www. chameleoncloud.org

PLATFORM FOR INNOVATING IN THE EDGE TO CLOUD CONTINUUM

Kate Keahey

Mathematics and CS Division, Argonne National Laboratory

CS Department, University of Chicago

keahey@uchicago.edu

OpenInfra Days at North America, October 2024


Chameleon
700+
Papers published
1,000+
Unique projects
11,000+
Users
9

CHAMELEON: AN EDGE TO CLOUD TESTBED

- Chameleons like to change – testbed that adapts to your experimental r
  - **From bare metal reconfigurability/isolation** -- KVM cloud – to containers for edge (**CHI@Edge**)
  - Capabilities: power on/off, reboot, custom kernel boot, serial console access, etc.
- From large to small – diversity and scale in hardware:
  - **Supercomputing datacenters** (UC/ALCF, TACC, NCAR) over 100G network – to **edge devices**
  - **Diverse:** FPGAs, GPUs, NVMe, NVDIMMs, Corsa switches, edge devices via CHI@Edge, etc.
  - **Distributed: CHI-in-a-Box** sites at **Northwestern and UIC** – and now also **NRP**!
- Based on mainstream open source – proud to be cheap!
  - 50% leveraging and influencing **OpenStack** + 50% "special sauce" (incl. fed id)
- Promoting digital artifact sharing
  - Integration with **Jupyter** for non-transactional experiment packaging
  - **Trovi** for experiment sharing and discovery, **Chameleon Daypass** for access sharing
  - Reproducibility and education: digital sharing killer apps!

NOT JUST A TESTBED, A COMMUNITY

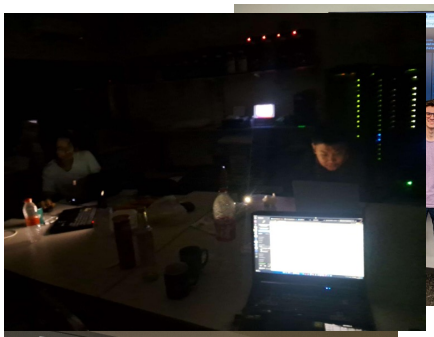

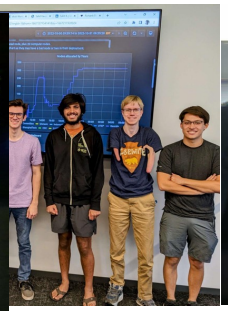

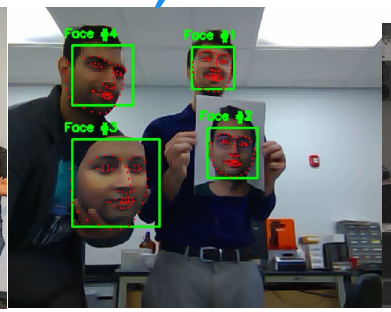

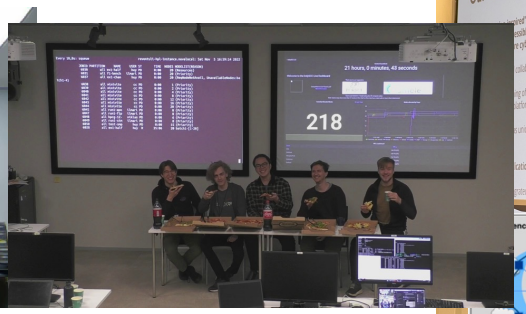

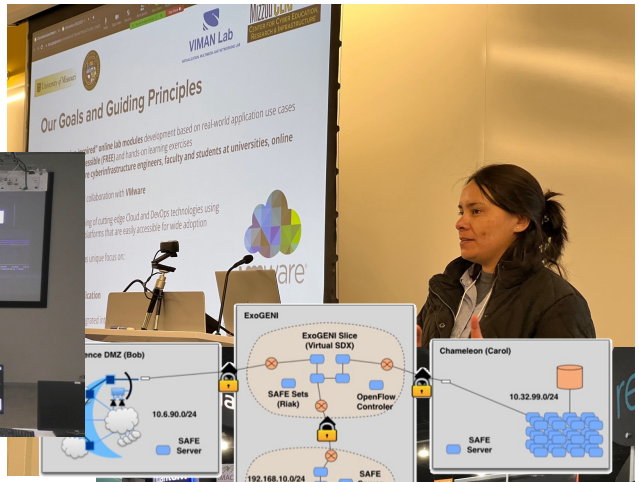

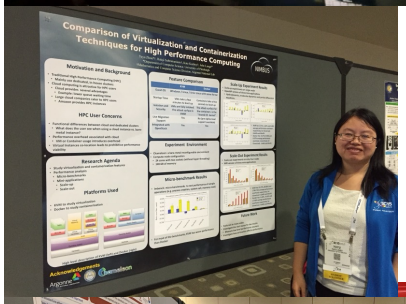

Supporting research projects in architecture, operating systems design, virtualization, power management, real-time analysis, security, storage systems, databases, networking, machine learning, neural networks, data science, and many others.

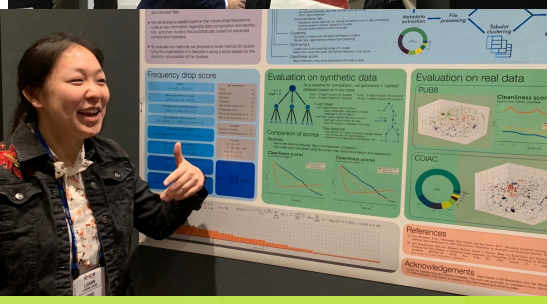

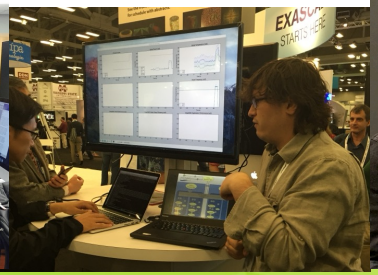

Chameleon www.chameleoncloud.org

Check out user experiment stories on our blog:
https://www.chameleoncloud.org/blog/category/user-experiments/

FROM CLOUD TO EDGE WITH CHAMELEON

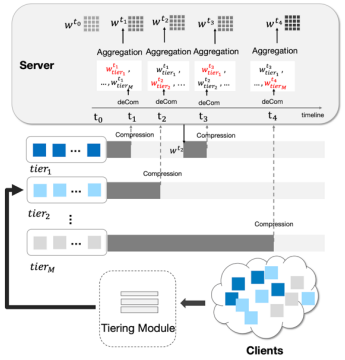

federated learning

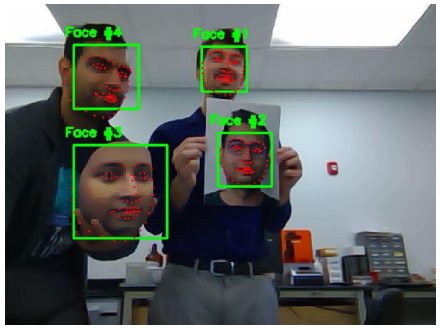

biometrics

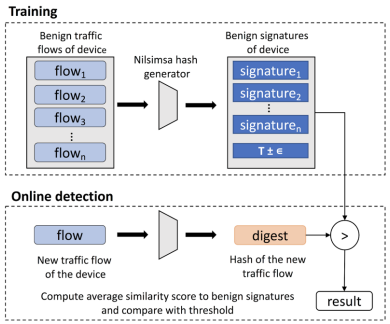

network traffic fingerprinting for IoT devices

- Increasingly more Chameleon project applications working on IoT/edge
- Simulation/emulation don't always provide the answer: What are the impacts of this approach on power management on edge device? How will the performance transfer to edge? Can we measure the impact of distribution/networking for edge/cloud applications?
- **Goal: "realistic edge to cloud experiments from one Jupyter notebook"**

HIGHLIGHT: CHI@EDGE

A lot like a cloud!
All the features we know and love – but for edge!
"Edge to cloud from one Jupyter notebook."

Not at all like a cloud!
Location, location, location!
IoT: cameras, actuators, SDRs!
Not server-class!
And many other challenges!

- CHI@Edge: all the features you love in CHI, plus:
  - Reconfiguration through non-prescriptive **container deployment** via OpenStack interfaces (using K3 under the covers)
  - Support for "standard" **IoT peripherals** (camera, GPIO, serial, etc.) + easy for you to add support for your own peripherals
  - **Bring Your Own Device (BYOD): Mixed ownership** model via an SDK with devices, virtual site, and **restricted sharing** – building on OpenBalena

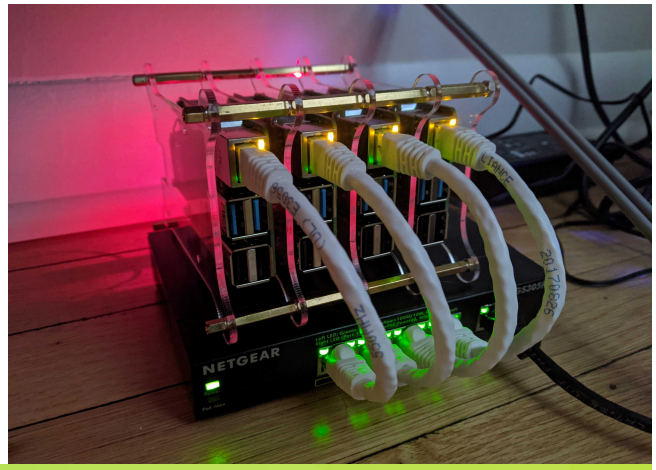

Paper: "Chameleon@Edge Community Workshop Report", 2021

FLOTO: GIVING BROADBAND MEASUREMENT AN EDGE

- Scientific instrument for measuring broadband
- Deploy 1,000 Pis nationwide (~500 so far)
  - Chicago, IL; Milwaukee, WI; San Rafael, CA
  - Marion County, IL; Beaver Island, MI -- and others
- Measurement Applications
  - Netrics; Measurement Lab's (MLab) Measurement Swiss Army Knife (MSAK) toolkit; RADAR toolkit; NetUnicorn; rural broadband tests (ARA) – and others
- Data
  - 11M data points, spanning 17 providers (national and local), across multiple different technologies
  - Publicly available on FLOTO website
- How powerful is this dataset?
  - Marion County: 32% of sampled households below 25/3 Mbps federal threshold
  - Beaver Island: area challenge to FCC -> reassessment of broadband coverage

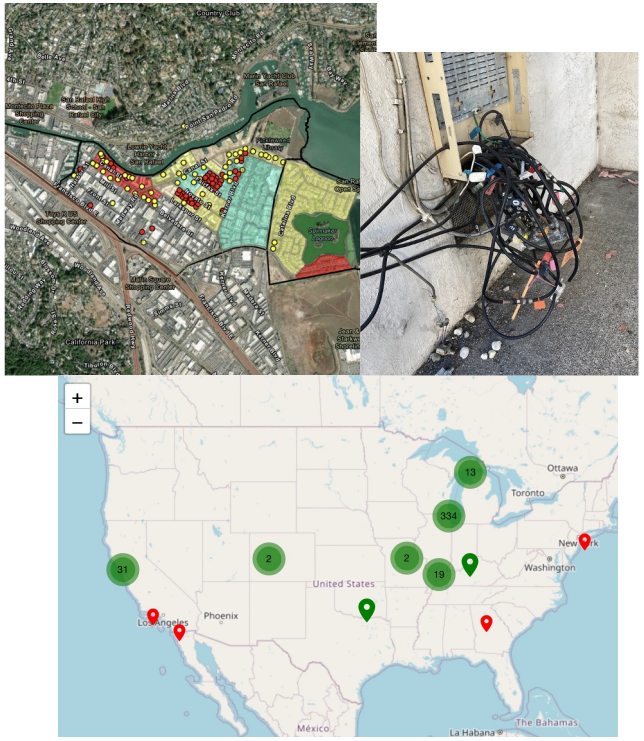

floto.cs.uchicago.edu

MEASURING RURAL WIRELESS

- Collaboration with ARA project
- Assessing the quality of rural 5G networks
  - Measuring device to device latency
  - Clock synchronization
  - Comparing over different network fabrics
- Deployed 6 Raspberry Pi devices with 5G connectivity in rural Iowa
- Latency measurements: GPS-based time synchronization for precise measurements (4000x more precise than NTP over 5G)
- Tested using Hadoop
- Hey presto: 5G networks can support distributed computing with performance comparable to wired connections!

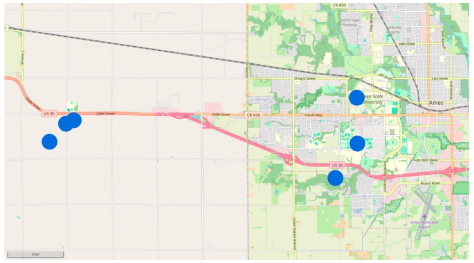

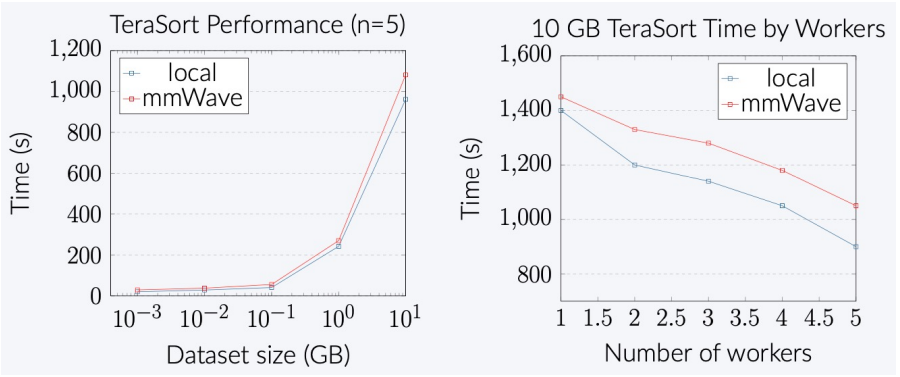

NCAR WEATHER SENSING STATIONS

- openIoTwx: NCAR 3D printed weather stations
- Richer continuum: IBIS SBCs connecting to openIoTws via LoRa
  - Exploring power (4x factor), connectivity (cellular vs aggregation via LoRa), sensing (additional camera sensors), and processing (to e.g., reduce size of data) trade-offs
- Future challenges
  - Image-based weather prediction methods, scaling up to create dense, high-resolution weather monitoring networks, and assessing long-term reliability in diverse outdoor environments

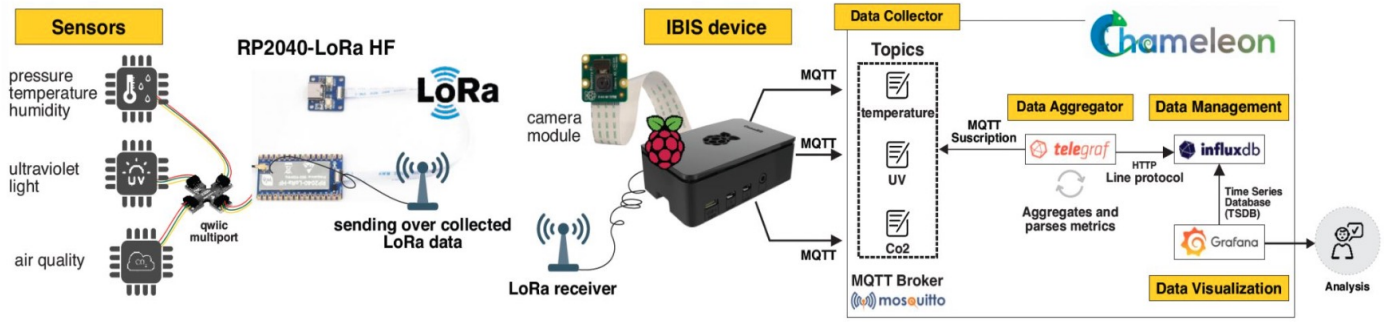

SENSOR STATIONS FOR MARINE AND COASTAL ECOSYSTEMS

- Smart buoy system: sensor stations for oceanic data collection (water quality, water movement, water levels, etc.)
- Collaboration with FIU
- Integrated multiple environmental sensors with IBIS infrastructure
- Demo deployment with real and simulated data
- Implemented cloud-based data visualization system
- Collaboration with FIU

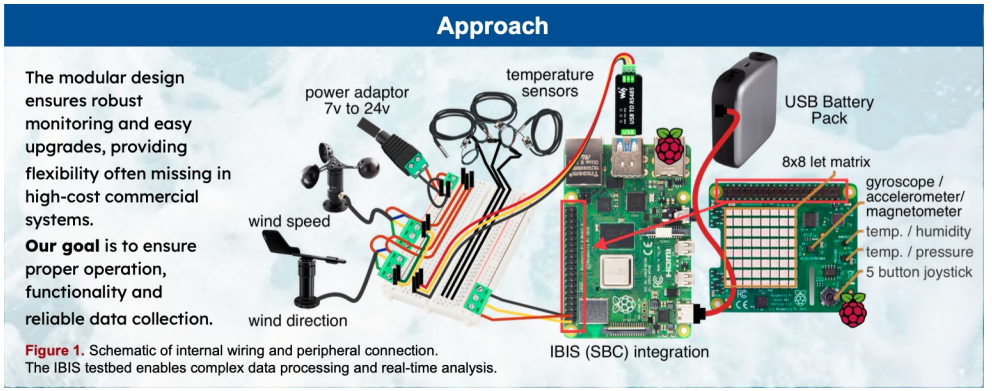

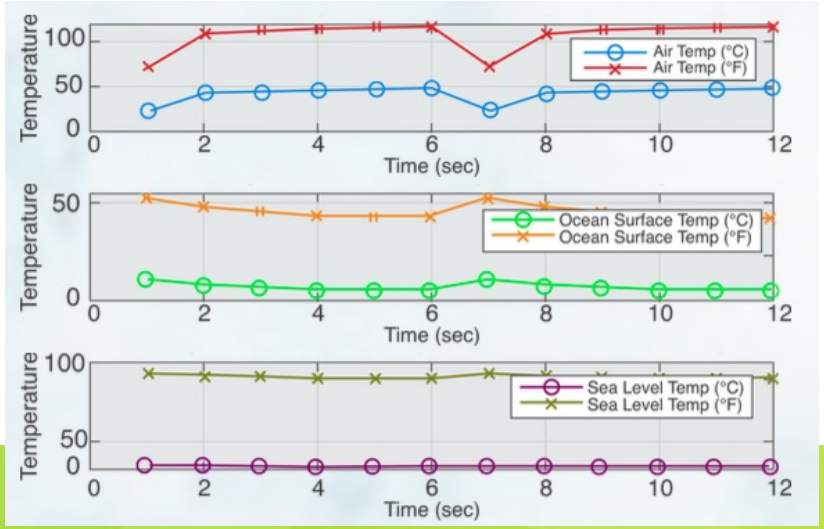

AUTO LEARN

- Can I experiment with self-driving cheaply? Can I teach edge to cloud AI concepts in a class in an engaging manner?
- AutoLearn Trovi artifact
  - Data collection (actual car versus simulator)
  - A library of ML modules
  - Verification via self-driving (actual car versus simulator)
- Different emphasis
  - Introduction to engineering
  - Machine learning with just the simulator
- Lots of scope for individual exploration

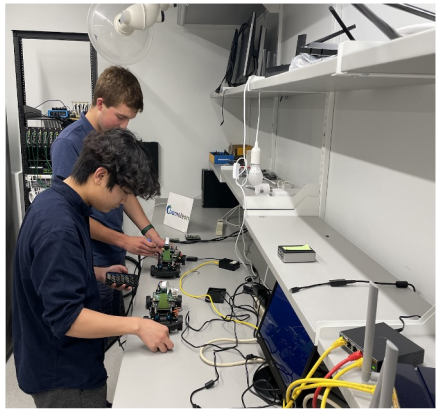

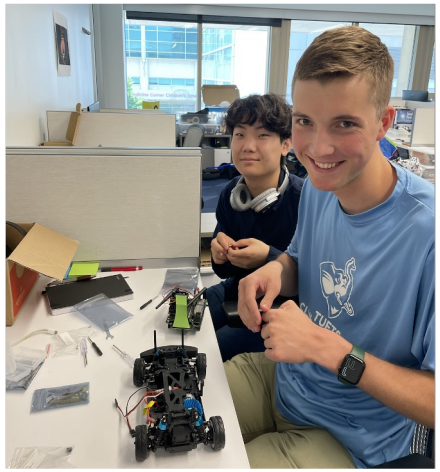

REU 2023 students working on hardware setup for autonomous vehicles

Paper: "AutoLearn: Learning in the Edge to Cloud Continuum", EduHPC'23

AND OTHERS…

- Predicting air quality with federated learning
- Soundscaping and forestry data analysis
- Precision agriculture: optimizing greenhouse environments
- Meteorologic monitoring system for ML-based weather forecasts
- And others…

SHARING SCIENCE DIGITALLY

- Can digital experiments be as sharable as papers are today?
- Can I simply integrate somebody's model into my research instead of reinventing the wheel and get to a new result faster?
- Can I discover something new through playing with somebody else's experiment?
- Can I develop exercises for my class based on most recent research results?
- Should our next Program Committee meeting be a reproducibility hackathon?

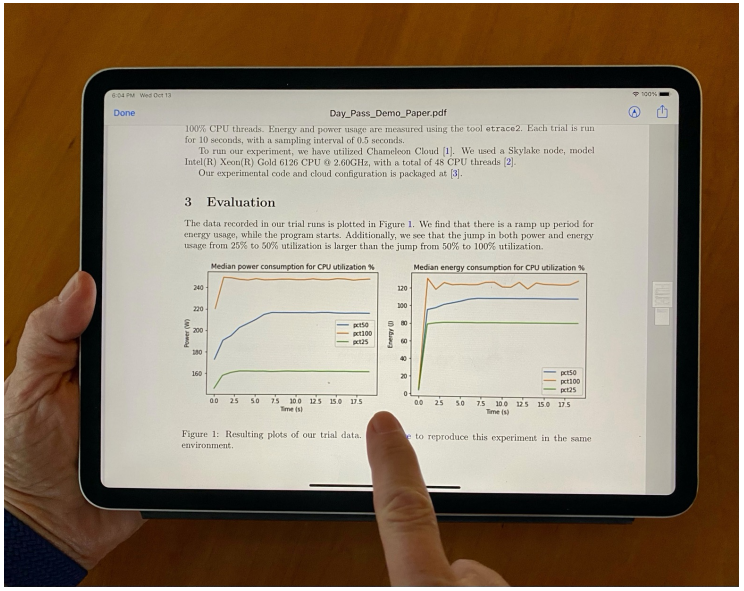

https://repeto.cs.uchicago.edu

5TH CHAMELEON USER MEETING (REPRODUCIBILITY)

- What is still a challenge for packaging and reproducing systems experiments?
  - Gather authors/reviewers from reproducibility initiatives relying on Chameleon (SC24 in particular)
  - Report and path forward on how to support reproducibility for computer science
  - Can we create a
- Date: November 18, 2024, Atlanta (same week as big SC conference)
- Keynote: Torsten Hoefler
- Details and registration in the announcement section of the Chameleon website or reproduciblehpc.org

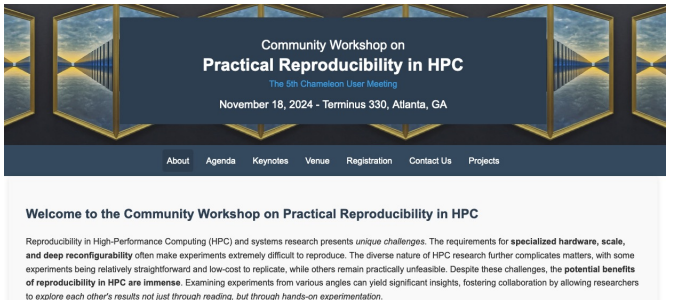

PARTING THOUGHTS

- Edge to Cloud: unprecedented access to observational data (and some analysis) combined with strong computational capabilities in the cloud
- It really takes a village
  - Partnerships with FABRIC, NRP, and PAWR projects to add a different dimension
- The different facets of continuum
  - Hardware/capability continuum, power continuum, configuration continuum, operational expertise continuum, connectivity continuum, processing continuum, etc.
- Configuration continuum: bare metal versus virtualization faceoff
- Working with the community on making great research reproducible!

We're here to change

www.chameleoncloud.org