

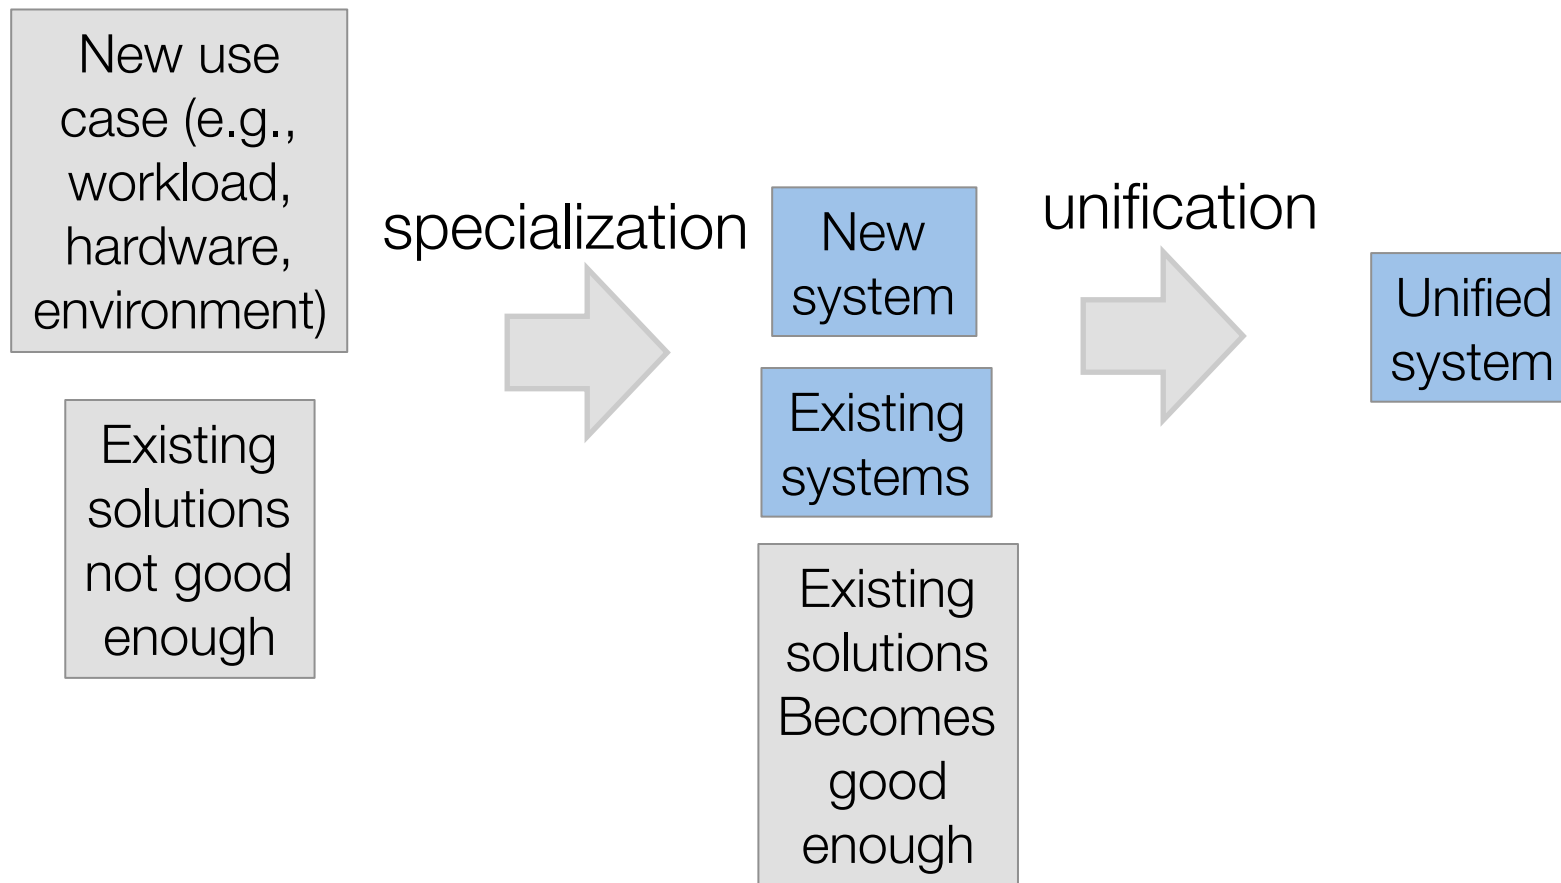
CS 294: A few patterns and techniques

December 2, 2015

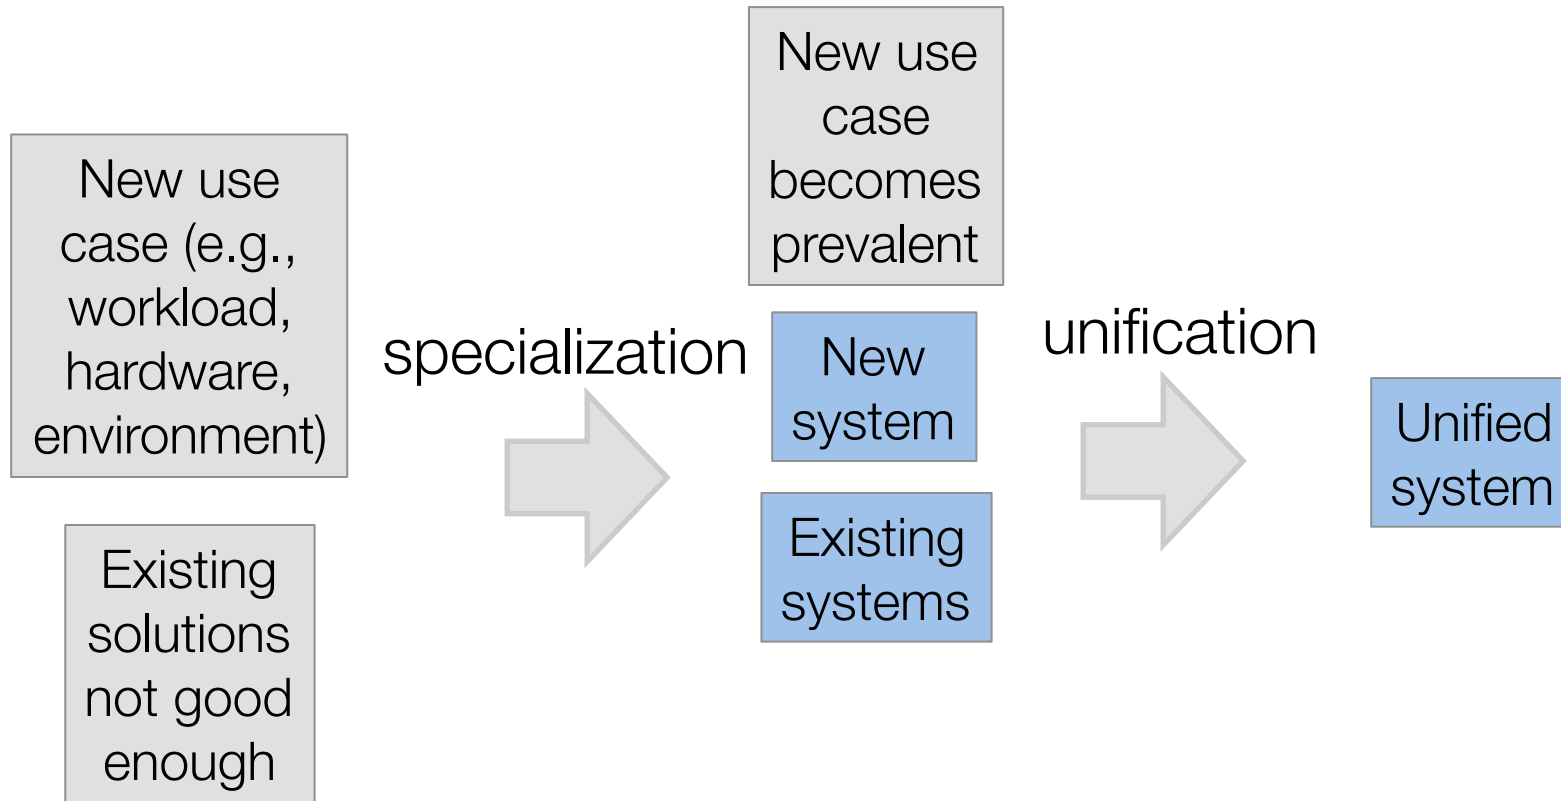
Ion Stoica

(<http://www.cs.berkeley.edu/~istoica/classes/cs294/15/>)

More on Spec. vs Unification



More on Spec. vs Unification



More on Spec. vs Unification




First cellular phones

Specialized devices

Unified device (smartphone)

More on Spec. vs Unification



More on Spec. vs Unification



More on Spec. vs Unification



Some random points...

Immutability

Dramatically simplifies design

- » No side effects (functional language)
- » Easy replication
- » Easy checkpointing

Coupled with deterministic operations

- » Easy fault recovery and straggler mitigation (through re-execution)

Challenges

Cannot support fine grain updates

So a bunch of papers gave up on immutability to better support new workloads

- » Parameter server
- » Asynchronous updates (also for speeding up ML)

Consistency

Strong consistency means coordination

Coordination means slow (hurts parallelism)

So what can you do?

- » Alleviate coordination by restricting the set of operations (e.g., CRTD)
- » Identify workload which really needs strong consistency (red-blue)
- » Provide more semantics from app or rewrite app (RAMP)

Performance

Congestion / high load

Slow algorithms, system overhead

Failures

Unpredictability (e.g., stragglers)

Networking Performance

Congestion / high load

Slow algorithms, system overhead

Failures

Unpredictability (e.g., stragglers)

Networking Performance

A network constantly congested is unusable

» Only solution is to upgrade it!

So the only interesting case is burstiness
causing congestion

» Only solution: prioritize important traffic

Insight

Short flows are typically latency sensitive

» Many short flows but few total bytes

Long flows are throughput oriented

» Fewer flows but majority of bytes

So prioritize short flows

» Little impact on long flows

Also true for jobs (see Dolly)

Flow Length?

If you do not know length

- » just prioritize each flow for at the beginning ...
- » ... decrease priority as flow continue to send bytes

Same idea in OS schedulers

- » Multi-feedback queue




Figure – Multilevel Feedback Queue Scheduling

Full Bisection Bandwidth

Common technique:

- » Use multiple paths
- » Route along different paths
- » Link layer vs. network layer

Doesn't solve all problems:

- » Edge points can be still congested
- » Dependence between senders and receivers makes things complicated (see faircloud)

Looking for Trends

Workload trends → new systems

» Yes, this is specialization!

Logs → append only → GFS

Analytics on big data → MapReduce

Graph queries → Graph databases

...

What Else?

Survey: Best Papers

35 different papers out of 55 papers!

Top 3:

- » MapReduce, Spark: 8 votes
- » CryptDB: 7 votes
- » GFS: 5 votes

Survey: Trends

Storage: non-volatile memory, SSDs: 13 votes

Faster networking: 7 votes

GPU & hardware specialization: 3 votes

Survey: Challenges

Easier to use tools (e.g., better visualization, faster prototyping, support non-developers): 6 votes

Auto-tuning systems: 3 votes

Security: 2 votes

Survey: Other Topics?

Huge variety...

More theory papers (e.g., dist systems, algorithms) : 3 votes

“Security”, and “Value in big data”: 2 votes

Others: search, IoT, visualization, e2e infrastructure, ...

Survey: Topics didn't care about?

Again, big variety...

Networks (fewer lectures or not at all): 6 votes

Coordination & consistency: 3 votes

Graphs: 3 votes

Others: ML, scheduling, fairness, security, ...

Survey: Other Topics?

Huge variety...

More theory papers (e.g., dist systems, algorithms) : 3 votes

“Security”, and “Value in big data”: 2 votes

Others: search, IoT, visualization, e2e infrastructure, ...

Survey: Suggestions

More discussions: 4 votes

Too many papers: 3 votes

More guest speakers: 3 votes