

# One, Two, Three.... Infinity: Lower Bounds for Parallel Computation

Faith Fich  
Prabhakar Ragde  
Avi Wigderson

## Abstract

In this paper we compare the power of the two most commonly used concurrent-write models of parallel computation, the COMMON PRAM and the PRIORITY PRAM. These models differ in the way they resolve write conflicts. If several processors want to write into the same-shared memory cell at the same time, in the COMMON model they have to write the same value. In the PRIORITY model, they may attempt to write different values; the processor with smallest index succeeds.

We consider PRAM's with  $n$  processors, each having arbitrary computational power. We provide the first separation results between these two models in two extreme cases: when the size  $m$  of the shared memory is small ( $m \leq n^c$ ,  $\forall \epsilon < 1$ ), and when it is infinite.

In the case of small memory, the PRIORITY model can be faster than the COMMON model by a factor of  $\Theta(\log n)$ , and this lower bound holds even if the COMMON model is probabilistic. In the case of infinite memory, the gap between the models can be a factor of  $\Omega(\log \log \log n)$ .

We develop new proof techniques to obtain these results. The technique used for the second lower bound is strong enough to establish the first tight time bounds for the PRIORITY model, which is the strongest parallel computation model. We show that finding the maximum of  $n$  numbers requires  $\Theta(\log \log n)$  steps, generalizing a result of Valiant for parallel computation trees.