Where's The Bear? -- Automating Wildlife Image Processing Using IoT and Edge Cloud Systems

Andy Rosales Elias, Nevena Golubovic, Rich Wolski, and **Chandra Krintz**

Dept. of Computer Science
UC Santa Barbara
ckrintz@cs.ucsb.edu

IOTDI - Apr. 18-20, 2017

Wildlife Monitoring

- Evaluating diversity, species, and habitat health
- Extracting patterns in activity and behavior of animals
- Monitoring change in land use
- Avoiding dangerous human/animal encounters & overlap
- Educational experiences
- Citizen science

UCSB

Digital Photography

- Alternative to labor-intensive observation & tracking
- Cost effective and scalable
- Safe and non-invasive
- Increasingly autonomous
  - Motion triggered
  - IoT devices

- Goal: Leverage IoT & cloud to improve wildlife monitoring

UCSB

Image Classification

- Enormous number of images (size and count)
- Classification and identification once done by humans
  - Now automated by data analysis tools in the cloud
    - Google Tensorflow, Caffe, Torch, Theano, CNTK...

UCSB

Automatic Image Classification: Training

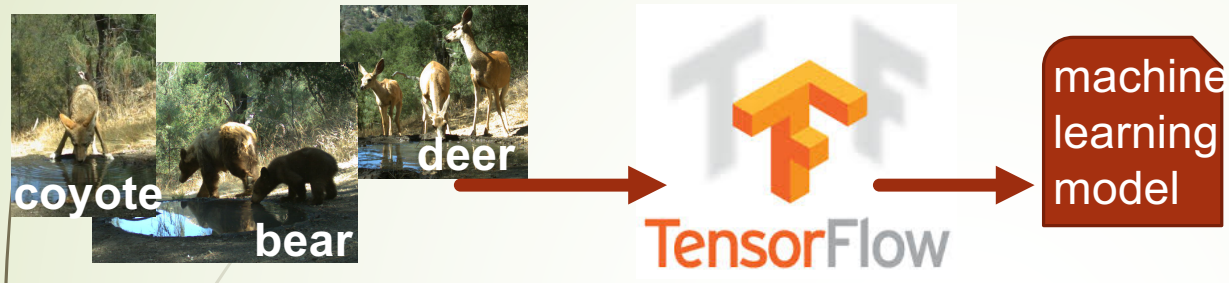

- Process: train a model then use it to classify images
  - Training set (large, ***manually*** labeled, ground truth)
    - Requires lots of compute power and GPUs (**thanks *public cloud***!)
    - Can take days to weeks; but can be incremental

UCSB

Automatic Image Classification

- Process: train a model then use it to classify images
  - Training set (large, ***manually*** labeled, ground truth)
    - Requires lots of compute power and GPUs (**thanks *public cloud***!)
    - Can take days to weeks; but can be incremental
  - Classification **is fast and uses many fewer resources**

UCSB

Challenges With IoT+Cloud for Wild Life Monitoring

- Enormous numbers of images (size and count)
  - Cameras limited storage & processing power
- Automatic classification requires labeling by humans
  - **Images must be moved to where they are processed**

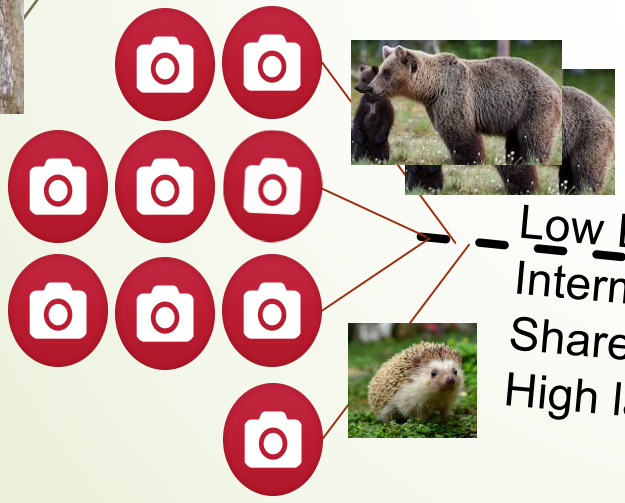

Challenges With IoT+Cloud for Wild Life Monitoring

- Enormous numbers of images (size and count)
  - Cameras limited storage & processing power
- Automatic classification requires labeling by humans
  - **Images must be moved to where they are processed**
  - **Extremely time consuming, tedious, and error prone**

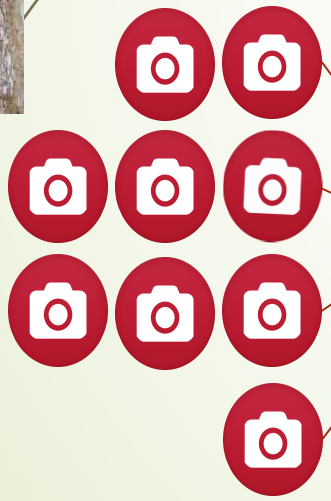

Low bandwidth Network
Intermittent connectivity
Shared, costly
High latency

Where's The Bear (WTB)

- Multi-tier IoT system
  - Move the ***code to the data (images)*** not vice versa
    - Via ***Edge Clouds:*** robust, self-managing appliances, on-site
      - Low latency, high-bandwidth direct connectivity to cameras
      - Local image classification

Where's The Bear (WTB)

- Multi-tier IoT system
  - Move the ***code to the data (images)*** not vice versa
    - Via ***Edge Clouds:*** robust, self-managing appliances, on-site
      - Low latency, high-bandwidth direct connectivity to cameras
      - Local image classification

WTB: Build Model Using Public Cloud With "Fake" Images

Public Cloud

Produce many 1000s of labeled training images automatically

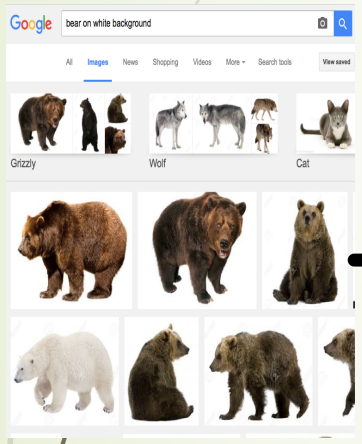

Google Images

Subtract background (OpenCV)

+

Empty Images

Random Images

Transmit a few "empty" Images **from edge to public** cloud (fast!)

UCSB

WTB: Build Model Using Public Cloud With "Fake" Images

Public Cloud

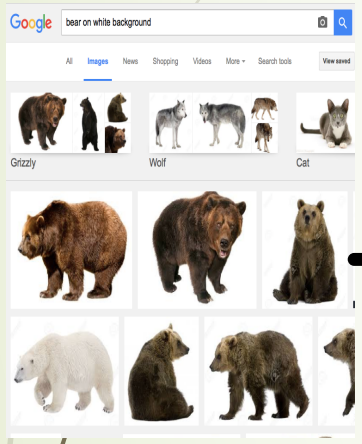

Google Images

Subtract background (OpenCV)

+

Empty Images

Random Images

Train the model

machine learning model

Transmit (tiny) model **back** to edge cloud (fast!)

UCSB

Where's The Bear (WTB)

- Multi-tier IoT system
  - Move the ***code to the data (images)*** not vice versa
    - Via ***Edge Clouds:*** robust, self-managing appliances, on-site
      - Low latency, high-bandwidth direct connectivity to cameras
      - Local image classification using model trained in public cloud

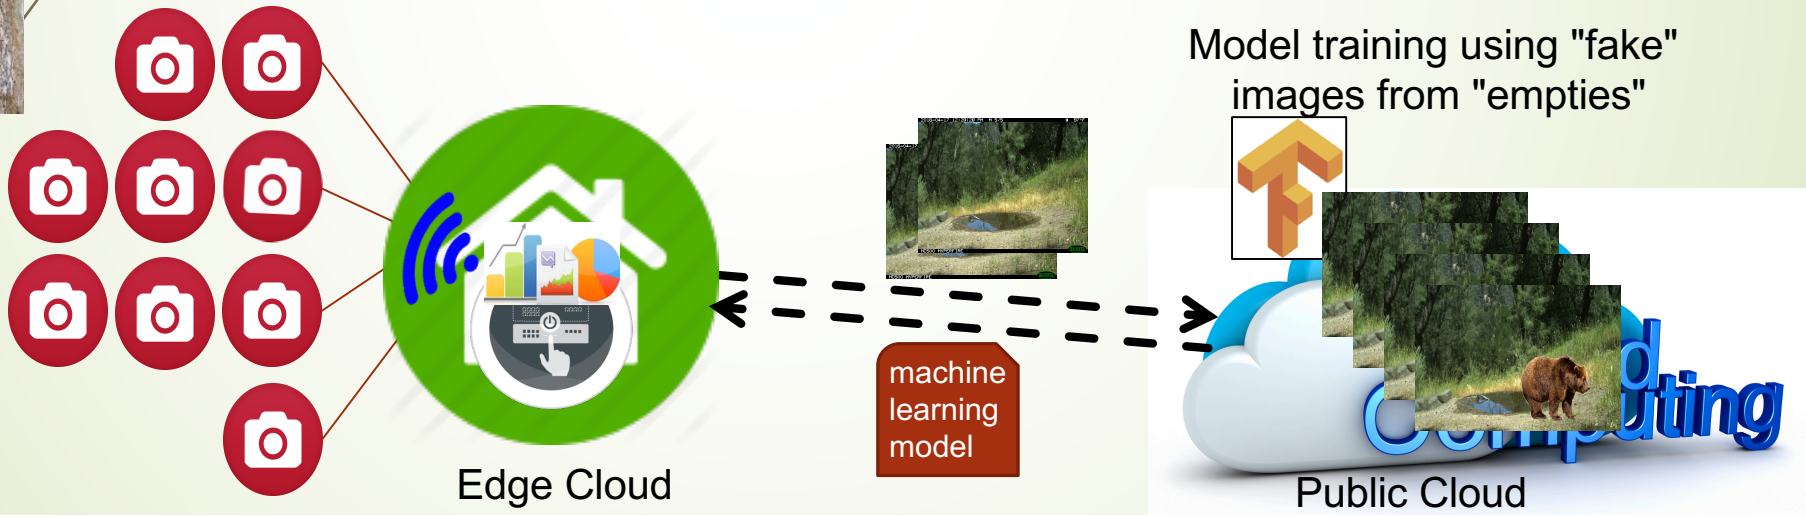

Where's The Bear (WTB)

- Multi-tier IoT system
  - Move the ***code to the data (images)*** not vice versa
    - Via ***Edge Clouds:*** robust, self-managing appliances, on-site
      - Low latency, high-bandwidth direct connectivity to cameras
      - Local image classification

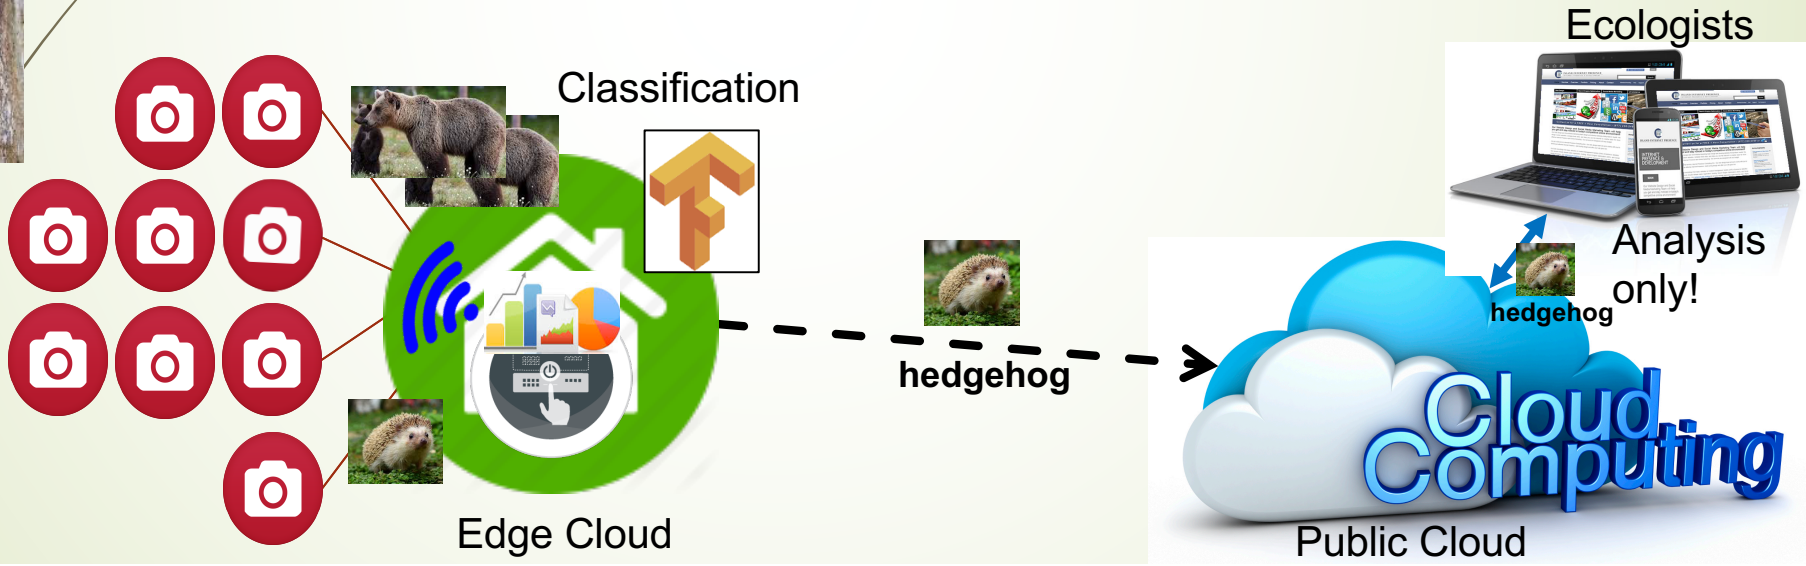

Where's The Bear (WTB)

- Multi-tier IoT system
  - Move the ***code to the data (images)*** not vice versa
    - Via ***Edge Clouds:*** robust, self-managing appliances, on-site
      - Low latency, high-bandwidth direct connectivity to cameras
      - Local image classification

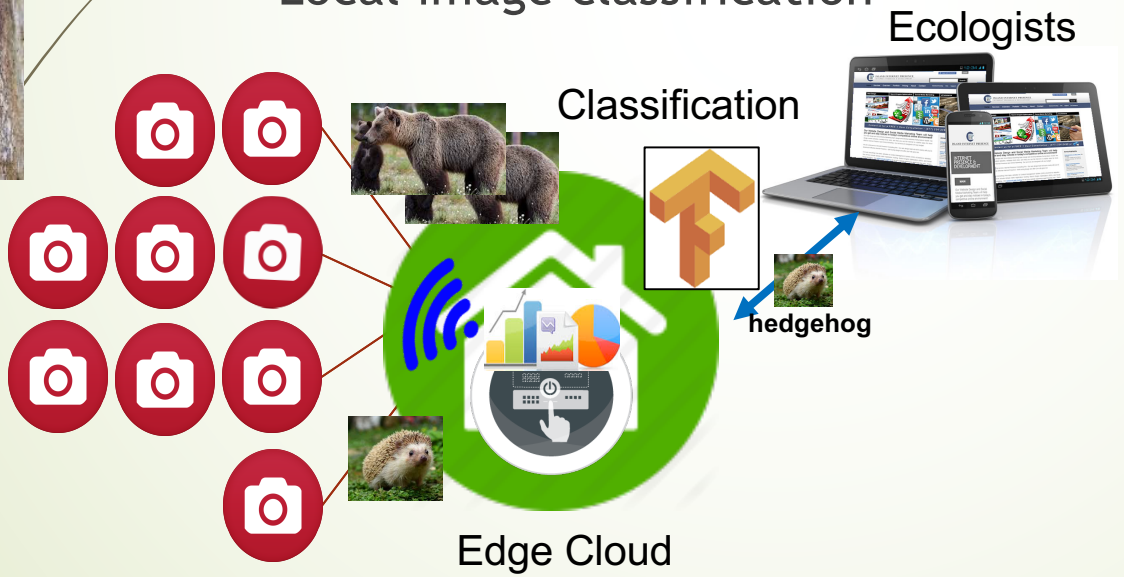

Deployment and Empirical Methodology

SEDGWICK RESERVE

UC SANTA BARBARA NATURAL RESERVE SYSTEM

UCSB

Total: 1.12 Million images: 714GB

This study: 238K images: 93GB

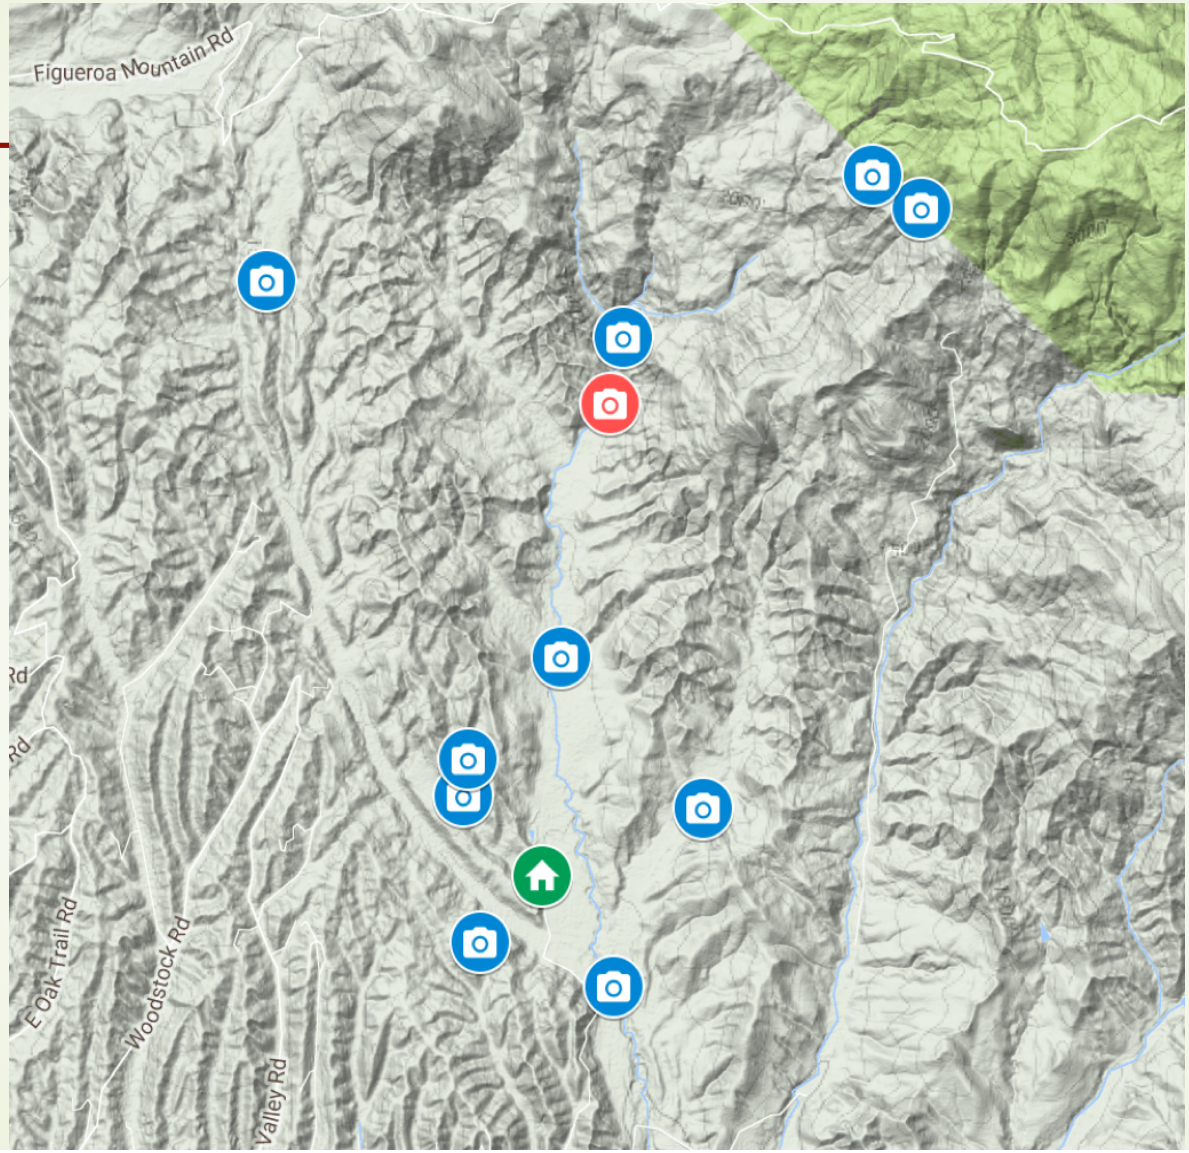

UCSB Sedgwick Reserve

Motion-triggered Wildlife Camera Traps

9 sq. miles

UCSB

Where’s the Bear?

Results: 4890 randomly selected images

Transmit only those of interest

1473 vs 4890

= 1.6hrs vs 5.3hrs

Which also saves researcher time (1250 images/hr)

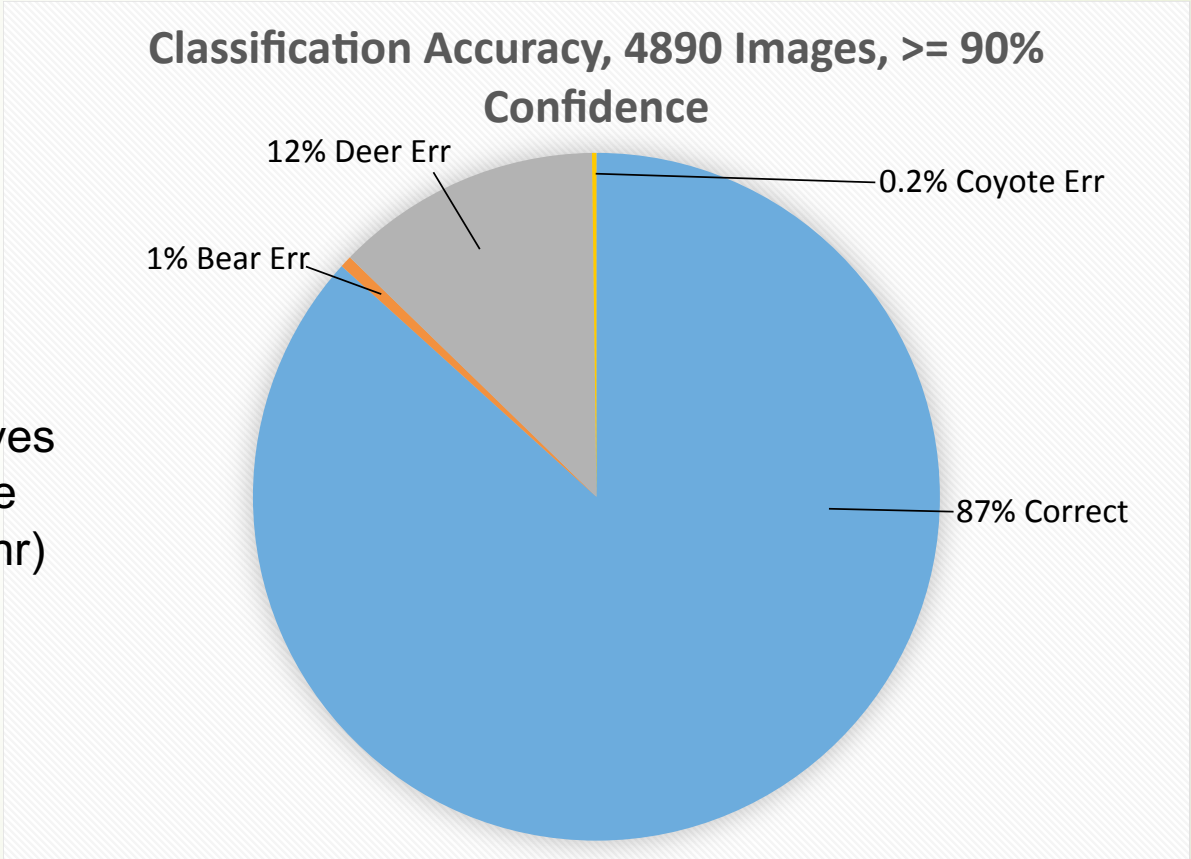

UCSB

Challenging Images

UCSB

WTB Findings and Future Work

- ***For IoT, the cloud model alone does not work***
  - App and deployment model must change to fit IoT
  - Can save significant network use (image transfer avoidance)
  - Can save significant researcher time (automatic classification)
    - While maintaining high accuracy
- ***Able to classify animals that only appear rarely***
  - In insufficient numbers to train with
- Next steps
  - Small animals vs empty images, improving accuracy
  - Identifying features, counting
  - Making edge cloud robust to outage and faults

UCSB

Thanks!

- Current Student Researchers
  - Will Berman, Kyle Carson, Stratos Dimopoulos, Jonathan Easterman, Angad Gill, Nevena Golubovic, Benji Lampel, Wei-Tsung Lin, Kevin Malta, Andy Rosales Elias, Michael Zhang
- Collaborators
  - Rich Wolski (UCSB)
  - Grant Canova-Parker, Kate McCurdy, Tony Nelson (UCSB Sedgwick)
- Support
  - Google, IBM Research, Huawei, NSF, NIH, California Energy Commission

http://www.cs.ucsb.edu/~ckrintz **ckrintz@cs.ucsb.edu**
http://www.cs.ucsb.edu/~ckrintz/racelab.html
http://www.linkedin.com/in/**chandrakrintz**

UCSB