

Optimal Caching and Greedy Exchange

In the optimal caching problem, our computer has a main memory of size n , a cache of size k , and we are presented with a sequence of data $D = d_1, d_2, \dots, d_m$ that we must process. When an item is not in the cache, we have a cache miss, and must bring the item into the cache and evict something else if the cache is full. We assume that the cache is initially empty. Our goal is to give an algorithm that minimizes the number of misses.

Determine a cache maintenance algorithm by coming up with an eviction schedule.

Example 1:	a, b, c, b, c, a, b	$k = 2$
Example 2:	$a, b, c, d, a, d, e, a, d, b, c$	$k = 3$

Definition 1. A schedule is *reduced* if it does the minimal amount of work necessary in a given step.

- If the requested item is already in the cache, nothing is done.
- Otherwise, the missing item is brought into the cache (after possibly evicting another item).

Lemma 1. For every non-reduced schedule, there is an equally good reduced schedule (that brings in at most as many items as the original schedule).

Prove this by construction.

Hint: You might *charge* bringing in an item in one schedule to bringing in an item in another schedule to show that it doesn't bring in any more items than the original schedule.

Observation 1. *For any reduced schedule, the number of items that are brought in is exactly the number of misses.*

Proof by Greedy Exchange

Step 1: Label. Label your algorithm's solution ($A = \{a_1, a_2, \dots, a_k\}$), and a general solution ($O = \{o_1, o_2, \dots, o_m\}$).

Step 2: Compare. Compare greedy with the other solution. Assume that they're not the same and isolate some difference.

Step 3: Exchange. Swap the elements in O without making the solution worse. Argue that swapping a finite number of times will result in A . Hence, greedy is just as good as *any* optimal or arbitrary solution.

The Exchange Step

Let A_{FF} be the schedule created by the farthest-in-future algorithm and let S be an arbitrary reduced schedule.

Lemma 2. *Suppose S is a reduced schedule that makes the same eviction decisions as A_{FF} through the first j items in the sequence for some j . Then there exists a reduced schedule S' that makes the same eviction decisions as A_{FF} through the first $j + 1$ items and incurs no more misses in total than S does in total.*

Prove this by constructing S' . *This is an **exchange** argument.*

a. What happens if the $j + 1^{st}$ item is in cache?

b. What happens if the $j + 1^{st}$ item isn't in cache, but S evicts the same item as A_{FF} ?

c. What happens if the $j + 1^{st}$ item isn't in cache, and S evicts a different item as A_{FF} ? What should S' do?

d. How can you get S' 's cache back to the same as S 's without incurring more total misses?

e. How do we know that S' is a reduced schedule?

f. Sanity check: Are all parts of the lemma true?

Theorem 2. *A_{FF} incurs no more misses than any other schedule O and hence is optimal.*