SUMMARY:

This manuscript (1) reports on hole-punch cloud observations that occurred during CLOUDLAB experiments, and (2) conducts an LES sensitivity study to see how the timing and extent of simulated hole punch clouds change with two different model parameters: the sub-grid turbulence parameterization, and the initial LWP of the cloud at the location where seeding is performed.

In all simulation runs, a 100% reduction in LWC is achieved in over 5% of the grid cells (unclear whether this is throughout the entire simulation domain or just at the seeding height), providing a great platform to study the limit of complete LWC depletion. When looking at the largest decrease in LWC with time, the investigators observe that a minimum is reached after a longer period of time in runs with larger initial LWP. The investigators also note that the area of the hole is larger with increasing initial LWP. The authors explain that higher LWP provides ample consumable cloud droplets for the hole to continue growing, compared to low LWP where the available water droplets are consumed, shutting off the continued growth of ice particles. The authors also note that varying the Smagorinsky constant produced no noticeable effect on the cloud microphysics.

The study offers a conceptual picture of the processes governing the hole punch clouds in the simulation. It does not use the observational data in the analysis. The results are thought provoking and will be of interest. Generally, the paper struck me as being very short, and there are many questions that arise as I read through this brief article. I believe that offering more complete details to the reader will fill out the paper, increase readability, and offer a more thorough analysis of the physical concepts under study. Beyond minor clarifications of methodology, there are two important scientific clarifications which I feel must be addressed, which are the first two comments below:

Relationship between LWC and LWP

The first clarification that should not be overlooked is the relationship between LWP and LWC, and how the differences between LWC and LWP impact the data interpretation. They are used nearly interchangeably in the paper, yet they are not the same, and are not necessarily required to scale with each other. (deep cloud with low LWC can have the same LWP as shallow cloud with high LWC)

According to Table 1, it seems that all simulations are conducted on the same cloud system just selected at different times of day during its evolution. Does the cloud depth change with time in your simulation, or does the LWC change with time? Or is it some combination of both? How would that impact how these results are to be interpreted?

For example, if LWC remains constant, but cloud depth is changing, then it could be more appropriate to interpret results as being a function of cloud depth.

Ln 133: Data is grouped and interpreted as “Low background LWC”, but simulations are classified by LWP, not LWC.

With LWC and cloud depth being the most important variables for LWP, it would be helpful and interesting context for the reader to see initial LWC and cloud depth values for each of these simulation cases. So much information about the cloud is hidden by not presenting these values. It could be included in the table, or as a time-series similar to Figure 3. It could optionally go into supplementary materials if the authors feel it is not critical to the main text.

Physical interpretation of the sub-grid mixing length scale.

I'll start with the caveat that I am not an expert in LES. This paper review, however, did send me down a rabbithole of sub-grid scale modeling. Please give this reviewer grace upon emerging with the following comments - I apologize in advanced for being long-winded.

I'll first mention that the rabbit-role journey began because the manuscript does not sufficiently describe the Smagorisky constant, and does not provide sufficient references. This reader feels that in order to motivate why the authors chose to vary the smagorinsky constant, a clearer explanation of its physical significance to the simulation is required. Especially so, since this turbulence parameter is central to the abstract and conclusions.

The phrasing of Line 46 "more or less mixing", and Ln 88 “Low values indicate more turbulent flows, while higher values yield more mixing”, is unclear to me and I can’t discern the intended meaning. Two references are provided: Lilly 1962, wherein there is no mention of the “smagorinsky constant/coefficient”; and Dipankar 2015, which only mentions the smagorinksy constant once, without defining it.

Before going further, I’ll mention that Dipankar as well as other references below use C_s (capital C) instead of lowercase c_s, so the authors may consider adjusting their notation to be consistent with the more predominant usage.

The lower value of 0.17 is attributed to Lilly 1962, but I was unable to find this minimum value in the paper. During my rabbit hole adventure it seems that most other literature references Lilly 1967 for this 0.16/0.17 value rather than 1962, but I was unfortunately not able to locate a version of the 1967 report to read it myself. It also may be possible that Lilly 1967 used lowercase c_s, but it seems that C_s has become more standard throughout literature.

Other literature provided me with some introductory understanding of the smagorinsky constant. Very helpful was Mason and Brown 1999. I suggest that this paper (and perhaps others as well) be referenced for the readers benefit, but more importantly the authors may find the discussion in this paper relevant towards interpreting why changing the smagorinsky constant did not impact their results:

Mason, P. J., and A. R. Brown, 1999: On Subgrid Models and Filter Operations in Large Eddy Simulations. *J. Atmos. Sci.*, **56**, 2101–2114, https://doi-org.cuucar.idm.oclc.org/10.1175/1520-0469(1999)056<2101:OSMAFO>2.0.CO;2.

My very entry-level understanding upon emerging from the SGS rabbithole is that there exists the filter length scale which is related to the smallest resolved length scale in the LES. Then in the subgrid, there is kinetic energy removed from the system by friction through the eddy-cascade process. The sub-grid model is needed to quantify this energy removal, with the smagorinsky constant playing a role in parameterizing this energy transfer.

The smagorinsky constant comes across in literature as a loosely-constrained empirical constant, and I have a difficult time relating it directly to a physical property of the simulated cloud. Sometimes it seems like the constant is treated as something that needs to be set appropriately in order to avoid non-physical results and not “break” the simulation. In the range of coefficients explored in this study, it’s possible the authors remain within a realm of validity where large impacts on the resolved scale are not to be expected.

However, to the extent that energy dissipation on the subgrid scale is expected to change the resolved-scale dynamics, I wonder if the investigators can confirm/quantify any changes to the turbulence in the resolved scale due to their modifications of the smagorinsky constant? Did or did not the changes made to the constant affect the resolved-scale turbulence in the cloud? If not, then it would indicate that changing this constant did not truly allow you to compare different cloud turbulence conditions.

Ln 117-119: By referencing other literature about sub-grid scale turbulence, I think these statements can likely move outside the realm of hypothesis and speculation. Depending on how one interprets the prior literature on this topic, one could go so far as to say that the results seen here are an expected outcome, and that the resolved turbulence is most definitely of higher importance.

I acknowledge that the authors will have greater expertise in this area than I, and do not wish to speak with absolute authority, but I do encourage the authors to investigate this a bit more thoroughly.

Questions about the microphysics parameterization

The results of this study are based solely on the simulation, so I think more emphasis can be placed on providing relevant background information about the model, particularly on the microphysics/ice nucleation parameterization. I understand that references are provided, but relevant info can be summarized, similar to the overview which was provided for CLOUDLAB, so that the reader is not required to retrieve and read each reference.

Ln 85: Can the authors provide a brief summary of the immersion freezing parameterization and its key dependencies?
Ln 88: Is ICNC in the simulation given the same >25um requirement as the observed ICNC?
Ln 107: What size exactly?

WBF mechanism: I am curious about the ratio of droplets converted to ice via immersion freezing compared to those lost to evaporation via WBF. Or do only the seeding particles experience immersion freezing, such that no pre-existing LWC is lost to immersion freezing?

Clarifications of methodology and interpretation

Abstract: “delay in the strongest reductions of liquid water content”.
Ln 135: “Hence, a delay in the appearance of strong LWC reductions is notable”.

I wonder if this “delay” can be interpreted as a longer duration/lifetime of the hole-punch cloud, as the “minimum” corresponds to a transition point when the LWC is no longer decreasing and starts to fill-in with liquid water again (i.e. the closing of the hole)?

Ln 130-131: “As soon as maximum reductions of −100 % LWC are achieved, no more liquid water is present to be consumed, such that the extent of the hole is limited.”

> This statement lacks nuance. Can the authors expand more on why this occurs, and frame it within the context of their ice nucleation parameterization?
>
> It’s not immediately clear to me why this would limit the extent of the hole. If seeded particles are dispersed and advected across an identical area in all cases, then why does achieving -100% LWC in one region of the cloud stop ice from growing in other regions where seeding particles are present? Based on the immersion freezing parameterization detailed in Miller et al 2025, I am curious about the interplay between the ambient saturation ratio, the size evolution of the seeded solution droplets and their water activity, and their subsequent ability to nucleate into ice crystals. Is it possible that ice nucleation itself is limited to a smaller lat/lon area in the lower LWP cases?
>
> Is the area of the hole limited in the along-wind direction? Or perpendicular to the wind (cross-wind) direction?
>
> One interpretation is that, since seeding is ongoing for 6 minutes, and is dispersed into an existing cloud, then every “wave” of seeding particles is met with an initial LWC from which to produce a concentration of ice crystals and a “hole” that is advected downstream. In this sense, it seems that the along-wind length of the hole will be the same in all cases (would depend primarily on the wind speed). In terms of the cross-wind width of the hole, this seems like it will depend on the horizontal mixing timescale relative to how long the seeded particles are “active” in terms of their ice nucleating abilities. Perhaps in low LWC conditions, by the time the seeded solution particles mix horizontally, they have already evaporated to a point where they can no longer nucleate new ice. Maybe in the higher LWC conditions, seeded particles can remain efficient at ice nucleation for longer, and this longer timescale relative to the horizontal mixing timescale allows ice particles to nucleate along a greater cross-wind width of cloud, before advecting downwind.
>
> Anyways, a lot is left open to speculation on the part of the reader, so it would be great if the authors can detail some of these nuances in their analysis.

Figure 2(a-c):

- In the caption, mention that a mask (ICNC>0.001cm-3) is used.
- The contour labels are not legible.

Figures 2(g-i) and Figure 3:

- At each timestep, are values of LWC in all grid cells of the simulation used to compute the 5th percentile? Or is it only calculated using the grid cells at the seeding level height?
- Or are LWC values calculated only for the masked (ICNC>0.001cm-3) grid cells?

Ln 123: For the optical thickness, at each timestep, am I correct in understanding that the optical thickness is calculated across the full vertical column of the simulation? Is r_e the effective cloud droplet radius?

Ln 127: How is the hole "area" calculated? Would the lat/lon extent of the hole be influenced by the wind speed and direction? Are the windspeed and direction identical for all simulated cases?

Ln 136-139: Can your hypothesis not be confirmed by looking at the simulation? Did indeed the plume encounter first lower and then higher LWC?

The use of percentiles can be a little tricky to interpret. The authors could improve readability by stating a physical interpretation of the percentiles: "5% of the grid cells had LWC reductions greater than the value shown", "In Figure 3(b), once the value reaches -100%, it indicates that at over 5% of the grid cells had a complete depletion of LWC."

Figure 2(g) 95th percentile ICNC → At first glance it seems like ice crystal number concentrations are constantly decreasing, despite active seeding occurring... but the interpretation of this 95th percentile metric requires some thought.
"Of the grid cells containing ice, 95% of them have ice concentrations larger than this value": If the value is decreasing, it means that more grid cells with lower concentrations are occurring. This can be seen in the pink contours, where the dark 0.2 line at first encompasses a high ratio of the masked cells, later on the 0.2 contour area shrinks and the others grow, such that 0.2 only is a small ratio of the masked cells. This could indicate dispersal of ICNC… where there are fewer regions of extremely high concentrations. This is overall tricky to interpret without a fair amount of pondering on the part of the reader. Did the authors have a specific intent for using the 95th percentile here?

Rather than 95th percentile, the total ICNC could be nice to see instead. If the total ICNC increases with time it means continued nucleation/production of ice, whereas if it decreases it would indicate precipitation or aggregation.

Figure 3(b): Once the relative reduction in LWC reaches -100%, it seems to remain fixed there, but in panel (a) the absolute reduction in LWC starts to go back up. For example, looking at L60, the absolute LWC reduction in panel (a) increases to zero at ~19 minutes (implying zero reduction relative to the initial LWC), but the relative LWC reduction remains flat at -100%. Can the authors explain how this happens?

Ln 138-139. The results on cloud optical thickness: Is this a true statement? In some cases it appears that LWC reductions start going back up to zero in panel (a), while the relative changes in optical thickness don’t change accordingly.

I am curious about the response of the cloud after seeding ceases. What changes in the cloud evolution are expected once seeding ceases after 6 minutes? In the figures, no noticeable changes happen at the 6-minute mark, does this align with what the authors expect?

Questions about the observations

I understand that the purpose of this paper is not to analyze the observational data, nor to perform a direct model-obs comparison. However, I do have a few general questions about the observations which could be clarified, as they are relevant to the study of LWP dependence:

For the observations, I'm curious if the LWP corresponds to a remote sensing measurement or is it derived from the in-situ LWC? Ln 70-72: How are the estimates of observed LWP and LWC obtained?

The magnitude of ICNC seems to be comparable between the observation and the simulated L25 case shown in Figure 2. However, the observed cloud shows decreases in LWC on the order of ~0.1 g/m3, while the simulation, despite starting with initial LWP of 25 g/m2 or higher, never exceeds more than 0.05 g/m3 of LWC reduction. Can the authors comment on this?

The simulated L25 cloud has comparable LWP to the observed case, but does it have smaller LWC? How, then, does the cloud depth compare between the simulation and observed case?

Other small clarifications

Ln 53: "This type of clouds" →"This type of cloud"
Ln 64-65: Did all three happen on the same day?

This is a very nit-picky language thing, but the use of negative values with the term "reduction" is not exactly precise, though the reader can still understand what you mean. The word "reduction" has a negative sign implied within it, thus adding another negative sign negates it, resulting in a positive. For example, "-60% addition" = "60% reduction", and likewise "-60% reduction" = "60% addition".

Some examples where this comes up (non-exhaustive):
Ln 111-112: "minimum of mean reduction" -> The words being said are "minimum reduction", however you're actually referring to the point of maximum reduction in LWC. The trouble arises from the fact that your "reduction" values are negative.
Ln 134: "minima of LWC reductions"
Ln 155: "-60% reduction"

Possible solutions:

- Instead of labeling the figures as "LWC reduction", the authors can use the label "LWC change". The term "change" does not have an implied sign as the term "reduction" does. Then, the authors are free to say "60% reduction" and "maximum of LWC reductions" in the text, which makes more sense when reading.
- The authors could also simply leave it as it is, just knowing that a bit of confusion can occur. It's not difficult for the reader to understand from context what the authors mean to say. However, without looking at the figure axes, the phrase "minimum reductions" in the text could definitely be misunderstood, and it would be challenging for others to quote this out of context.

Ln 127: Capitalize after the colon.
Ln 128: I suggest "there are two processes" instead of "we have two processes".