We thank the reviewers for their careful reading of the paper and their insightful feedback. Please find below answers to the questions that were raised. For clarity, we sometimes use blue text to quote from the reviews.

Review #1.

R: The result expressions for term t and the running time all ignore the dependency for α, it is better to give an illustration for about the dependency on α, to show the differences to previous works more clearly.

We quickly remark that the dependencies on α both in the exponent of the running time and the additive error are all polynomial. We will write out the dependencies more clearly in the revision. Thank you for the suggestion.

R: Authors show that the computational complexity of these two problems [Densest Ball and 1-Cluster] are essentially the same by the binary search. I've some questions about this binary search. Authors don't show the data universe and the precision of the binary search, then what is the time for this binary search? Intuitively, it will depend on them. Authors are expected to explain that before saying they are the same.

There are two directions in the equivalence. When reducing from 1-Cluster to Densest Ball, we binary-search on the target radius. In this case, the number of iterations needed for the binary search depends logarithmically on the ratio between the maximum possible distance between two input points and the minimum possible distance between two distinct input points. This is explained in more detail in Appendix F of the Supplementary Material.

Conversely, when reducing from Densest Ball to 1-Cluster, we binary-search on the number of points inside the optimal ball. Here the number of iterations will be logarithmic in the number of input points.

We will add a remark regarding these in the main body of the revision. Thank you for pointing this out.

Review #2.

R: The paper is a solid theoretical contribution and I vote for acceptance. I would have loved to see an experimental analysis – so I am curious, why no experimental analysis? especially for the k-means, there are so many standard implementations, why not pick one and apply the generic framework of the second result?

Thank you for your suggestion. The main issue facing an experimental evaluation of our algorithms is that they would rely on constructions of lattices and solvers for the Closest Vector Problem on lattices. Unfortunately, these are currently not efficient at scales that would be interesting from a practical clustering point of view. Nevertheless, obtaining a differentially private clustering algorithm that is also practical is one of the directions that we are currently pursuing.

Review #3.

Thank you very much for the careful reading and helpful comments. We will incorporate the low-level/stylistic suggestions.

R: Line 420 (above): Shouldn't this be a less than or equal, rather than an equal?

Yes, the expression before line 420 in the Supplementary Material should indeed be $\leq$ instead of $=$.

R: Eq. (23): in the last derivation, I am not sure how eq. (21) has already been applied

The last line in the derivation of (23) should indeed be removed.

Review #6.

R: Additive approximation terms, perhaps unavoidable.

Yes, the additive approximation term is indeed inevitable under differential privacy requirements. We will emphasize this in the text.