HPTS 2022

~30 Students in ~30 mins

Achilleas Benetopoulos (UCSC)

- 2nd year PhD student. Advisor: Peter Alvaro

- Pre-PhD:
  - Did some work on early versions of PaSh (EuroSys ‘21, ICFP ‘21, HotOS ‘21, OSDI ‘22)
  - Worked on Harp (CCS’21)
- Currently focused on programming paradigms for distributed systems (come watch my talk later!)

- My full name was too long for twitter until at least 2015.

Allen Aboytes (UC Santa Cruz)

- 4th year PhD student. Advisor: Darrell Long

- Operating Systems support for Large-Scale Data-Intensive Applications
- Working on the Twizzler Operating System
  - Exploring scheduling and distributed aspects of its design

- Left handed and have a bearded dragon: Flambo

Anna Herlihy (EPFL)

- 2nd Year, advised by Anastasia Ailamaki and Martin Odersky
- My two great loves: Databases + Programming Languages
- Working on: Datalog, metaprogramming, static analysis, any system that consumes code! Compile & runtime optimizations [CIDR22]
- Previous life: Staff engineer at MongoDB

Harry: The 2nd best dog in these slides

Audrey Cheng (UC Berkeley)

- 3rd year PhD student, advised by Natacha Crooks and Ion Stoica

- Worked on adding transactional guarantees at large-scale (billions of QPS)
  - RAMP-TAO (VLDB ‘21)
- Developed new benchmark based on Meta’s social graph workloads
  - taobench.org! (VLDB ‘22)
- Working on caching algorithm for transactions and distribution-aware concurrency control

TAOBench

- Snorkeling enthusiast, bird lover

Aunn Raza (EPFL)

- 5th year PhD student, advised by Anastasia Ailamaki
- Enabling efficient operational analytics on fresh data through Adaptive HTAP
  - Tuning performance-freshness trade-off in HTAP systems, at runtime

- Adaptive HTAP in Proteus (proteusdb.com): open sourcing very soon!
  - Modelled workload-isolation in HTAP as a scheduling problem [SIGMOD ‘20]
  - Worked on CPU-GPU based hardware isolation in HTAP [CIDR ‘20]
  - Saw (MVCC-) garbage and proposed OneShot GC for OLTP [SIGMOD ‘23]
  - Working on workload-driven snapshot layout for efficient access to fresh data

- I have spent more time with performance profilers than actual programming systems

Austen Barker (UC Santa Cruz)

5th year/recently graduated, advised by Darrell Long

Dissertation topic: deniable and steganographic storage systems

Artifice: A Design for Usable Deniable Storage Informed by Adversary Threat

Interests:

Operating Systems, storage systems, and cryptography

Starting at Sandia National Laboratories

I like working on old computers and Volvos

Daniel Bittman (UC Santa Cruz)

(he/him) ; @danielbittman ; @sysminuscontext

- 5th year PhD student, advised by Peter Alvaro and Ethan Miller, graduating this year
- Interests
  - Operating systems (ATC ‘20)
  - Distributed systems (HotNets ‘21, HotCloud ‘19)
  - Programming models (PLOS ‘19)
  - Non-volatile Memory (FAST ‘18)

- ~~I’ve been writing kernels since I was 14~~

I play 3 musical instruments (2 of them badly) and I love backpacking

Devashish R. Purandare (UC Santa Cruz)

- 5th Year PhD Student, Advisor: Ethan L. Miller

- Log-Structured Data Management on Log-Structured SSDs (CIDR ‘22)
  - Using zoned SSDs to prevent inefficiencies from Log on Log (on Log sometimes!)
  - Updating filesystems, databases, other log-structured systems
  - Talk to me about storage!
- Large scale workload characterization (CHEOPS ‘22, SIGOPS OSR ‘22)

- Was refused dev@ucsc.edu address because of security policy

Esteban Ramos (UC Santa Cruz)

- 2nd year PhD student. Advisor: Peter Alvaro

- Interested in data management problems, although I’m increasingly interested in networking problems as well
- Currently working on distributed TwizzlerOS

- I biked to HPTS!

Hamish Nicholson (EPFL)

- 2nd year PhD student. Advisor: Anastasia Ailamaki

- Working to tame the complexity of the modern storage hierarchy
  - Hear more at my talk on Tuesday
- Previously a software engineer @ Ocient
- Talk to me about hardware

- I love mountains and deserts: disappearing into the wilderness after HPTS

Jack Waudby (Newcastle University)

- 4th year PhD student. Advisor: Paul Ezhilchelvan
- Thesis on High Performance Concurrency Control & Commit Protocols in OLTP Databases
  - E.g., Mixed Serialization Graph Testing (TPCTC'22), Epoch-based Multi-commit (SRDS'22)
- Linked Data Benchmark Council Benchmark (LDBC) Social Network Benchmark (SNB) Task Force Member
- Distributed systems engineer at Neo4j
- Host of "Disseminate: The Computer Science Research Podcast"
- Hobbies: football (soccer), skiing, and (hot) yoga

Jun Zhang (UC Santa Cruz)

- 4th year PhD student. Advisor: Peter Alvaro

- Focus on general purpose state collection and analysis
- 3MileBeach (Socc 2021) provides real-time control options on network communication patterns in distributed applications.
- Currently moving the concept to a lower level (record-replay)

- Badminton & Tennis player. Beach & Trail lover.

Kamala Ramasubramanian (UC Santa Cruz)

Defended in June. Advisor: Peter Alvaro

Thesis: Addressing distributed systems problems such as fault-tolerance and troubleshooting by *querying* observations of executions

I did my first pull-up recently after 4 years of attempts!

Konstantinos Kallas (University of Pennsylvania)

- 5th year PhD student. Advisor: Rajeev Alur

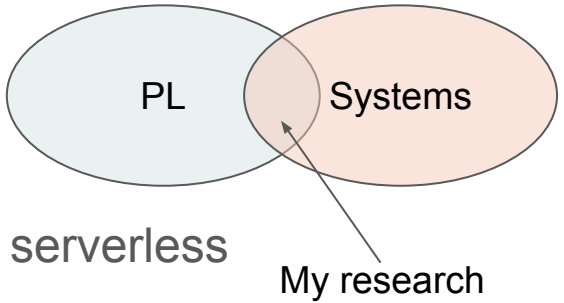

- Currently focused on correct and efficient stateful serverless
  - Durable Functions (OOPSLA 21)
  - Netherite (VLDB 22)
  - Microservices on Serverless (POPL 23)
- Also work on improving the shell (parallelization, distribution, etc)
  - PaSh (EuroSys 21, ICFP 21, HotOS 21, OSDI 22, hosted by LF)
- And Stateful Stream Processing (OOPSLA 20, PODS 21, PPoPP 22)

- I enjoy most human activities, currently a lot of climbing, reading

Matt Butrovich (Carnegie Mellon University)

- 4th year PhD student. Advisor: Andy Pavlo

- Embedding DBMS logic in kernel-space using eBPF:
  - Training data collection for Self-Driving DBMSs (SIGMOD 2022)
  - Application Layer (i.e., L7) wire protocol logic for proxies (Soon…)
- NoisePage Self-Driving DBMS
- BusTub Educational DBMS

- Owner of The Best Dog: Scout

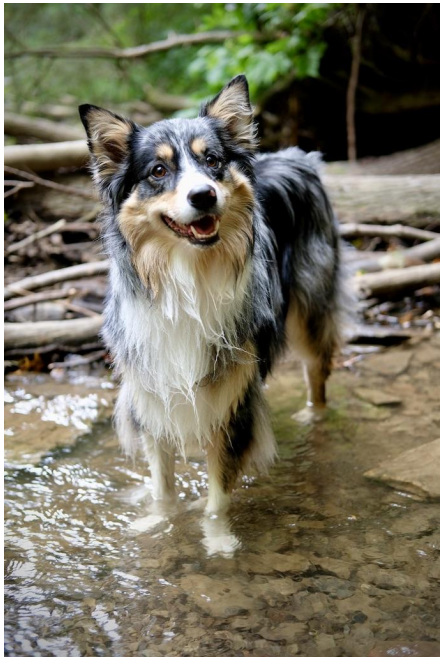

Michael Freitag (TU Munich)

- 5th year PhD student, advised by Thomas Neumann

- Working on the flash-based Umbra HTAP database
  - Buffer management (CIDR'20)
  - Transaction processing & MVCC (VLDB'22)
  - Cardinality estimation (CIDR'19)
  - Worst-case optimal join algorithms (VLDB'20)
  - Spent an unreasonable amount of time on our B+-tree implementation

- I share an office with Rudolf Bayer

Peter Kraft (Stanford)

- Final year PhD student, advised by Matei Zaharia and Peter Bailis

- Working on the DBOS project
  - DBMS-oriented serverless computing (Apiary)
  - Polystore transactions
  - Provenance/observability

- Also work on analytics systems (parallel query serving, data lakehouses)

Puneet Mehrotra (University of British Columbia)

- 3rd year. Advisor: Margo Seltzer

- Designing a system to store & process large graphs that eliminates ETL pipelines
  - Flexibility to change the on-disk representation based on workload/statistical properties of the graph
  - Supports both Edge-centric and Vertex-centric APIs
  - Is *not* pig-dog slow
- Previously: solving the problem of combed degree distributions in Kronecker graph generators (GRADES 2020)
- Benchmarking crimes in graph processing system evaluations

- I’ve been attacked by birds in 3 continents, once inside an Airport

Qian Li (Stanford)

- Final year PhD student advised by Christos Kozyrakis

- Working on DBOS (CIDR'22, VLDB'22)
  - Apiary: Serverless environment for DBOS
  - Transaction oriented debugging
  - Polystore transactions
- Also work on serverless/microservices
  - INFaaS: Model-less ML inference serving (ATC'21).

- Love designing logos and swags.
- Easily distracted by birds.

Saba Jamilan (UC Santa Cruz)

- 5th year PhD student. Advisor: Heiner Litz
- Profile-guided compiler optimization for data center workloads
  - Utilizing run-time profiling information for emitting timely software prefetches for indirect memory access patterns
    - APT-GET: profile-guided timely software prefetching (EuroSys 22)
  - Profile-guided compiler optimization for non-temporal data management
  - Speculative Common Path Optimizations for Datacenter Workloads
- I enjoy outdoor activities: biking, hiking

Sidhartha (Sid) Agrawal (University of British Columbia)

- 2nd year PhD . Advisor: Margo Seltzer

- General Purpose Isolation Mechanisms in Operating Systems
  - Develop a Model for Resource Isolation
  - Develop a capability-based framework to realize the model
  - More at the Gong Show tonight
- Recently worked on integrating seL4 with CHERI capabilities at ARM

- Unusually obsessed with Star Trek TNG

Sujaya Maiyya (UC Santa Barbara)

Tianyu Li (MIT)

- 4th year. Advisor: Sam Madden

- Resilient building blocks for the modern cloud
  - Speculative execution across failure domains (SIGMOD ‘21)
  - Fault-tolerant service composition with resilient message-passing steps (Coming soon)
  - Building resilient applications (e.g., transaction processing) that are both simpler and more efficient/performant using these building blocks (Coming soon-ish)

- Plays guitar in the MIT DSG (and friends) rock band, Impedance Mismatch

Ties Robroek (IT University of Copenhagen)

- 2nd year. Advisor: Pınar Tözün

- Resource-Aware Machine Learning
  - Deep Learning models have grown increasingly hungry for resources
  - Efficient Deep Learning is pivotal for sustainable and accessible research
  - Currently investigating resource consumption patterns of popular models

- Wasting most of my free time developing a remarkably poor game engine

Xiling Li (Northwestern University)

- 2nd year Ph.D. student. Advisor: Jennie Rogers
- Interests on building secure, private and verifiable DBMS
- Current work: ZKSQL = Zero Knowledge Proof + Query Evaluation
  - Clients can get faithful query answer while DB owners preserve data secrecy
- Want to make more friends in DB area

Personal Website

Xinjing Zhou (MIT)

- 2nd year Ph.D. student. Advisor: Mike Stonebraker

- Working on the DBOS project.
  - Transaction Processing
    - Better distributed transactions for H-Store/VoltDB style architecture. (VLDB 22)
    - Polystore transactions.
  - Storage
    - Improving buffer pool memory utilization of index and storage structures when there are skews .

- I have hitchhiker’s thumb

Zheng Zhang (Northwestern University)

- 2nd year Ph.D. student. Advisor: Jennie Rogers
- Interested in improving security and privacy guarantees of DBMSs.
  - Zero-Knowledge Transaction Processing
  - Adversarial Robustness in ML Systems

- Cats person, proof: ⟹

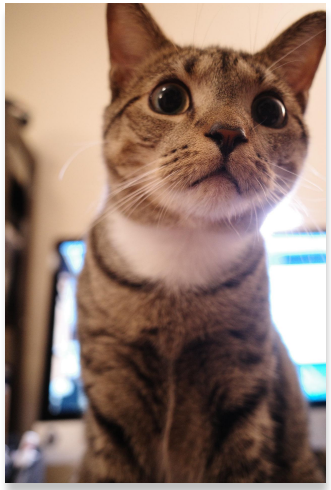

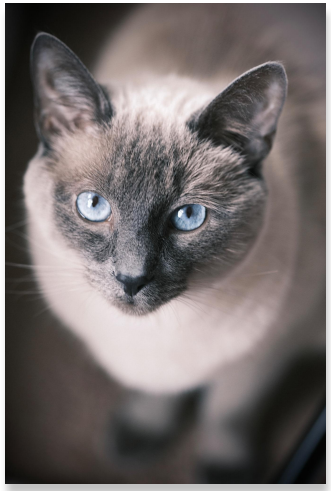

Zhihan Guo (University of Wisconsin - Madison)

5th year Ph.D., Advisor: Xiangyao Yu

- Transaction Processing
  - **Atomic Commit** on Disaggregated Cloud Storage (VLDB'22)
  - **Concurrency Control** for High Contention (SIGMOD'21)
- Others
  - How Good is My **HTAP** system? (SIGMOD'22)
  - **Caching** on Modern Storage Hierarchy (FAST'21)
  - ML-Driven Data Cleaning (SIGMOD'20)

Basketball, Climbing (Bouldering)

Now go and mingle! …

Happy HPTS’22