Python lists:

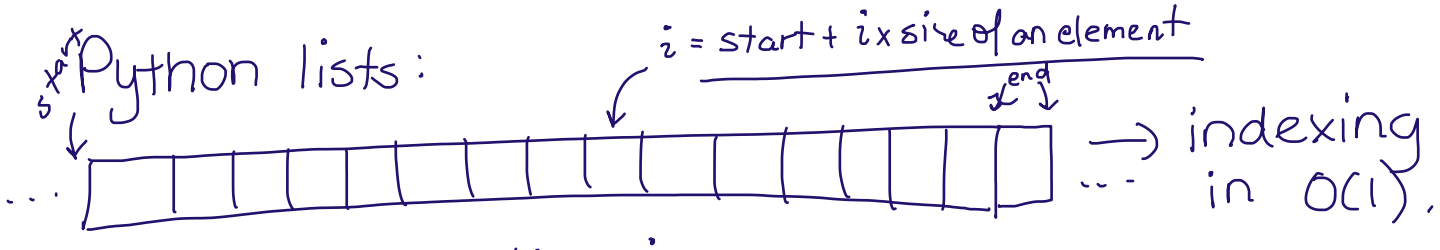

→ consecutive in memory.

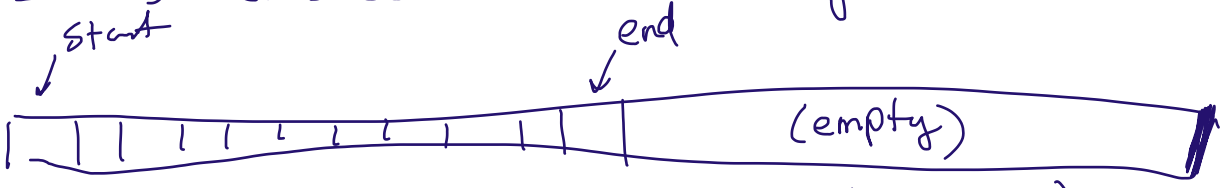

append is usually O(1)

append only slows when we run out of empty spots

What if objects could tell us which index to find them at?

→ hash function. object → int

x → 9 we can look at index 9 for x

↑ hash value.

Desireable Properties

1. fast (otherwise, just use a list)
2. deterministic (always same for same object) → immutable keys.
3. well-distributed.

hash value % size of list → 0 .. len - 1 valid indices.

what if 2 things give same index?

collision

Collisions will happen,
and often!

Strategies:

1. Chaining
   - store a collection of objects with same hash %table size instead of single object
   → problem: too many collisions at same spot, turns into linear search.
2. open address probing
   - look for the next spot that is available.
   → could turn into linear search too
   → more collisions!