Modern Fast Streaming Data

Todd L. Montgomery
@toddlmontgomery

QCon NEW YORK

- □ *Why Should We Care?*
- □ *Myths & Misconceptions*
- □ *You can’t escape the Math*
- □ *Technologies & Techniques*

☑ *Why Should We Care?*

*Human Knowledge is now doubling every year**

** by discipline, 12-18 months*

Fueled by Technology

Middle Ages - 1500 yrs
Renaissance - 250 yrs
Industrial Revolution - 150 yrs
WWII - 25 years

Buckminster Fuller - Critical Path 1981

IoT

~~IoT~~

Ubiquitous Computing

In the near future,
Human Knowledge could
double every 72 hours

What this could mean for our systems…

Either ingest or streaming.
2x for Request/Response

Updates/Sec

=

Devices * Frequency * Market Share

Updates/Sec

=

Devices * Frequency * Market Share

9 Billion (Today)

50 Billion by 2020 (Cisco)

26 Billion by 2020 (Smartphone/Tablet - Gartner)

75 Billion by 2020 (Morgan Stanley)

Updates/Sec

=

50 Billion * 6/min * 1%

=

50 Million/sec

Bandwidth

=

50 Billion * 6/min * 1% * 200 bytes

=

9.3 GB/s (74.5 Gb/s)

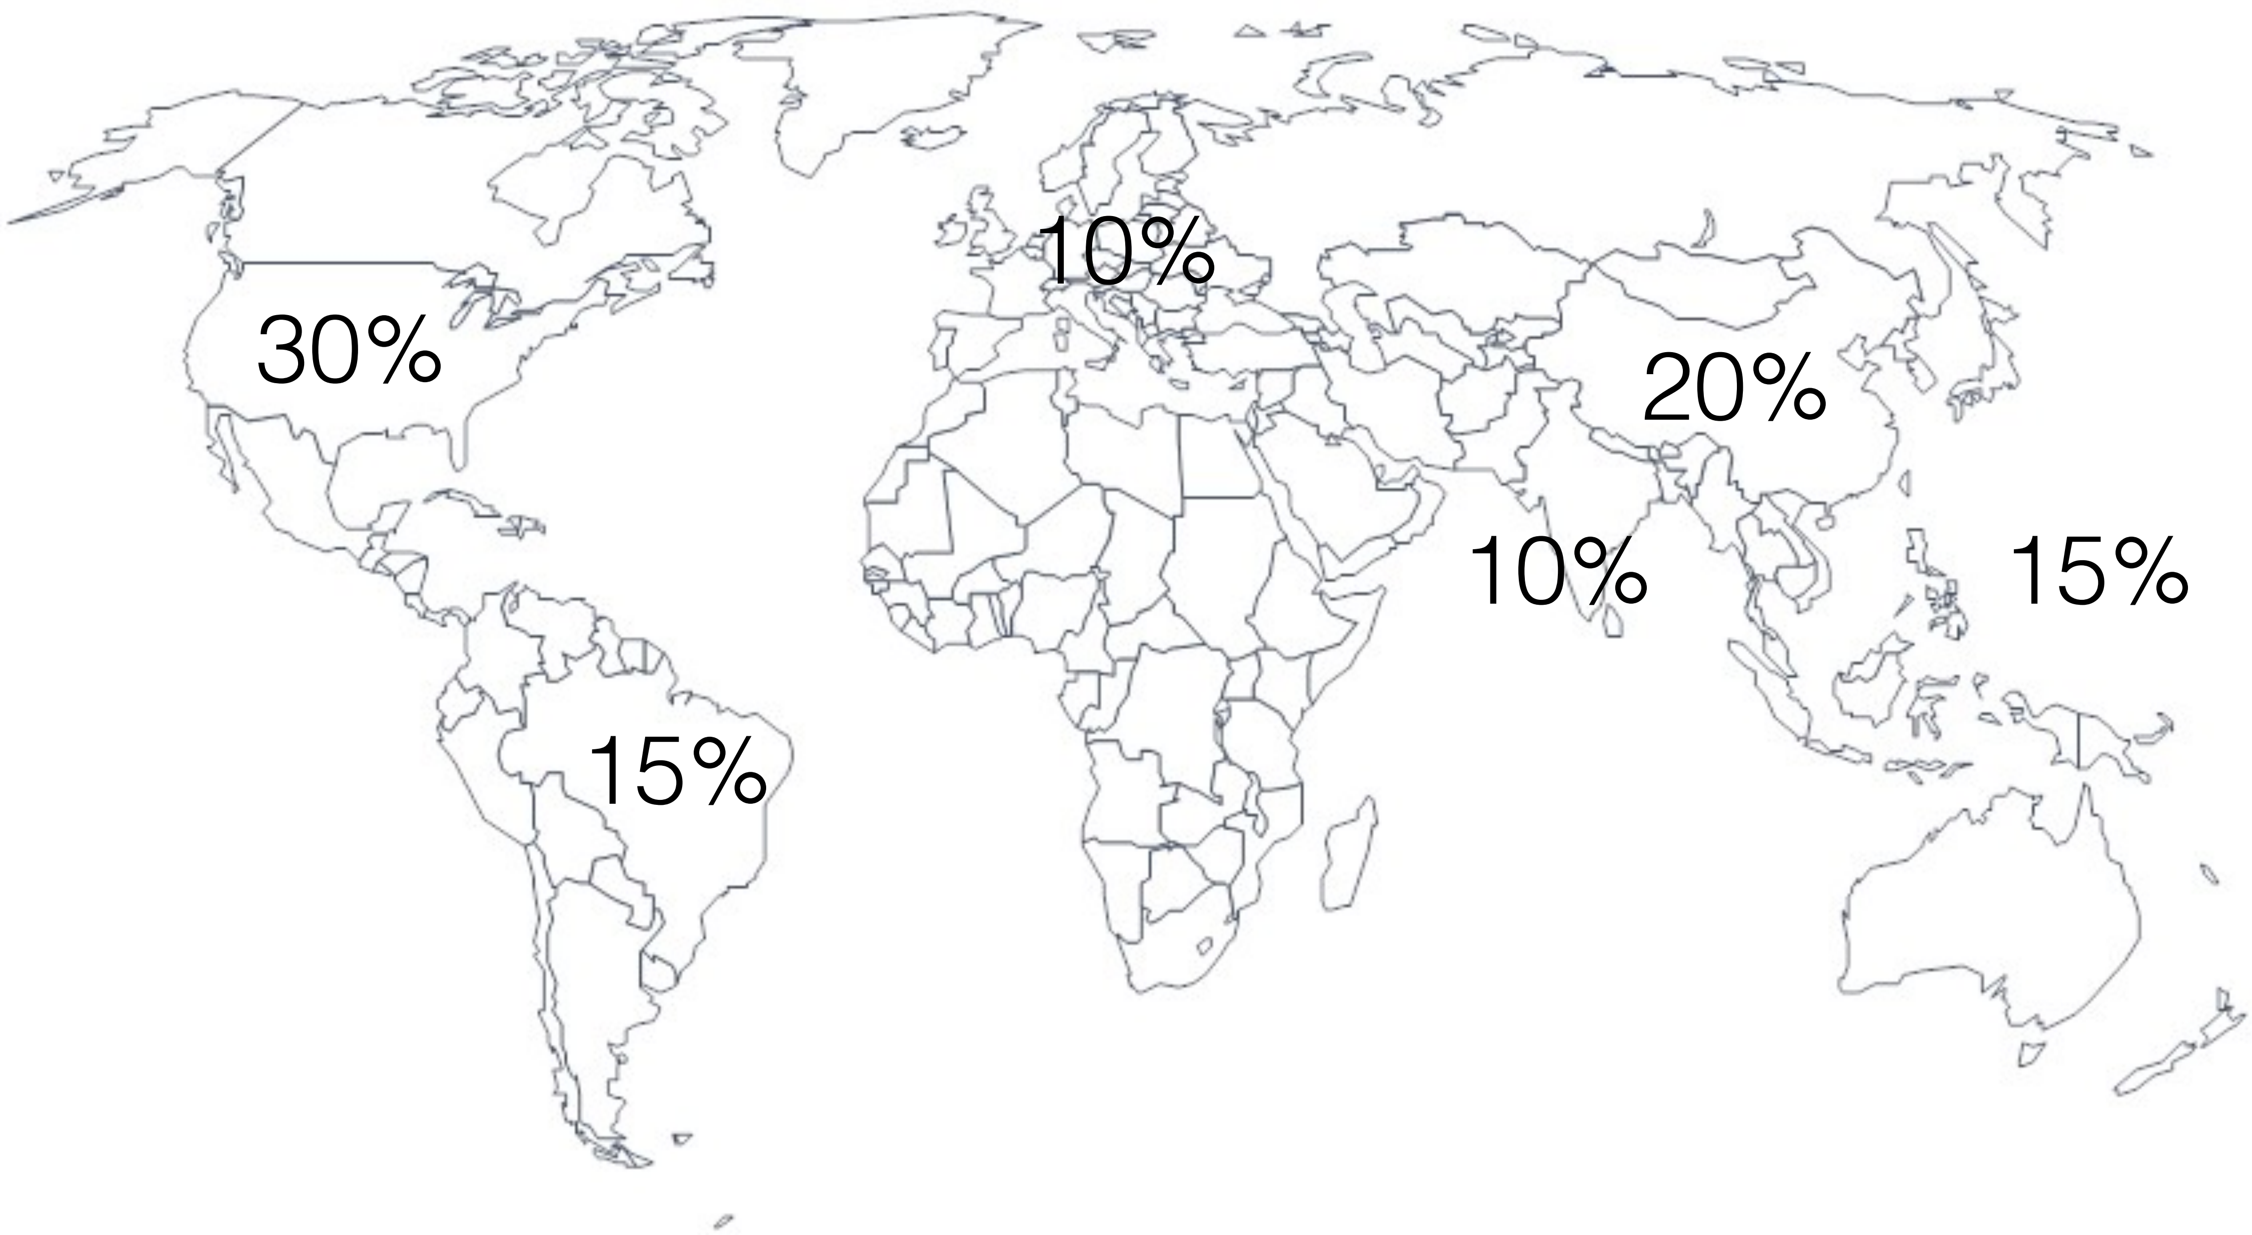

And…
Geographic Distribution

Social & Societal demands
will require
processing an intense stream
of data in real-time

☑ *Myths & Misconceptions*

Excuses, Excuses!

Myth

*(CPUs, Storage, Networks) are not capable of processing in real-time**

** for some unknown, unquantified data volume*

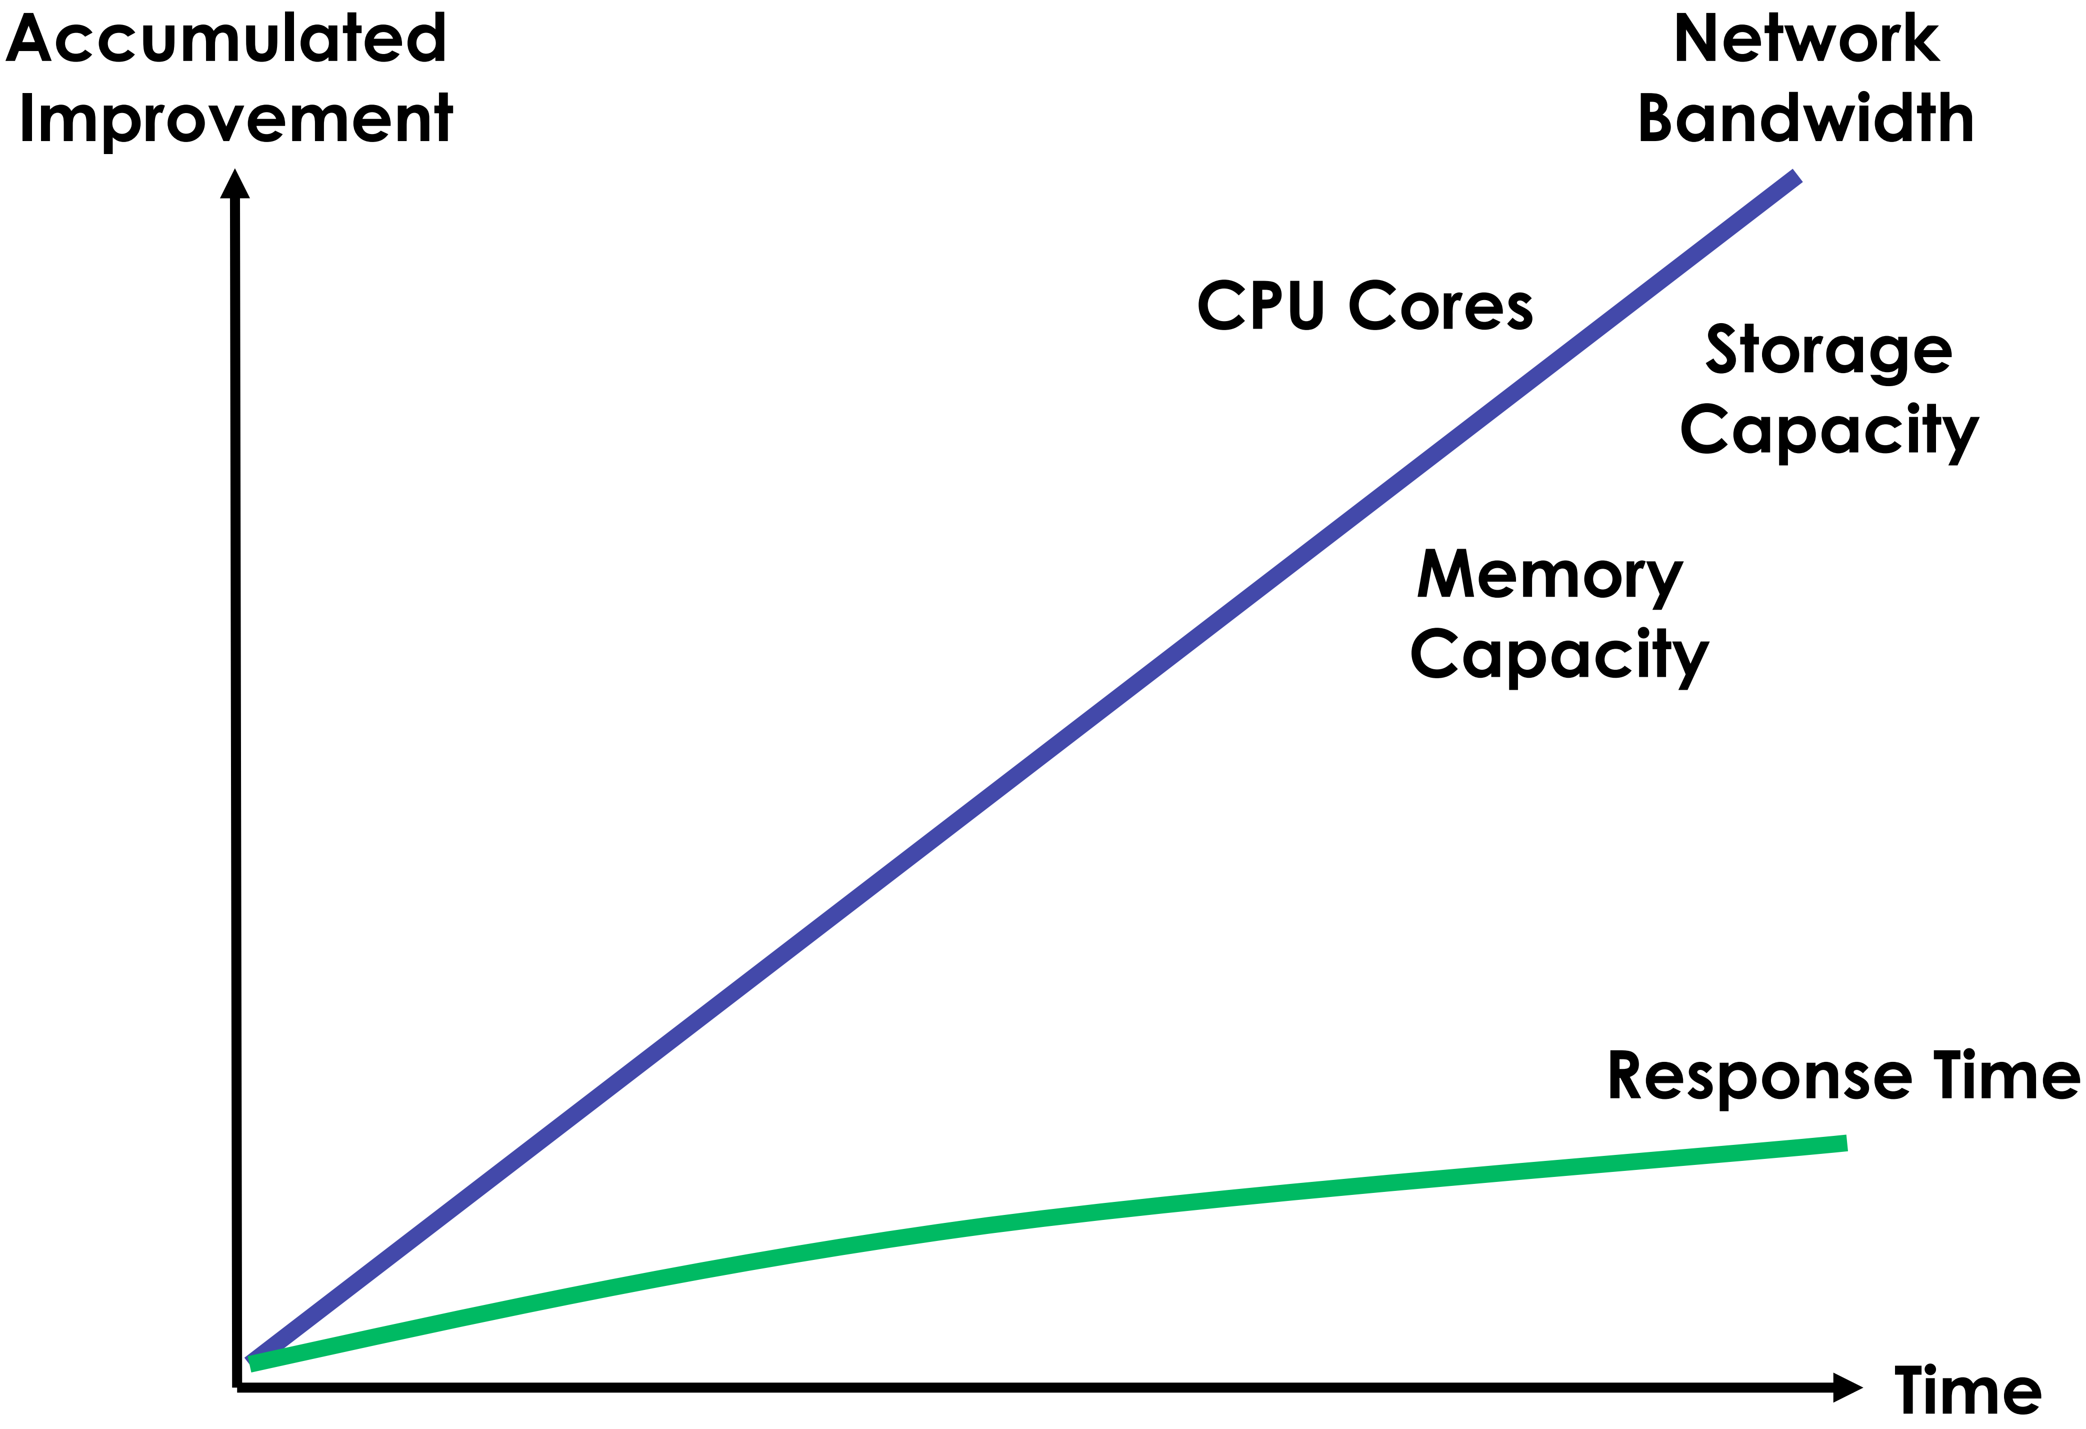
Accumulated Improvement
Network Bandwidth
CPU Cores
Storage Capacity
Memory Capacity
Response Time
Time

Year	Processor	MIPS
1974	Intel 8080	0.29
1982	Intel 286	1.28
1993	PowerPC 601	157
2003	Pentium 4 Extreme	9,726
2008	Intel Core i7 920 (Quad)	82,300
2011	Intel Core i7 2600K (4/8) Sandy Bridge	128,300
2014	Intel Core i7 5960x (8/16) Haswell	298,190

http://en.wikipedia.org/wiki/Instructions_per_second

Raspberry Pi 2
(Quad)

1,186 MIPS!

[*http://en.wikipedia.org/wiki/Instructions_per_second*](http://en.wikipedia.org/wiki/Instructions_per_second)

DDRSSD
PCIe - 3
100 GbE
…
OmniPath

1 thread of awesome

>

128 cores of so-so

http://blog.acolyer.org/2015/06/05/scalability-but-at-what-cost/

http://www.frankmcsherry.org/graph/scalability/cost/2015/01/15/COST.html

Misconception

Data needs to come to rest to be processed

Data at rest is a liability

ETL

Warehouse

Data at rest is a liability

ILM

MDM

Why would you want transient data at rest?

But what about sort?

…REALLY?!

☑ *You can’t escape the Math*

… and you <u>don't</u> want to

The Math will guide you to a solution

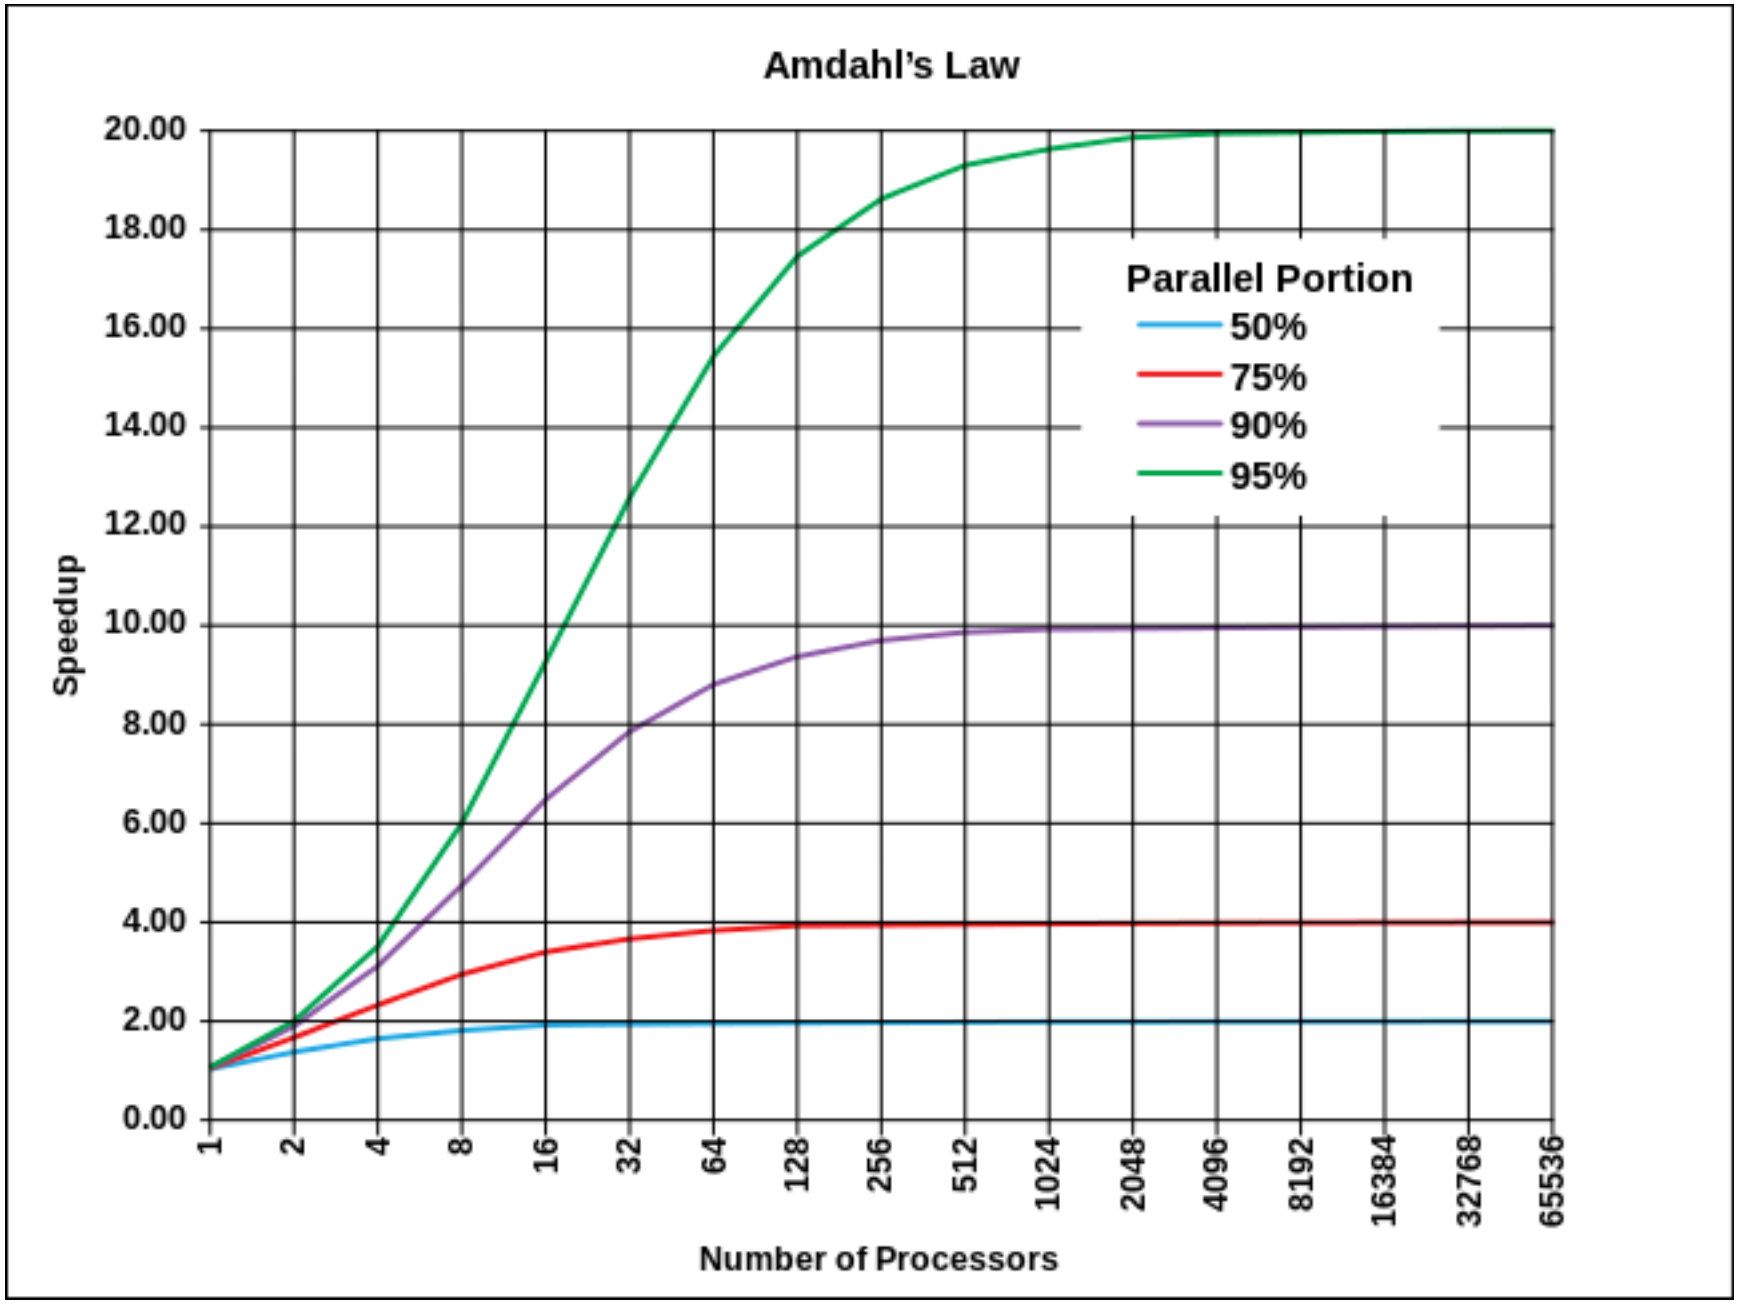

"AmdahlsLaw" by Daniels220 at English Wikipedia - Own work based on: File:AmdahlsLaw.png. Licensed under CC BY-SA 3.0 via Wikimedia Commons

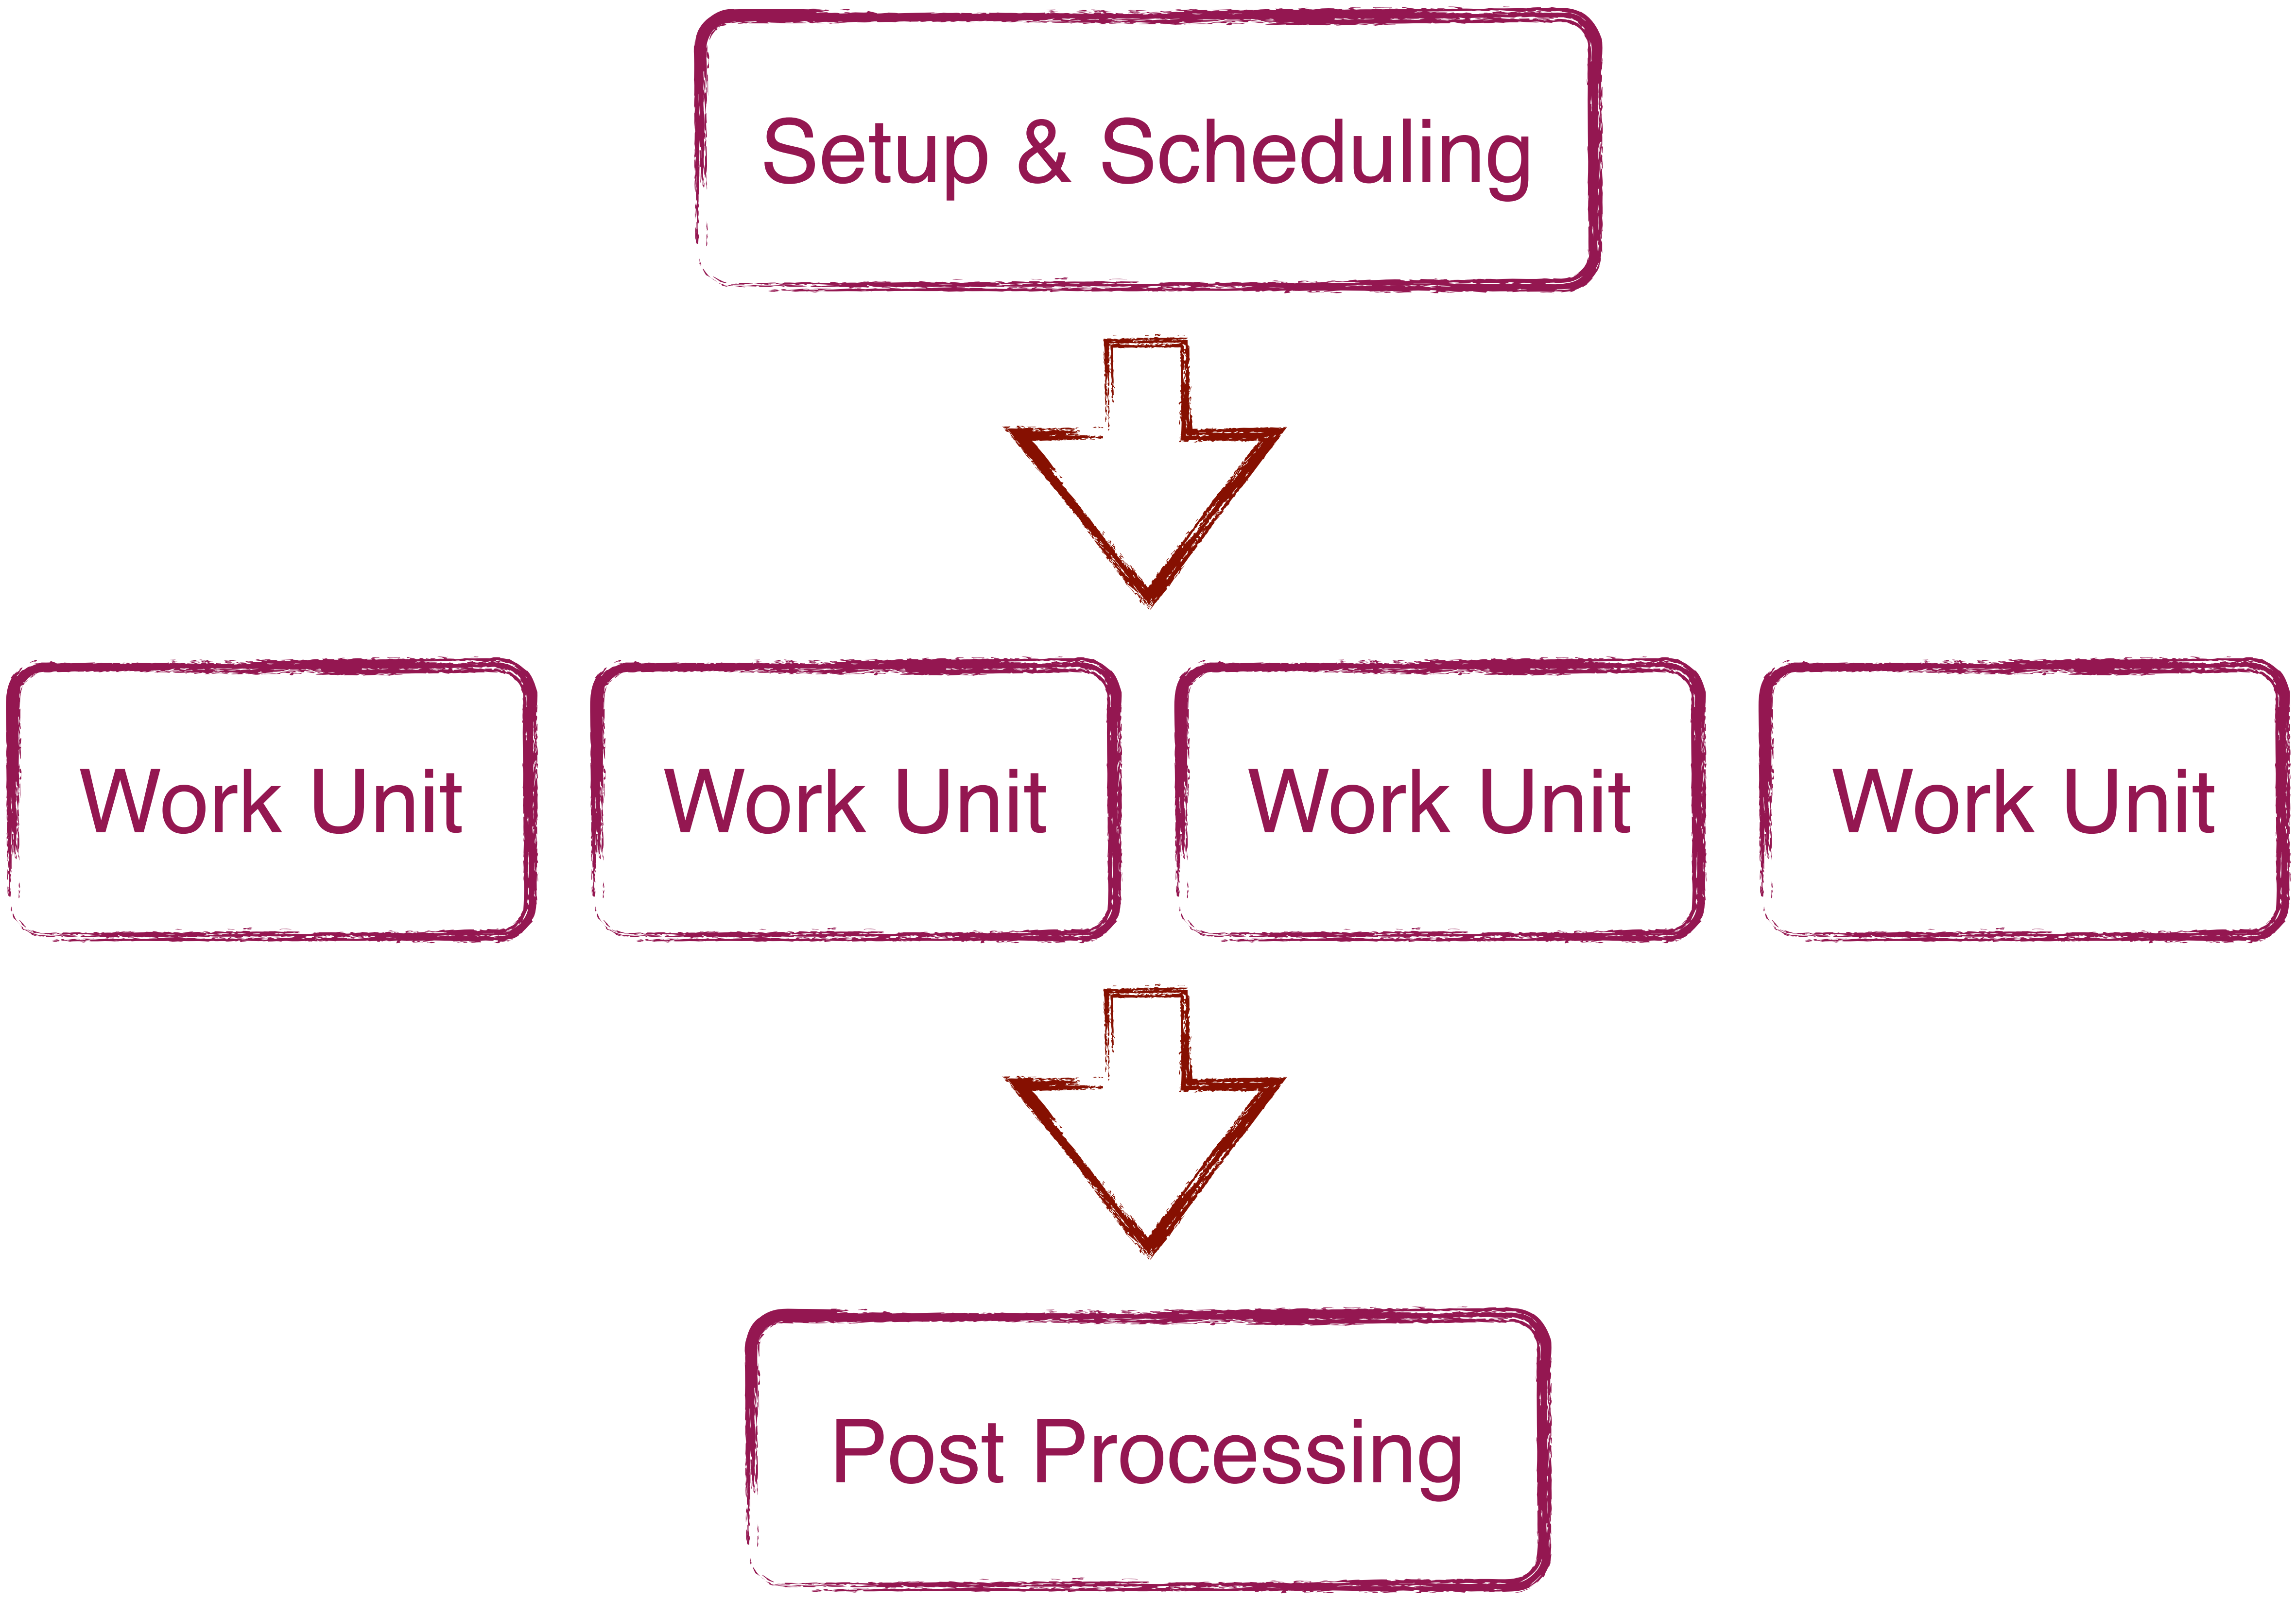

Setup & Scheduling

Work Unit

Work Unit

Work Unit

Work Unit

Post Processing

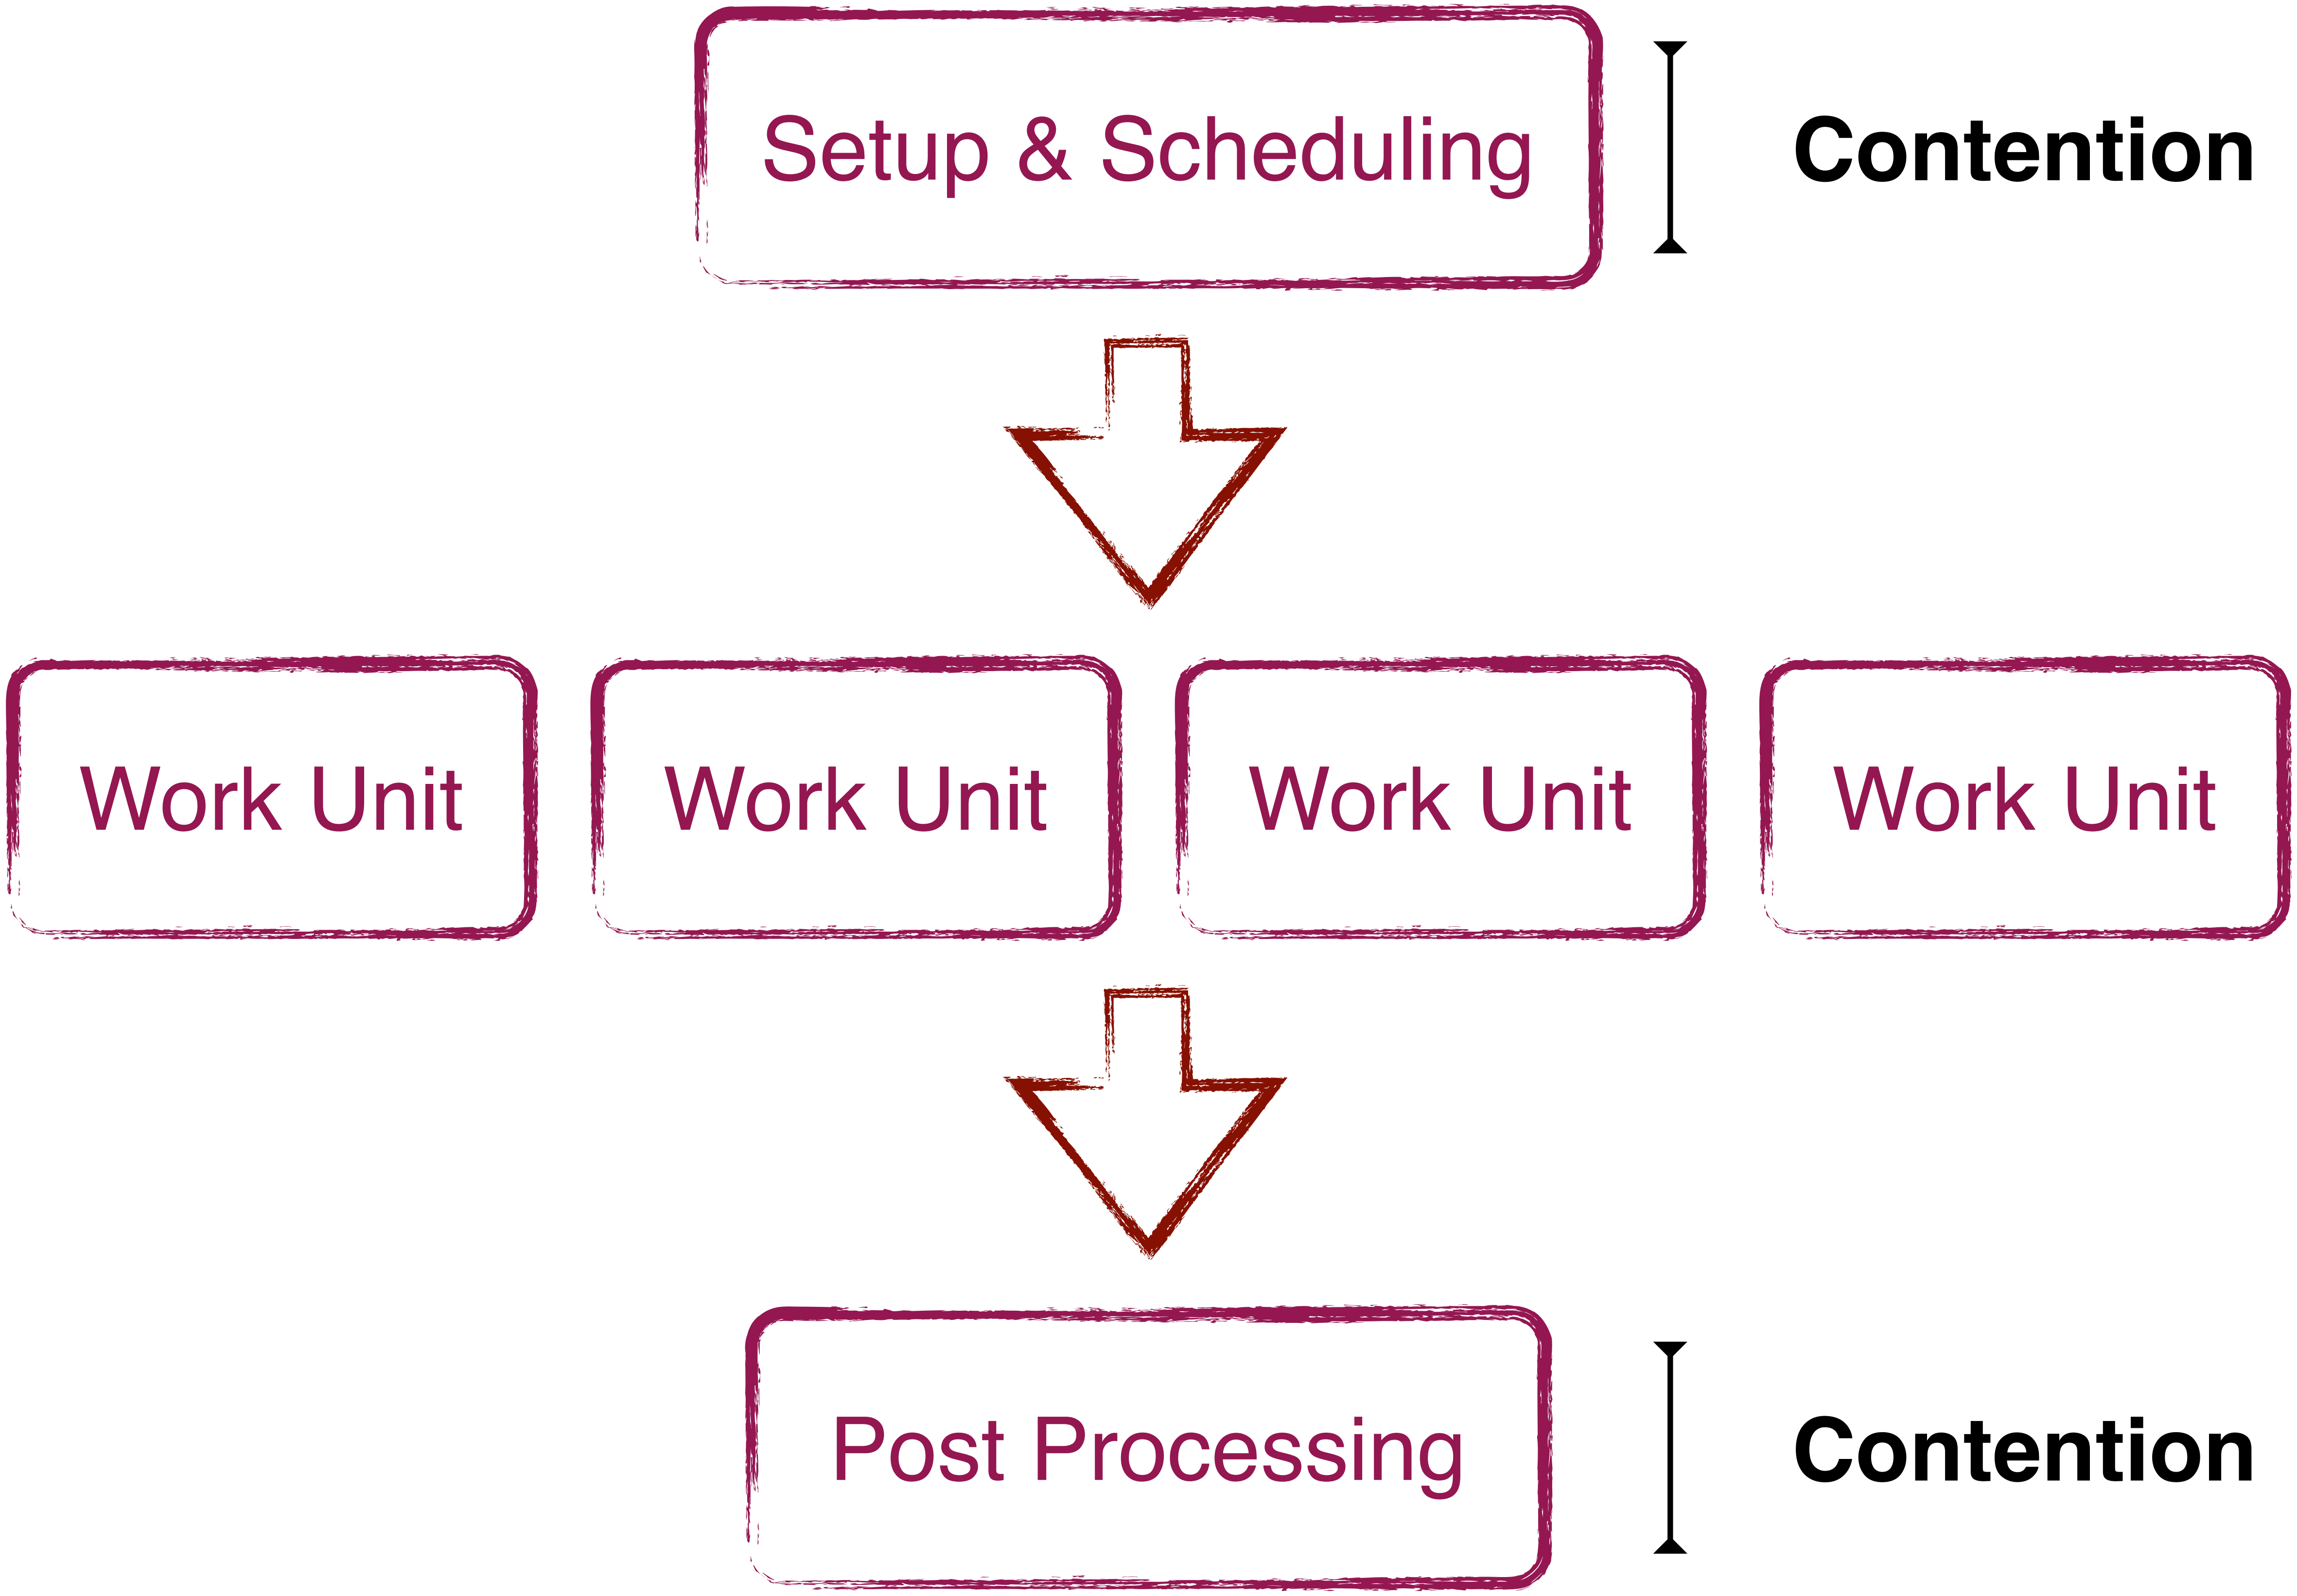

Setup & Scheduling

Contention

Work Unit

Work Unit

Work Unit

Work Unit

Post Processing

Contention

Contention isn’t the biggest enemy

Coherence is!

Universal Scalability Law

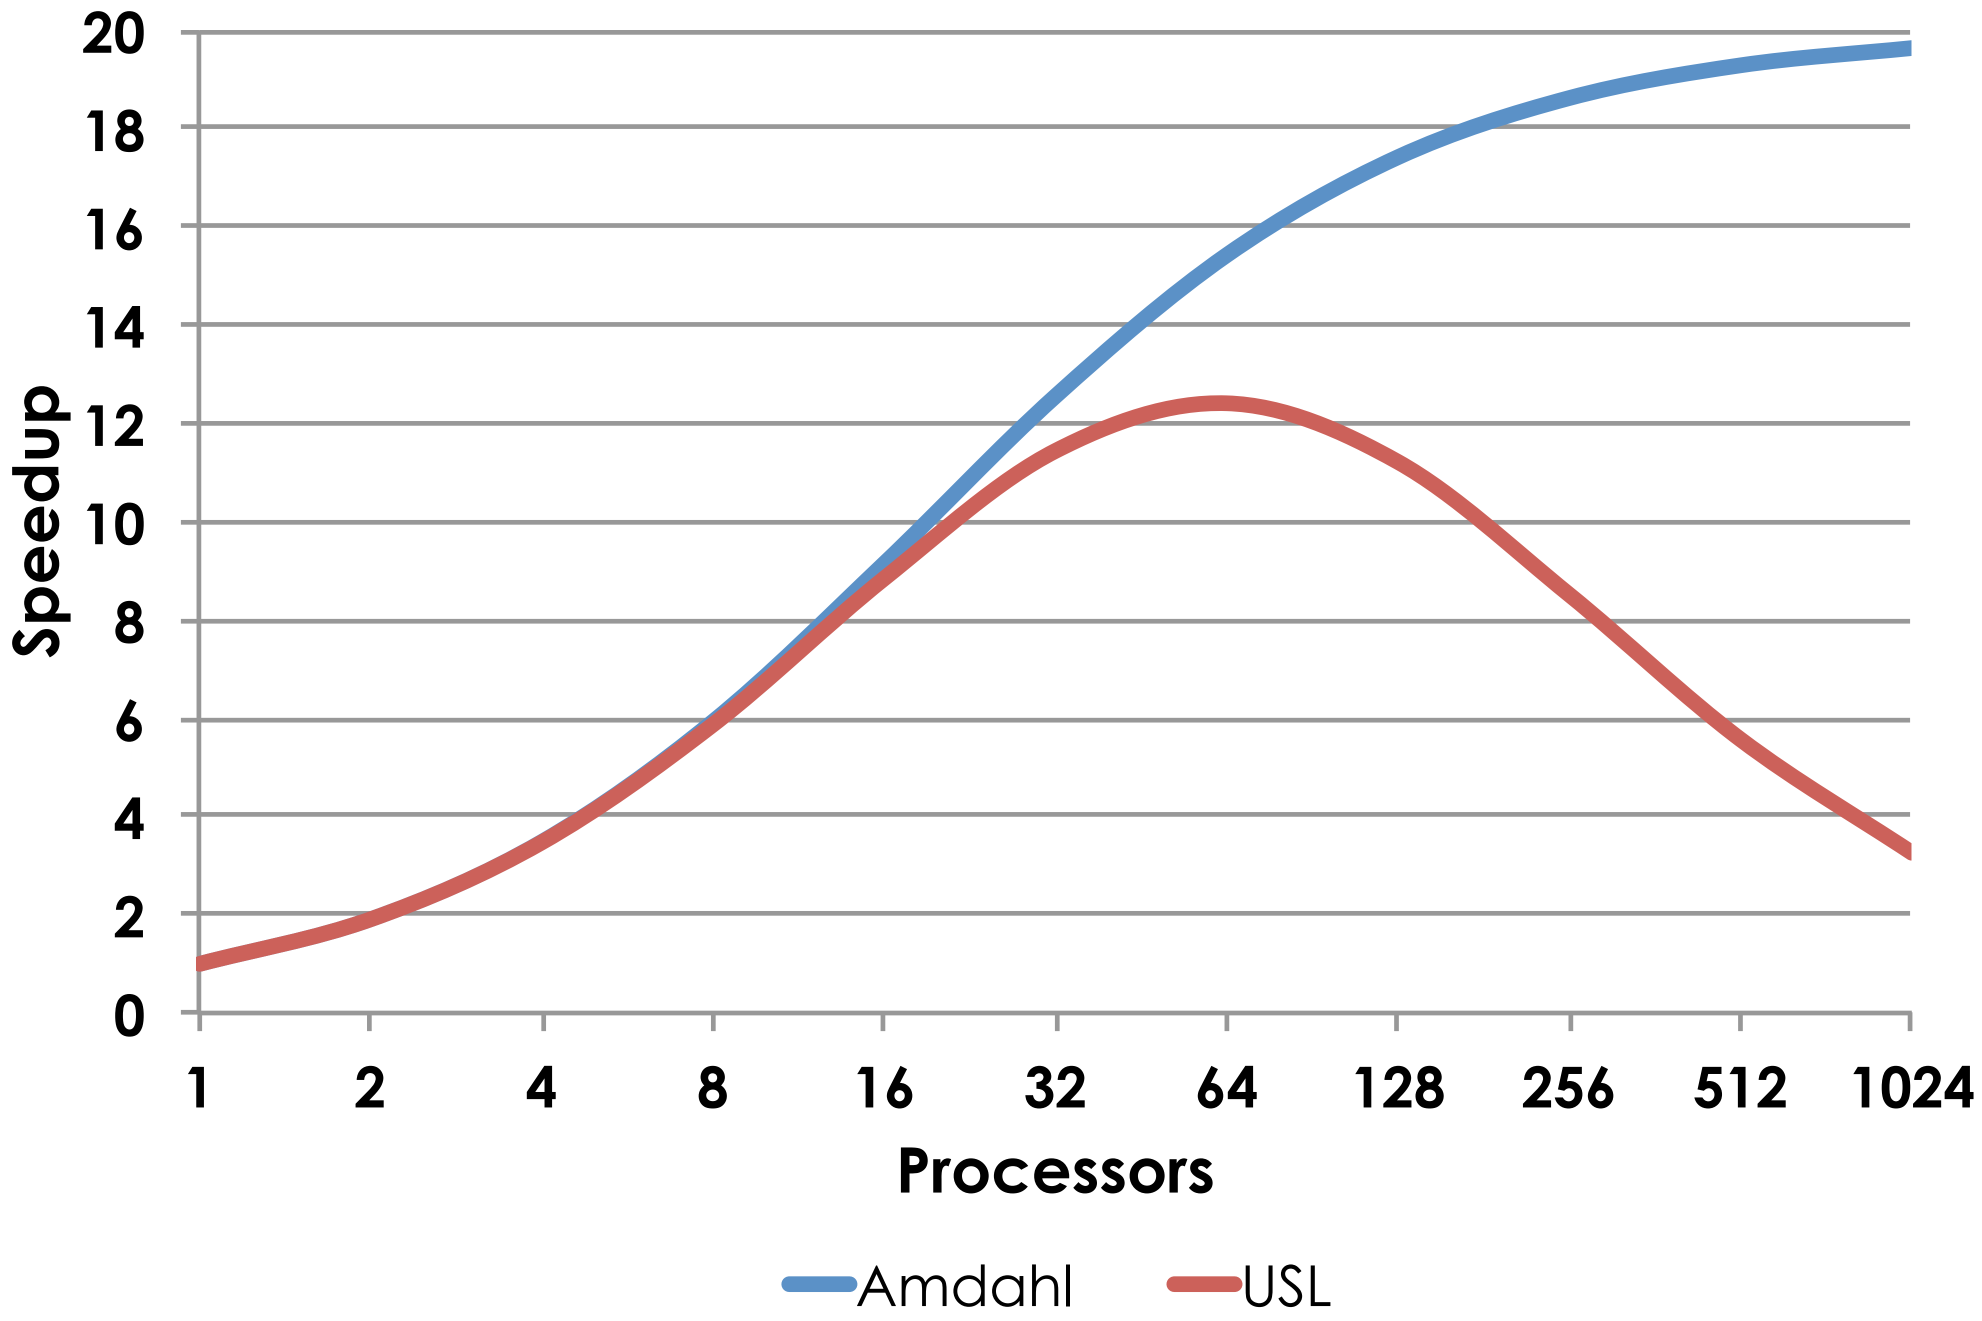

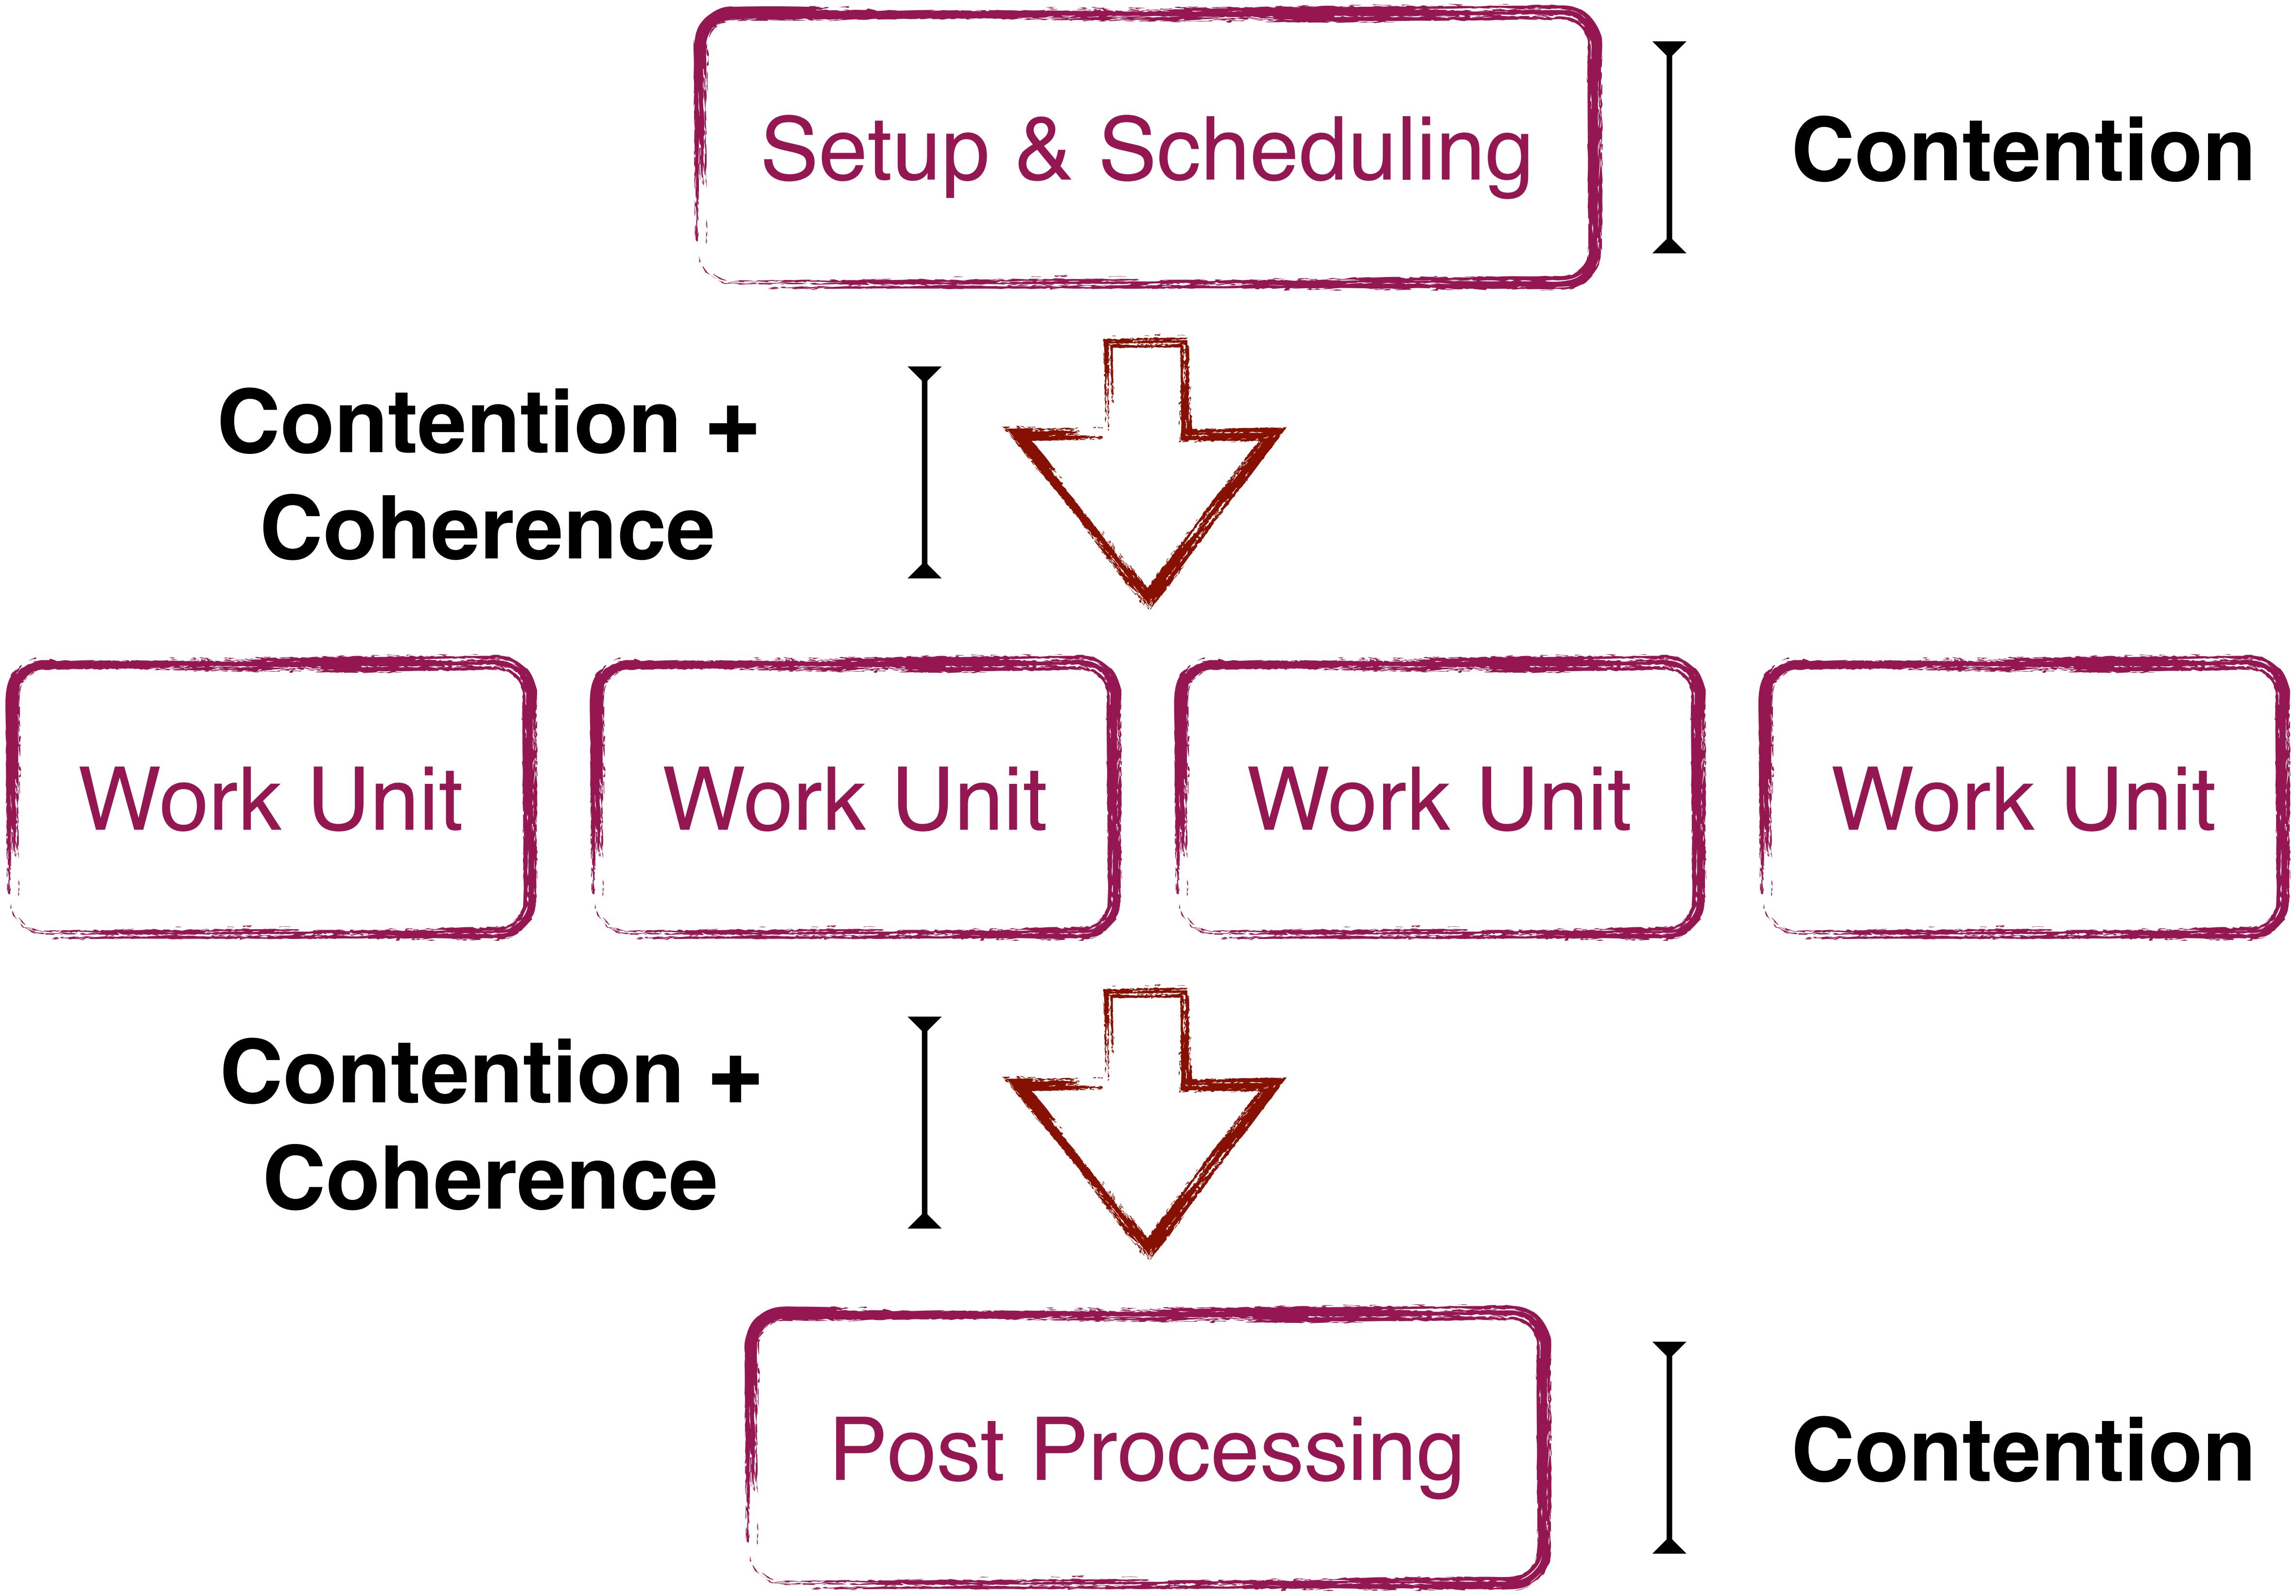
Setup & Scheduling

Contention

Contention +
Coherence

Work Unit

Work Unit

Work Unit

Work Unit

Contention +
Coherence

Post Processing

Contention

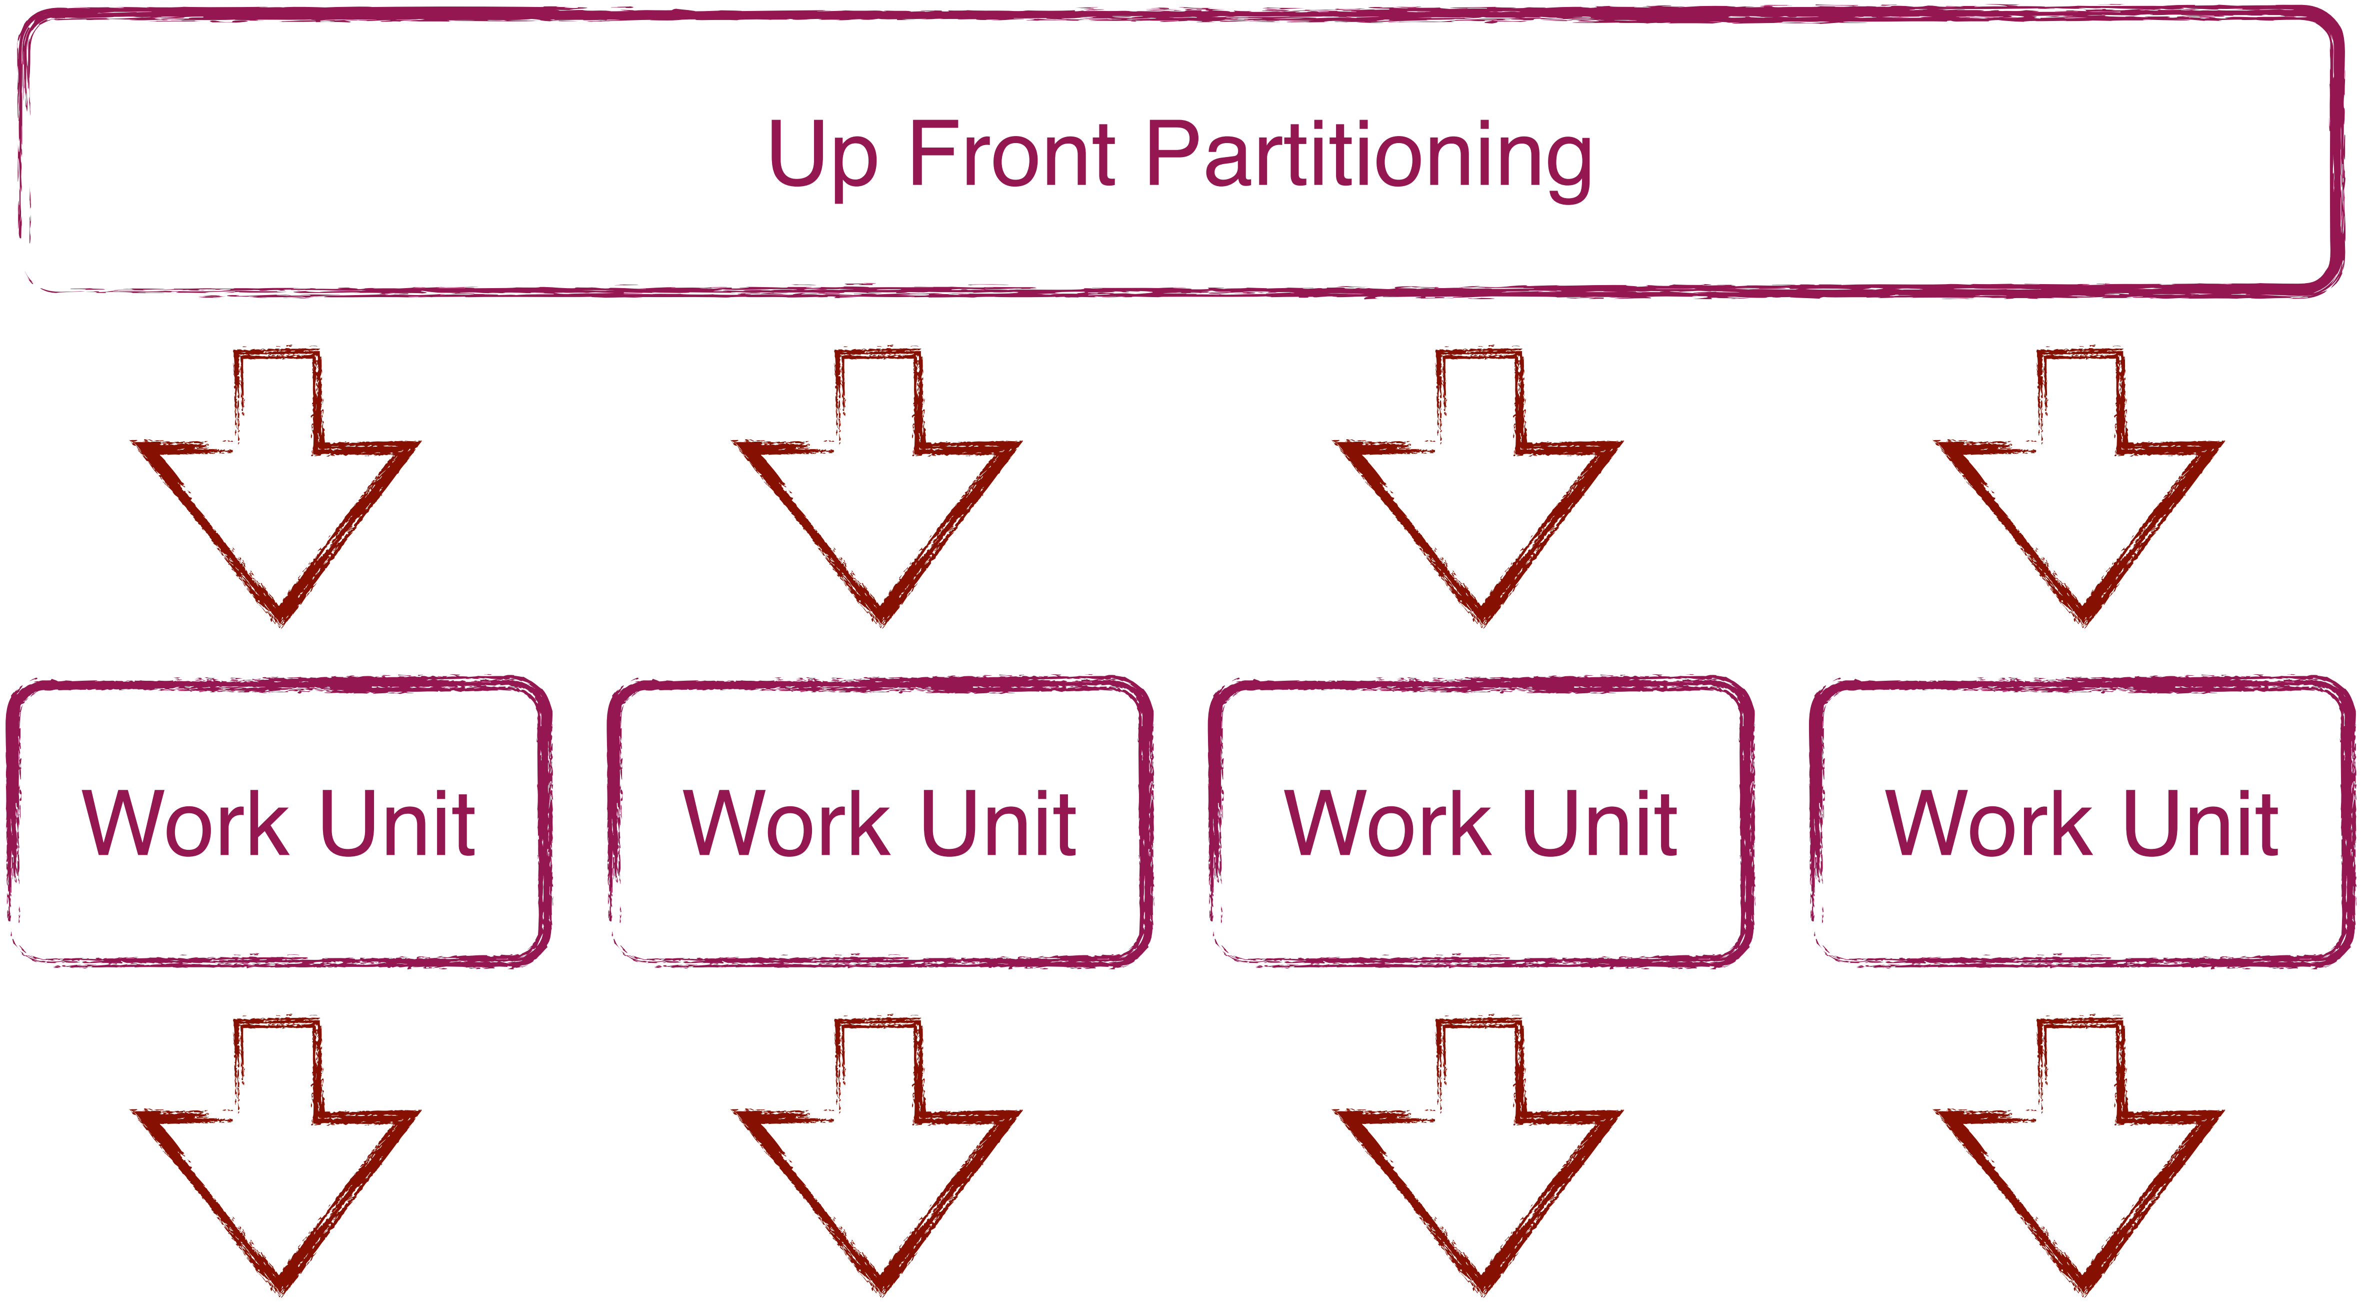

Up Front Partitioning

Work Unit

Work Unit

Work Unit

Work Unit

… and more

Queuing Theory ⭐
Complexity Theory
CAP Theorem

☑ *Technologies & Techniques*

The Essence of Architecture

Data Structures
Protocols of Interaction
Mechanical Sympathy

Understanding is essential

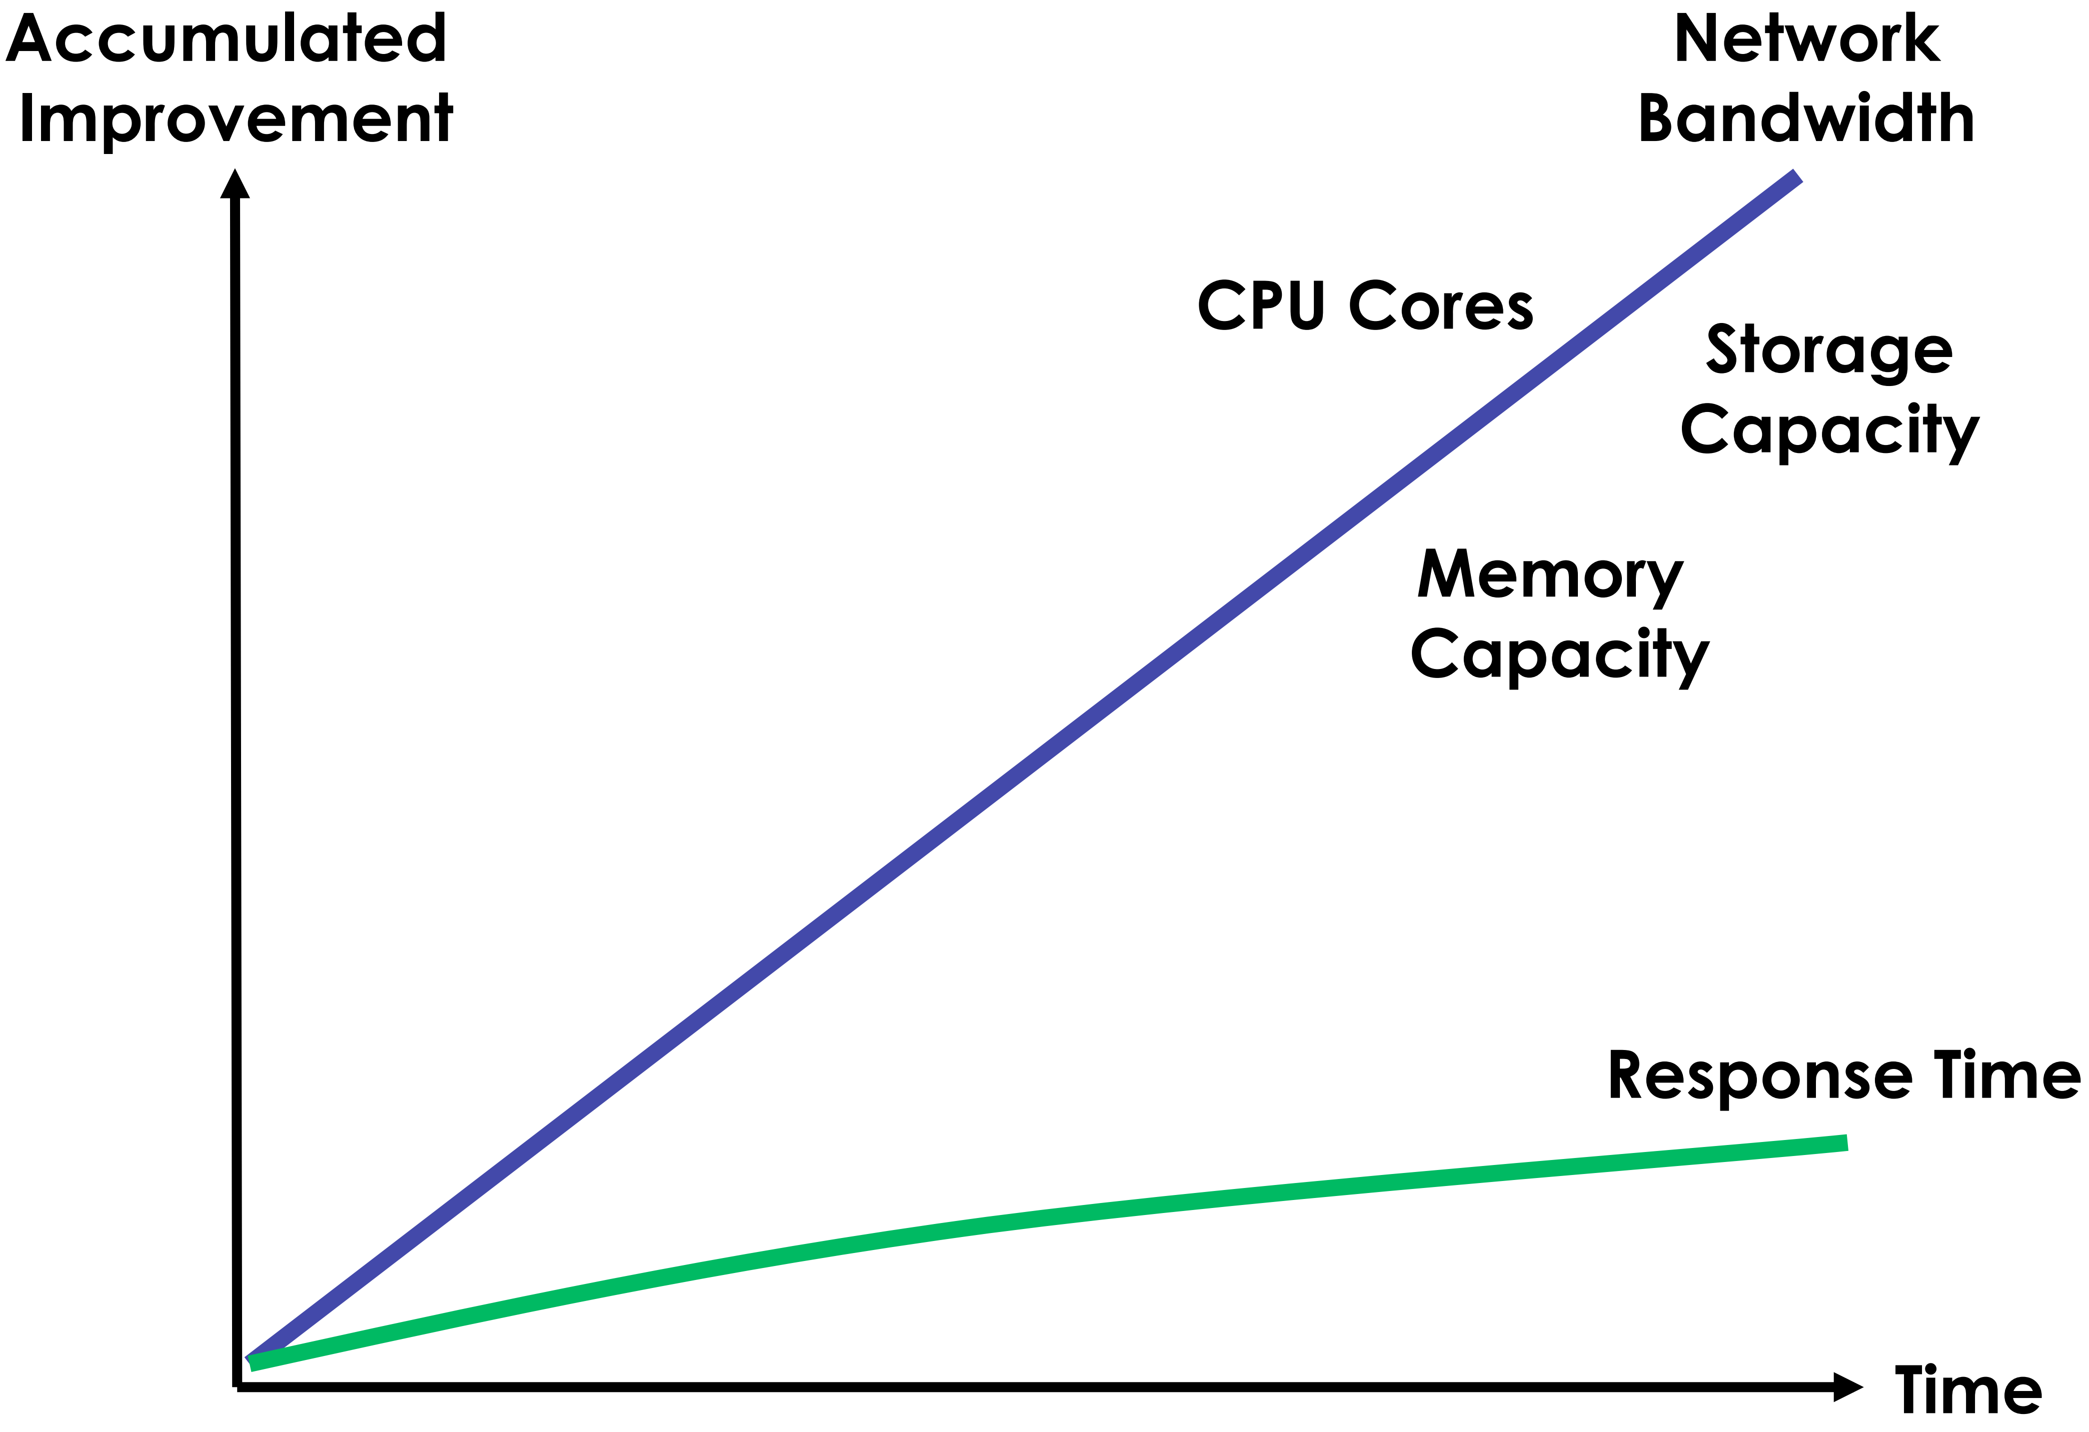
Accumulated Improvement
Network Bandwidth
CPU Cores
Storage Capacity
Memory Capacity
Response Time
Time

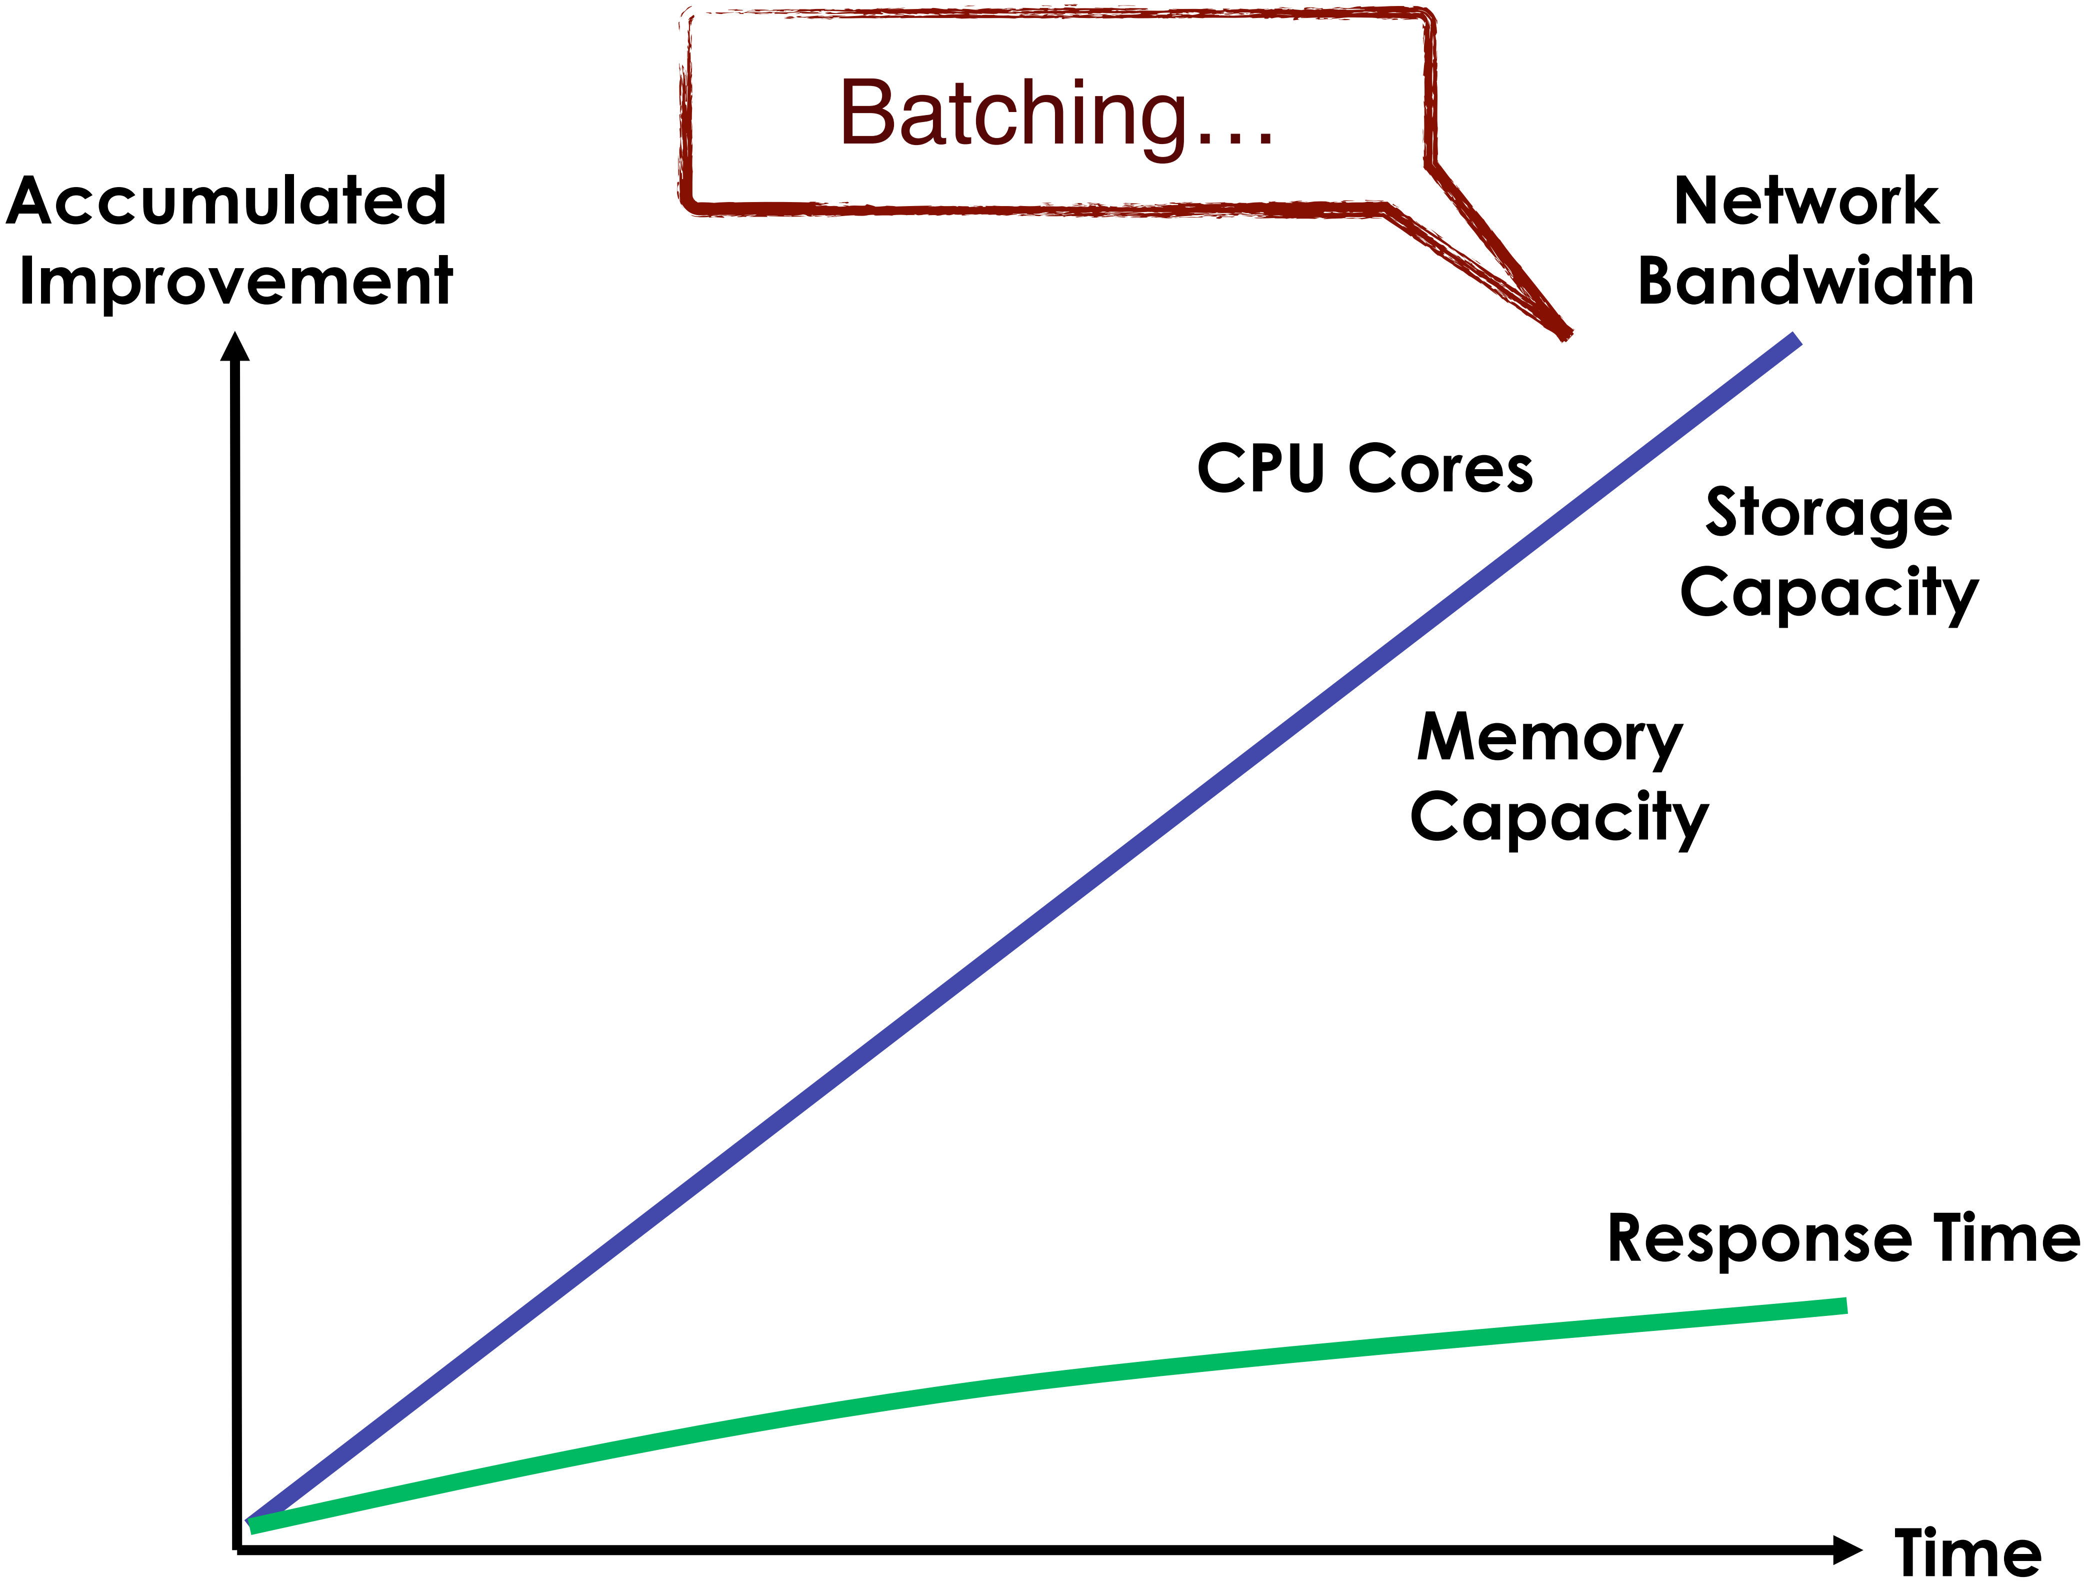

Batching…

Accumulated Improvement
Network Bandwidth
CPU Cores
Storage Capacity
Memory Capacity
Response Time
Time

Technique

Smart Batching (Natural Batching)

http://mechanical-sympathy.blogspot.com/2011/10/smart-batching.html

Resource

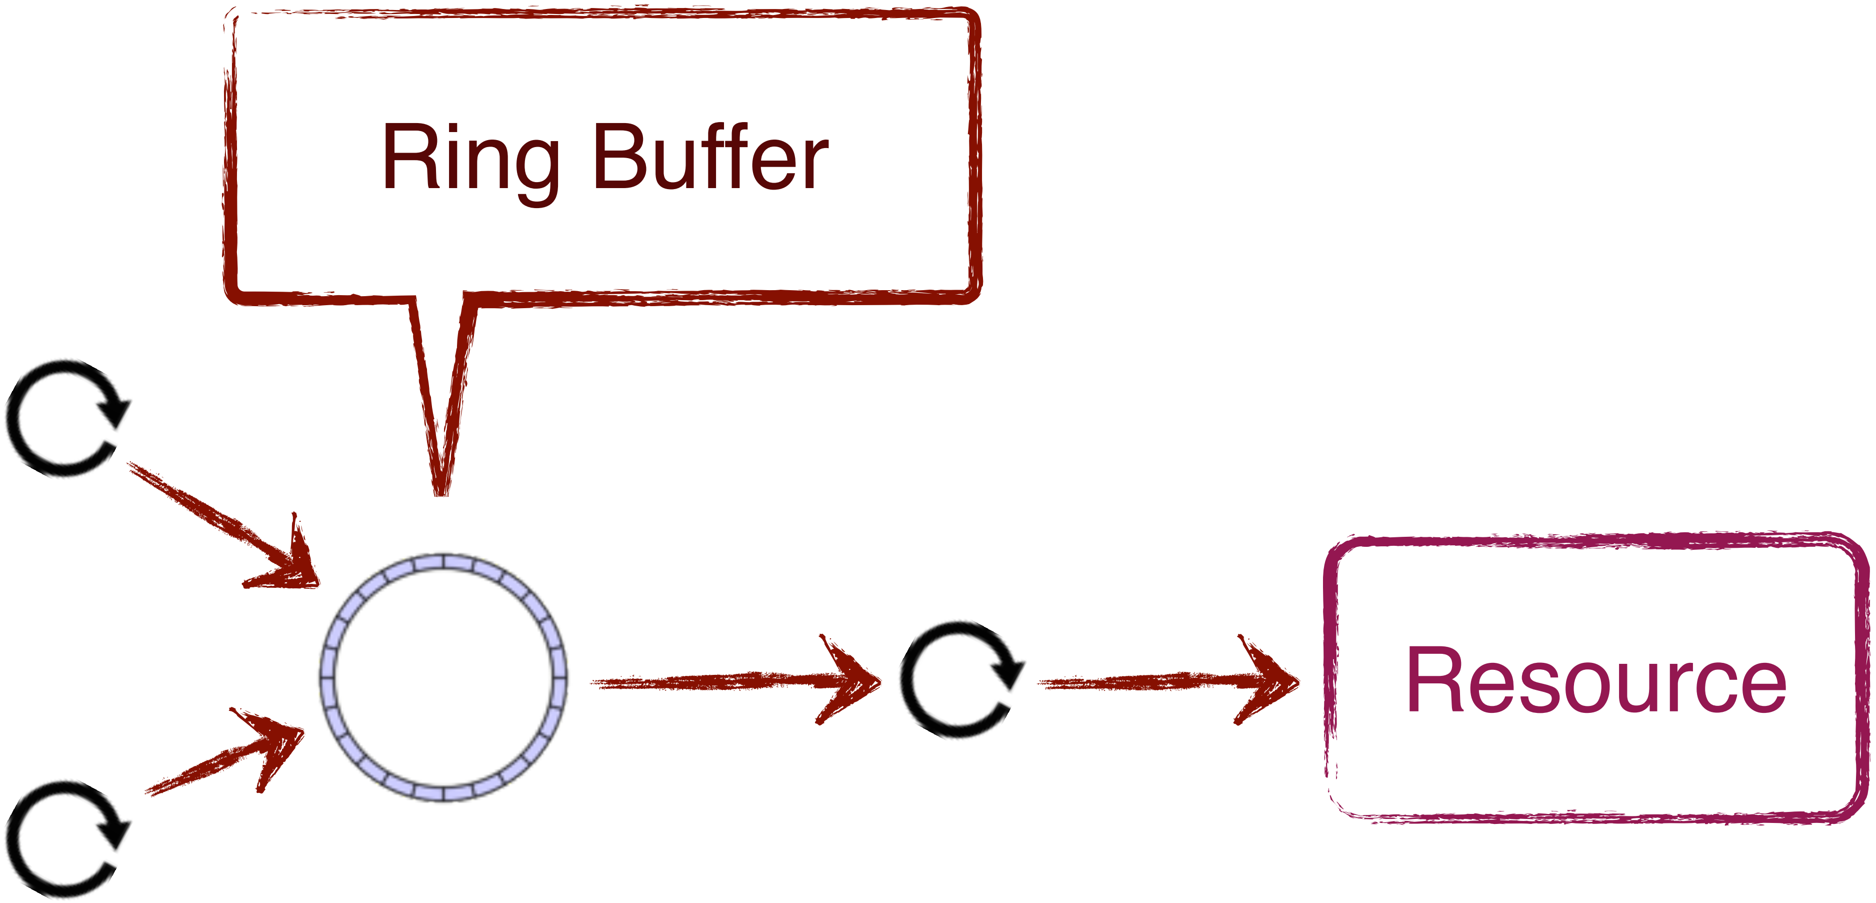

Ring Buffer

Resource

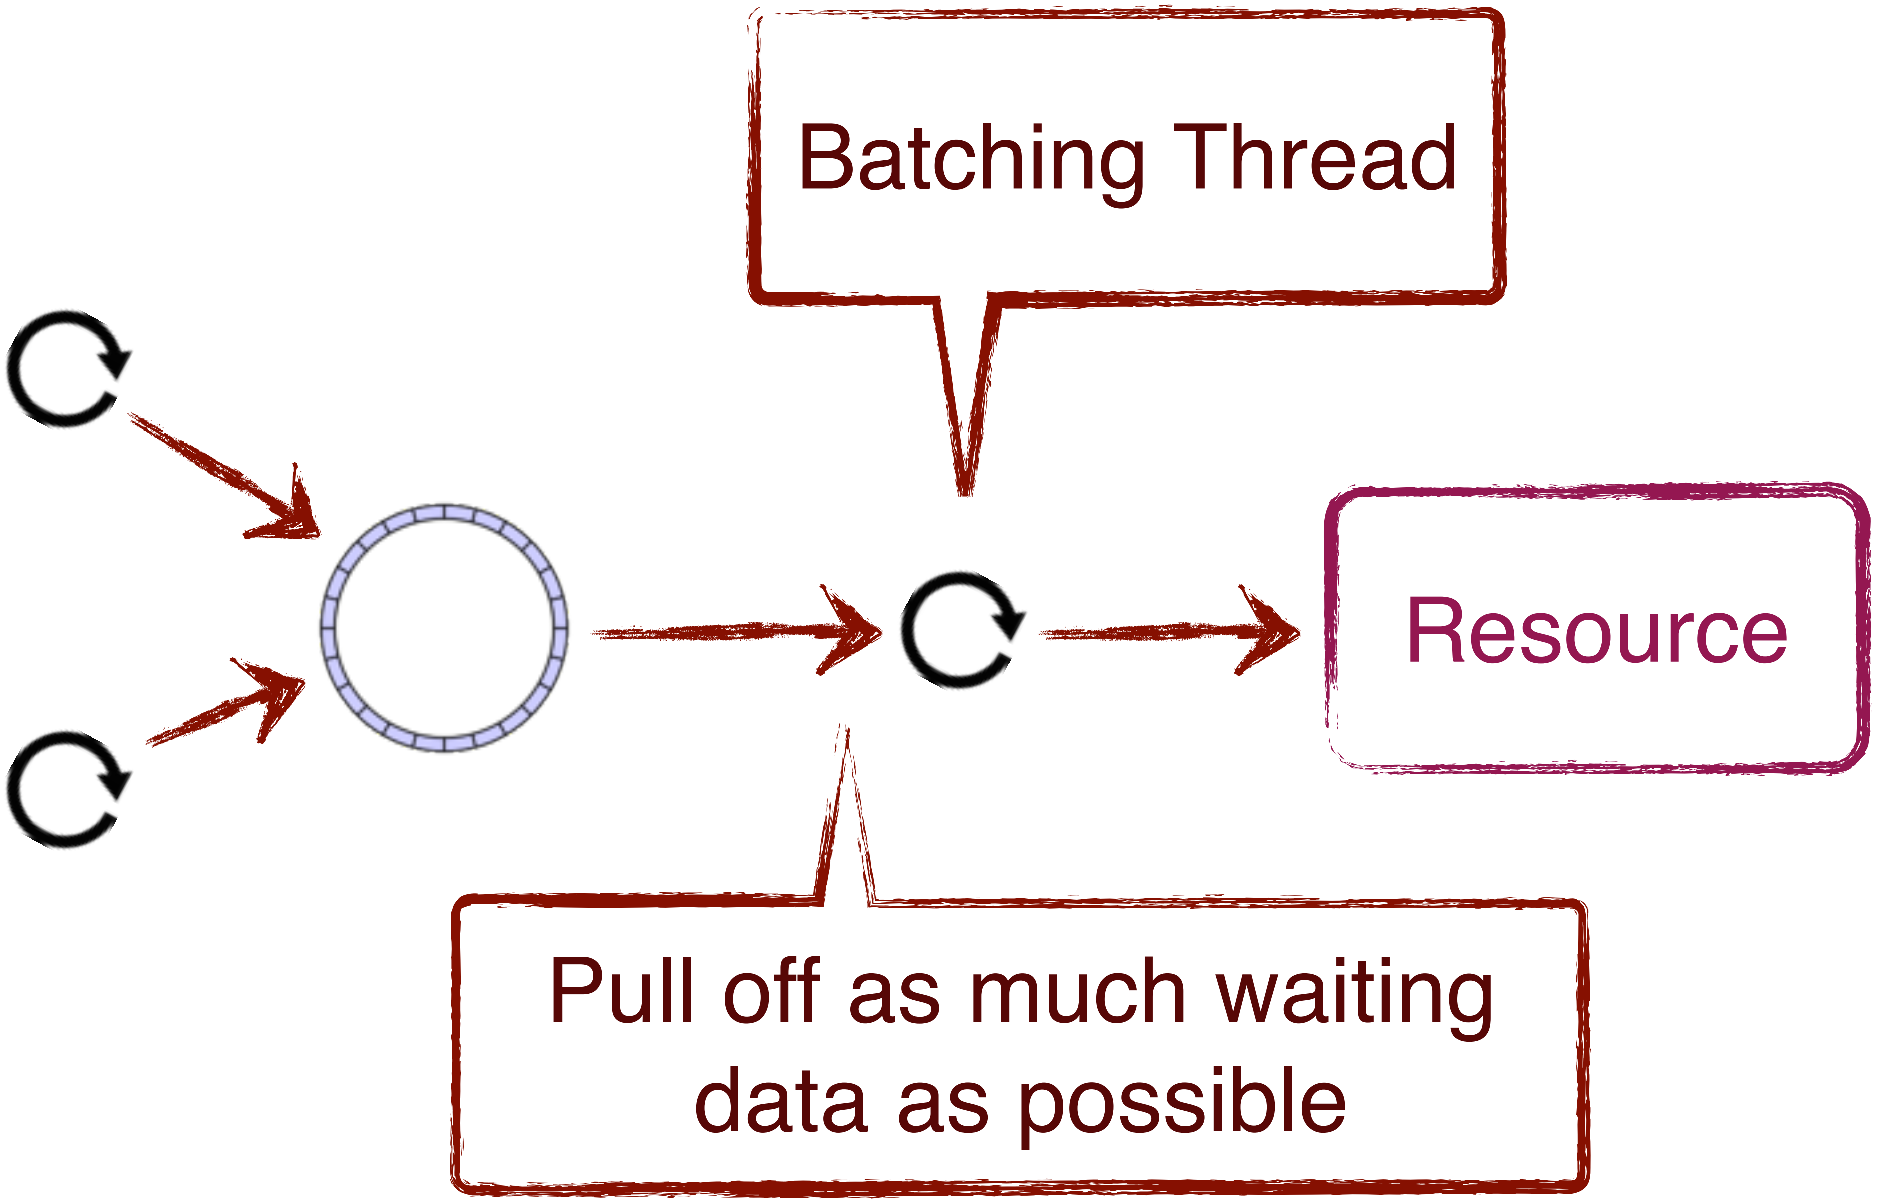

Batching Thread

Resource

Pull off as much waiting data as possible

- ☑ *Single Writer Principle*
- ☑ *Avoid Resource Contention*
- ☑ *Batching only when needed*
- ☑ *Rate Decoupling*
- ☑ *Back Pressure*

Techniques

Freedom!
Lock-Free, Wait-Free

http://en.wikipedia.org/wiki/Non-blocking_algorithm

FREEDOM!

Words Matter

Obstruction-Freedom

Partially completed operations aborted & changes made rolled back

Lock-Freedom

Individual thread may starve, but guaranteed system-wide throughput

Lock-Free is Obstruction-Free

Wait-Freedom

Starvation free and guaranteed system-wide throughput

Wait-Free is Lock-Free

These properties are awesome!

Who wouldn’t want them?

System-wide properties start at the lowest level

Essence

Just because we could take an action right now, doesn’t mean we should

PATIENCE

Because you know that someday, you'll be able to beat the #$%^ out of that cat.

Technology

CRDTs

http://en.wikipedia.org/wiki/Conflict-free_replicated_data_type

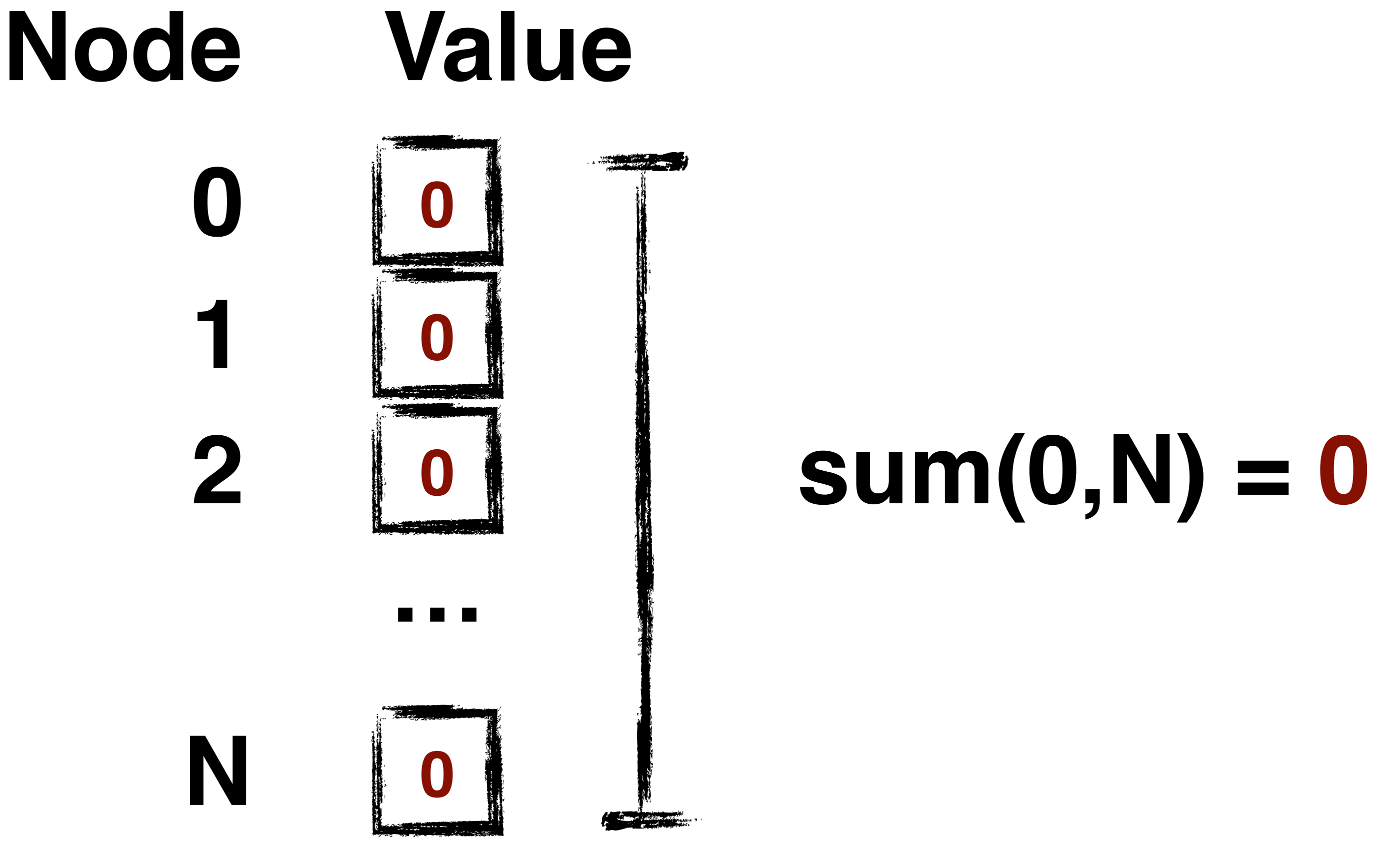

Node	Value
0	0
1	0
2	0
	…
N	0

sum(0,N) = 0

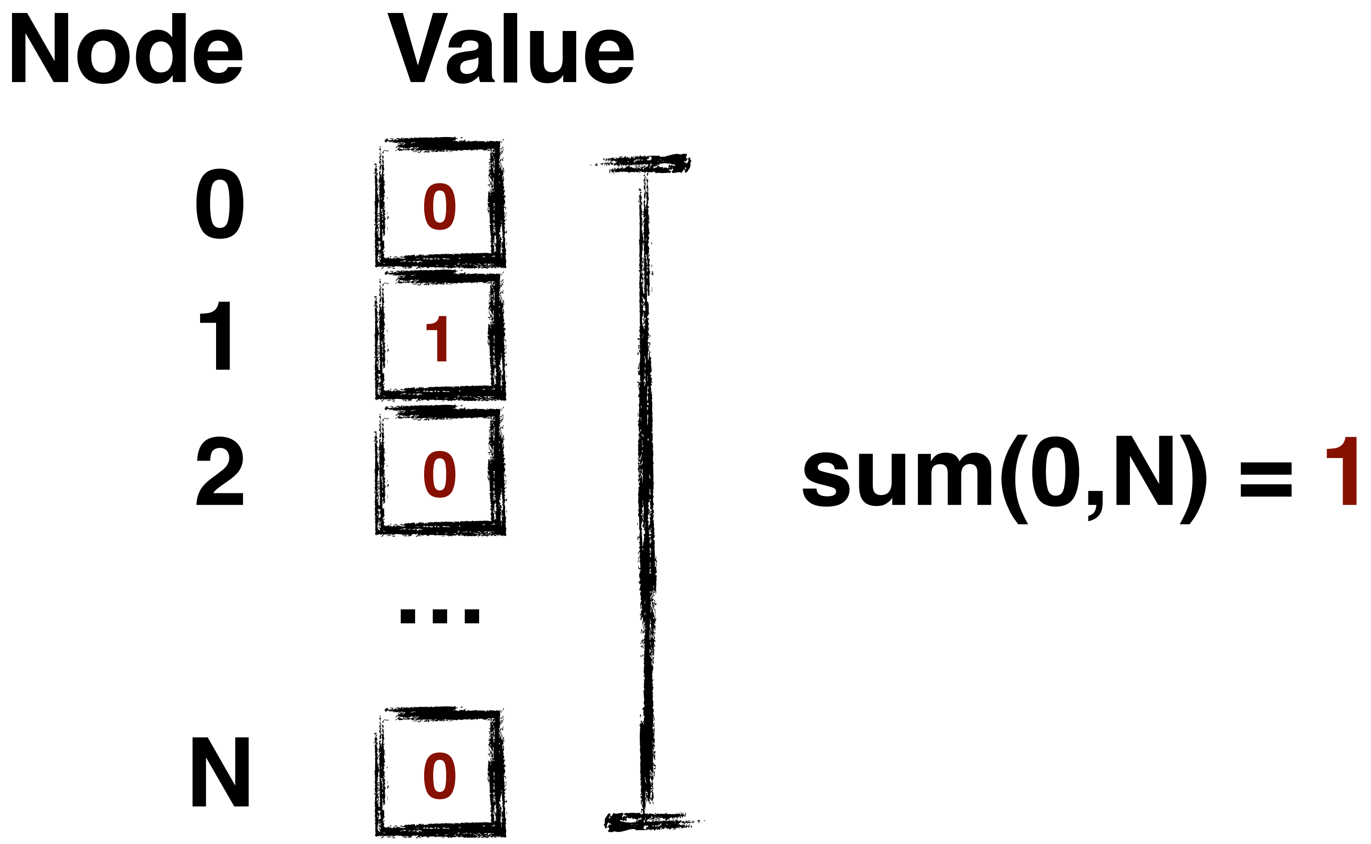

Node	Value
0	0
1	1
2	0
	...
N	0

sum(0,N) = 1

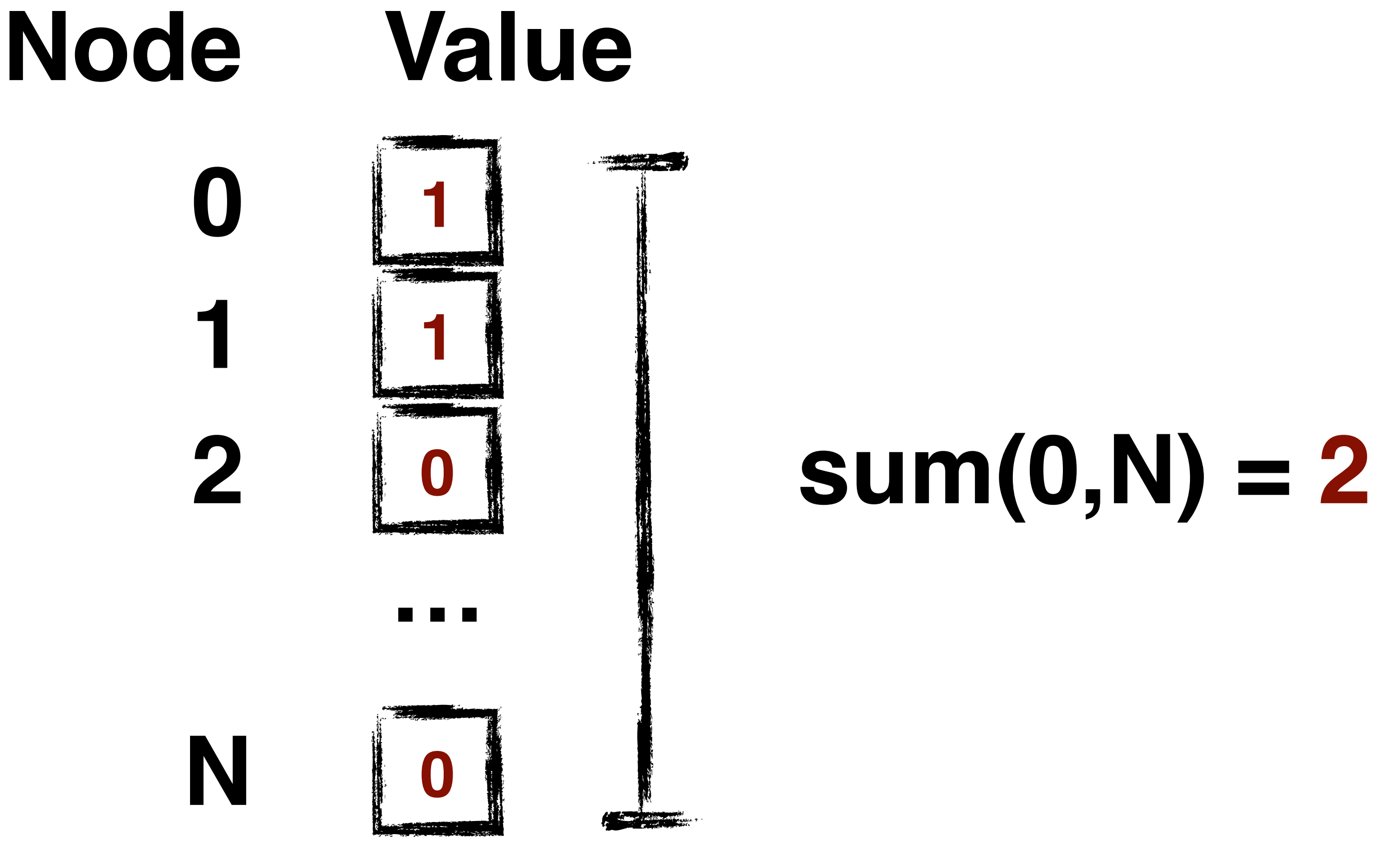

Node	Value
0	1
1	1
2	0
	…
N	0

sum(0,N) = 2

Gossip for visibility

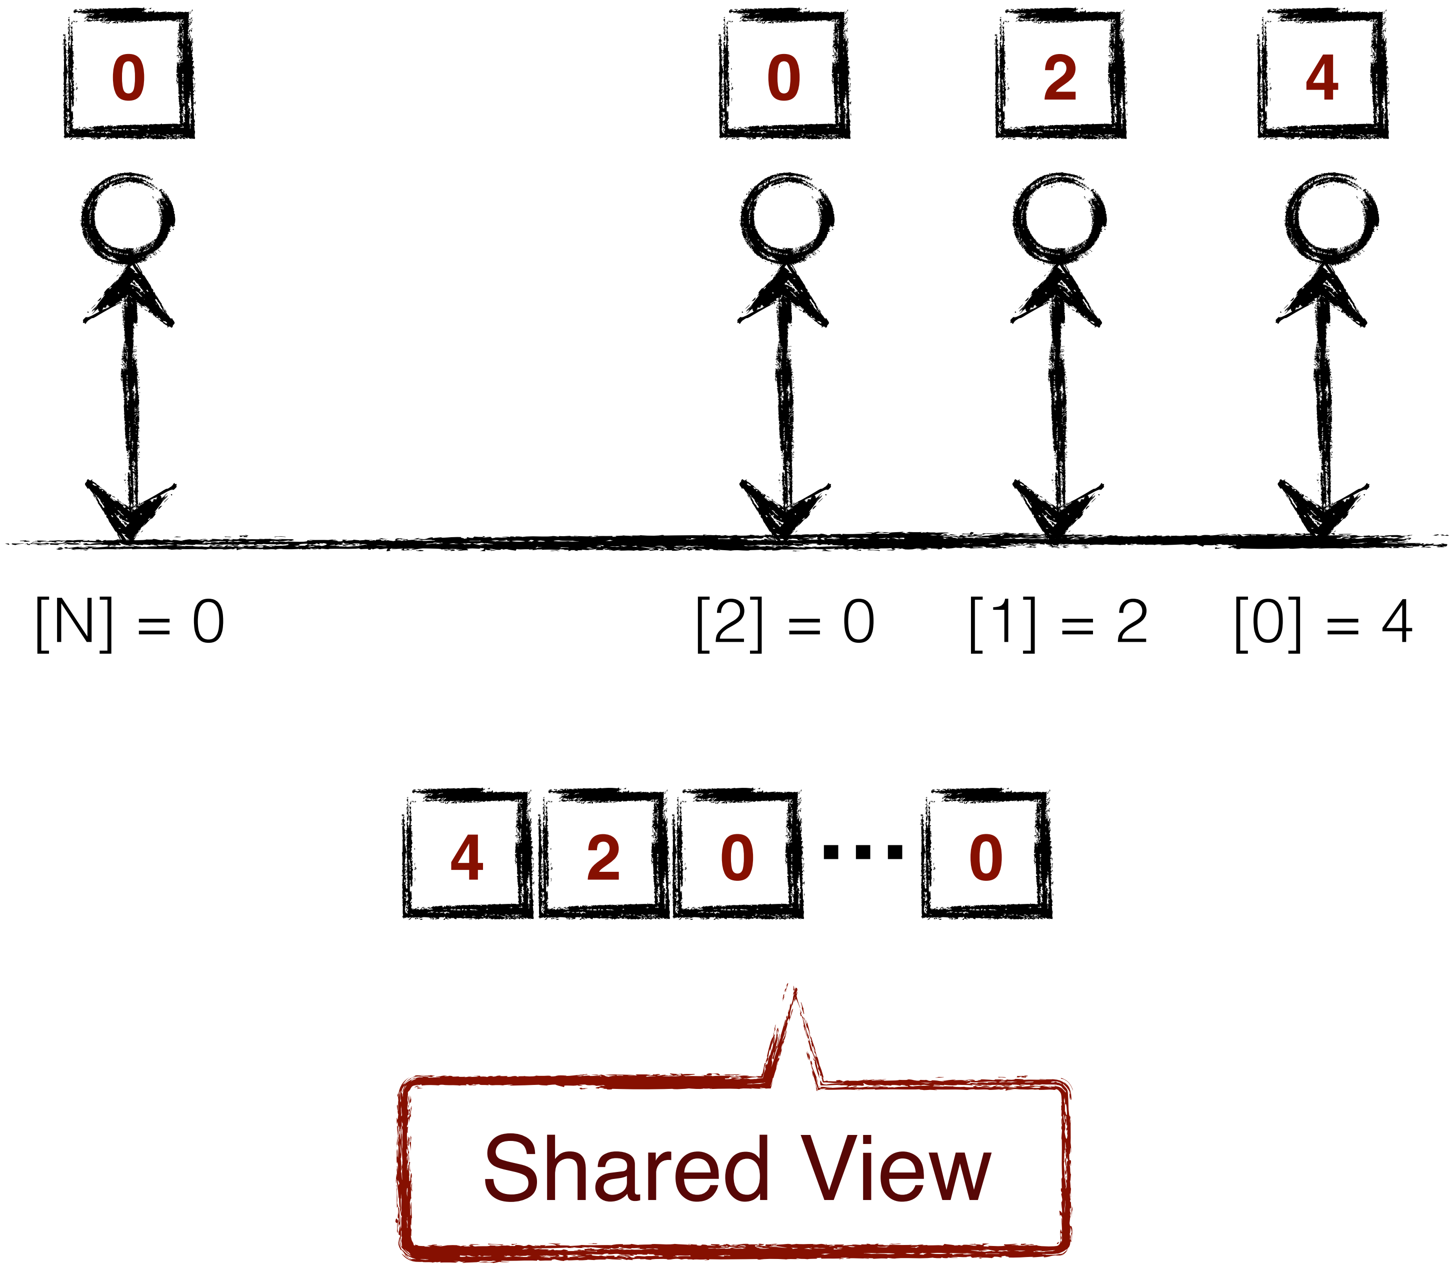

Shared View

Technology

Append-only Data Structures

https://github.com/real-logic/Aeron

Log

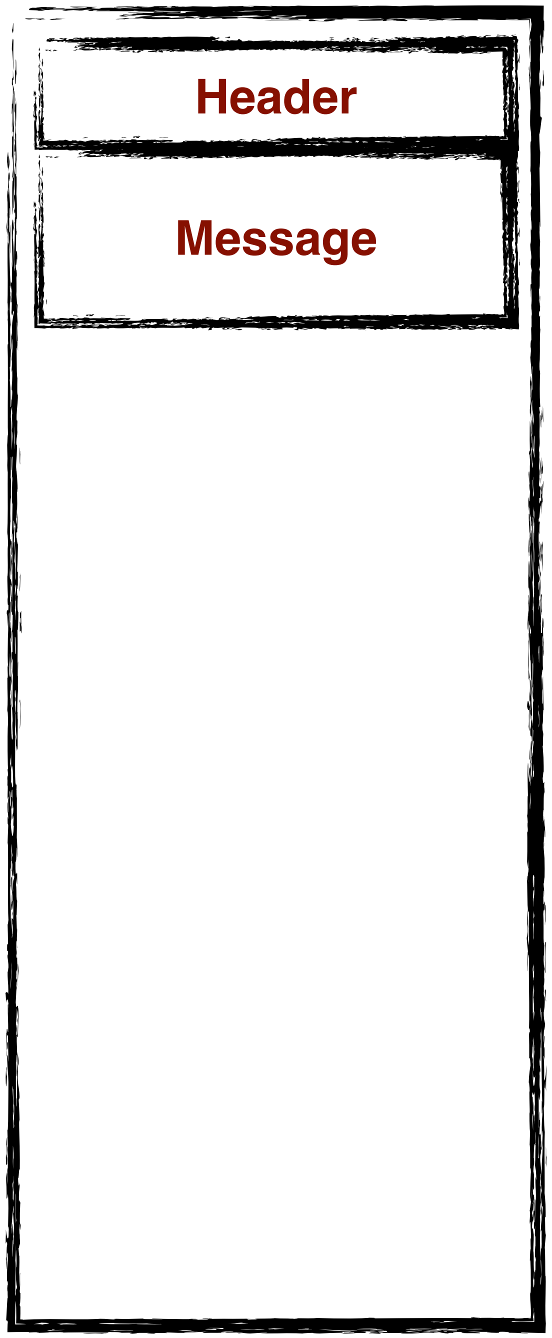

Header

Message

Log

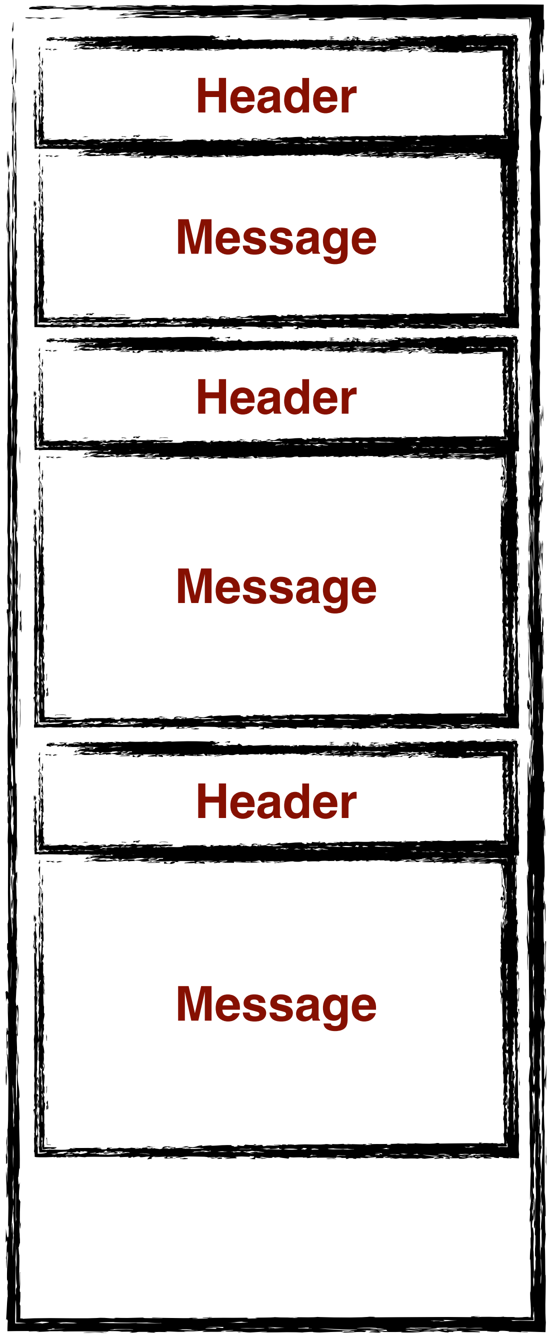

Header

Message

Header

Message

Header

Message

Efficiently Replicating an Append-only Log

What If…?

The Data Structure could be directly sent to the “network”?

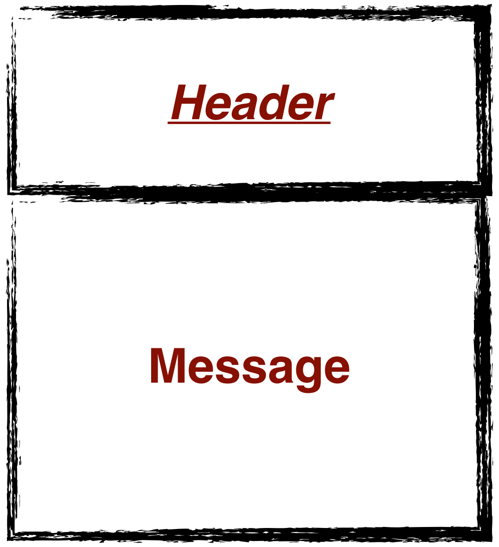

Header

Message

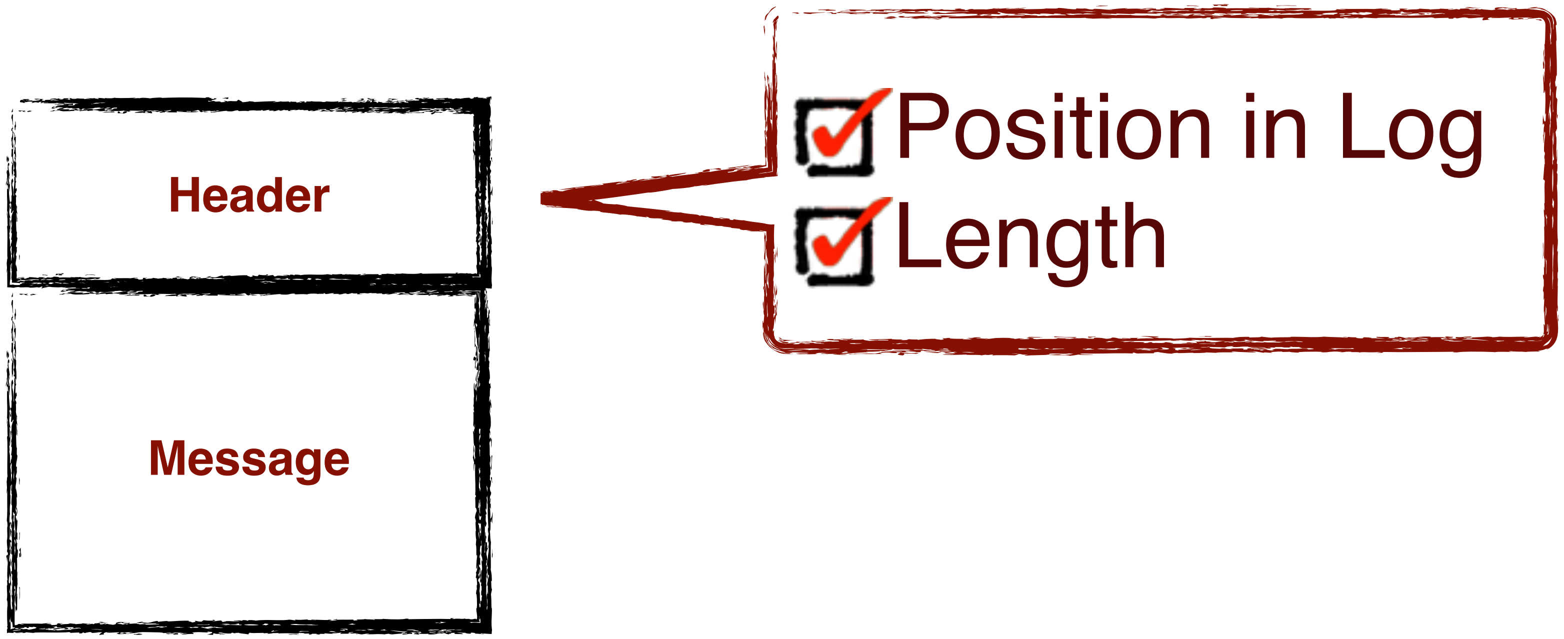

Header

Message

☑ Position in Log

☑ Length

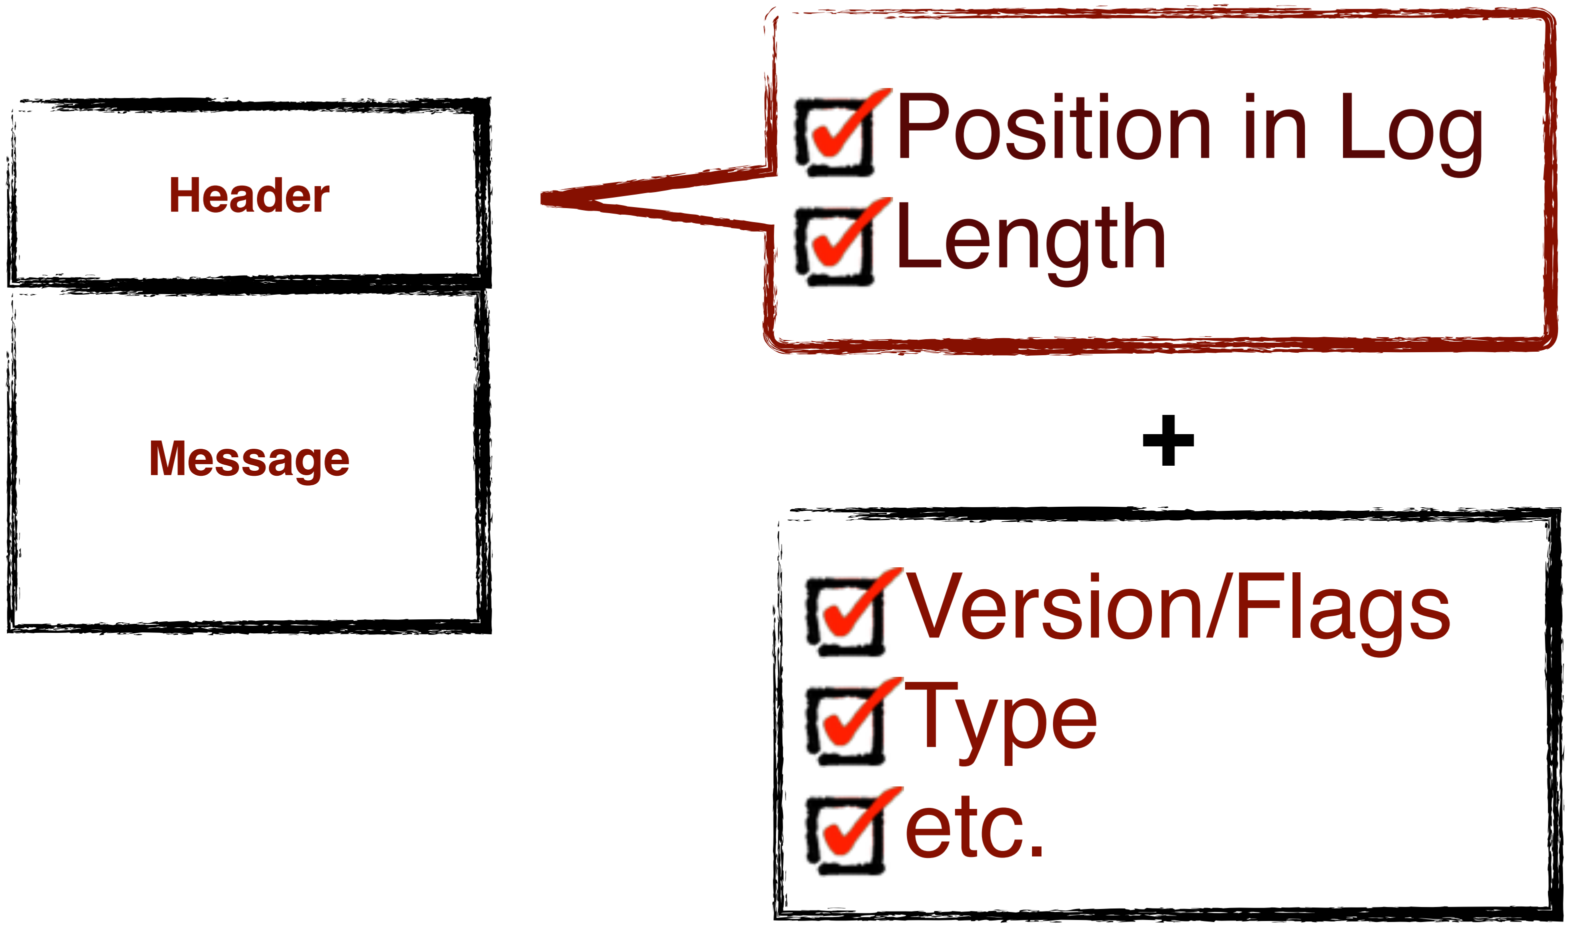

Header

Message

☑ Position in Log
☑ Length

+

☑ Version/Flags
☑ Type
☑ etc.

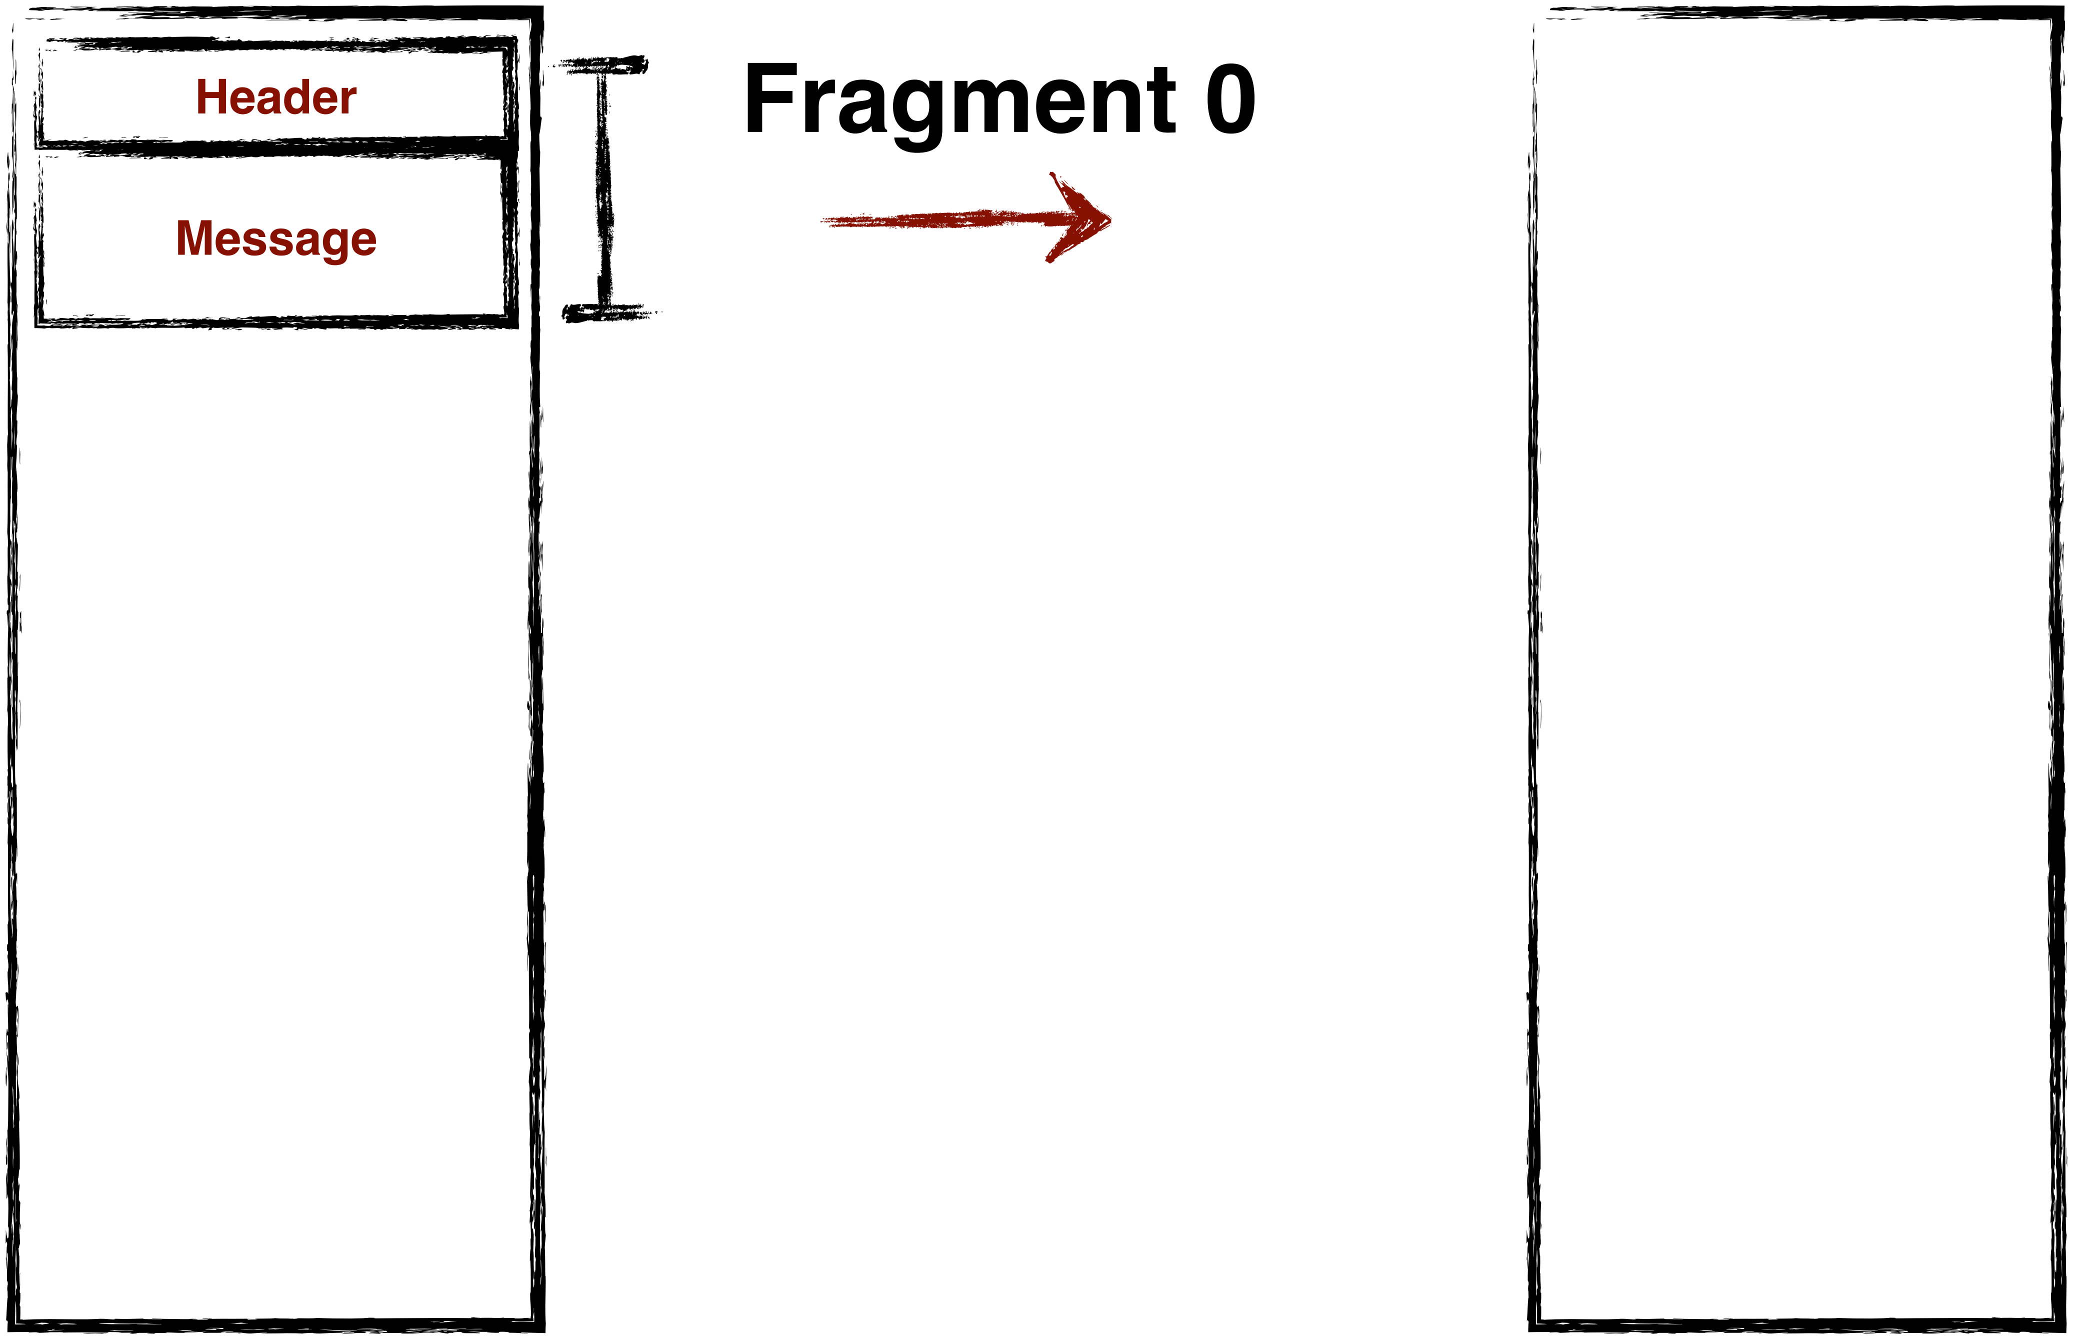

Header

Message

Fragment 0

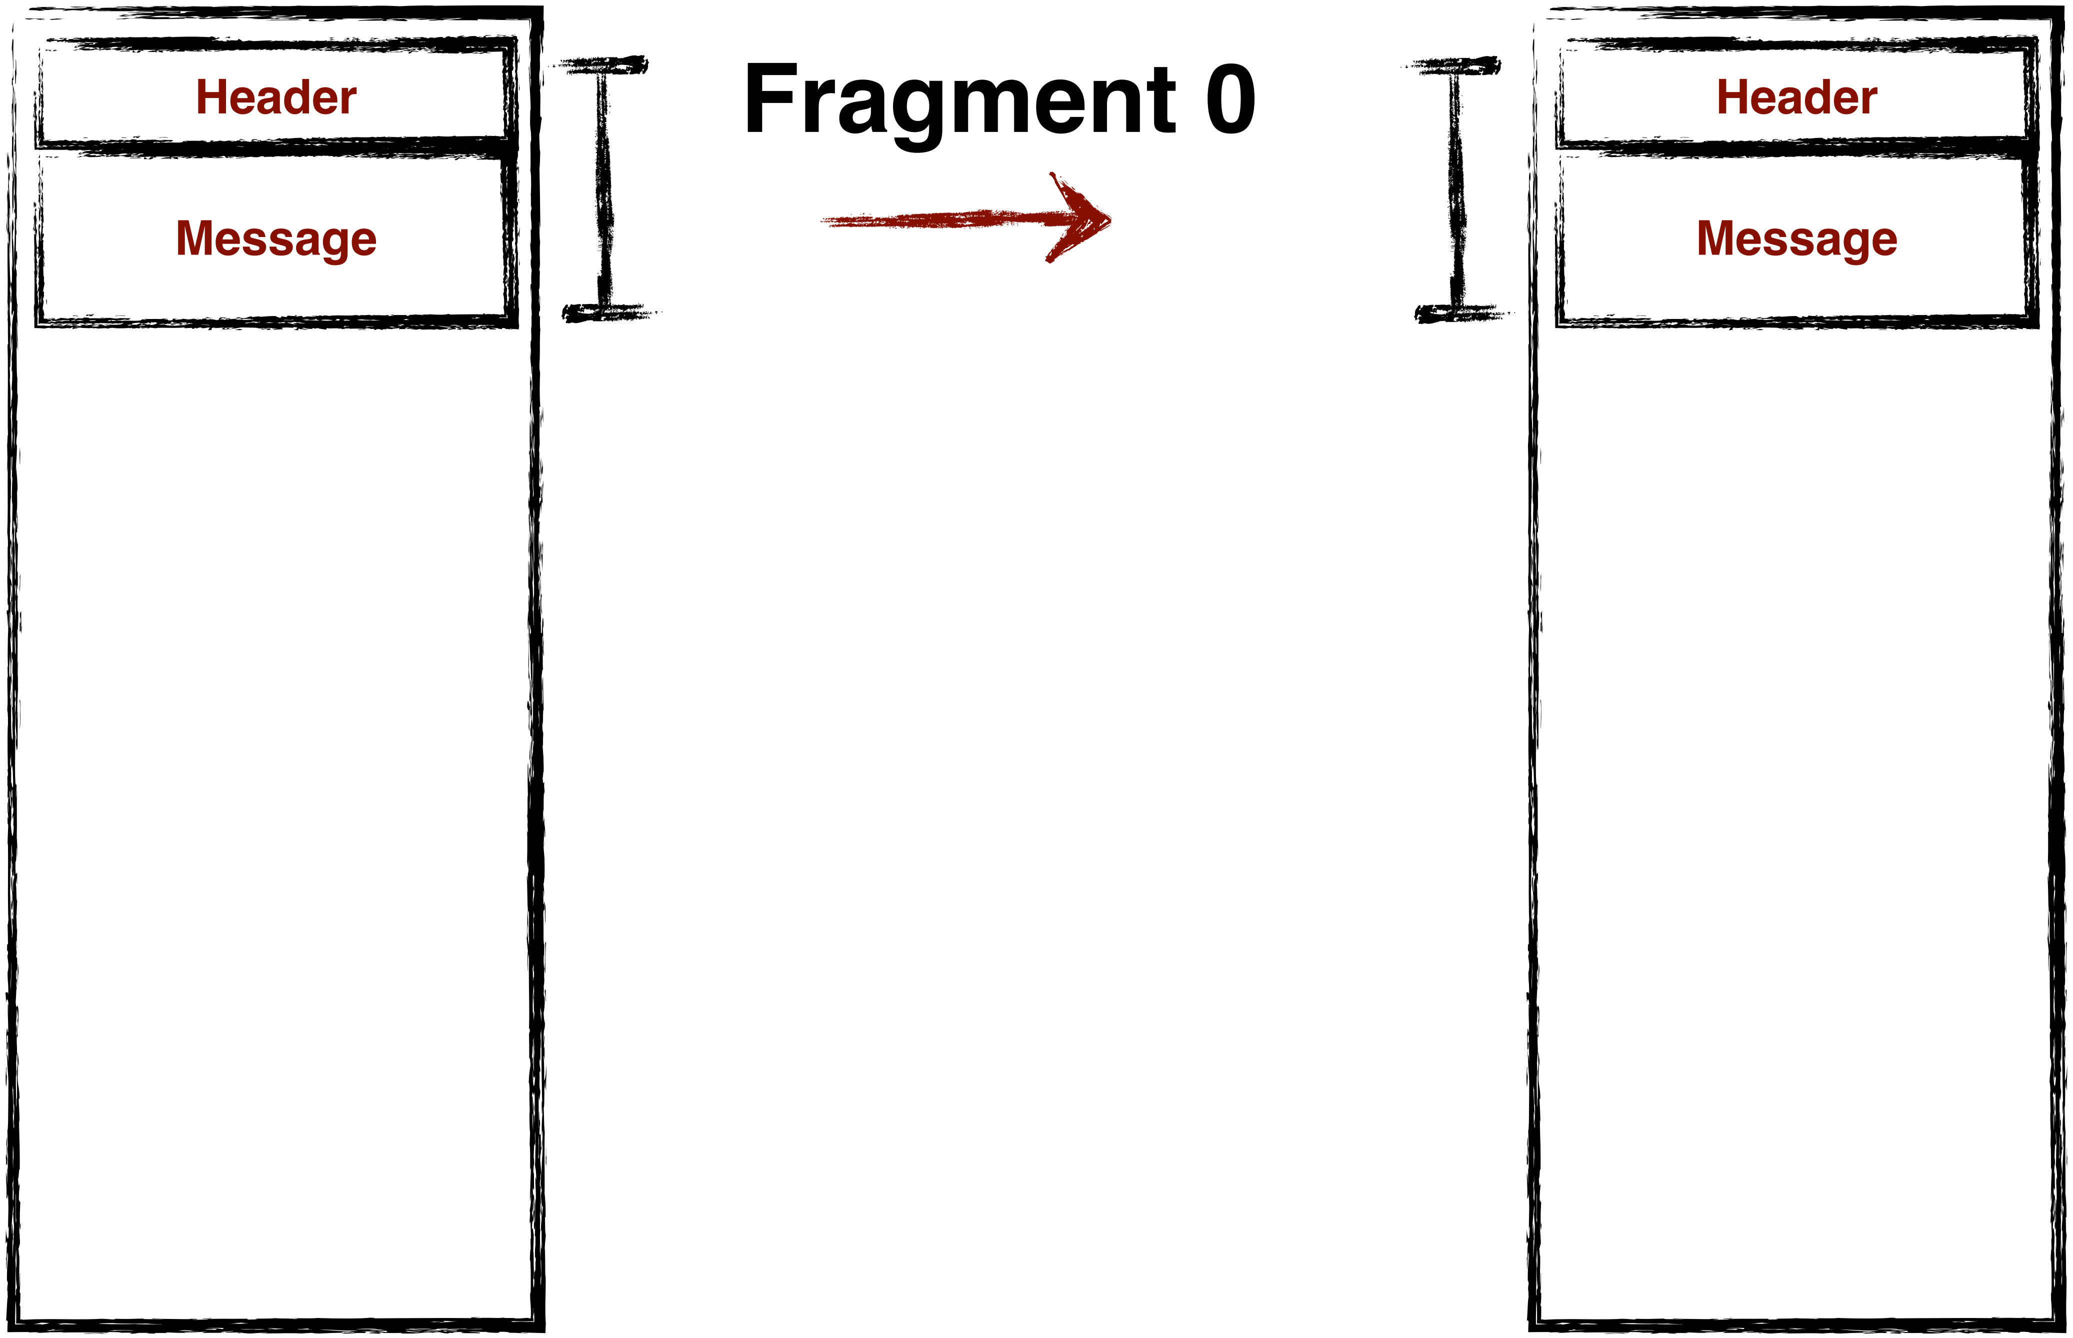

Header

Message

Fragment 0

Header

Message

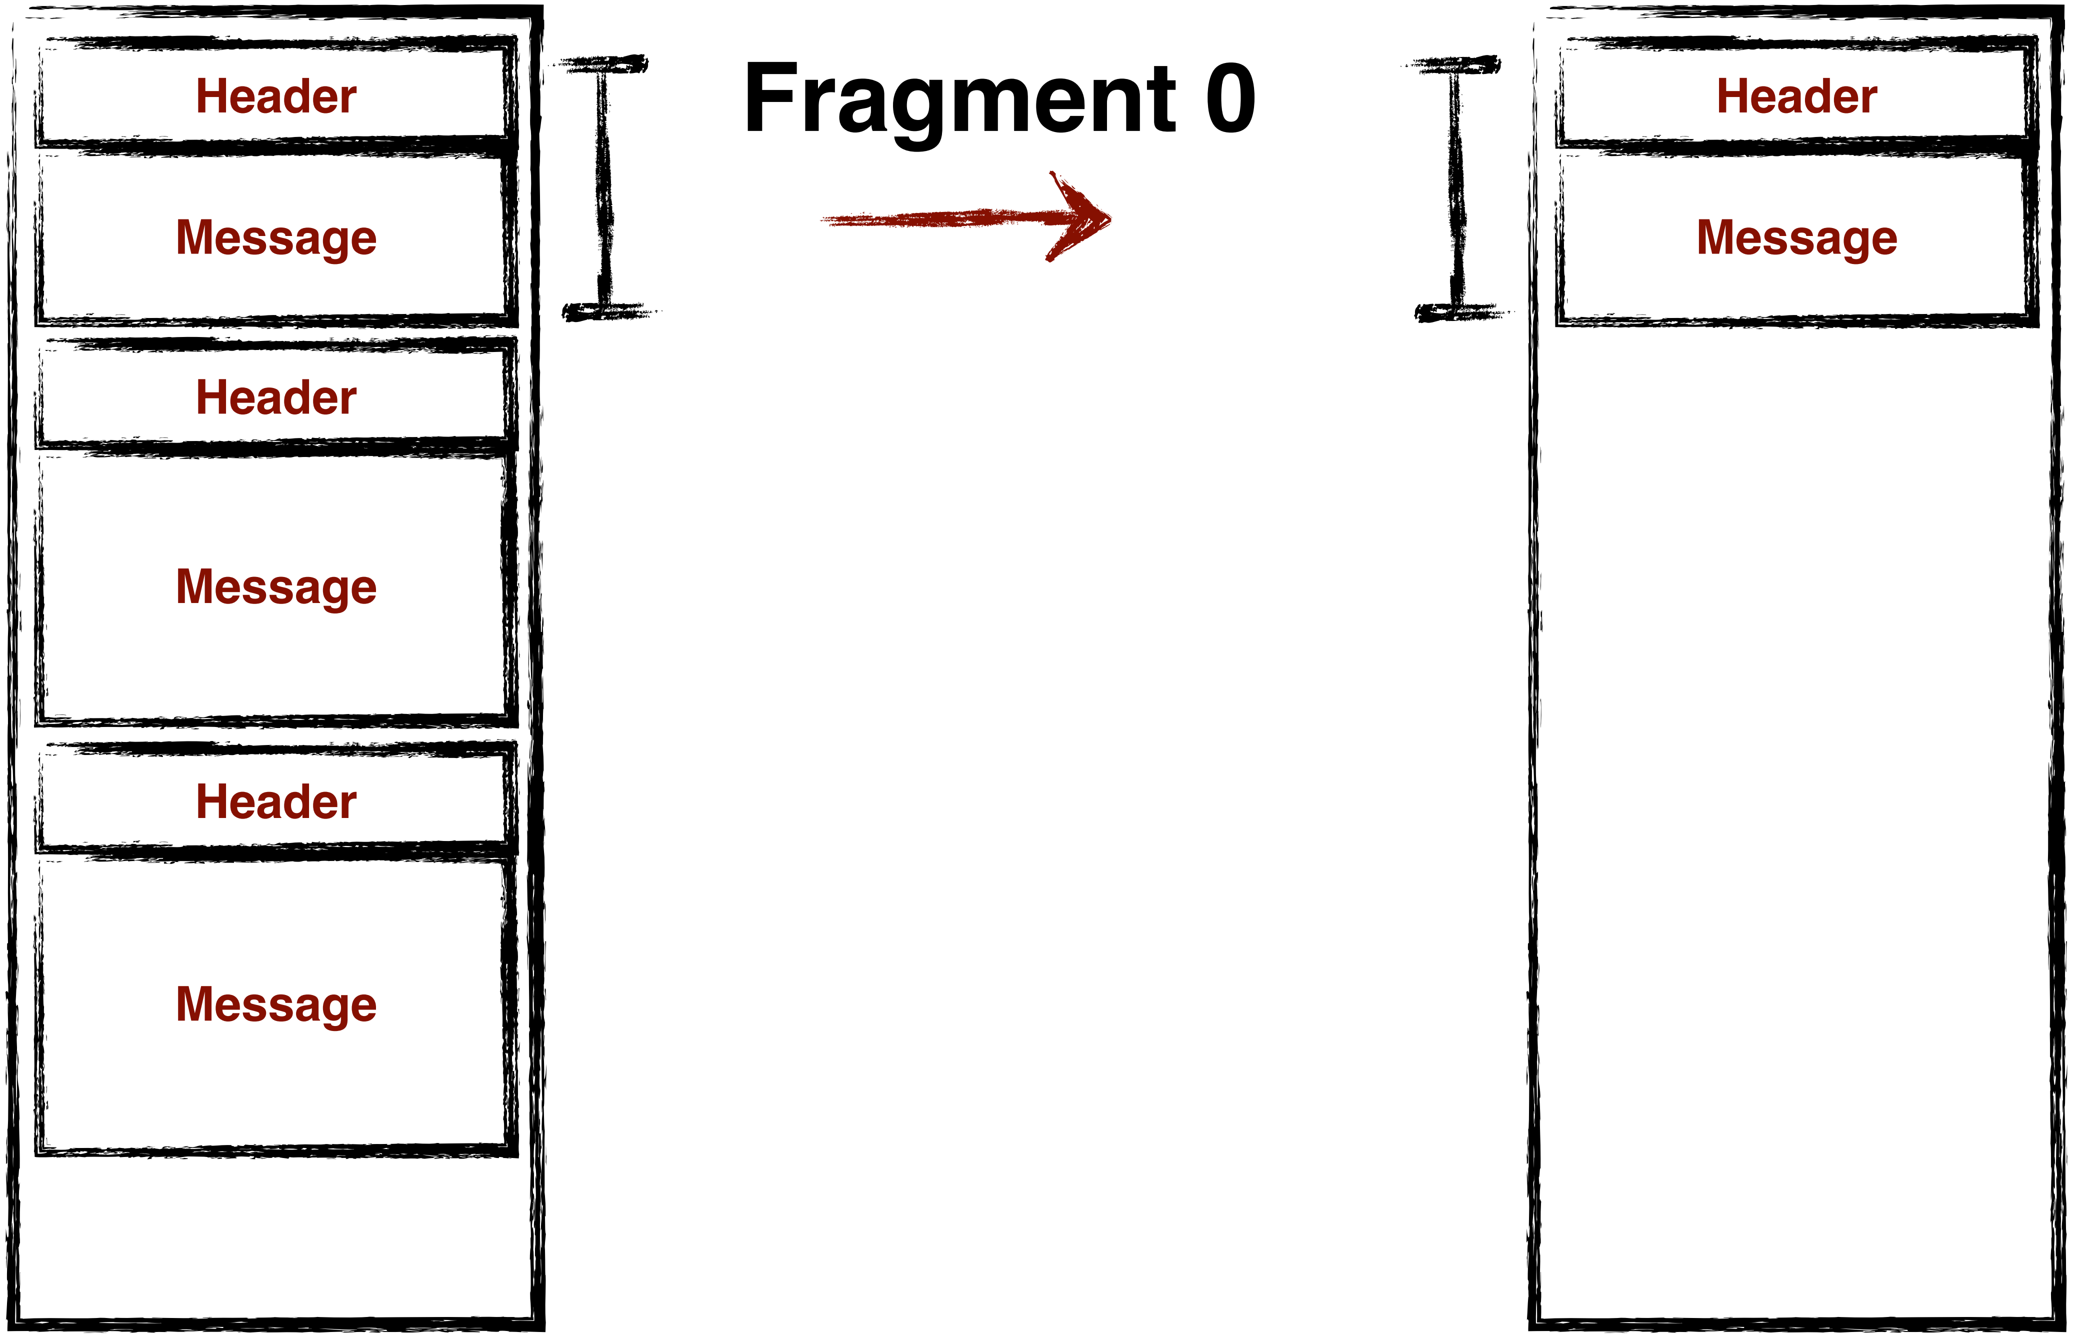
Header
Message
Header
Message
Header
Message
Fragment 0
Header
Message

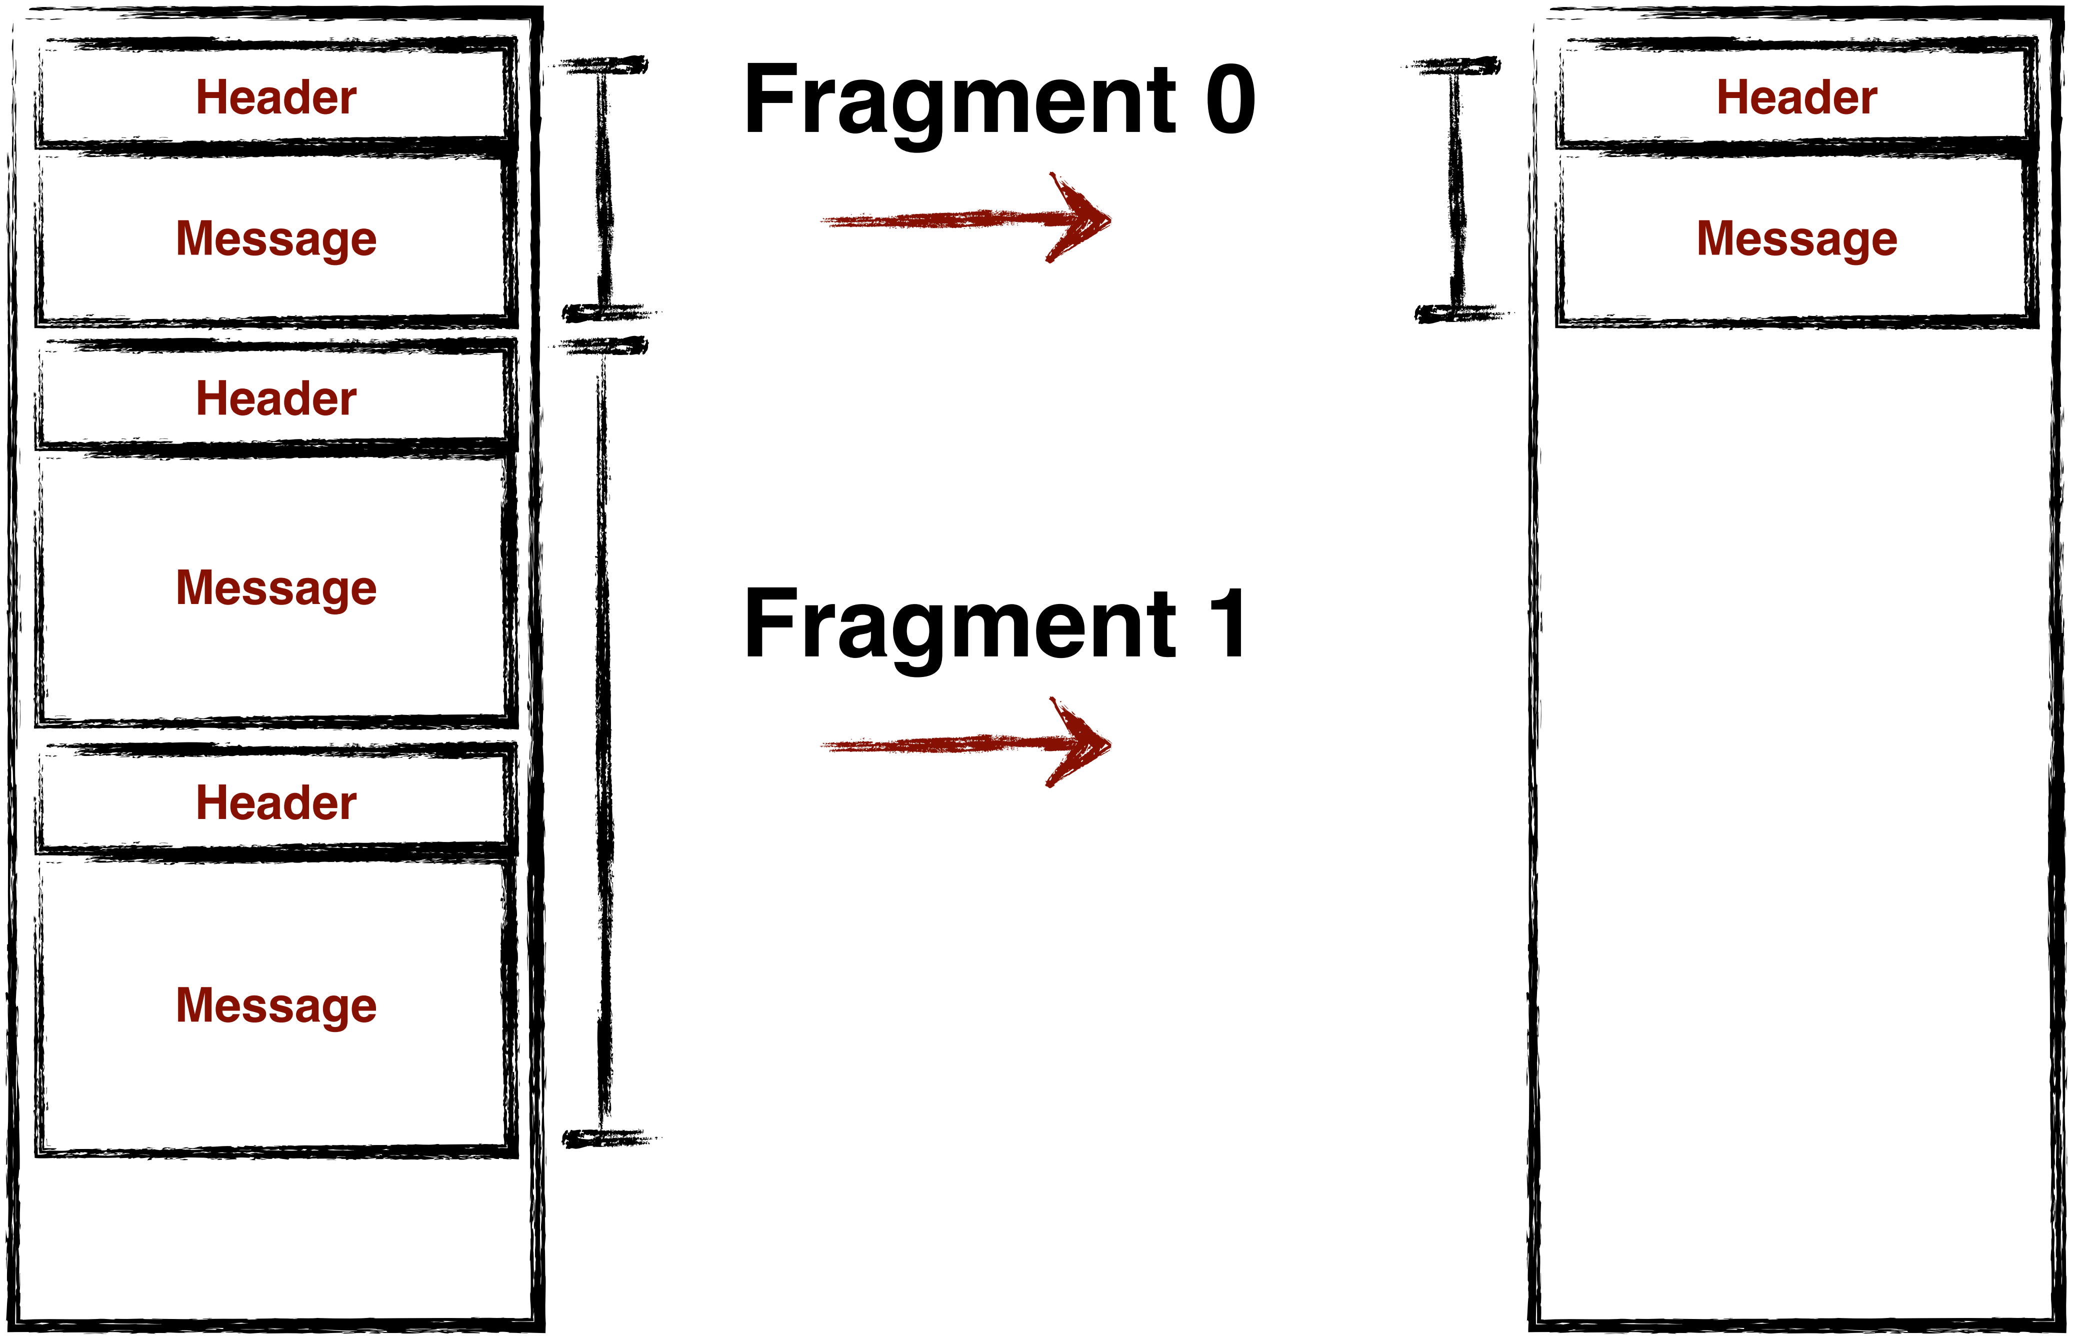

Header

Message

Header

Message

Header

Message

Fragment 0

Fragment 1

Header

Message

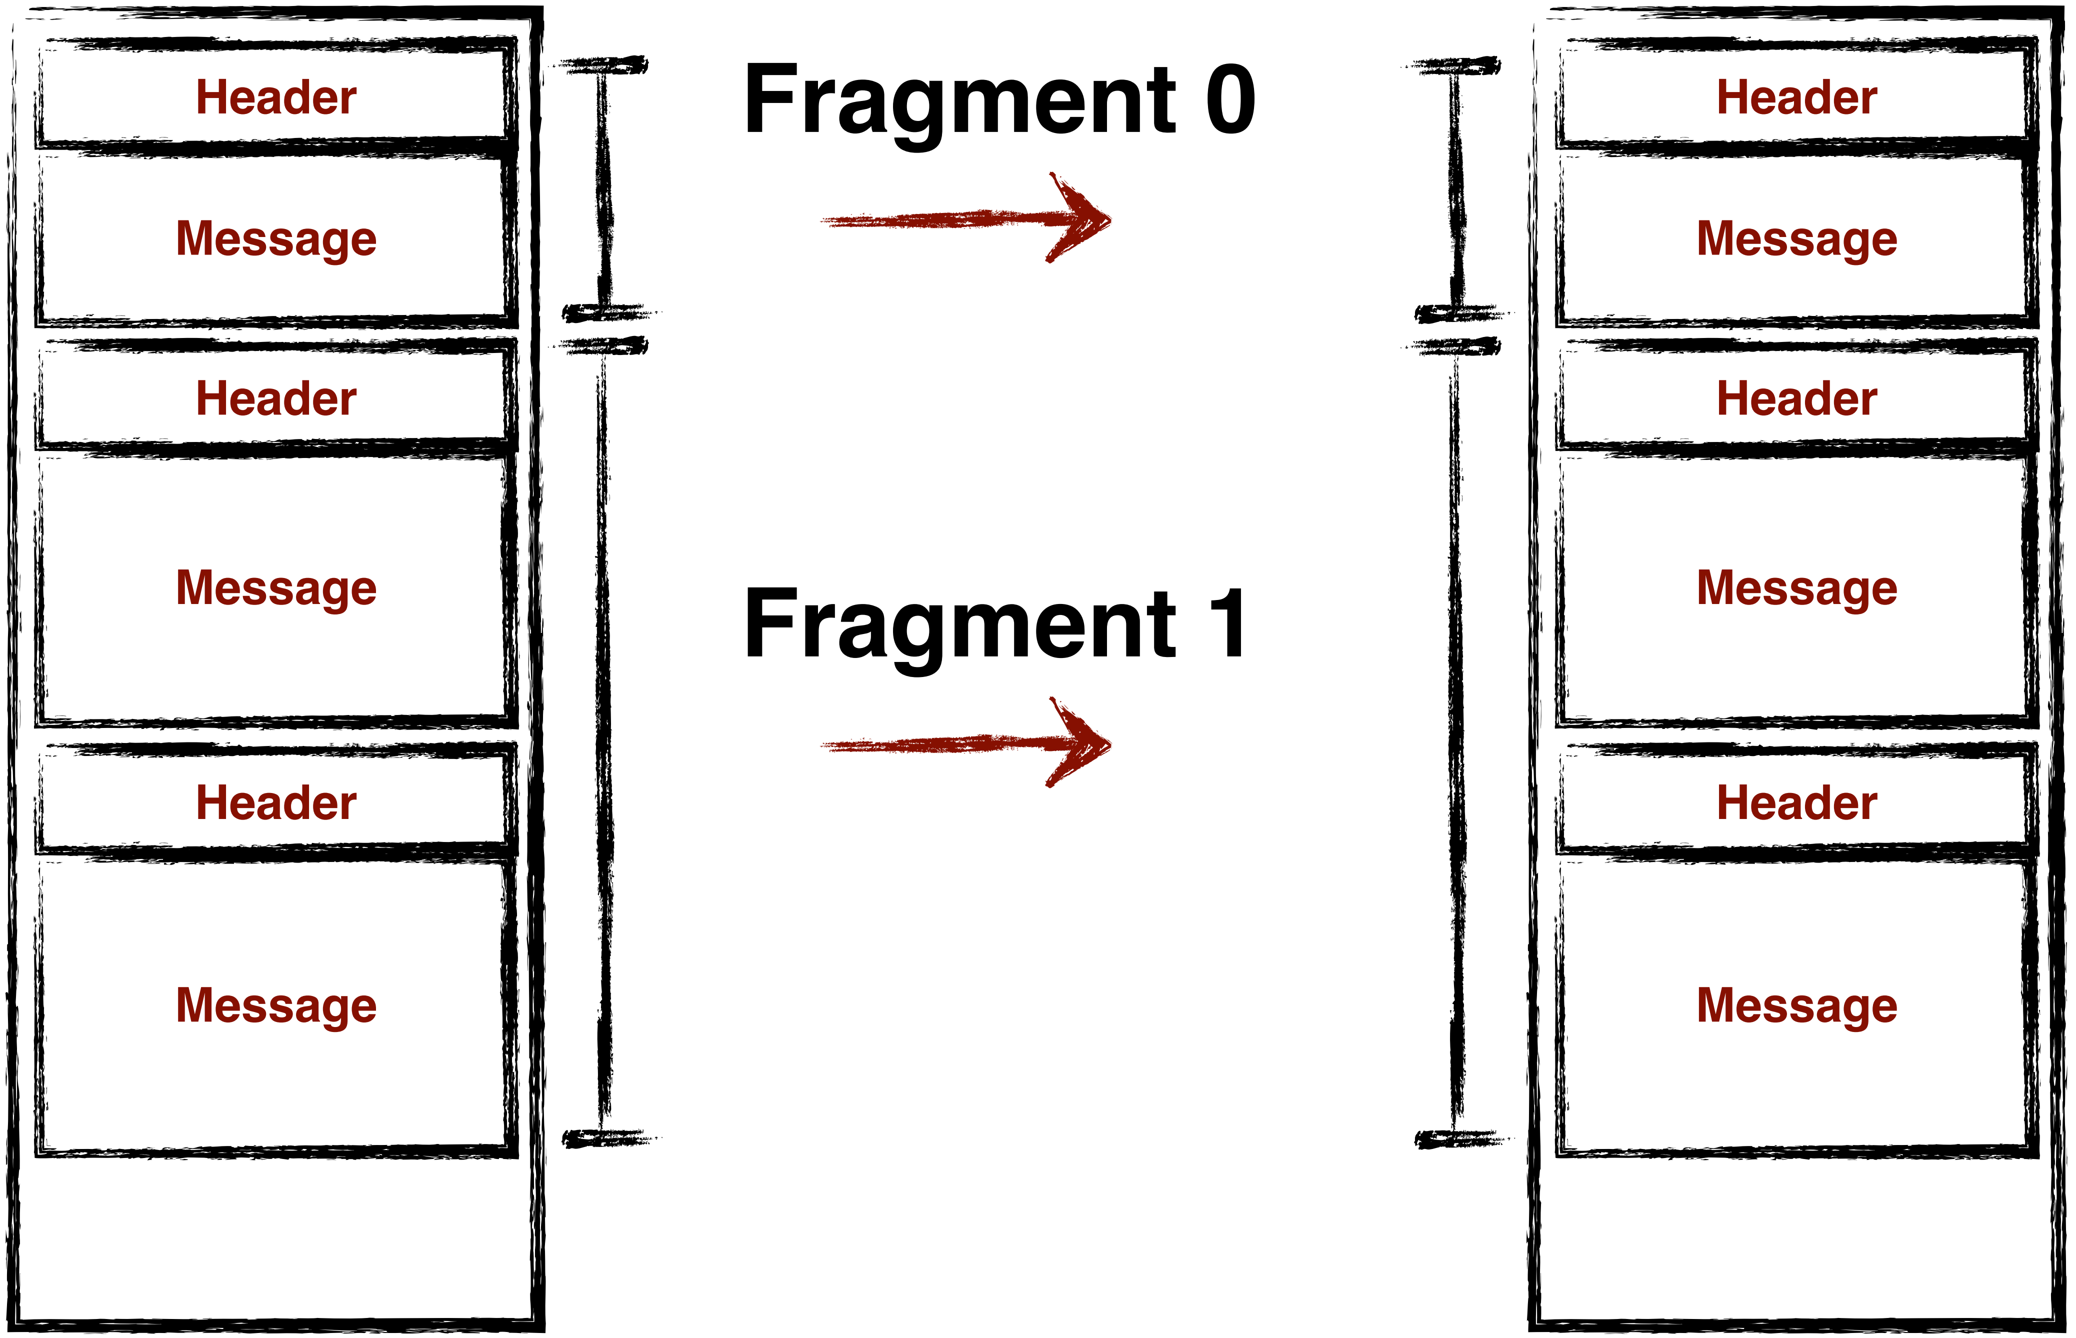

Header

Message

Header

Message

Header

Message

Fragment 0

Fragment 1

Header

Message

Header

Message

Header

Message

Natural for broadcast replication

In Closing…

A flood of data is coming,
Many say it is already here,
How will you deal with it?

Ever seen a PitBull drink from a Fire Hydrant?

It won’t give up…

Be the Pitbull at the Fire Hydrant!

Questions?

- Aeron *https://github.com/real-logic/Aeron*
- SlideShare *http://www.slideshare.com/toddleemontgomery*
- Twitter @toddlmontgomery

Thank You!