"MicroPulse DIAL (MPD) – a Diode-Laser-Based Lidar Architecture for Quantitative Atmospheric Profiling" by Spuler et al.

Referee #1 (05 Mar 2021)

General remarks

In this article, Dr. Spuler and his colleagues detail the technical advances made for the new generation of the unique micro-pulse DIAL lidars they have been developing at NCAR and the Montana State University for more than 15 years now. These network-able systems may well represent the future of operational high-resolution atmospheric sounding for water vapor, temperature and aerosols. The 5th generation of instruments is here described in greater detail than in previous articles, and many technical difficulties are addressed and solutions are found that are highly relevant to the operational capabilities of these systems. The new water vapor measurements are validated, with a focus on the lowest altitudes (below 500 m), where accuracy is clearly improved.

Therefore, this work is of great interest for the lidar community, and for the validation of an MPD network as a reliable tool to the greater atmospheric science and meteorology communities. The science and technical aspects presented are perfectly sound and clearly part of a staggering expertise gathered by Dr. Spuler's team on these systems, requiring mastership of cutting-edge aspects both in electro-optical systems and lasers on the one hand, and in atmospheric physics on the other hand.

As a lidar scientist, I can testify to the fact that because of this, it is never straightforward to find the right platform or means of expression to satisfy readers of both backgrounds when writing such a paper. In the case of this article, my thirst for technical details on the lidar is well-satisfied except on a couple aspects that I would ask be developed. But I wonder if the average reader from the atmospheric science background would find the article too dense as it is. Maybe another reviewer from another background will confront this aspect: a lot of very specific technical aspects are addressed with no real common thread leading from one to the next. The importance given to some developments is not proportional to their relevance to the whole. And mostly there seems to be a lack of structure that would allow a reader to skip certain specific parts and get to what interests them. Some sections are unclear and must be made much easier to understand for the non-expert reader. The overall schematics of the system is given too late in the paper, and there is no outline, yet the reader definitely needs a map to understand where they are in this complex system.

The main path I see out of this issue is to rewrite the introduction with a clear statement of the focus of the paper, an outline, and to separate the introduction sub-sections in a Section 2 that will also include the lidar main schematics. Only then the authors may delve into more details with the transmitter and receiver. Also in those sections there could be more structure, maybe with sub-sections. I am indeed sorry to ask this because these recommendations just amount to reshaping the paper into the orthodox, classical, almost catalog-like structure for a lidar validation paper, whereas the current structure was more out-of-the-box, but I must confess the current shape does not work. This, and the several new technical aspects I think were not developed enough or unclear, and would maybe entail a new appendix, requires what would be called "major" changes in the article before I would state it is ready for publication. However this should not require too much work, and I am certain the authors would even be able to submit a suitable version in a matter of days.

We sincerely thank the referee for their time to provide this thorough and valuable review. They have provided helpful and constructive criticism that will greatly improve the readability of the paper. As a result of the feedback, we have restructured the paper, adding a paragraph outlining that structure and introducing the instrument schematic upfront (delving into details later). The changes are not dramatic but it should provide a framework that makes it much easier for the reader to navigate. As part of the restructuring we removed figure 4. That figure provided a zoomed in version of the transmitter, and is no longer needed now that the overall schematic has been introduced earlier. Finally, responding to the referee's concern that some details may be too dense, or not proportional to their relevance to the whole, we moved the description of how the spectral purity tests were done (including Figure 7) to an Appendix.

Specific remarks

I have provided all of my specific remarks in the shape of an annotated PDF manuscript. I also copy the remarks that are not about typos or syntax below, so that they can be more easily addressed by the authors in their reply. Please see the annotated PDF for other corrections.

L12: Please specify at the end of the abstract that the paper will validate the capabilities of the new generation on water vapor measurements only. Maybe give some summarized results.

The following sentences were added. "The work presented focuses on general architecture changes that pertain to both the water vapor and the temperature profiling capabilities of the MPD. However, the specific subcomponent testing and instrument validation presented are for the water vapor measurements only. A fiber-coupled seed laser transmitter optimization is performed and shown to meet all of the requirements for the DIAL technique. Further improvements such as a fiber-coupled near-range receiver, the ability to perform quality control via automatic receiver scanning, advanced multi-channel scalar capabilities, and advanced processing techniques are discussed. These new developments increase narrowband DIAL technology readiness and are shown to allow higher quality water vapor measurements closer to the surface via intercomparisons within the MPD network itself and with radiosondes."

L47: This intro is a bit unorthodox, with no outline and Figures in the introduction; this is not a problem in itself, but at this point I do not see where the paper is going. There is indeed a lot of information and it would help if an outline was specified somewhere in the introduction. I would suggest here, because the introductive subparts are rather long. And maybe separating them as Part 2. Please also specifiy that the focus of the qualification part of the article is WV measurements and that the HSRL/T capabilities are not developed.

The following paragraph was added. “This paper presents the latest details of our diode-laser-based lidar architecture which is being developed to provide quantitative atmospheric profiling. In section 2 we start with a discussion of technical approaches that currently seem promising toward addressing the need for a ground-based thermodynamic profiling network and attempt to highlight where Micropulse DIAL is a unique approach. In section 3 we provide an overview of the MPD Gen 5 lidar architecture. In section 4 we take a step back to provide some background and history of the MPD development, highlighting some field test results which occurred during the transition from Gen 4 to Gen 5 to provide context for why certain design choices have been made. In Section 5 and 6 we delve into the technical details of the Gen 5 transmitter and receiver architecture; respectively. We explain design trades and show laboratory test results validating the design choices. Section 7 provides an overview of the data processing techniques. In section 8 we discuss the validation of the latest instrument design which includes initial field testing and intercomparisons with radiosondes. And finally, we end with a summary and provide conclusions.”

L185: It is not mentioned whether the previous part of the setup is made with PM fibers. But if up until then none of the fibers are polarization maintaining, how is it that the polarization alignment required to the passing axis of the Farady isolator / the slow axis of the TSOA does not fluctuate with temperature / fiber flexing?

The following was added within the new Instrument overview section. “The transmitter uses single-mode (SM) fiber and the correct polarization for subsequent components is maintained using paddle-style controllers that use stress-induced birefringence produced by wrapping the fiber around a spool to create half-wave plates. The fiber is secured to the bench to avoid movement and maintain the alignment.”

L208: Meaning unclear.

In hindsight, this was perhaps an unnecessary clarification. Since it was unclear, and there was a comment about the details being too dense at times, we chose to remove the sentence to improve the readability of the paper.

L216: by attenuation or by modifying the current on the TSOA?

This section was moved to Appendix A. And the sentence was modified to clarify that the sampled beam was attenuated. “The amplified detector has a 2 mW damage threshold, so the sampled pulse was attenuated so the peak power was below that threshold.”

L216: I would not dare criticize the English of a native speaker. But 's possessives are seldom used in academic writing for inanimate objects or concepts. Common exceptions being of course eponyms such as "Student's t-test" or "Parkinson's disease", or the phrase "according to the manufacturer's instructions". Please consider whether this should be modified throughout the article.

This is a fair point (and it is not unusual for a non-native English speaker to have a greater awareness of grammar rules than native speakers). The possessives allowed for more compact sentence construction, but we agree that it is uncommon for inanimate objects in technical writing. We have removed this construction throughout the paper.

L221: I had much trouble understanding this paragraph even though I am familiar with the technology, which make me think this should be made much clearer for other readers.

1) It is unclear what etalon we are talking about here. I do not think the etalon filter at the reception has been mentioned let alone sufficiently described at this stage of the article, nor that several transmission peaks remain within the IF bandpass. But actually the etalon is not what's important here, it is given too much focus. => Consider omitting it and talking about the receiver bandpass that necessarily allows both ON/OFF wavelengths in.

2) It took me three readings to understand that what you were referring to was the fact that due to leakage in the SWITCHING SETUP (not mentioned), there could be a bi-modal seed.

=> Consider mentioning this instead.

3) This part comes between the description of the setup to measure spectral purity and the results of the experiment. It should definitely be put somewhere else, either before or after.

To address these concerns, this paragraph was moved to the end of the spectral purity discussion. And, thanks to the referee's suggestions, hopefully more clear in the new format. It is as follows.
"A critical element of the MPD architecture is using the same amplifier for online and offline seed lasers and the same detector to measure both atmospheric profiles for DIAL measurements. Therefore another aspect of spectral purity, that is important to the MPD design, is measuring leakage in the optical switches used to alternate the online and offline seed lasers. Cross-talk in the switches could result in a bi-modal seed, which would not be noticeable in the test setup described so far, and since the receiver bandpass is designed to pass both online and offline wavelengths it is a possible spectral purity contamination source. So, as described in the receiver section below, the instrument was designed for automatic receiver scans to measure both the receiver bandpass and the isolation between the online and offline wavelengths. To ensure the spectral purity specification of 99.9% is maintained, greater than 30dB isolation between online and offline is monitored and verified"

L233: Consider adding a very short sentence about the impact (or lack thereof) of the pulsed output of the TSOA on the observed spectrum. Basically, it means convoluting the CW spectrum by a frequency comb with 8 kHz FSR and 1.6 MHz envelope, right? So, no impact here.

The referee is correct. But, in the context of spectral purity, and since we measure the actual end result (including any nonlinearities and noise amplification), we chose not to add this comment.

L246: Please expand on why. In case there is a loss of purity/TSOA output power over its lifetime?

Added the parenthetical clarifier "(e.g., in the case of component power degradation over its lifetime)"

Fig10 legend: O2 has only been mentioned once before in the article, and it may not come to the unaware reader why it is addressed and that it is associated with temperature measurements (please mention it here). I think that despite the citation of Stilwell 2020 in the introduction, there should be a summarized explanation of the measurement principles before, in a third part of a new section 2 using also the current 1.1 and 1.2. Maybe the schematics of Figure 11 and the associated text could be put there.

L266: I think this problem could be solved by first describing the instrument as suggested hereabove.

Agreed. Adding the instrument description above helped clarify the O2 in the caption. Additionally we added "which is used to measure temperature" in the caption.

Figure 11: The polarization controllers are not described in the text and may respond to a previous question of mine. Please consider adding a word about them in the adequate section.

This was added in the new section as discussed above

L282: Could be expressed more clearly.

Changed to "For example, as will be discussed in more detail later, a narrow field-of-view receiver is inefficient at capturing light from ranges close to the surface."

L294: between the emitted beam and the receiver field of view (overlap between what? undefined yet)

Changed to "The range where the 'object' (i.e. the footprint of the transmitted beam scattering from the atmosphere), can be imaged onto the receiver without clipping at the field stop is known as the region of `full overlap' in the lidar discipline."

L295: Loose statement. Not as much typical as perfectly theoretical. The authors could be more precise and even give the textbook formula for hyperfocal distance H = D/FOV = 450mm/150µrad = 3km, which helps discuss what follows.

The hyperfocal distance is the simple geometrical optic relationship we were referring to. We attempted to make this more clear as shown below, and added a footnote about the practice of adjusting the receiver focus to reduce the range to full overlap.

"The MPD primary receiver does not achieve a full overlap until approximately 3~km range (see Fig. \ref{fig.Gen5_overlap_funtion}). When ignoring any secondary mirror effects, simple geometrical optics define the range to full overlap -- using the hyperfocal distance H ≈ D/FOV, a distance of 400mm/{115μm} ≈ 3.5 km is expected"

If I may delve deeper into that matter, if one focuses at 3 km, one gets both 1.5 and infinity in focus, and thus full overlap after 1.5 km. Was it achieved in practice?

This is correct and what we do in practice. A footnote stating this and referencing our earlier work where this was discussed, Repasky et al. 2013, was added, "the MPD receiver focus is adjusted to decrease that distance to about 1.5 km."

L307: It unclear here why adding a separate smaller receiver with a larger FOV, compared to another channel with larger FOV on the same receiver, changed this aspect. But this can be explained using the formula above. Using the same receiver with Kfov times larger FOV you get H/Kfov full overlap range for the WFOV channel, but Kfov2 more background noise, meaning the near range to reach a threshold SNR is not so good. Using a separate Kd times smaller receiver (independent of FOV), you get H/Kd full overlap range, but no more background noise relative to the signal. This could be better explained in Annex A. Of course, this formula may be an oversimplification, but then an extra figure plotting Signal & Noise or SNR vs range for NFOV, WFOV, separate WFOV could be used, supporting the discussion about the noise figures of Appendix A. This brings another question: Isn't the smaller receiver equivalent to adding a smaller pupil stop in the WFOV channel? I suppose the concentric emitter/receiver beam configuration prevents that. In any case, I would be interested in seeing a small discussion in Appendix A.

Yes we agree. Repeating the formula in the Appendix makes this relationship easier to see. We added the following text: "As mentioned in the main text, when ignoring any secondary mirror effects, a simple geometrical optical relationship can be used to define the range to full overlap – hyperfocal distance H ≈ D/FOV. This relationship makes it easy to see that increasing the receiver field-of-view reduces the distance to full overlap and thereby improve the efficiency of collecting light closer in range (and thereby closer to the surface). The wide field-of-view also collects more background light, so the WFOV receiver diameter must be balanced. However reducing the receiver diameter reduces the signal photon counts as function of the receiver area (just like the background count).

L308: Unclear. Same sky background photocount, as explained in Appendix A, but not the same for the signal photocount, right?

Correct. We added that clarification "the signal counts, just like the background count, is a function of the receiver area. So the signal will decrease by decreasing in the collection area." and more explanation and a plot in the Appendix.

To address the specific comment on L308, the Gen 4 design used the shared telescope for both the WFOV and NFOV. We added the following clarification "Each dataset had to be processed independently, and the Gen 4 WFOV channel used the same primary receiver telescope and had a significantly higher background rate..."

L325: Thanks for pointing this out. Narrowband filtering (even in Raman lidars) is so terribly dependent on angles of incidence, and this effect of non-common path between Raman or DIAL channels so often overlooked.

This is a change from the previous Gen 4 version where the field stop was after the filter stage. So we have made that clear.

L335: unnecessarily heavy

Changed to: "The combination of filters passes an ultra-narrow band (<50 pm) centered on the absorption lines of interest."

L357: Unclear. Indeed, this works perfectly. But in the Gen5 MPD, there isn't a field stop, right? Why not? Do I understand correctly that you prefer to validate thereafter that this is not a problem?

There is a field stop. In the new section added earlier in the manuscript – providing an overview of the instrument – we now make this point more clearly. The following was added "As discussed previously, the MPD has a field stop limiting the angles into the filter to <3 mrad, so concern raised is not applicable".

L362: Despite the experimental results, this statement is a strong one considering the opposing claims of Spath et al. 2020, and could be substantiated a bit more. I really think a theoretical/simulation approach would help here.

- What is the etalon CWL shift (or peak broadening maybe) associated with 3 mrad? Then what is the resulting change in the observed WV absorption cross-section? What is the WV concentration bias that ensues?

More generally, knowing the FOVs, diameters and focal lengths of the receivers, it is possible to roughly calculate or more precisely simulate how the etalon peaks would shift/broaden at short ranges. At what range do we only get xx% change in effective CWL, yy% change in observed cross-section, and zz% bias in WV concentration?

We understand the referee is suggesting this analysis only if a field stop was not used (which was not clear before, but hopefully that has been corrected). The argumentation in Späth is based on an analysis in Wulfmeyer that assumed the system did not contain a field stop. With the clarification about the stop added, we would argue that development of a theoretical/simulation is unnecessary.

S3.4: This whole section is interesting and very important for operational deployment of these lidars. But it is very confusing as it is. Please consider rewriting it entirely, taking into account my comments below.

L366: heavy.

L369: The spectral/modal behaviour of the TOSA is known to change tremendously between CW & pulsed operation, if the thermal load varies. How is this mitigated/taken into account?

L370: For this scan, one needs to make sure the backscattered light follows the exact same path as light gathered from the beam's propagation in the atmosphere. On the main lidar schematics I see it would be in the focal plane of the telescope. Please state this plainly and refer to the figure as this is quite important.

Also, laser light reflected back into the TSOA could be damaging. Please explain why it is not.

L373: This is very unclear and yet crucial. The central sentence states that the scan is simultaneous, which is, I believe, incorrect. How then could the respective transmission of the online and offline channels be separated? The detector would add both of their contributions.

The only way I would understand the results is if two scans are not done simultaneously but in a sequence, both with the online laser turned off and offline laser emitting. The first (red) scan is done with the switches adressing the online laser. Because it is turned off, we only see the leaks through the switches. The second (black) scan is done with the switches adressing the offline laser. Then we see the full power. So the red curve is basically the black curve attenuated by the switches and with a noise floor.

If that is correct, this section is extremely confusing and should be rewritten from scratch, detailing the steps rigorously. The two channels online/offline turned on and off are inherently tricky to talk about in a clear way; please indulge the reader.

We think the confusion stems from two things. First, we never mentioned the switches are still running during the scan, and second we said "The scan measures the transmission of both the online and offline wavelengths simultaneously". We are not measuring both emitted wavelengths at the same time, however it only seems that way from the perspective of the data collection (MCS) from which it was written. So we have added some sentences to clarify what is going on. The TSOA is turned off during the scan, so there is no concern about back reflections or that it would behave differently in pulsed vs cw mode. The scan is simultaneous. The paragraph has been rewritten to clear up the confusion that what is being measured was crosstalk in the optical switches and not the actual other seed wavelength. The following clarification is made
“Although there is only one laser active, the 1x1 and 2x1 optical switches used to combine online and offline seed lasers are still alternating. Therefore, during this offline scan, any light that ends up being recorded as the ’online’ occurs when the switches are set to pass the online laser, and what is being measured is the offline light leaking through the optical switches. So the red trace in Fig.10 is the measurement of the offline to online optical switch crosstalk”

L376: These two results come too soon, they should come at line 399, and as already done, line 401.

Corrected the figure location and the transition from the receiver scan to the discussion about the fitting

L382: What are the n channels? This could be at least three things: the range channels i.e. the range bins as defined later in the paper, the online/offline channels, the NFOW and WFOV channels. I first thought it was range channels. But then at line 395, I saw g1/g2, so there are two of them and I was wrong. After much pondering, I still do not know.

The n subscript indicates the detection channel designation. For example, if we scan the online laser over its etalon feature, the offline channel will still collect data. We can arbitrarily designate n=1 for the detection channel corresponding to the scan laser and n=2 for the detection channel opposite (e.g. offline in this example). Note that the only difference between the two channels is the efficiency. Ideally g_2 should always be zero, but because the optical switches have some leakage, we must account for this.

We have clarified the text below eq. 2 in accordance with the above explanation
"where $s_n(\lambda)$ is the mean number of photon counts as a function of wavelength λ and the subscript n is the detector channel designation (e.g. if the online laser is scanned, the online channel has designation $n=1$ corresponding to the scanning laser and the offline detection channel has designation $n=2$)"

L398: State here the results in terms of peak effective width, the distribution of w(x) that was found, and the associated standard deviation for x.

We have updated the caption in Figure 10 to summarize the incident angle weights. Standard deviation is probably not the right metric, as the incident angle is not particularly Gaussian-like. However the fitting process selected weights that are bounded less than 3 mrad.

"The incident angles weights are nonzero between 0 and 3 mrad with a peak at 2 mrad."

L403: Confusing, please rephrase. I am confused by "transmitted" as it could refer to the laser transmitter as you named it, and by the "location" of the wavelength, rather than its value; please avoid those terms and 's when rephrasing. Also, detail which wavelength: "(the value of) the wavelengths transmitted by the etalon"

Is that correct? Or did I get it wrong and it actually refers to: "(the values of) the emitted wavelengths"

On a side note, I have trouble with "Transmitter" (instead of "Emitter") in the whole document. I assume the former is more correct in English, but it certainly is confusing when we are talking about the transmitted wavelengths sent into the atmosphere by the laser source, and the transmitted wavelengths through the etalon. In other languages, a transmitter conveys a signal that is already there, does not create it.

Rewritten as “One example of this utility is remotely setting the online and offline wavelengths and confirming the receiver's passbands are properly centered on those wavelengths. As described in the introduction, changing the wavelength is done to alter the MPD's sensitivity to water vapor. “

L406: I see transmit*ted* pulse also below, line 427. Which is best? I would argue the one with the participate.

Changed to transmitted

L411: I believe MCS is undefined until 3.6.

This is correct. The new overview instrument section added now introduces it much earlier in the paper.

L415: corresponding to what range? This seems useful as a lower bound for Rmin, independent of laser pulse duration. Can it be chosen?

It seems the referee is asking about the 250 ns accumulation time. The following text was modified “As a balance of performance and file size, the Gen 5 MPD instrument typically uses a 250ns photon accumulation time and \Delta R = 150 m but these values are adjustable.”

L415: How is this chosen? Can it be reduced?

It seems the referee is asking how the absorption path length Delta R is selected. Using shorter delta R will increase the noise. We say at the end of the paragraph (L 437) that these values can be adjusted and now that the line 417 and 434 are reunited this discussion is hopefully more coherent.

“As a balance of performance and file size, the Gen 5 MPD instrument typically uses a 250ns photon accumulation time and \Delta R = 150 m but these values are adjustable.”

L418: This paragraph lies right between the discussion on pulse length and its conclusion, i.e. the choice made for the Gen 5 instrument. As a result, we do not understand where the line of thought of line 417 is leading. Please consider reuniting line 417 and line 434. The discussion about after-pulses is interesting, but it should be put elsewhere. Maybe just after line 442.

Reunited line 417 and 434.

Fig13: I wish I could see Rmin in Figure 13. For a while I thought there had been a confusion between R3 and Rmin. Also, what is R0, shouldn't it be just 0? And R3, just a random range?

We added an R_min and R_0 was relabled to 0. The main reason for having an R_3 was to easily add the label "Tau x c" for the pulse length. It was challenging to add the label cleanly with only 3 pulses, so we chose to keep the pulse but relabeled it to R_n for generality. Finally, the distance from T_1 to T_2 is now equal to that of R_1 to R_2.

Also I don't understand why the photon accumulation time doesn't just create a wider rectangle.

Because the pulse is moving during the accumulation time they would have to be drawn as trapezoids. That doesn’t seem to add much clarity (at least to us) so we chose to leave it as a separate rectangle.

L438: SNR previously undefined.

Corrected.

L454: Please choose transmit, transmitted or transmitter throughout the document.

The use of transmitter pulse power seems appropriate here.

L457: Still, this requires photodetectors that are unspecified here, and not on Figure 11. What type of detectors? Are they calibrated?

This is addressed in the new section centered around the previous Figure 11. Yes it still requires photodetectors. They are not calibrated, so only a relative power level is given.

L460: This long paragraph detailing the advantages of a custom FPGA system and Ethernet communication does not seem so important here and could be shortened to two sentences.

Noted. Moving away from USB to ethernet communication has been an important step for instrument reliability. So although we could shorten this, and are sympathetic to the referee’s point, we have left this alone.

L471: there or here?

“Here” seem appropriate

L490: It intuitively seems sub-optimal to the uninformed reader to perform smoothing in two steps. If it had to be done at 10 min, 170m resolution, why not do it in the first place (line 473). Maybe add a short sentence here to explain that.

We assume that the referee understands the reason for this, so we have added clarifying language to the manuscript.

The size of the smoothing kernel applied to the photon count profiles, and therefore the amount of random noise that can be suppressed, is restricted by the high frequency backscatter structure from clouds and aerosols. Additional smoothing would impart biases on the backscattered signals by smearing these structures. However after we calculate the water vapor field from those backscatter profiles, the high frequency terms from backscatter structure are canceled while the remaining water vapor field is still noisy. Therefore we apply a second smoothing stage to the water vapor field using a Gaussian kernel where it's standard deviation is set to the desired time and range resolution. In the data presented here, the time kernel standard deviation is 10 minutes and the range kernel standard deviation is 170 m.

L493: What is said here is very true, linear propagation of errors does not work at all and another method is required. This bootstrapping method, which in all honesty I did not know of, could work. However I do not find enough information here to properly understand it and evaluate if it does, given the problem of finding enough actually independent samples, but of the same atmospheric column, in the lidar data. Variability is certain at the 10-min / 170 m scale. So I do not understand which samples or profiles are compared. I could not procure the book Hastie et al 2001, but it seems to deal with random samples in general and not signals. In what I could find online, examples involve stationary random processes, but the atmosphere is all but stationary.

From the referee's comment, it appears that we need to clarify how we obtain independent samples. We employ a process called Poisson thinning (described in detail in Hayman 2020) that allows us to take one Poisson distributed profile and split it into two independent Poisson distributed profiles. This thinning process allows us to ensure that the two separate water vapor profiles are then statistically independent. We calculate their mean-square difference, then we repeat. The bootstrap iterations are not themselves independent of each other. They are resampled over and over from the original profile.

So there are two reasons we are not concerned with a non-stationary process.

1.) Poisson thinning allows us to split profiles without making separate measurements
2.) We only need two independent profiles per bootstrap iteration and the iterations are not independent.

Note that the concern over whether the process is stationary or not is of no more concern than it was with the original photon counting profile.

We have updated the text on bootstrapping to read
"To implement bootstrapping on the MPD data, we split the photon count profiles into two statistically independent sets with the identical mean values using Poisson thinning (for details on the Poisson thinning method see Hayman et al. (2020)). We calculate the water vapor number density, as detailed above, for both of the independent photon count sets. We then estimate the statistical variance as the difference squared of the two statistically independent water vapor profiles. This thinning and processing is repeated so that the final variance is approximated as the average of all the squared-difference results given by..."

Also more generally in this paragraph, explanations are very short. For instance, I only understood what is called bootstrap iterations by looking at other resources. How do we know that 50 resamplings are enough? I see B>= 1000 is a recommended minimum for estimating variance on a mean value.

We always monitor for convergence by plotting the maximum change in variance at each bootstrap iteration. Typically, this value starts to level off at about 20% by 50 iterations.

(This largest change is occurring at extreme values where the difference is not particularly important). The selection of 50 iterations is based on practicality. More is better, but it takes longer to process.

We have updated the text on the number of iterations to read
"We have found that using B = 50 generally provides a reasonable balance between convergence and processing time. We monitor the maximum change in variance (generally corresponding to the noisiest data points) over iterations. This is typically below 20% after 50 iterations. This bootstrapping method is much more effective at capturing the statistical errors of the final water vapor estimates and includes all processing steps that may be prone to statistical uncertainty (smoothing optimization, background subtraction, etc). Because it is a numerical technique, bootstrapping can capture errors that are otherwise difficult to propagate through an analytical calculation of first-order derivatives."

In the end, I am hesitating on what the best solution would for this section. Either explain the method in more detail, which would make it too lengthy and should be put in another Appendix, or find a better reference than Hastie et al., that would apply the method directly to non-stationary signal processing.

The bootstrapping is not complicated to implement, but perhaps difficult to follow when the methods are unfamiliar. Our hope is that our clarifications will alleviate the need to add large amounts of content. The method of bootstrapping is well established and widely used in statistics and signal processing. In addition to the Hasti reference, we have also included a reference to Willem Marais' PhD thesis which also used bootstrapping to estimate uncertainty in photon counting lidar data.

My recommendation is: If this is the first time it is applied to lidar, then it should be developed in a full appendix. If not, please find a suitable reference closer to the field where the reader can inquire about it. In either case, please explain where these independent samples with the same mean value are found despite atmospheric variability, and why the chosen B works best.

Hopefully, we have addressed this. Hayman 2020 gives details on Poisson thinning -- the method by which we obtain independent samples with the same mean value. The number of bootstrap iterations, B, is selected based on monitoring convergence and the practical limit of processing time.

Fig17: I would suggest plotting also the mean error, which is informative about potential biases. Although that information is contained within RMSE, it is convolved with random error.

Agreed, we have included the mean error in the plot

Please also note the supplement to this comment:
https://amt.copernicus.org/preprints/amt-2021-41/amt-2021-41-RC1-supplement.pdf

Thank you. We have corrected all the typos and syntax errors pointed out in the supplementary pdf. These changes can be seen in the marked-up revision

Referee #1 (15 Mar 2021)

After pondering on the overlap results for the new configuration proposed in the Gen 5 MPD by Spuler et al., I have another major comment on Figure A1. Because the resulting overlap functions of the combined channels, from the large reflector (935 cm2) and the small reflector (14 cm2), combined with ratio 10/90 as described in section 3.2 must be :

> Ocombined = Oreflector + Arefractor/Areflector * coupling ratio * Orefractor = Oreflector + 14/950 * 10/90 * Orefractor
>
> Ocombined = Oreflector + 1/601 * Orefractor
>
> This ratio is much smaller than the 1/10 ratio implied by the black curve on Figure A1.
>
> Where is the mistake?

We have made modifications to the Appendix (now Appendix B) to clarify this issue and show how the receiver area affects the signals received. First we removed the combined overlap function from Figure B1. Second, we rewrote equation B2, the relative background, to be more complete and added the relative aperture normalized signal as equation B3. The relative signals for the primary near/WFOV, and combined receivers are then shown as Fig. B2.

Referee #2 (29 Mar 2021)

General comment

The study “MicroPulse DIAL (MPD) – a Diode-Laser-Based Lidar Architecture for Quantitative Atmospheric Profiling” by Scott M. Spuler, Matthew Hayman, Robert A. Stillwell, Joshua Carnes, Todd Bernatsky, and Kevin S. Repasky provides a thorough description of the architecture of the MDP Gen 5 Lidar.

The level of technical detail provided is simply outstanding, the interested reader and the lidar expert can find so much information in this article to satisfy their thirst of lidar engineering thoroughly.

The last part of the study dealing with the validation of the humidity data is somehow less robust and less convincing than the technical part. The comparison with the radiosounding measurements is based on very few cases and does not allow to assess quantitatively the accuracy of the MDP Gen5 data. Despite this not fully satisfactory part of the analysis (which should be improved probably in a future validation study) the scope of this paper is perfectly adapted to the requirements of AMT and the study’s objectives are largely met. This study will contribute positively to the state of the art of DIAL (and in perspective HSRL) lidar technology. I have a short list of minor comments and curiosities that could be easily addressed by the authors.

We sincerely thank this referee for their helpful and insightful suggestions and comments. We are pleased to hear that the level of technical detail was useful. Our hope was to provide enough information for others to build on this concept and make further advances in this field. We agree that the field assessment of the Gen 5 data is preliminary at this point. We have added language to help clarify this point.

Technical comments

Abstract: the abstract results a little too generic, deprived of all quantitative results. The general reasons leading the authors to develop a diode-based to address the observational need are provided. The cost-effectiveness and the network aspect are also provided, but the information regarding the field testing and the radiosonde comparison are missing.

The following sentences were added to the end of the abstract to address this comment and a similar suggestion from referee #1.

"The work presented focuses on general architecture changes that pertain to both the water vapor and the temperature profiling capabilities of the MPD. However, the specific subcomponent testing and instrument validation presented are for the water vapor measurements only. A fiber-coupled seed laser transmitter optimization is performed and shown to meet all of the requirements for the DIAL technique. Further improvements such as a fiber-coupled near-range receiver, the ability to perform quality control via automatic receiver scanning, advanced multi-channel scalar capabilities, and advanced processing techniques are discussed. These new developments increase narrowband DIAL technology readiness and are shown to allow higher quality water vapor measurements closer to the surface via intercomparisons within the MPD network itself and with radiosondes."

Abstract, ln 4: replace "based on a diode-laser-based lidar architecture" with "based on a diode-laser technology".

Done

Introduction: it gives the context in which the authors decide to develop an active remote sensing technology capable for deployment in a network of instruments at the national scale, but the provided structure of the study, nor its scope are provided clearly.

The structure of the study and the scope was clarified (also to address the similar comment from referee #1 above) by adding the following paragraph.

"This paper presents the latest details of our diode-laser-based lidar architecture which is being developed to provide quantitative atmospheric profiling. In section 2 we start with a discussion of technical approaches that currently seem promising toward addressing the need for a ground-based thermodynamic profiling network and attempt to highlight where Micropulse DIAL is a unique approach. In section 3 we provide an overview of the MPD Gen 5 lidar architecture. In section 4 we take a step back to provide some background and history of the MPD development, highlighting some field test results which occurred during the transition from Gen 4 to Gen 5 to provide context for why certain design choices have been made. In Section 5 and 6 we delve into the technical details of the Gen 5 transmitter and receiver architecture; respectively. We explain design trades and show laboratory test results validating the design choices. Section 7 provides an overview of the data processing techniques. In section 8 we discuss the validation of the latest instrument design which includes initial field testing and intercomparisons with radiosondes. And finally, we end with a summary and provide conclusions."

Sect.1.1, Pg 3, ln 63-65: while it is technically true that by calibration against radiosounding, a possible systematic error in the reference becomes almost undetected, the radiosounding suffer typically an uncertainty of 0.1 K and 1%-RH in the troposphere, which is almost negligible when compared to the typical uncertainty of a DIAL or HSRL.

In the manuscript we are not commenting on the magnitude of calibration errors or uncertainty. However, it is not negligible. The reviewer's assessment of sonde accuracy omits representation error between the volume observed by the lidar and the volume observed by the sonde. Our claim in the manuscript is (1) the operational dependency of Raman lidar on sondes increases costs and complexity (both for personnel accessing remote sites and with extra infrastructure required), (2), comparison of one instrument to the other is unable to uncover any bias (regardless of magnitude) if they share common calibration.

Sect.2, Pg 6, ln. 140-142: could you provide a percentage value generally speaking for the needed power of single-frequency laser diodes compared to narrow-band solid-state lasers?.

The "needed power" is a matter of temporal and spatial resolution. We have added the following sentence " The ≈100x reduction in power requires longer averaging times to achieve adequate SNR levels, and therefore lidar instruments based on laser diodes are (currently) only suitable to observe atmospheric phenomena at longer timescales."

Sect.2.3, Pg 10, ln. 218: which etalon? Not sure what are we talking about here.

This is the etalon in the receiver used as the first stage filter. We have modified the sentence to be more clear. "The optical filters in the MPD receiver limit the bandpass to approximately <1.3GHz FWHM, so this measured frequency spectrum is sufficient to ensure the instrument's spectral purity." Note, this section was moved to an appendix in response to the concern from referee #1 that too much technical detail was contained in this section.

Sect. 3.3, pg 17, ln 329-335: the double filtering provides indeed an excellent narrow band, 50 pm is quite an achievement. The first stage is housed in a temperature-controlled mount (what is the precision of the controlled temperature), what about the second stage? why there is no temperature control on the Alluxa interference filter?

The first stage filter (the etalon) is temperature-controlled to <0.02ºC. This value has been added to the text. The thin film interference filters (fabricated by Alluxa) do not require temperature control beyond the environmental control of the field enclosure (+/- 5ºC).

Sect. 4 pg 22, ln 468-476: before removing the background why is not performed a dead-time correction?

Deadtime correction is performed, but we did not mention it. Deadtime is corrected using the standard non-paralyzable correction formula which assumes that range bins are statistically independent (deadtime << bin width) and the photon arrival rate is constant over the accumulation time. In order to minimize the second assumption, deadtime correction is applied at the base MCS histogram resolution (2 seconds), though this is still likely inadequate for the accurately capturing cloud structure.

We have updated the text to include this...

Photon counts are binned onboard the MCS at approximately 2 second resolution. Prior to any processing, a non-paralyzable deadtime correction is applied to correct some of the resulting nonlinearity in the captured signal (Müller1973). This standard correction assumes that photon counts between range bins are uncorrelated (deadtime is much less than range bin width -- in this case approximately a factor of 10) and that the photon arrival rate is constant over the bin accumulation interval. it is for the second assumption, that the correction is applied at the base acquisition resolution, though it is unlikely this is sufficient to accurately recover photon arrival rates from clouds (where water vapor signals are not retrieved).

Is the detection unit of the Gen 5 MPD following the non-paralyzable hypothesis? In any case (data corrected or not) it would be useful to add this information.

Yes we are using a non-paralyzable detector model. We are employing APDs with approximately 30 ns deadtimes (depending on model). Deadtime correction is applied based on the prior stated assumptions.

Sect. 5.1: the overall comparison between the Gen4 and Gen5 datasets is hardly meaningful. The Gen4 comparison dataset is made of up to 60 RS41-MDP Gen4 cases, while the RS41-MP Gen5 counts a maximum of 10 cases. Moreover, the 2019 comparison spanned across Spring and Summer, while the 2020 comparison is limited to the Winter. I agree on the concept, ”better than nothing”, but this validation is not at the level of the rest of the manuscript. It is merely a qualitative information about the better performances of MDP Gen5 compared to Gen4.

We agree that the 2020 radiosonde comparison provides only a preliminary and partial validation of the design. Yet, we believe it adds credibility, as part of the larger cumulative case given, for validating the new design. We have noted that these are initial results. A more complete study is planned for the summer of 2021.

I do not see the real advantage in putting Fig. B1 in the Appendix B, this figure brings at least a visual comparison of each profile. I would suggest to include it directly in section 5.1 after Fig. 17.

This figure was moved into the main body as suggested