# A  RECOMPUTING NODES WITHIN A STAGE

In Section 4.2, we allowed each operation in a computation graph to be evaluated in multiple stages, though a given operation can only be evaluated once per stage. Results are cached if needed later within a stage. However, the schedule could be partitioned into any number of stages to allow more evaluations of a given node, which may be useful for highly connected graphs. For example, one frontier-advancing stage in our formulation could be split into $k$ stages to allow for up to $k$ reevaluations. Note that prior work (Chen et al., 2016b) actually checkpoints a subset of results for the entire schedule after they are computed for the first time — these checkpoints can only be evaluated and cached once, while our ILP formulation allows for up to $n$ evaluations even without splitting stages.