



Two Elementary Instructions make Compare-and-Swap







Pankaj Khanchandani, Roger Wattenhofer
ETH Zurich - Distributed Computing Group (DISCO)

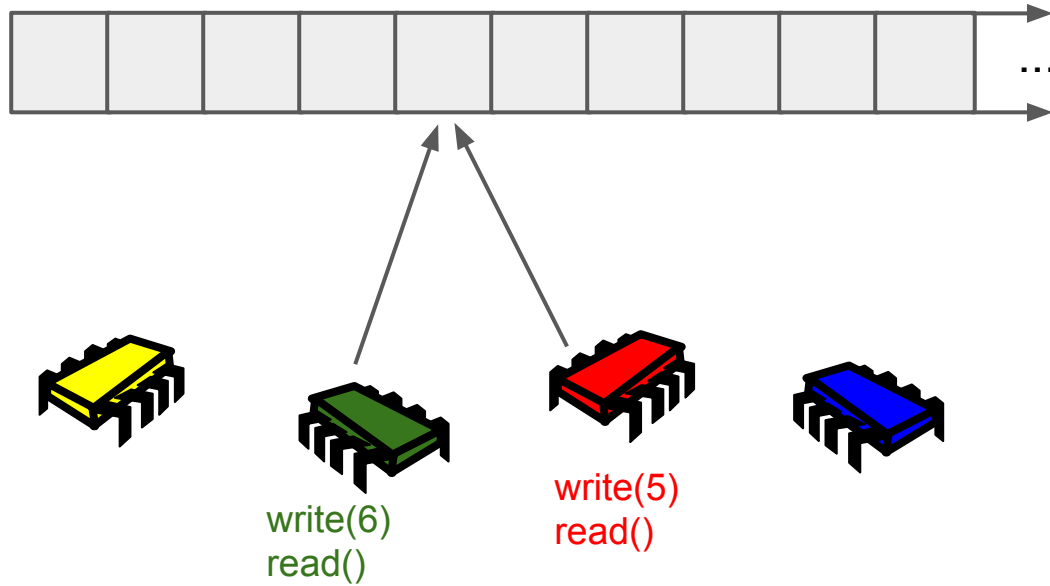


Compare-and-Swap


Shared Memory Model



Shared Memory Model




Shared Memory Model




Atomic Registers

fetch-and-add(x)




Atomic Registers

fetch-and-add(x)




Atomic Registers

fetch-and-add(x)

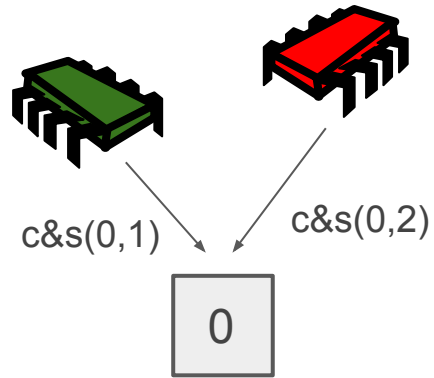


Atomic Registers

fetch-and-add(x)




compare-and-swap(x,y)

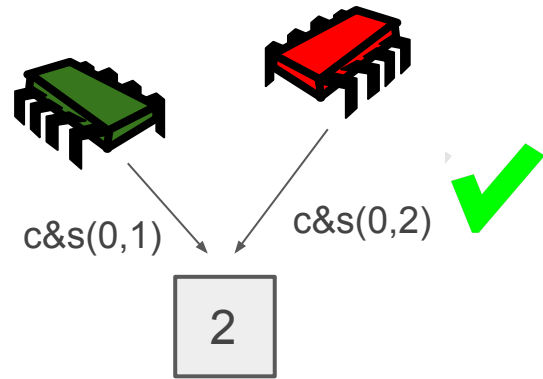


Atomic Registers

fetch-and-add(x)

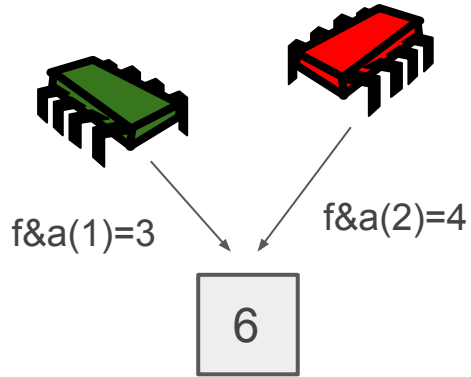


compare-and-swap(x,y)




Atomic Registers

fetch-and-add(x)



compare-and-swap(x,y)



Best Atomic Registers?

Consensus Numbers [Herlihy 1991]

compare-and-swap = ∞

...

n-register assignment = $2n-2$

...

fetch-and-add = 2

read/write = 1

Best Atomic Registers?

[Ruppert 1997]

compare-and-swap = ∞


...

n-register assignment = $2n-2$

...

fetch-and-add = 2

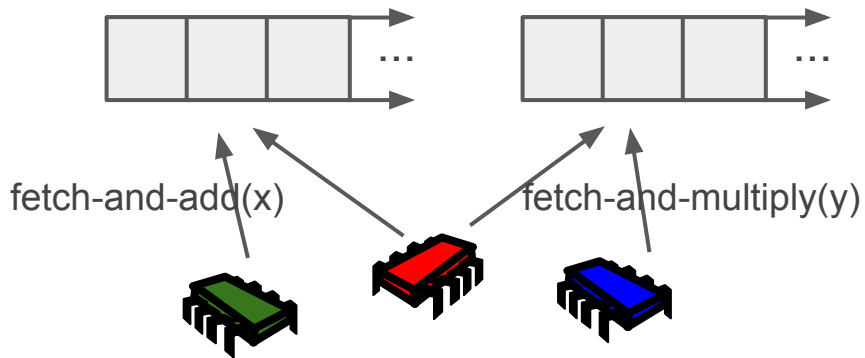
read/write = 1



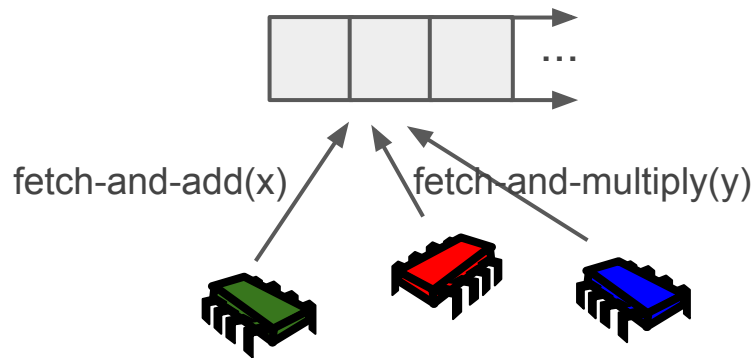
Best Atomic Registers?

[Ellen et al. 2016]

Assumption




Reality




Best Atomic Registers?

[Ellen et al. 2016]

Assumption



Reality

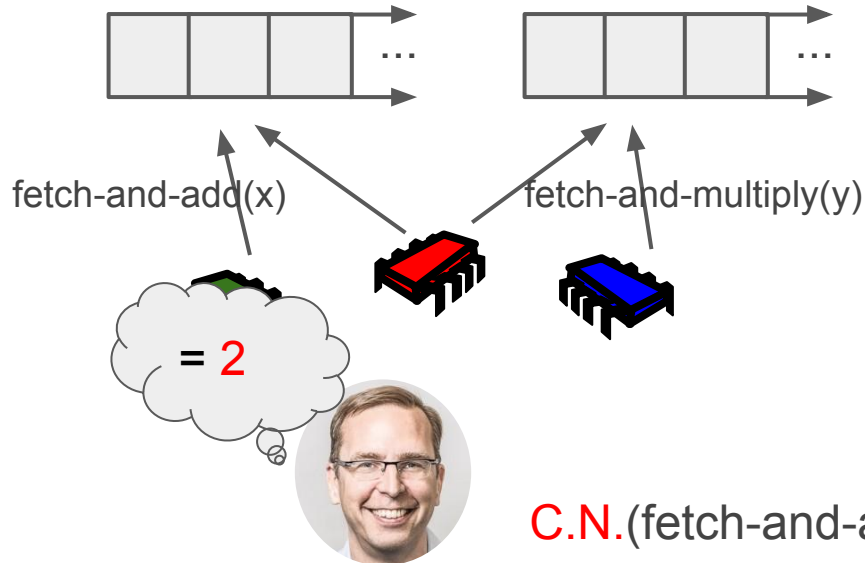


C.N.(fetch-and-add, fetch-and-multiply)

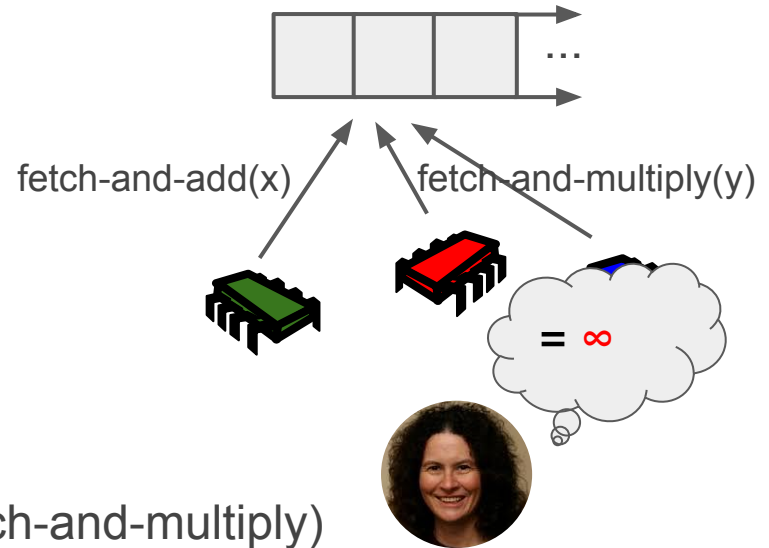
Best Atomic Registers?

[Ellen et al. 2016]

Assumption

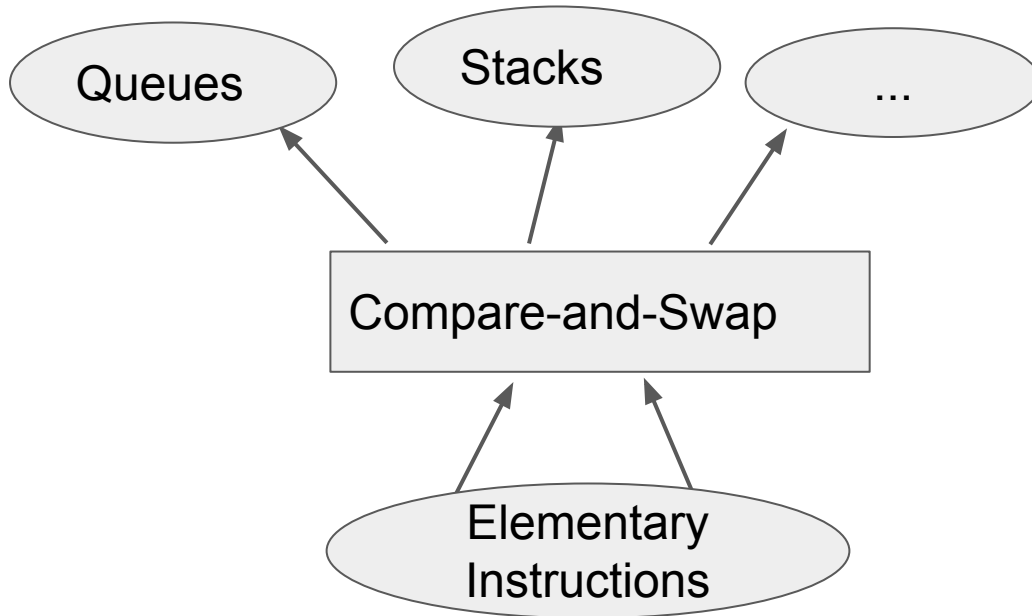


Reality



C.N.(fetch-and-add, fetch-and-multiply)


Best Atomic Registers?




This Work

Half-max and max-write operations can simulate Compare-and-Swap in $O(1)$ steps, in linearizable and wait-free manner.


This Work: half-max, max-write



This Work: half-max, max-write




This Work: half-max, max-write




Simulation Idea

Announce Array



Return Values Array



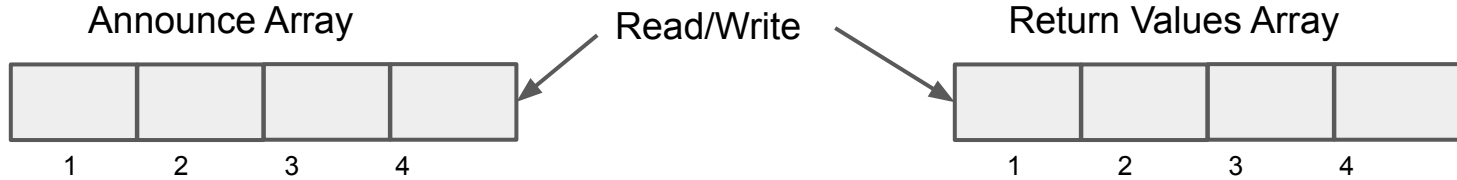
(seq., winner)



(seq., value)



Simulation Idea



(seq., winner)



← half-max/max-write


(seq., value)




← max-write

Simulation Idea


Announce Array




Return Values Array





(seq., winner)




(seq., value)




Simulation Idea



(seq., winner)




(seq., value)




false

Simulation Idea


Announce Array




Return Values Array







(seq., winner)



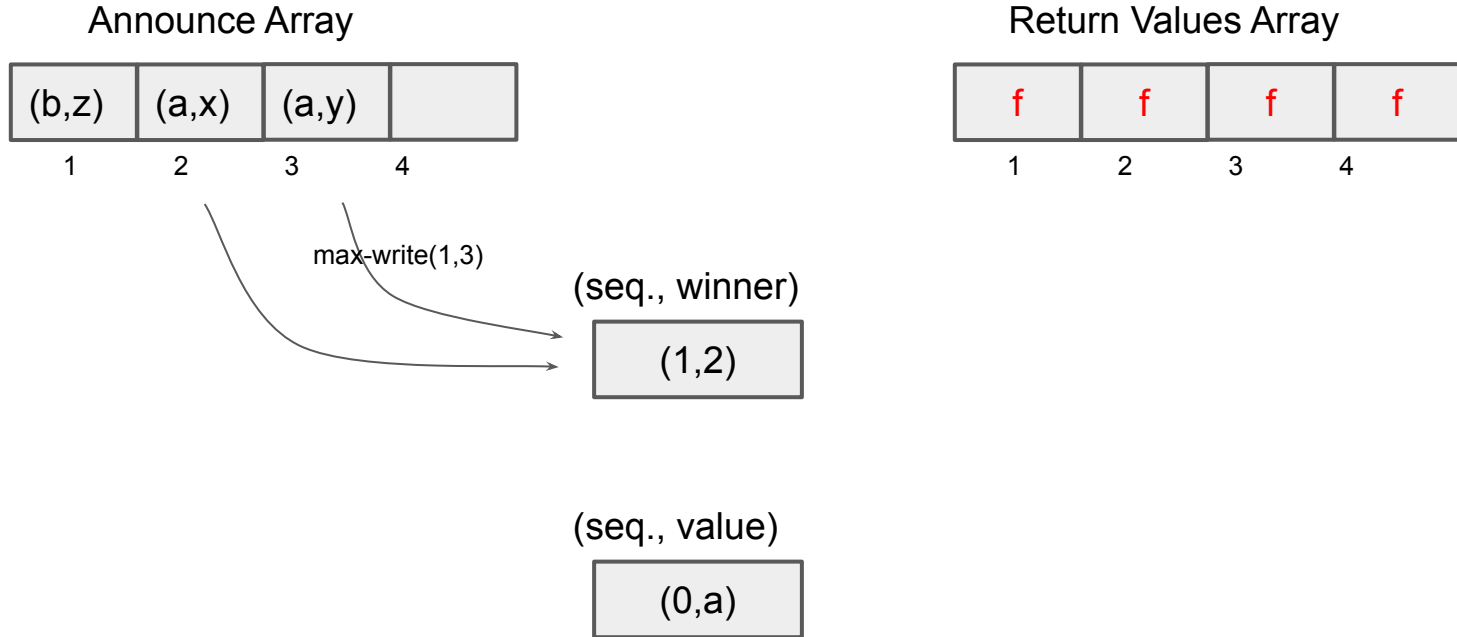
(seq., value)



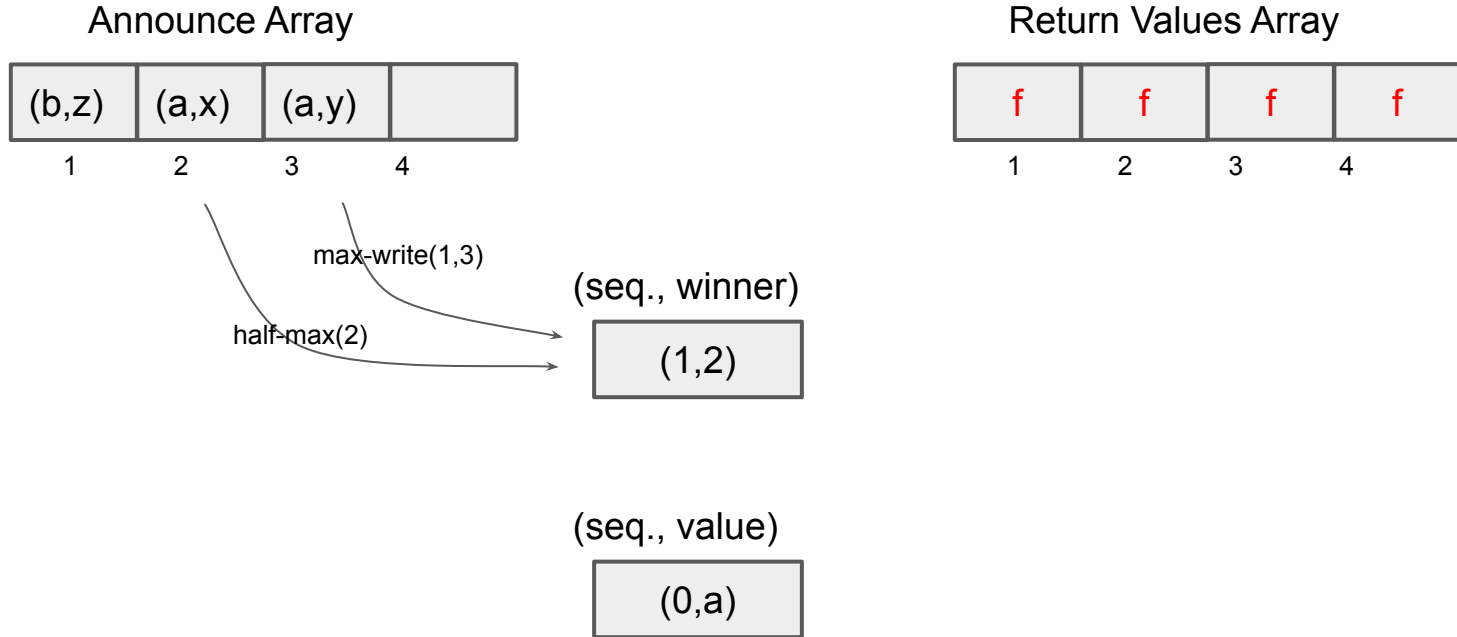
Simulation Idea







Simulation Idea





Simulation Idea



Simulation Idea



Simulation Idea



(seq., winner)

(2,2)

(seq., value)

(0,a)

Simulation Idea

Announce Array

(b,z)	(a,x)	(a,y)	
1	2	3	4

Return Values Array

f	f	f	f
1	2	3	4



(seq., winner)

(2,2)

(seq., value)

(0,a)

Simulation Idea






(seq., winner)

(2,2)

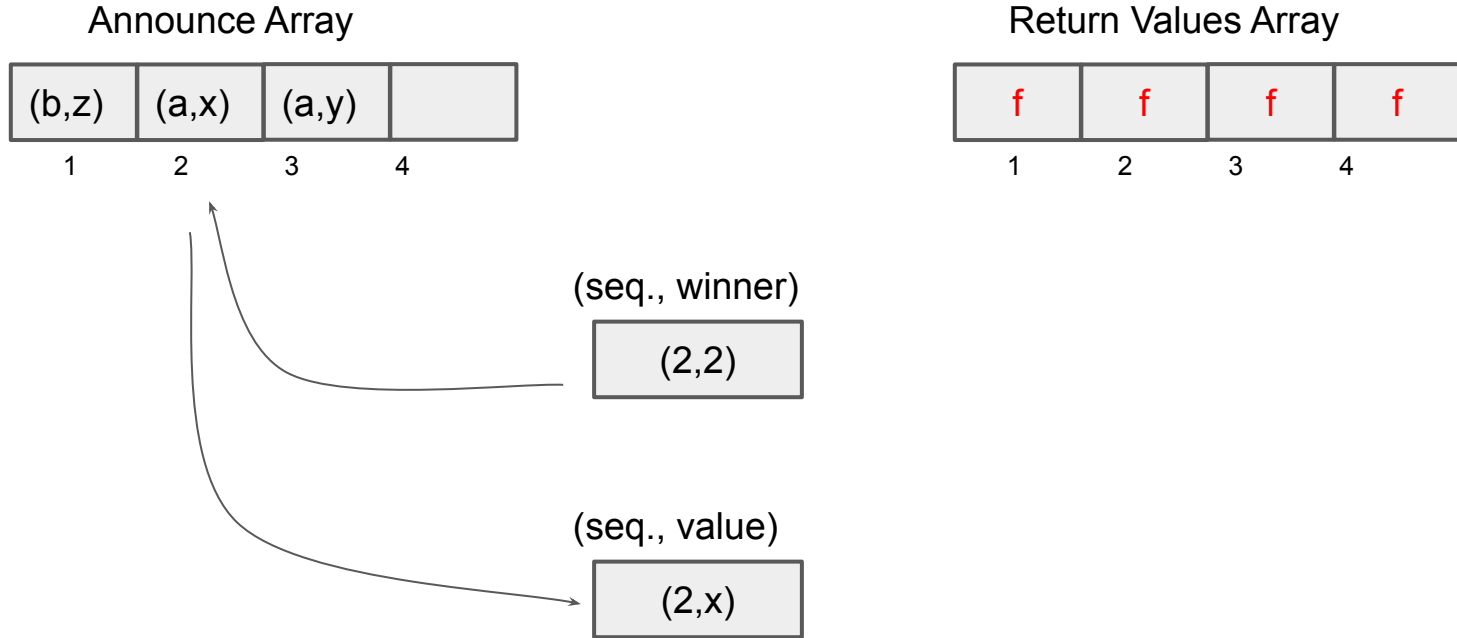
(seq., value)

(0,a)

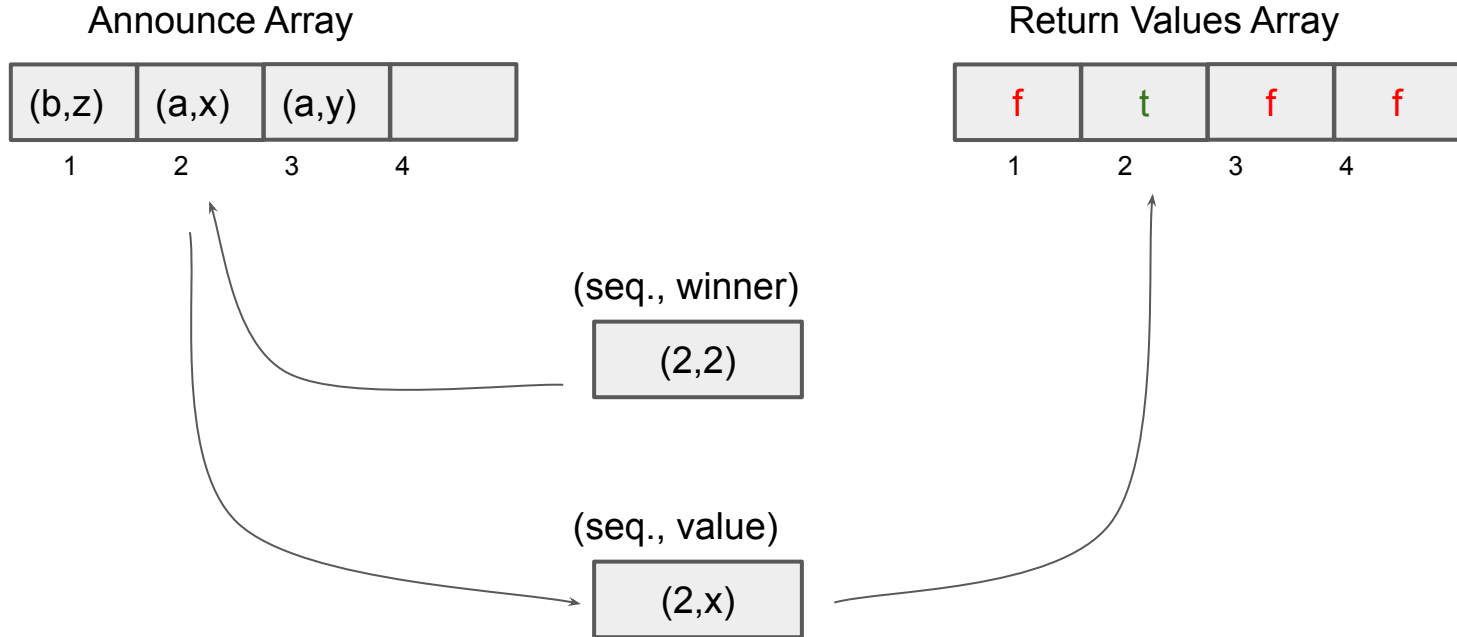
Simulation Idea



Simulation Idea



Simulation Idea




Consensus Numbers

Consensus: Each process has an input. The processes agree on an input of a single process and output that input.


Consensus Numbers

Consensus Number is the maximum number of processes n for which the instruction can solve consensus.

C.N.(compare-and-swap) = ∞

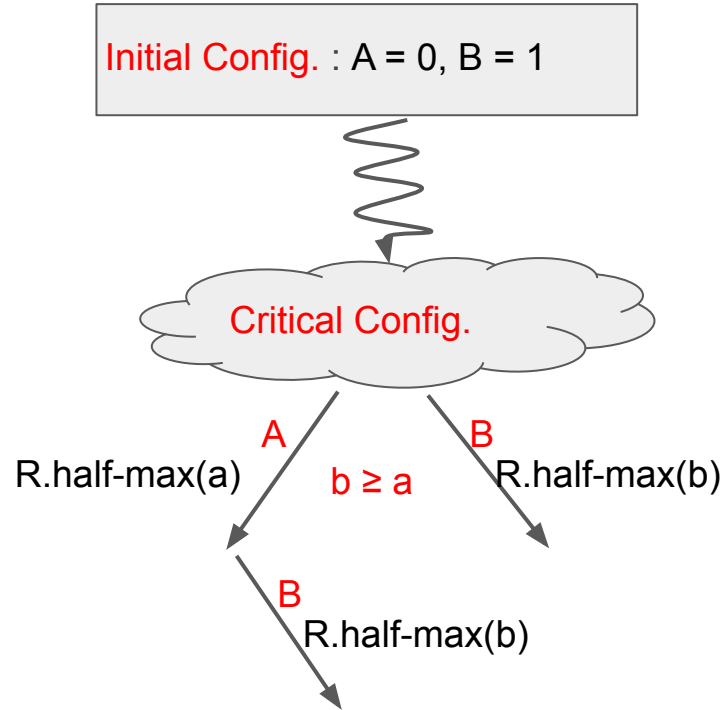


C.N.(compare-and-swap) = ∞




$$\text{C.N.}(\text{half-max}) = 1$$


C.N.(half-max) = 1




Are they better?




Are they better?




Are they better?




Are they better?



Are they better?



Are they better?







Are they better?

max-write(3,1)

max-write(4,2)

max-write(8,5)


max-write(9,4)



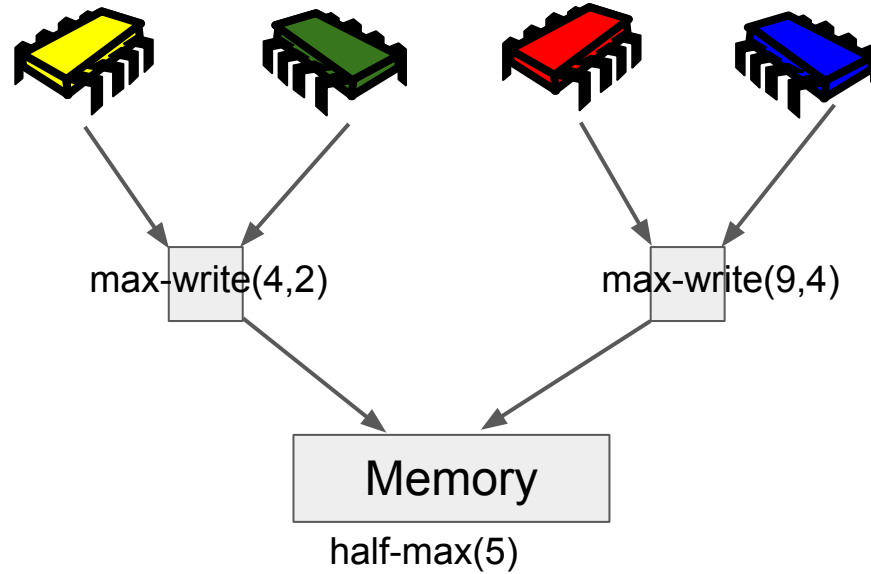
half-max(5)

half-max(2)

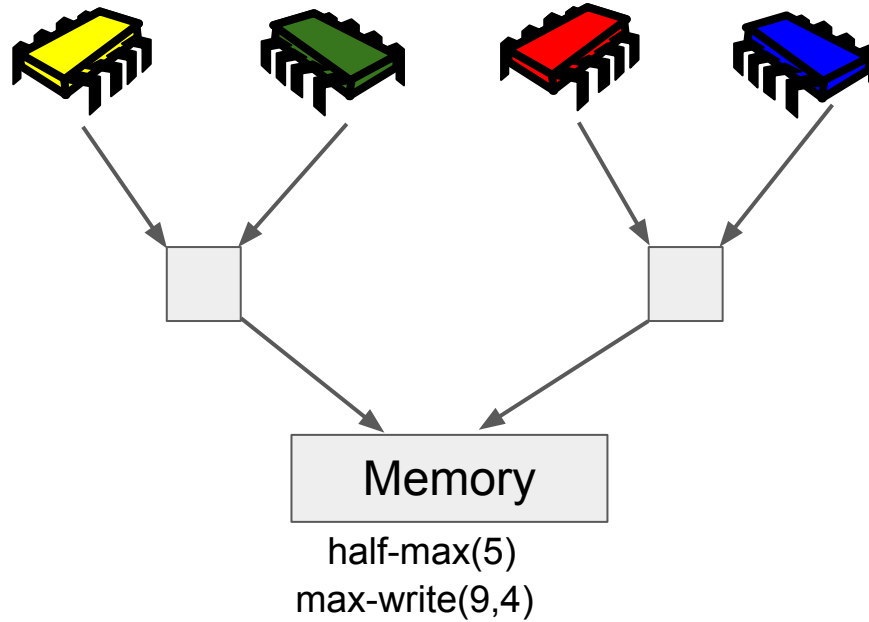
Memory



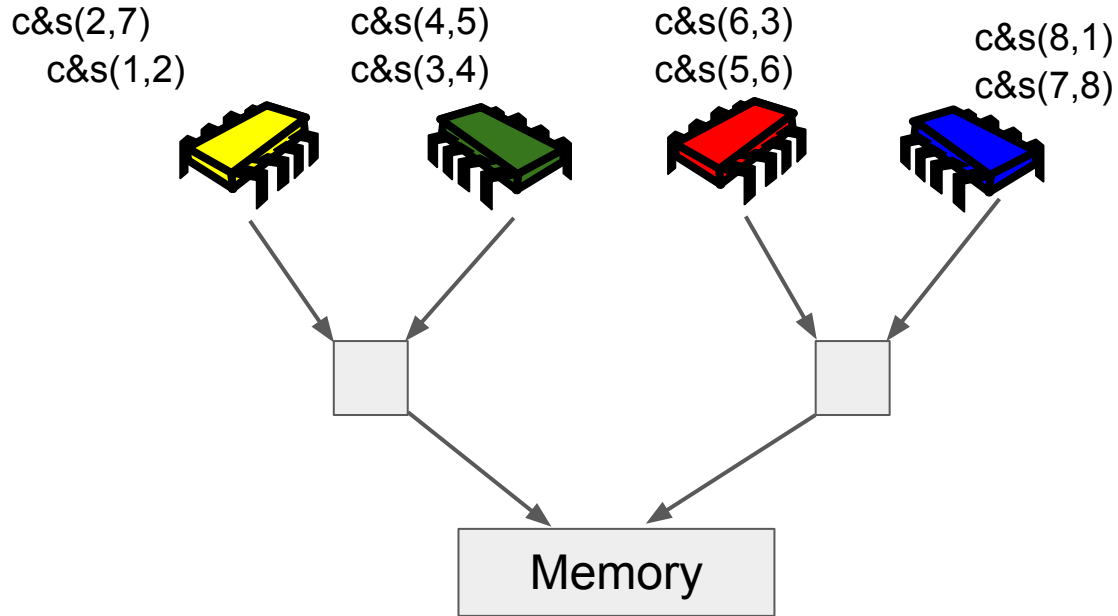
Are they better?



Are they better?



Are they better?







Are they better?

c&s(2,7)

c&s(4,5)

c&s(6,3)

c&s(8,1)




c&s(1,2)

c&s(3,4)

c&s(5,6)

c&s(7,8)

Memory








Are they better?

c&s(2,7)

c&s(4,5)


c&s(6,3)

c&s(8,1)




c&s(1,2) c&s(7,8)
c&s(3,4) c&s(5,6)

Throughput




Throughput



For c&s

$$\frac{n}{t_b}$$

Throughput



For c&s

$$\frac{n}{t_b}$$

For half-max + max-write

$$\frac{n}{6+t_b/m} \approx O(n)$$



kpankaj@ethz.ch

<https://disco.ethz.ch/members/kpankaj>