

Multi-Stage Checkpointing

ANDREAS KOWARZ <kowarz@math.tu-dresden.de>

ANDREA WALTHER <awalther@math.tu-dresden.de>

The computation of adjoints for large applications based on the Reverse Mode of Automatic Differentiation either require the storage of the full forward trajectory in large hard disk files or require complex recomputation of this data in time. As a compromise of these two extremes the checkpointing approach was developed and features recomputation from previously stored and properly selected intermediate states. All strategies designed so far are based on the assumption that a constant effort is needed for reading and writing the so-called checkpoints. However, in some cases even the set of all checkpoints does not fit into the main memory. Then, the underlying assumption of fixed costs for accessing checkpoints becomes invalid.

As an extension of the known checkpointing approaches we will present first results on Multi-Stage Checkpointing. The later is designed for the above-addressed tasks as well as applications running in environments of hidden memory access levels (e.g. distributed shared memory) with non constant costs indicated. Starting from the known and proven statements for constant cost checkpoint access we will present a cost function for this case and the more complicated non constant cases. We will also address the problems within the current state of our work as well as open questions to be answered on the way towards an optimal and general Multi-Stage Checkpointing solution.