
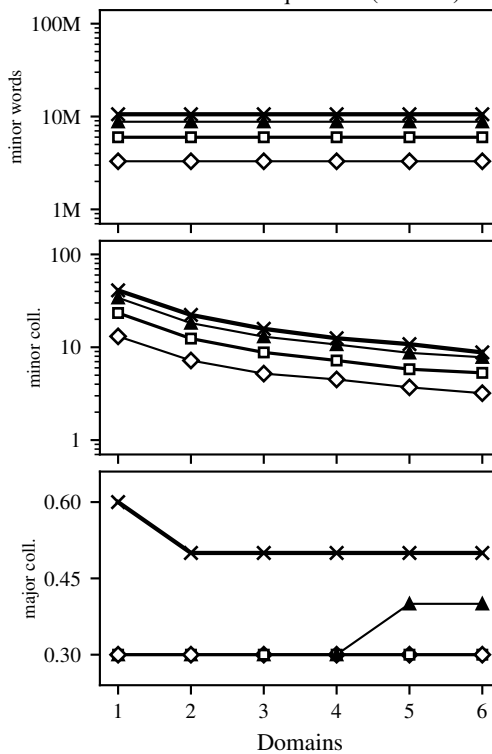


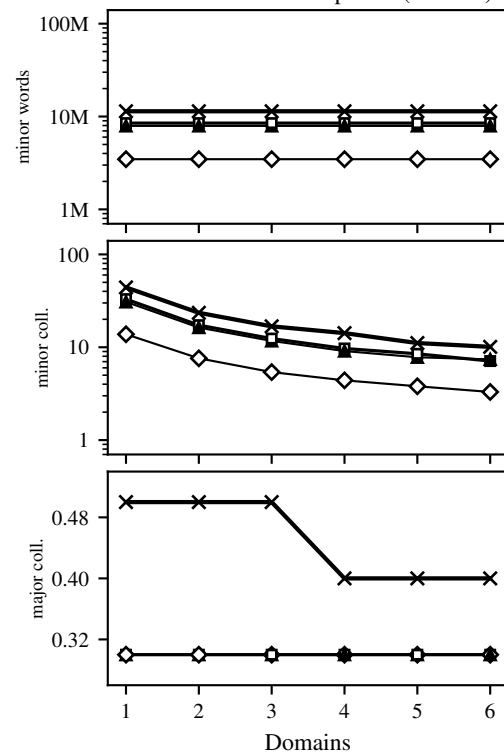
GC cost: stdlib vs par (1 domain)
(IC1, $n = 1M$, $b_{\max} = 128$)



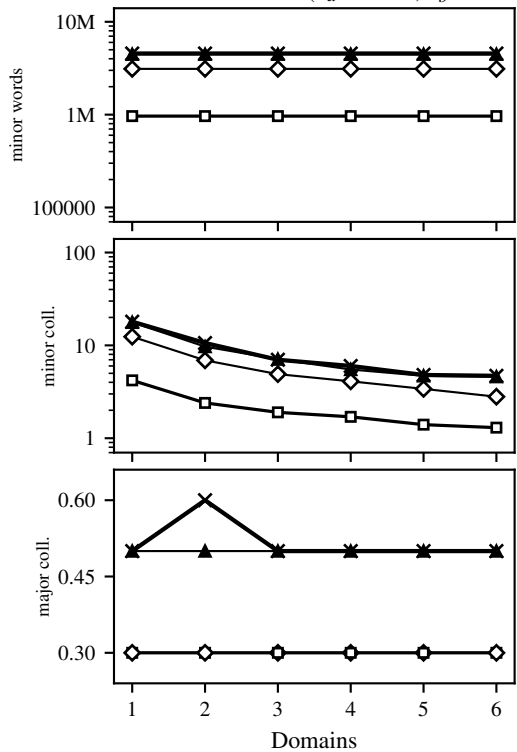
IC1: random equal-size ($n = 1M$)




IC2: controlled overlap 50% ($n = 1M$)



IC3: skewed ratio 9:1 ($n_a = 900k$, $n_b = 100k$)



IC4: sorted disjoint ($n = 1M$)



IC5: clustered keys ($n = 1M$, 20 clusters)

