Deep Patch Visual SLAM Supplemental Material

Anonymous ECCV 2024 Submission

Paper ID #272

A Experiment on Global Bundle Adjustment Efficiency

In Fig. A, we show the cost of the large-scale bundle adjustment on the KITTI dataset as a function of the number of pose variables. We compare the dense implementation from DPVO with our block-sparse implementation. As expected, the block-sparse implementation is much more efficient for moderate-to-large patch graphs. For very small graphs (i.e., in the VO frontend), the dense implementation is better since the patch graph is relative dense and it lacks the additional indexing overhead from the block-sparse version. The number of depth variables and edges remains constant in these experiments.

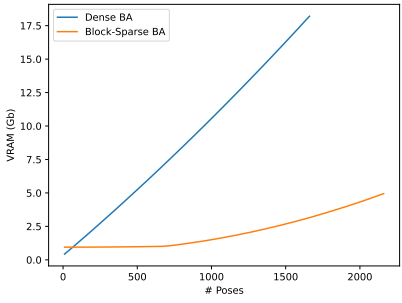

(a) BA VRAM on KITTI.

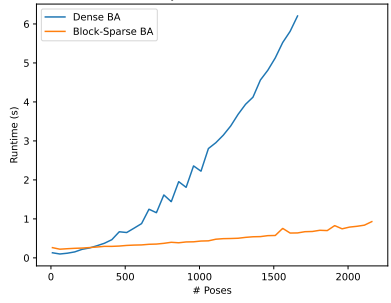

(b) BA runtime on KITTI.

Fig. A: The cost of running bundle adjustment on patch graphs, comparing our block-sparse implementation to the dense implementation. For large patch graphs, our block-sparse implementation is significantly faster and cheaper. In contrast, the dense implementation is faster/cheaper only for very small patch graphs, so it is used only for the visual odometry frontend.