

We propose two protocols for distributed priority queues called SKEAP and SEAP. SKEAP realizes a distributed heap for a constant amount of priorities and SEAP one for an arbitrary amount. Both protocols build on an overlay, which induces an aggregation tree on which heap operations are aggregated in batches, ensuring that our protocols scale even for a high rate of incoming requests. As part of SEAP we provide a novel distributed protocol for the k-selection problem that runs in time $O(\log n)$ w.h.p. SKEAP guarantees sequential consistency for its heap operations, while SEAP guarantees serializability. SKEAP and SEAP provide logarithmic runtimes w.h.p. on all their operations.