


Multiprocessor Scheduling




Co scheduling



[All (or almost all) threads of the
Same process should be scheduled
at the same time]


Multicore

TCB



Cosched ✓ ✓ ✓ ✓ ✓

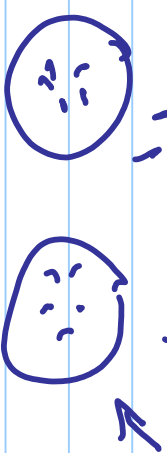
Cluster



- Gang Scheduling

→ All threads @ same time

- Matrix Coscheduling $\boxed{P_i}$

- Process Aware R.R 

- Dynamic Partitioning

- Hand off sched — $\rightarrow T_1$ was on

- Affinity Scheduling \rightarrow core 1
later T_1 should be on core 1

→ Dist Operating System

- Cloud

- Cluster


- Grid

- Fog

- Workstations

- Servers

- RPC / API




Dists

Tannenbaum

cooperation
of processes

→ A distributed operating system

is a collection of autonomous computing
element that appear to be users
as a single coherent (or integrated)
system.




Users

calendar

u_1
 u_2

$c_1 \rightarrow$
 c_2

(date, time of meeting)



← everybody's
schedule



↓
Anyone can
update

↘


↓

↘

Simple → centralized



implementation of a distributed calendar



← server

X process

Negatives

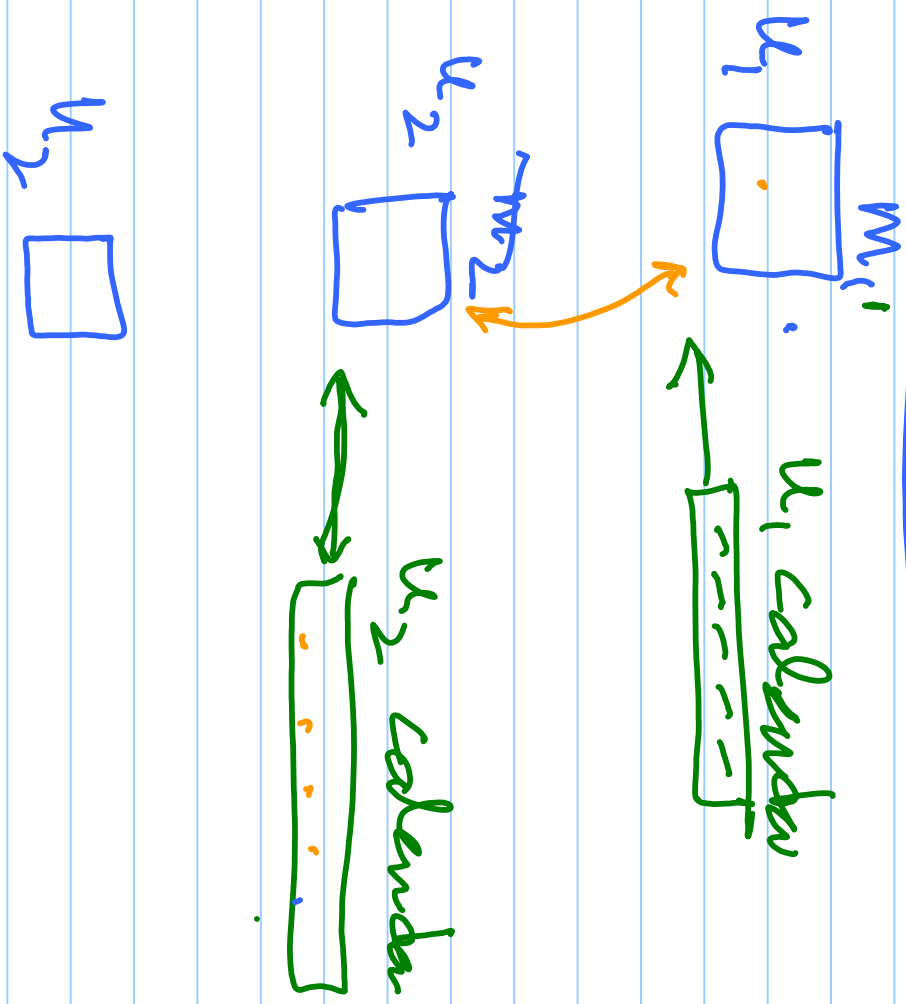
- single point of failure
- congestion
- compute
- network

calendar

positions

- ✓ - simplicity
- ✓ - administration

fully distributed (but per user)



No server

- no single pt of failure

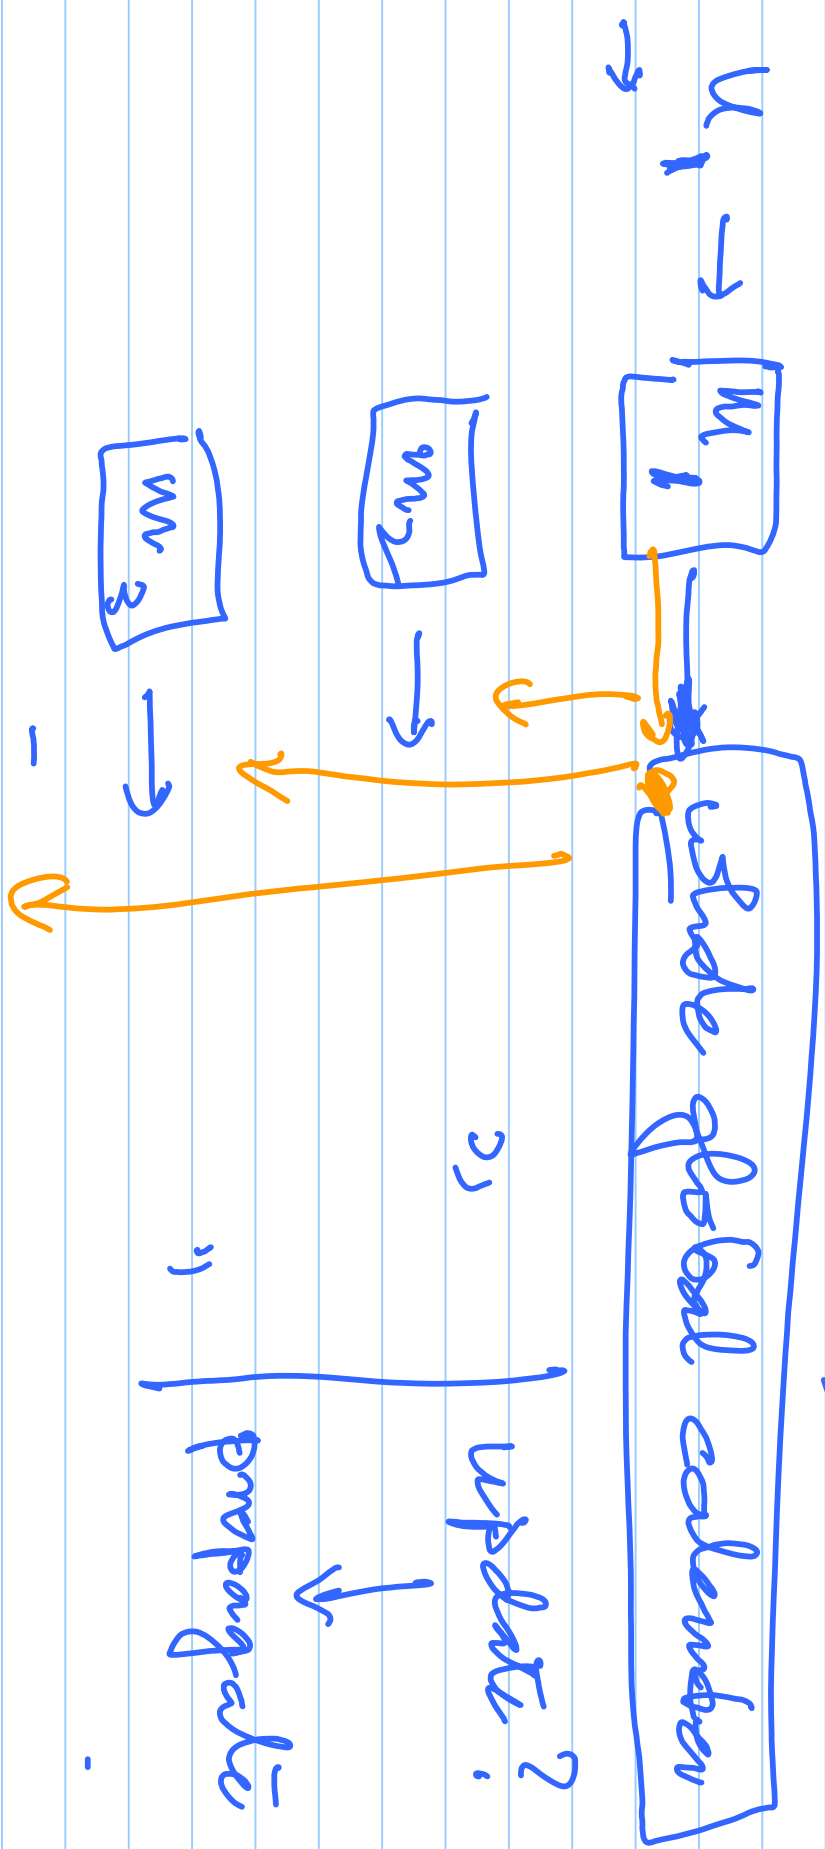
- no congestion

- global queries




- partial

fully distributed - replicated

(everyone has everything)



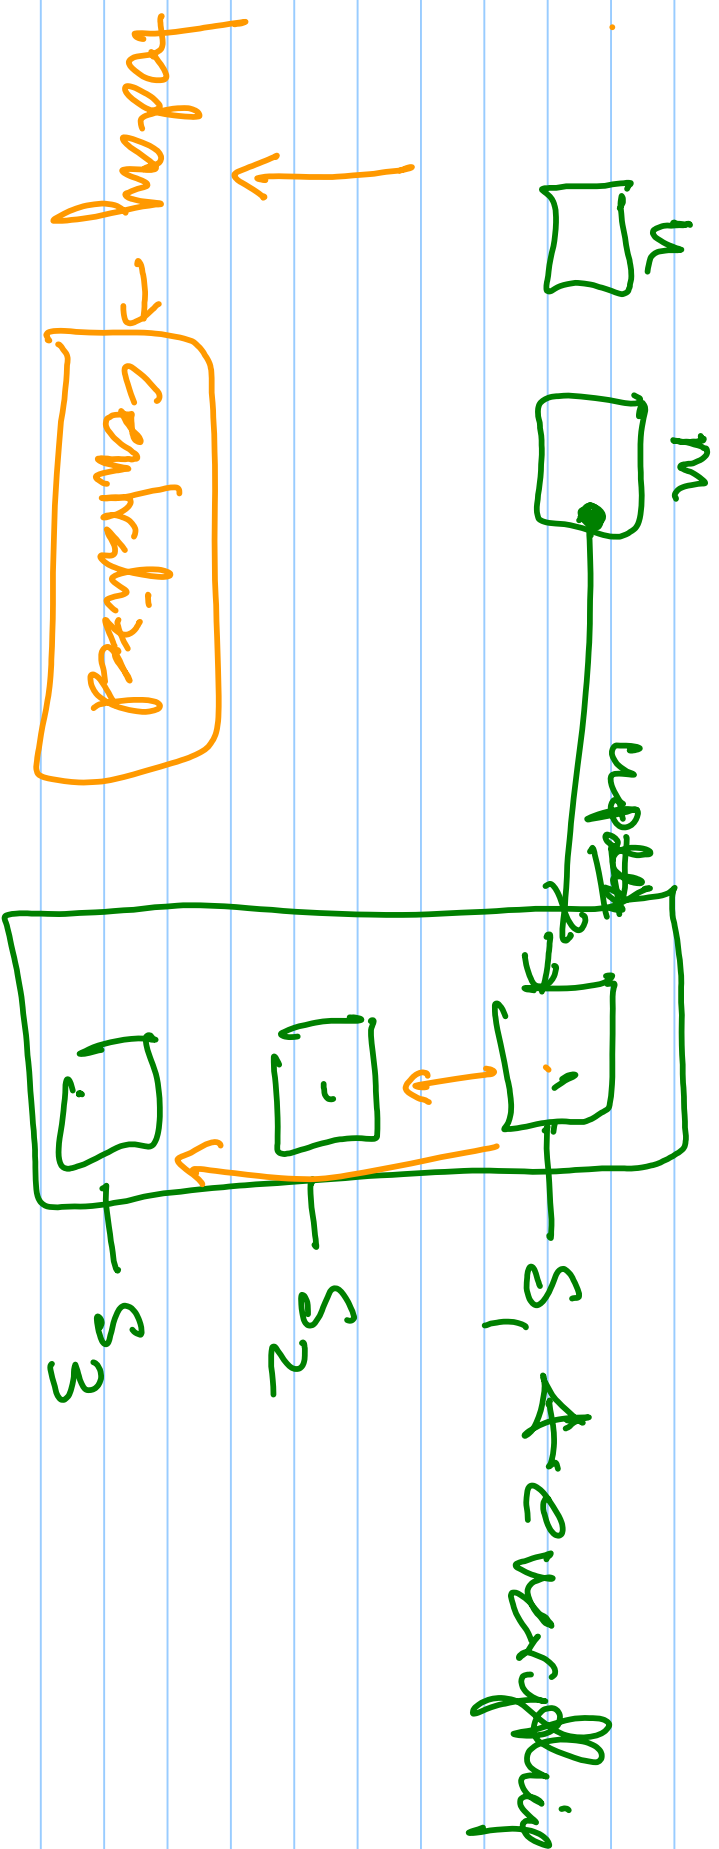
partitioned



← one
single
server
(But not)

↳

Replicated Servers

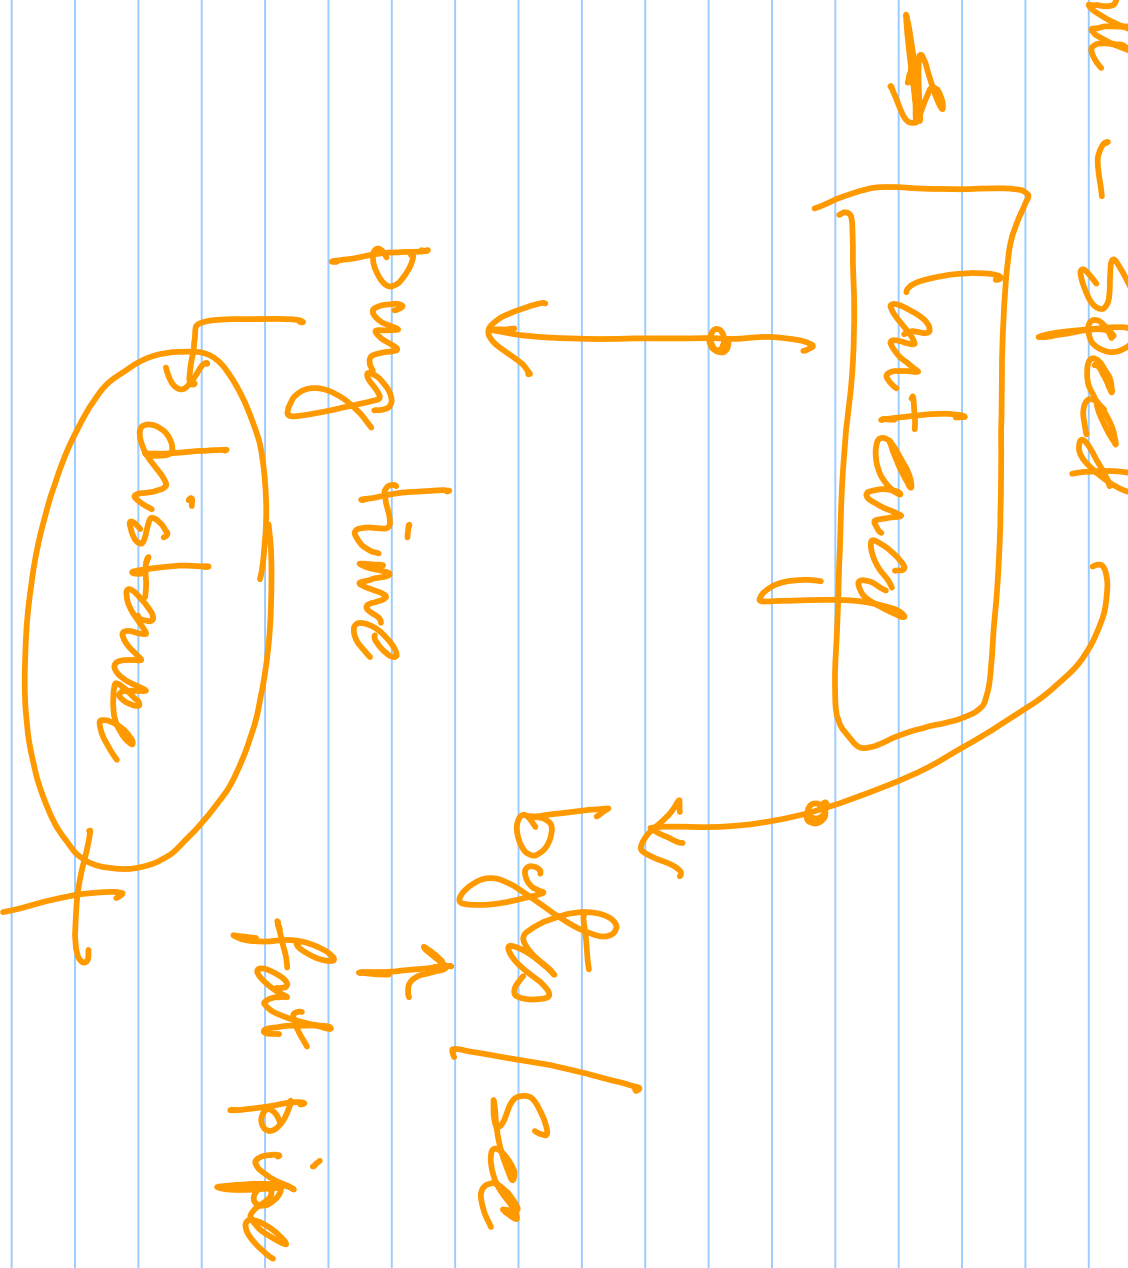


update problem


Goals

- reduce server computation
- client side scripting
- reduce network traffic
- client caching
- reduce overhead (?)
- network] head + latency

Network - Speed



Dist OS model



Properties

- cooperation
- autonomy
- Efficiency
- Availability
- Reliability
- Scalability