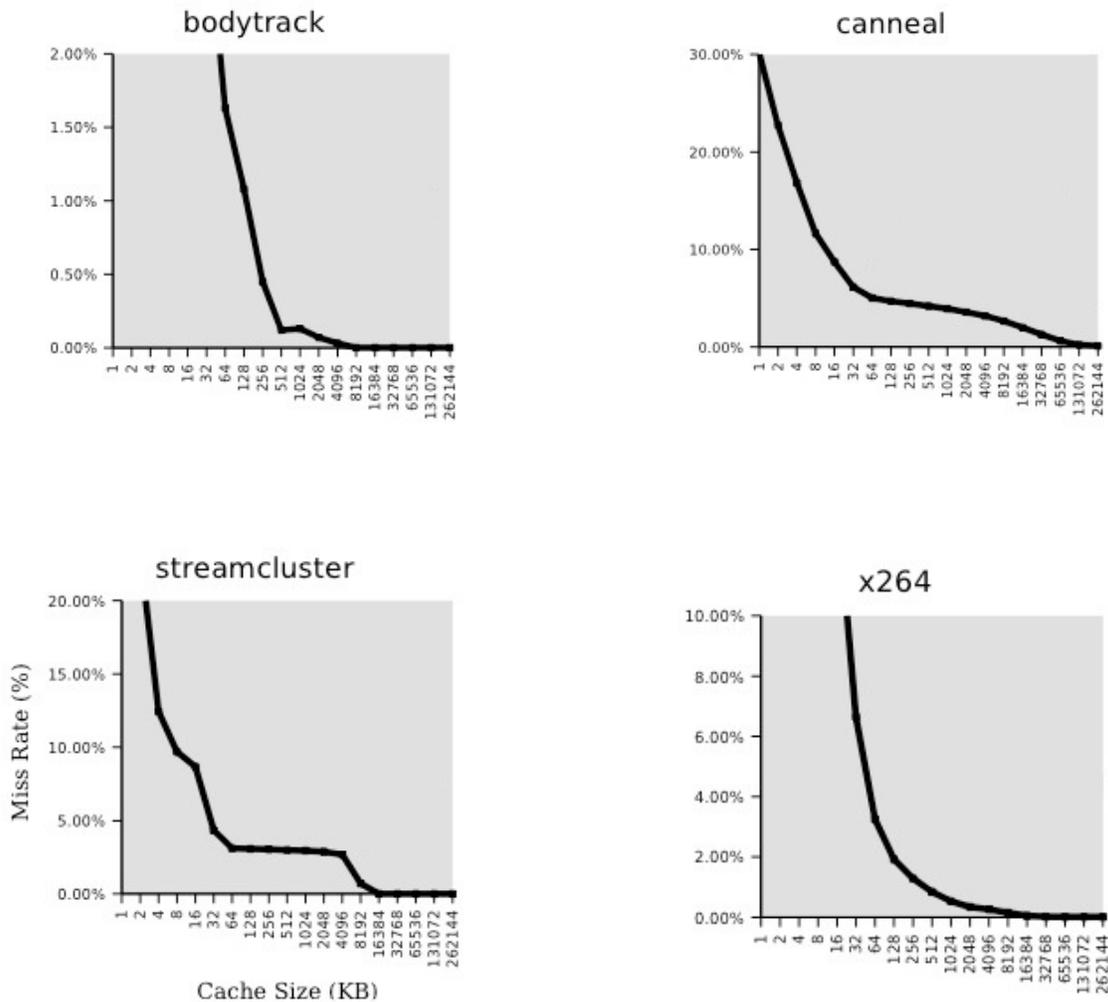


Problem 1: Working Set Size

Below are some plots of cache miss rates from four applications in the PARSEC benchmark suite. The cache modeled has 64B lines, is 4-way set associative, and its capacity is varied.



a) For each of the plots above, determine the application's working set size.

b) Why might bodytrack's miss rate have gone up with a bigger cache?

Problem 2: Short Answer

a) Cache A has a 90% hit rate and cache B has a 95% hit rate. Numerically, how much better is cache B?

b) Can you have a no-write allocate, write-back cache?

c) If you are concerned about memory bandwidth (to DRAM), should you have a write allocate cache?

Problem 3: Write Buffers

For this problem we are considering adding a write buffer to a single issue in-order CPU with one level of cache. The cache is write back, no write allocate, and for this problem you don't have to worry about lines ever getting evicted. By adding a write buffer, if a write misses in the cache, it can go directly into the write buffer rather than stalling the system for the miss.

	Quantity	Time (cycles)
read	Hit time	2
	Miss penalty	200
	Hit (after adding WB)	$3 + 2\log_2(\text{WB depth})$
write	Hit time	4
	Miss penalty	240
	Hit/miss time (after adding WB)	6
	Cache read hit rate	98%
	Cache write hit rate	70%

a) What is the average memory access time (AMAT) for reads before and after a write buffer of depth 4 is added?

b) What is the AMAT for writes before adding the write buffer?

c) Assume a program averages a write every 24 cycles and that they are decently uniformly distributed. What is the worst average write miss rate the program can have and expect a finite sized write buffer to not overflow?