

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Nov. 8, 2022

Agenda for Today

- Announcements
- Research tips: how to get started
- TPUs overview
- TensorFlow discussion

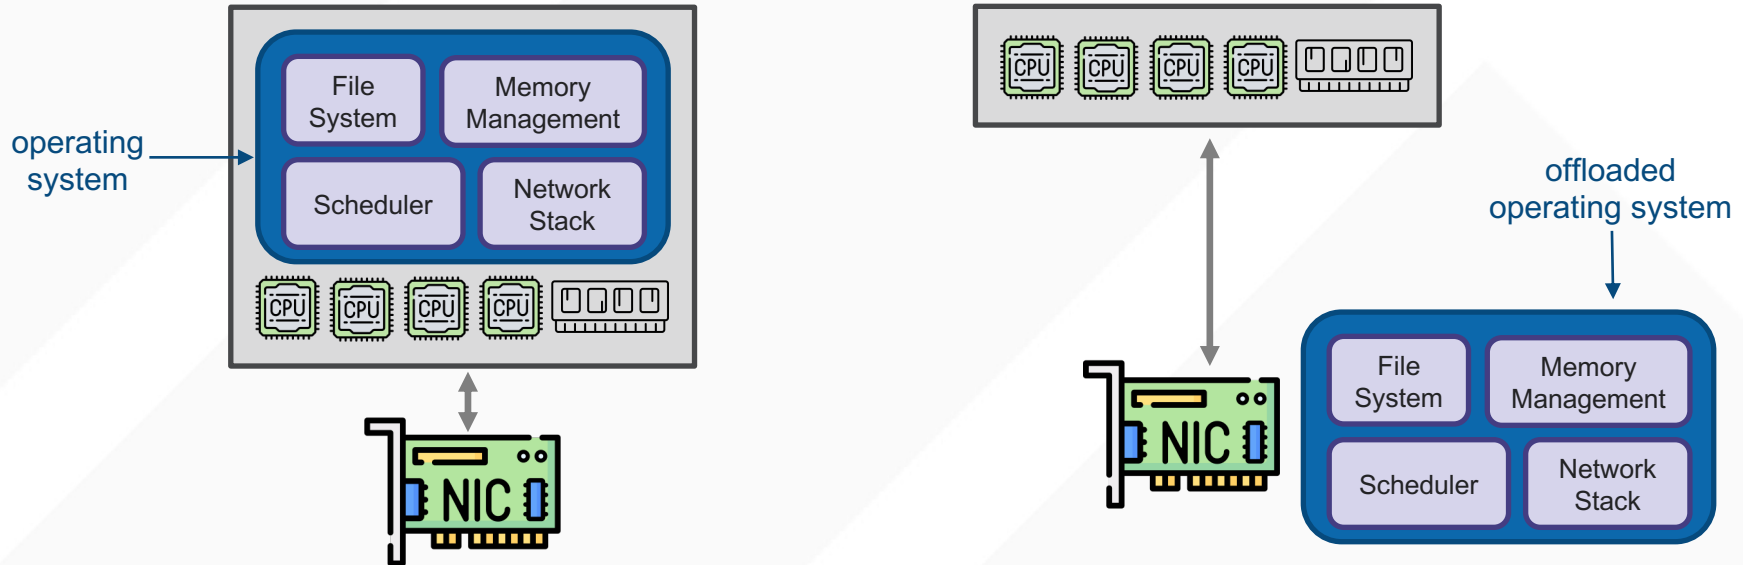
Announcements

- Project check-ins next week
 - Sign-ups will be posted later today
 - Time slots:
 - Tuesday 11/15, 2-3 pm
 - Wednesday 11/16, 10-11 am
 - Thursday 11/17, 2-3 pm
- Be prepared to talk about your progress so far
 - What have you learned?
 - What are you struggling with?
 - You can use the whiteboard or show diagrams or graphs
 - No need for a formal presentation

Research Tips: How to Get Started

How do you start answering a research question?

- Once you have a research question/problem/idea in mind, what do you do next?
- Example research question: is it possible to offload an OS to a SmartNIC?



Approach 1

- Identify workloads for testing such an OS-on-a-SmartNIC
 - Database, browser, computer game
- Offload an existing OS to a SmartNIC, one piece at a time
 - Network stack 😊
 - File system 😊
 - Scheduling 🤔
 - Memory management 😞
- Test our system after each piece to make sure it still works
 - Use the workloads we identified

testable milestones ✅

} it may take a while to discover if this won't work ❌

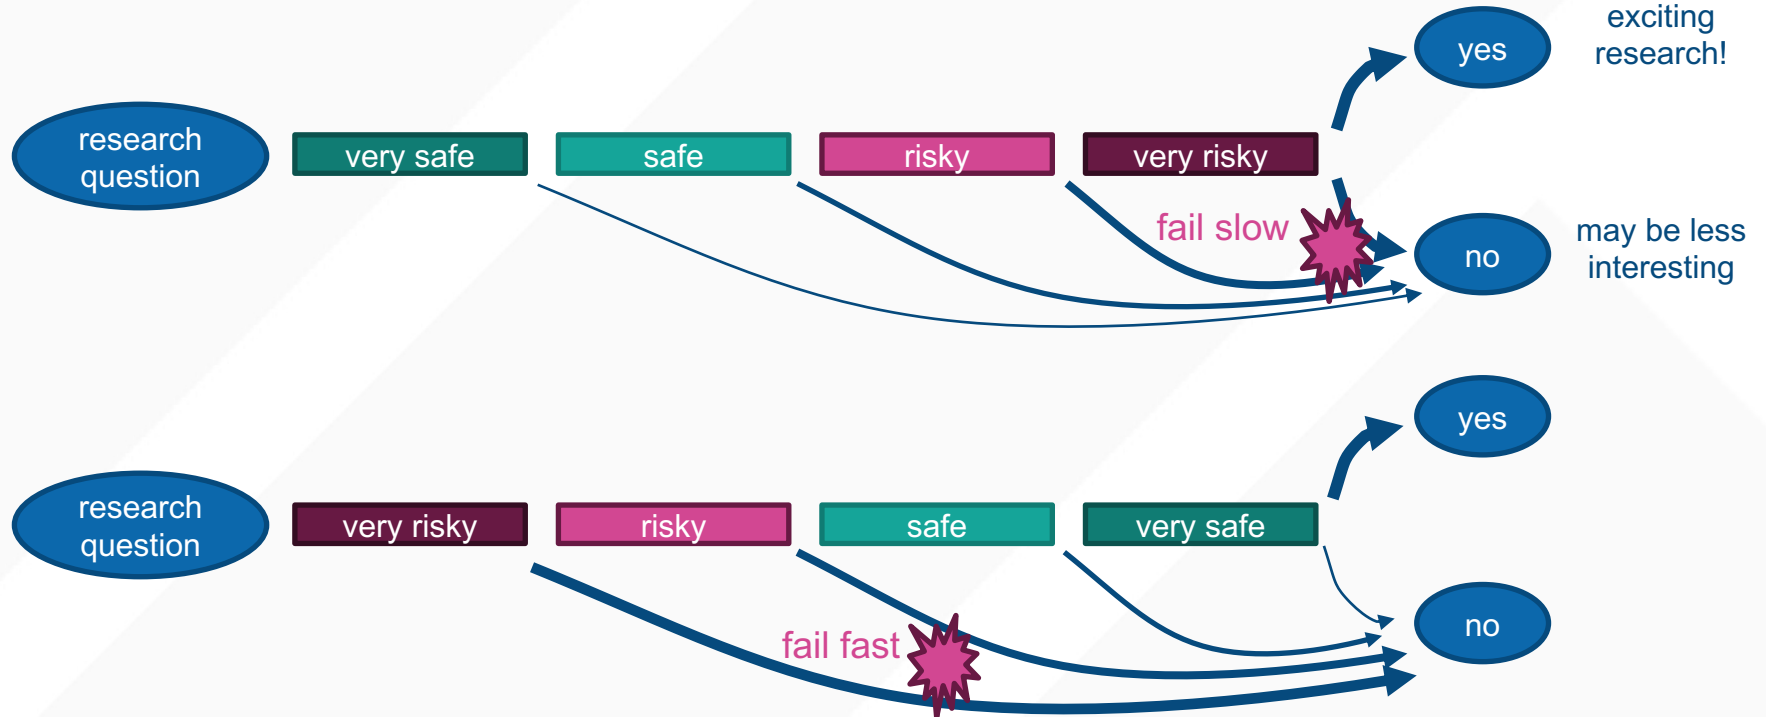
have to support a lot of features ❌

Techniques for Getting Started

- Create small milestones that can be tested
- Start simple
 - What is the simplest workload/setup that allows you to explore your idea?
- Try to fail fast
 - Start with the parts that are least likely to work

Failing Fast

- Fail fast - start with the parts that are least likely to work
 - Often these are the most interesting anyway!



Approach 2

- Identify the pieces that are least likely to work
 - Memory management and scheduling?
- Identify the simplest workloads that will test if these functionalities work
 - Scanning memory and handling page faults on the SmartNIC?
 - Sending timer interrupts from the NIC?
- Start by implementing these pieces and testing them on these simple workloads
- Add more complex workloads and functionality until you can confidently answer the research question

Getting Started Summary

- Create many small milestones
- Start simple
- Try to fail fast
- What is the most interesting aspect and how can I focus on that?

TPUs

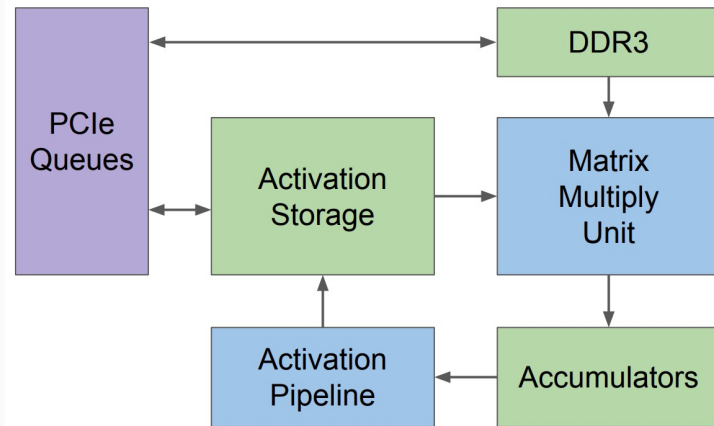
The Creation of TPUs

- 2013: machine learning was consuming more and more CPU cycles
 - Especially Deep Neural Networks (DNNs)
 - Very expensive
- Google set out to create a custom chip
 - Domain-Specific Architecture (DSA)
- Created the Tensor Processing Unit (TPU)
 - Used internally starting in 2015
 - Announced publicly in 2016
 - Codesigned with TensorFlow


Year	Version	Training?	Inference?
2015	TPUv1		✓
2017	TPUv2	✓	✓
2018	TPUv3	✓	✓
2020	TPUv4i		✓
2021	TPUv4	✓	✓

What is a TPU?

- Coprocessor connected via PCIe
- Primarily matrix multiplication and activations
- Optimized for 99th % performance
 - No caches, context switching, out-of-order execution, etc.
- Lower precision than CPUs
 - E.g., 8-bit multiplication
- Used for: improving search results, AlphaGo, etc.
- 30-80x better performance/watt than CPUs and GPUs (TPUv1)



TPUv1



Research on TPUs

- How to use TPUs for large-scale machine learning?
 - TensorFlow, OSDI 2016
- How to improve TPU performance?
 - “In-Datacenter Performance Analysis of a Tensor Processing Unit”, ISCA 2017
 - “Ten Lessons From Three Generations Shaped Google’s TPUv4i”, ISCA 2021

TensorFlow Discussion