Armstrong, D. A. McGinnis, and R. S. Pulwarty, "Characteristics of the western United States snowpack from snowpack telemetry (SNOTEL) data," *Water Resour. Res.*, vol. 35, no. 7, pp. 2145–2160, Jul. 1999. doi:10.1029/1999WR900090

[7] D. M. Gray and D. H. Male, *Handbook of Snow – Principles, Processes, Management and Use*. Pergamon, 1981, p. 776.

[8] W. Veatch, P. D. Brooks, J. R. Gustafson, and N. P. Molotch, "Quantifying the effects of forest canopy cover on net snow accumulation at a continental, mid-latitude site," *Ecohydrology*, vol. 2, no. 2, pp. 115–128, Jun. 2009. doi:10.1002/eco.45

[9] J. I. López-Moreno and J. Latron, "Influence of canopy density on snow distribution in a temperate mountain range," *Hydrol. Process.*, vol. 22, no. 1, pp. 117–126, Jan. 2008. doi:10.1002/hyp.6572

[10] E. E. Small, K. M. Larson, and J. J. Braun, "Sensing vegetation growth with reflected GPS signals," *Geophys. Res. Lett.*, vol. 37, no. 12, p. L12401, Jun. 2010. doi:10.1029/2010GL042951

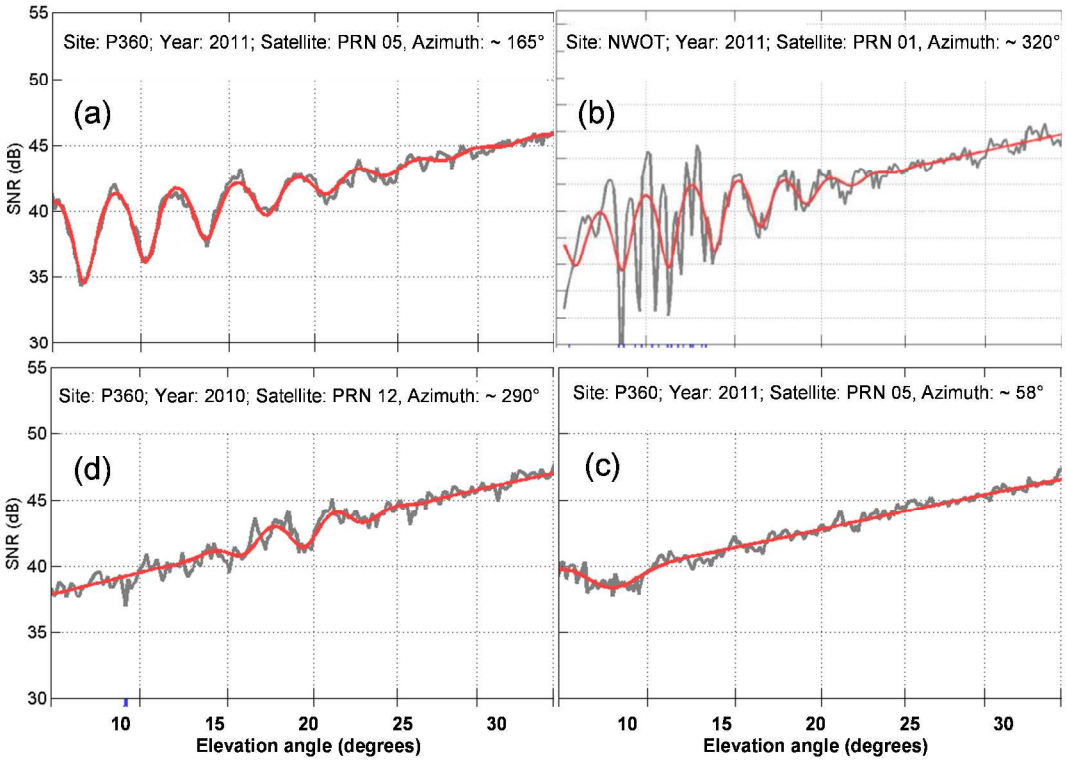

Figure 1: Examples of observations: (a) good fit; (b) presence of secondary reflections; (c) vanishing interference fringes; (d) atypical interference fringes

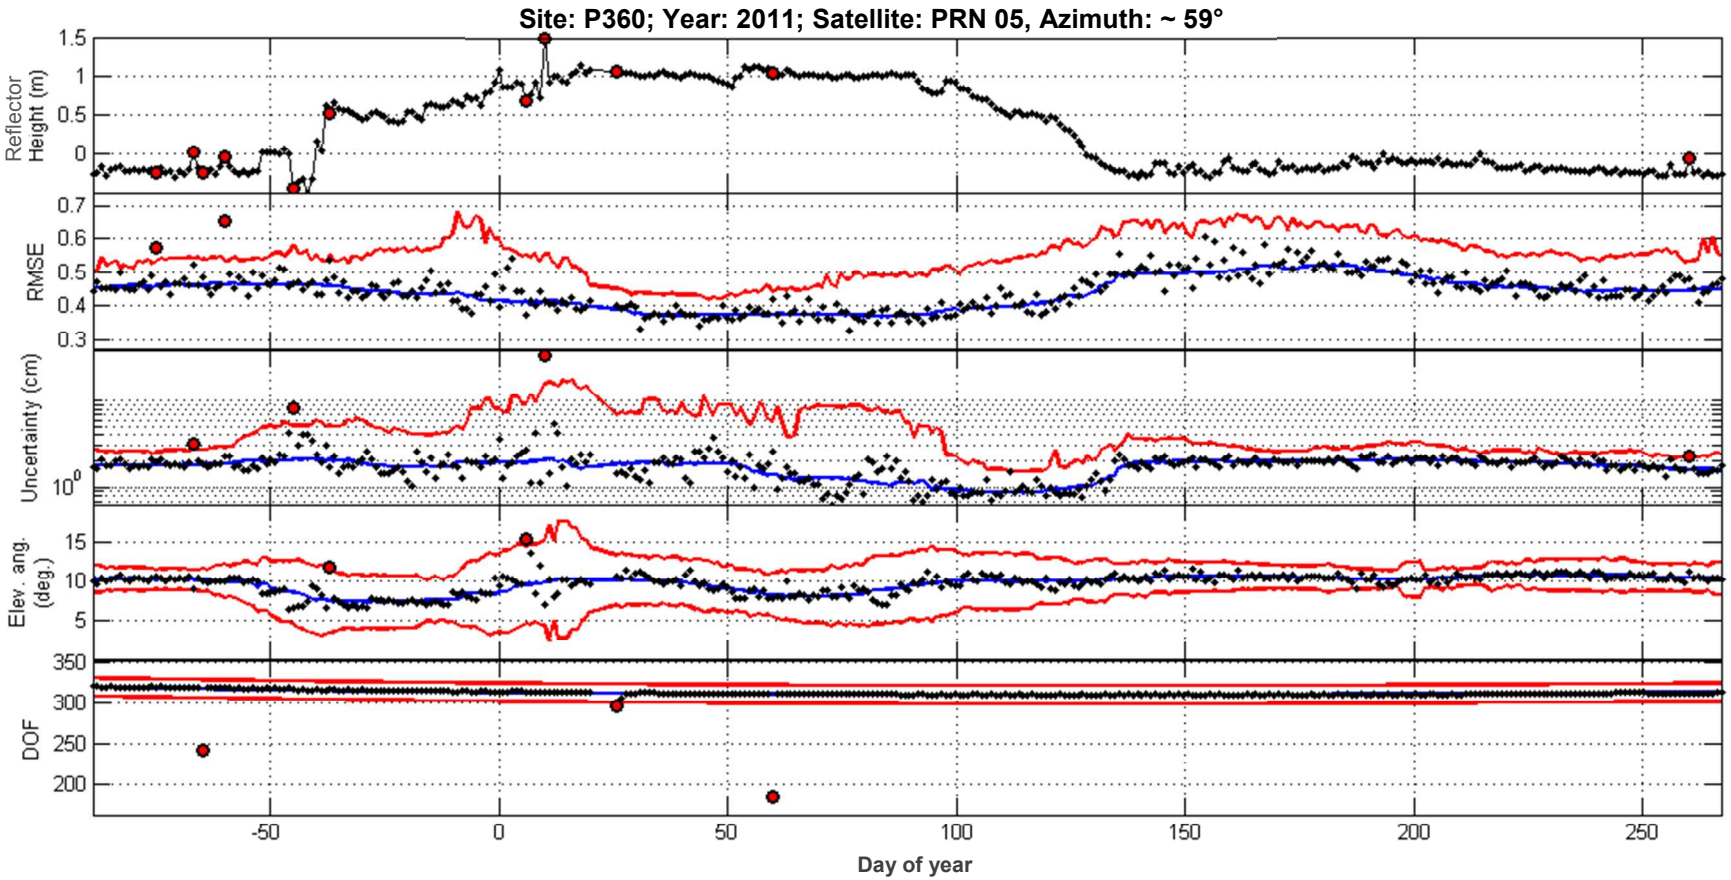

Figure 2: Time series of quality control tests for a single track cluster (satellite PRN 05, azimuth ~ 59°, ascending), as observed at the grasslands site (P360) in the water-year 2011. In each panel, black dots are independent day-to-day track retrieval statistics. Blue and red lines represent, respectively, the expected tendency and dispersion, both based on a 15-day moving average. The bounds are two-sided for normally-distributed variables (DOF and peak elevation angle) and one-sided for χ^2-distributed variables (uncertainty and RMSE). Red circles are retrievals deemed to be outliers, separately for each test in the lower four panels, and in conjunction for all tests, in the top panel. RMSE is the same as $\hat{\sigma}_0$ and is expressed in observation units (decibels); degree of freedom (DOF) is a number count.

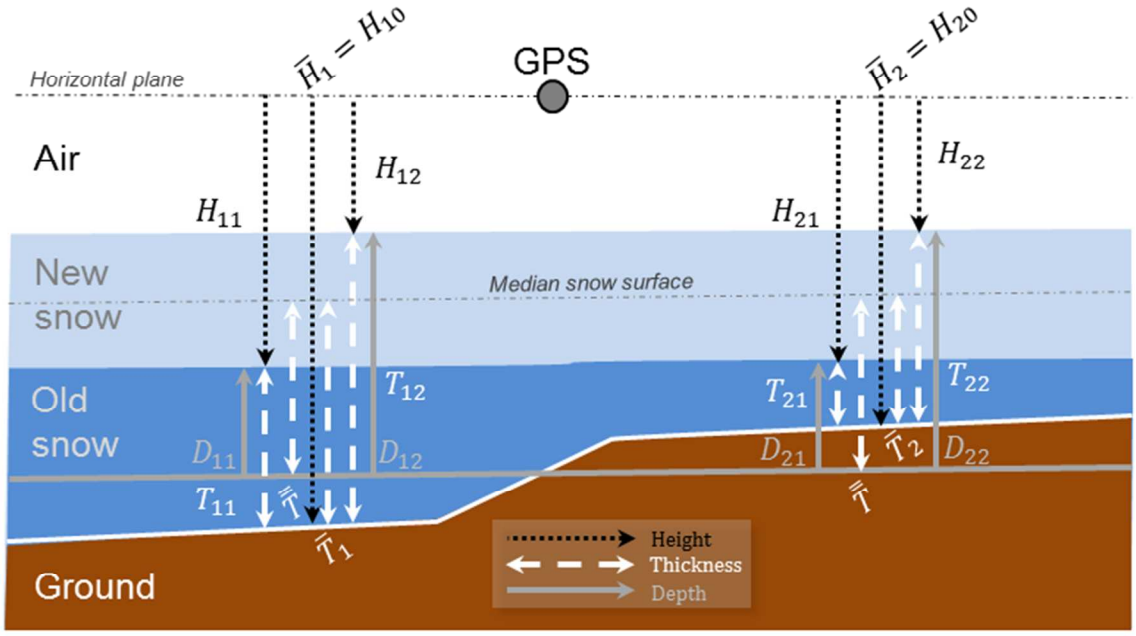

Figure 3: Snow reflector height, snow thickness, and snow depth as different vertical coordinates; for two track clusters (left and right) and three days (denoted 0, 1, 2).

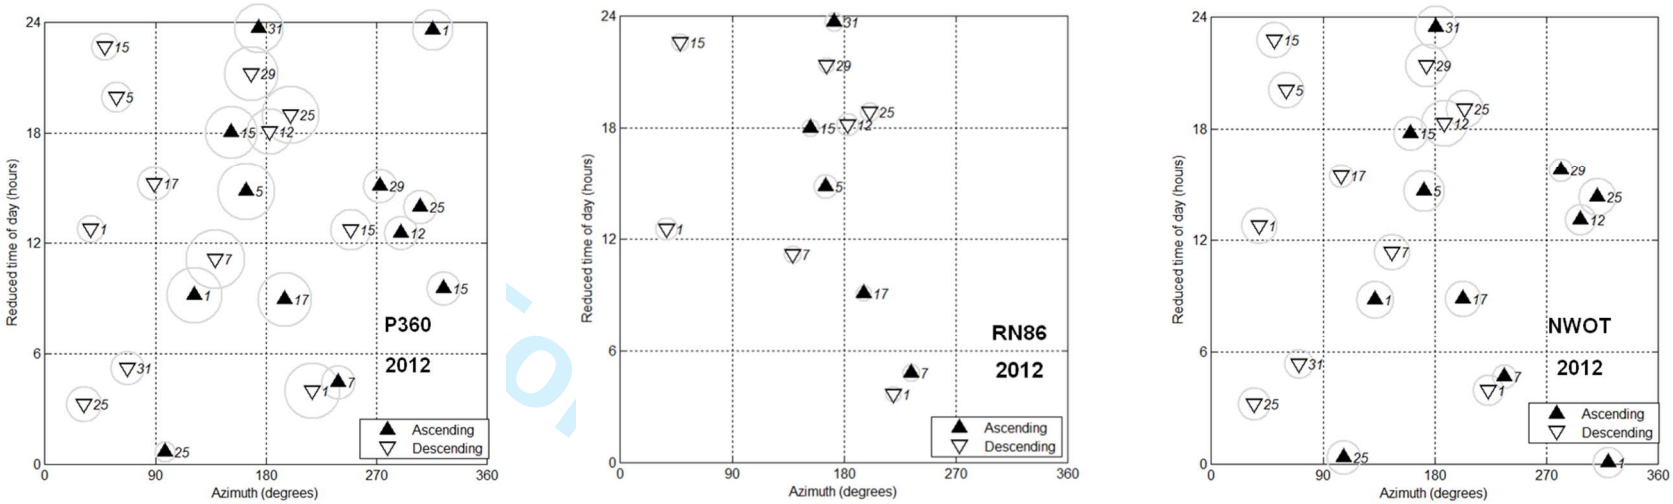

Figure 4: Diagram of satellite track clusters available at least 50% of the water-year 2012 at the grassland site (P360, left), forested side (RN86, middle), and alpine site (NWOT, right); italic labels denote satellite PRN number.

Figure 5: Aerial view of the forested site (RN86) around the GPS antenna (marked with a circle).

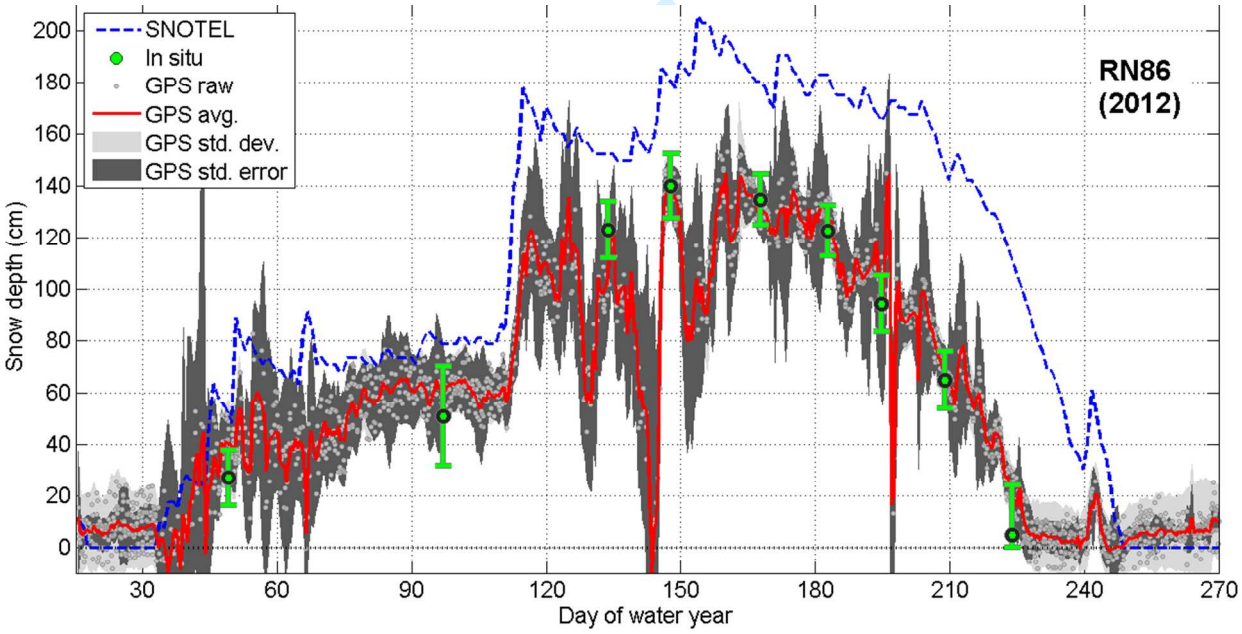

Figure 6: Snow depth measurement at the forested site (RN86) for the water-year 2012; see text for discussion and description of details; SNOTEL is ~ 350 m distant and same altitude.

Figure 7: Ground conditions in the vicinity of the GPS antenna at the grassland site (P360).

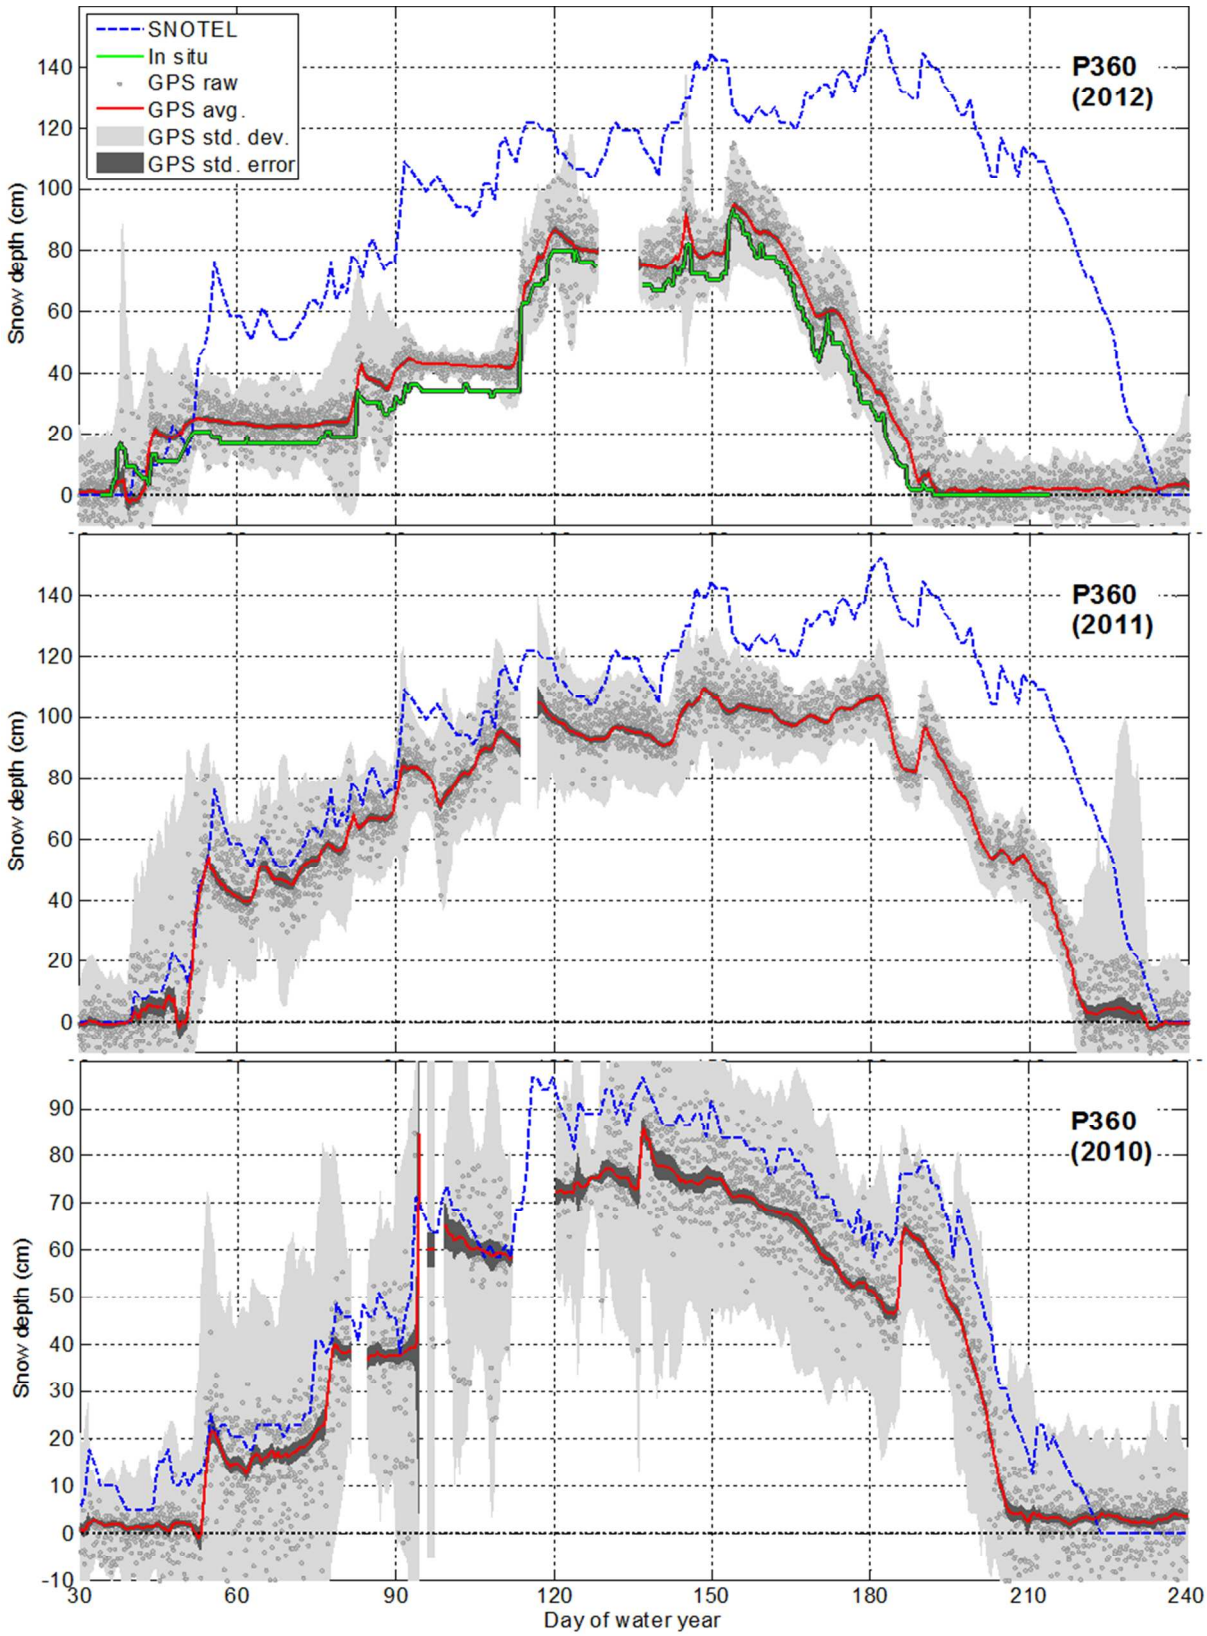

Figure 8: Snow depth measurement at the grassland site (P360) for three water-years; SNOTEL is ~ 12 km ± 4 km distant and 60-m higher altitude.

Figure 9: Ground conditions in the vicinity of the GPS antenna at the alpine site (NWOT).

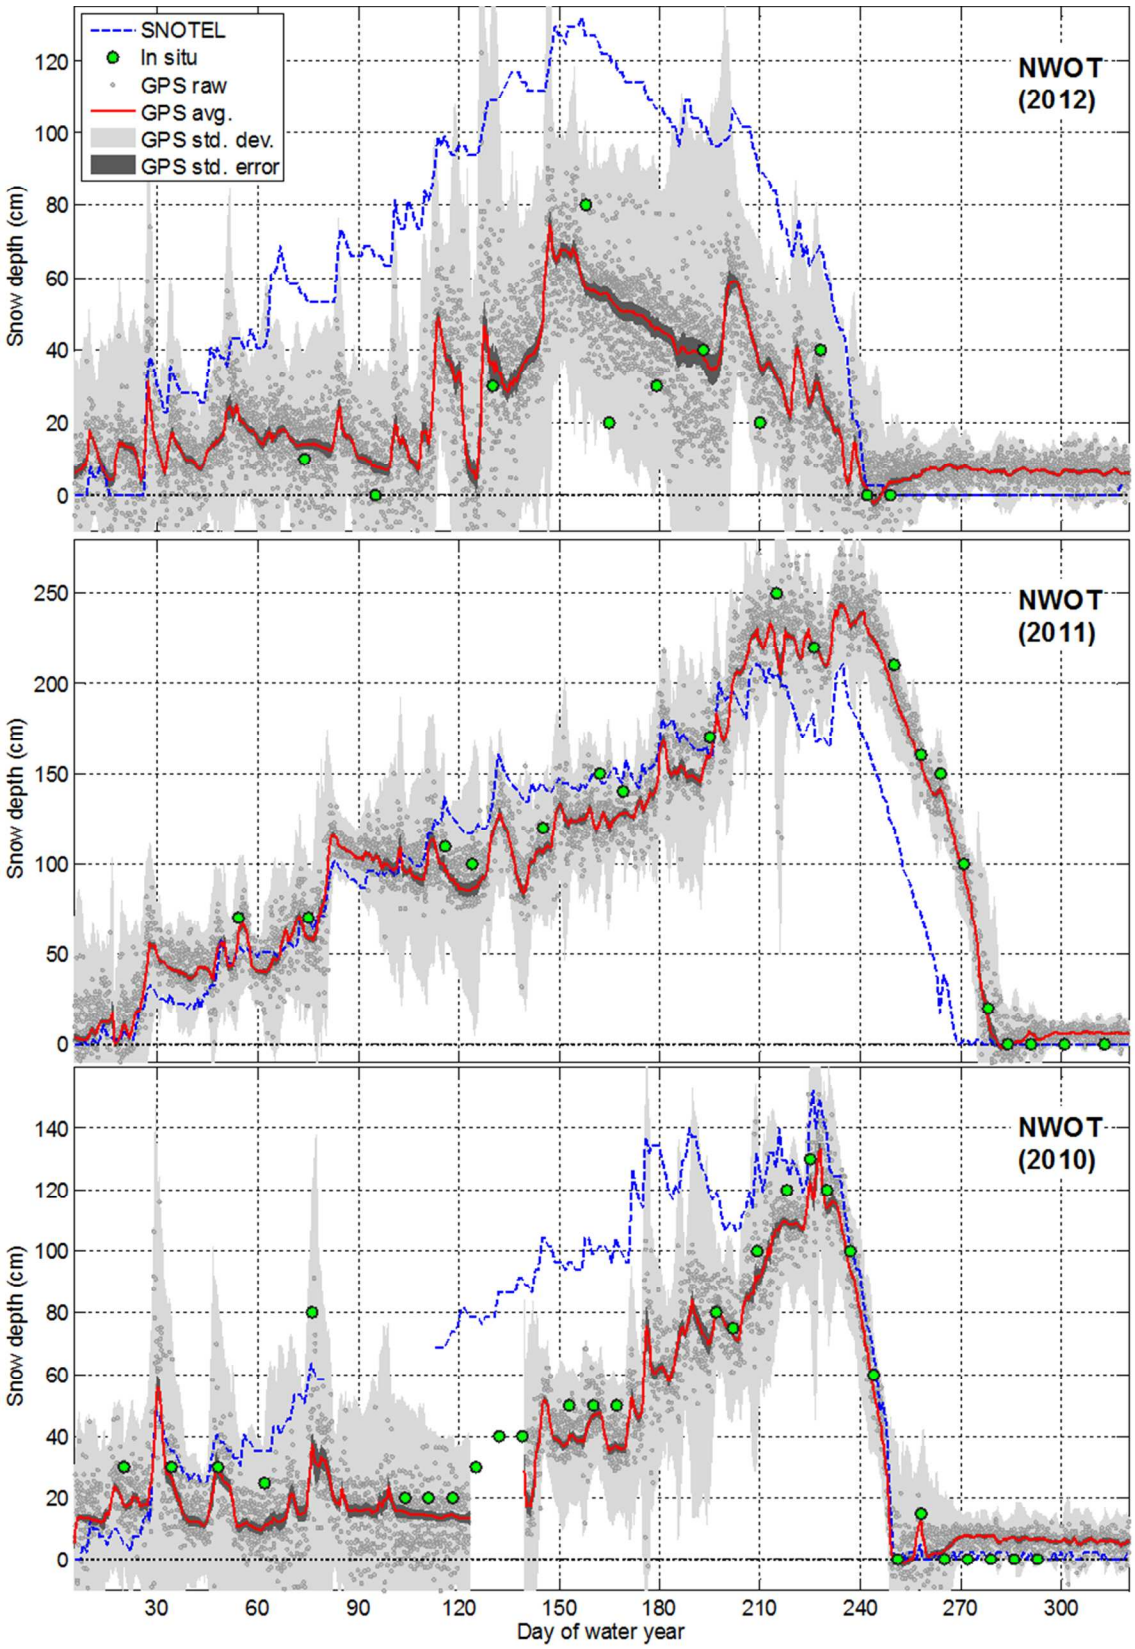

Figure 10: Snow depth measurement at the alpine site (NWOT) for three water-years; SNOTEL is ~ 4 km distant and ~ 380-m lower altitude.

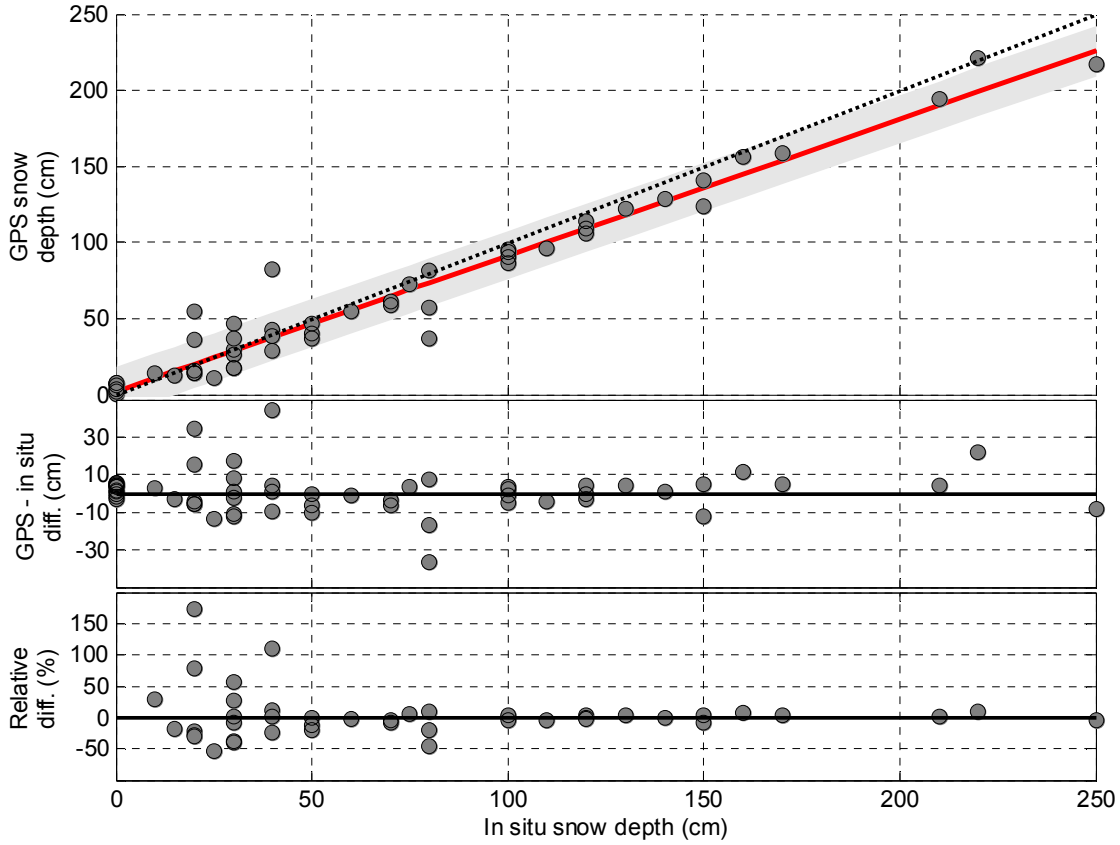

Figure 11: Scatterplot of GPS vs. *in situ* snow depth for all three years at the alpine site (NWOT). A simple linear regression is shown in red, with its 95% observation prediction interval shown as a light-gray band. Post-fit residuals (with respect to the red line) are shown in the middle panel; the bottom panel shows residuals normalized by the *in situ* snow depth value. The ideal 1:1 diagonal is shown as a dotted black line, for comparison.

BIOGRAPHY

Dr. Felipe G. Nievinski is a post-doctoral researcher at UNESP, Brazil, where he works in the field of GPS multipath reflectometry. He earned his PhD (Aerospace Engineering Sciences) from the University of Colorado Boulder in 2013, M.Sc.E. (Geodesy) from the University of New Brunswick in 2009, and B.E. (Geomatics) from UFRGS, Brazil, in 2005.

Dr. Kristine M. Larson is a Professor of Aerospace Engineering Sciences at the University of Colorado. Her current research focuses on GPS reflections.