

Semesterarbeit

Scheduling with Batching

This semester thesis deals with the following situation: a service must store a sequence of log entries in a high throughput environment. For this, it has a storage device available, which can store only one batch of data at a time. The duration of storing a batch depends only on its size. Log requests arrive over time, and have different sizes.

In this environment, we are interested in finding good online algorithms for grouping the arriving logs into batches and for scheduling these batches. The goal is to minimize the average flow time of a log request, that is the time difference between its arrival and its safe storage.

The aim of this semester thesis is to investigate some aspects of the described problem, for example finding online algorithms, or to show for some variants that achieving a good competitive ratio is not possible. It would be interesting to see if one obtains different results for slightly different definitions of the problem. Some possible variations are:

- The latency function, i.e. how the duration of storing a batch depends on its size (linearly, quadratic, etc.).
- In addition to the arrival times, one could consider *deadlines* assigned to each log request.
- Is *preemption* allowed, i.e. a store operation that is in progress can be cancelled?
- What if the number of log entries that can be grouped into a batch is limited?

In the beginning of the work, my Master's thesis "Message Bundling in a Storage Service" could serve as a starting point, as it considers some variants of the described problem.

Requirements: Interest in online algorithms.

For more information please contact: Beat Gfeller, CAB H16, gfeller@inf.ethz.ch