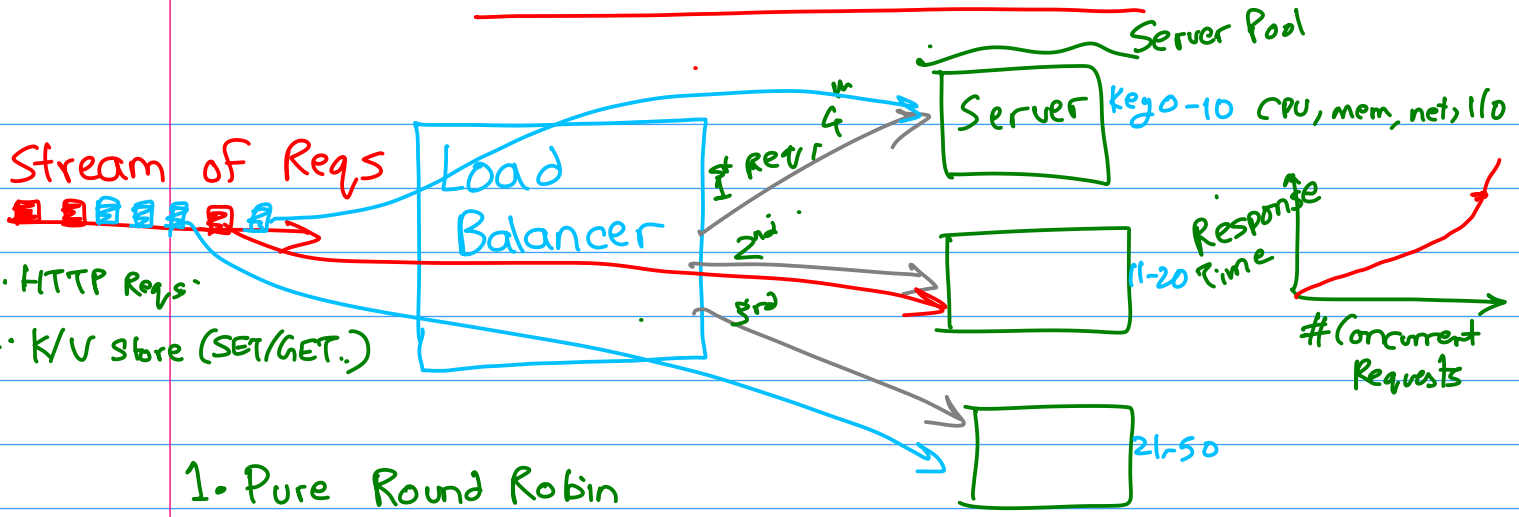



LOAD BALANCING



1. Pure Round Robin
2. Randomized

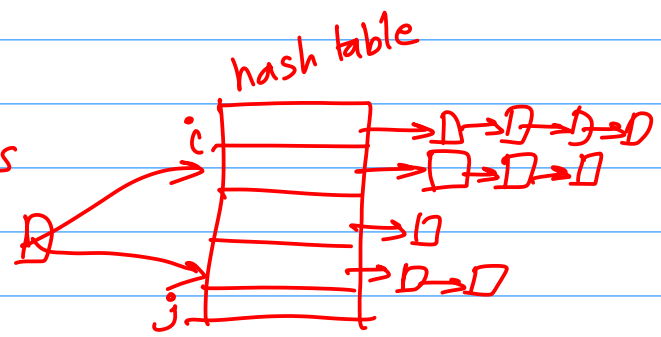
STATEFUL/PERSISTENT/"STICKY" → Locality of requests from users/IP/Req. id



Load imbalance

↳ distribution of the req. ids/keys

Skew



Random


"Power of 2 choices"

$\star \wedge$

$2^{Pow} hash(r)$ → i, j $j \in [1, k]$

Least loaded of i, j

Tuple output




hash(r) → i

hash(r + salt) → j

↑ fixed random number

XOR

Servers Pool can grow & shrink dynamically.



→ lose persistence & not stateful

Pools: $10 - 10^3$ servers

$\text{hash}(r) \% k + 1$

Consistent hashing

- hash all servers (IP addr) to a Circular Ring
- hash all reqs on

