This lecture is about the state-of-the-art in sorting and priority queues on a word RAM. An equivalence by Thorup shows that any sorting algorithm can be transformed into a priority queue with operations taking 1/$n$th the time to sort. So these are really one and the same problem.

The best results we know for sorting in linear time (and thus for constant-time priority queues) is when $w = O(\lg n)$ and when $w = \Omega(\lg^{2+\varepsilon} n)$. The first result is just radix sort. The second result is the main topic of the lecture: a fancy word-RAM algorithm called signature sorting. It uses a combination of hashing, merge sort, and parallel sorting networks.

The range of $w$ in between lg and $\lg^{2+\varepsilon}$ remains unsolved. The best algorithm so far runs in $O(n \sqrt{\lg \lg n})$ expected time.

# 6.851 Advanced Data Structures

Spring 2012