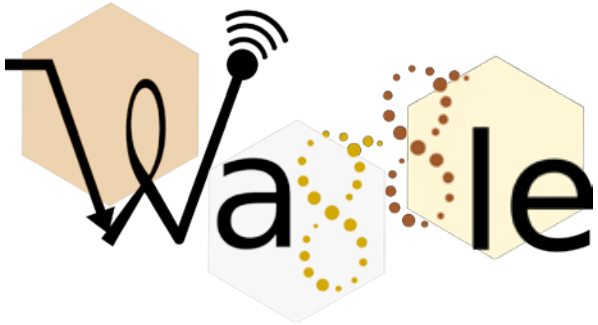

Parallel Computing at the Edge: Deploying Parallel Computers and Sensors on Chicago Street Poles

Pete Beckman, Charlie Catlett, Rajesh Sankaran, Nicola Ferrier, Rob Jacob, Mike Papka, and more….


Argonne National Laboratory
$675M /yr budget
3,200 employees
1,450 scientists/eng
750 Ph.D.s

Argonne's Next Big Supercomputer: Aurora

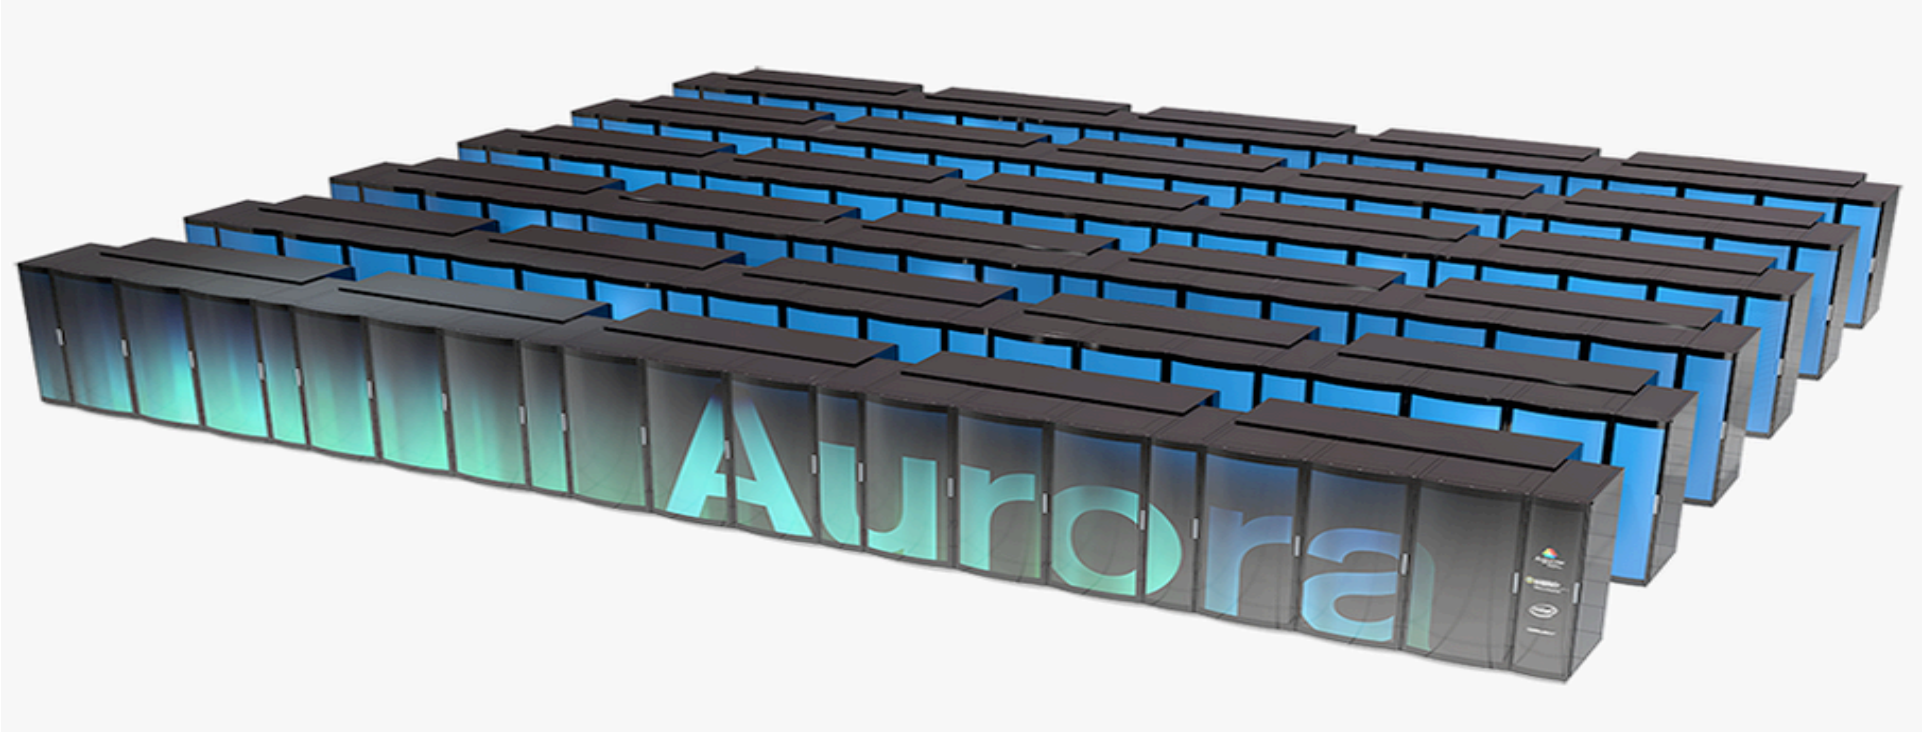

Big Sensor Science

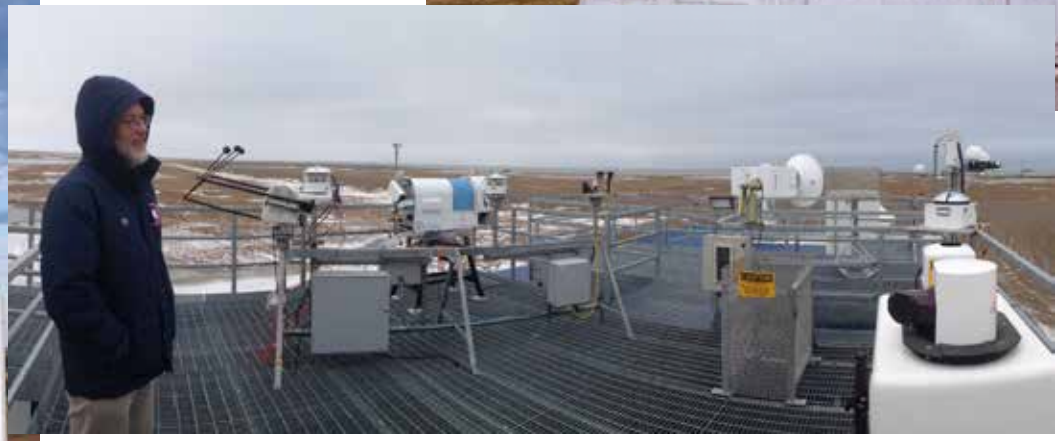

Big, Expensive, Precise, Sparse

Little Sensor Science

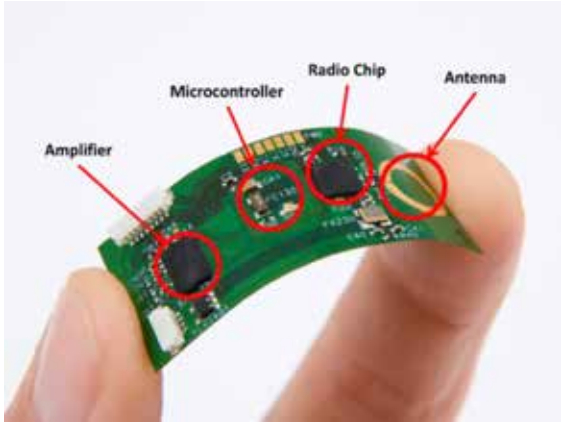

Small, Cheap, Imprecise, Dense

(almost no on-board processing)

Waggle: An Open Platform for ***Intelligent*** Sensors

Exploiting Disruptive Technology, *Edge Computing*, Resilient Design

Machine Learning

Computer Vision

Novel Sensors

Nano / MEMS

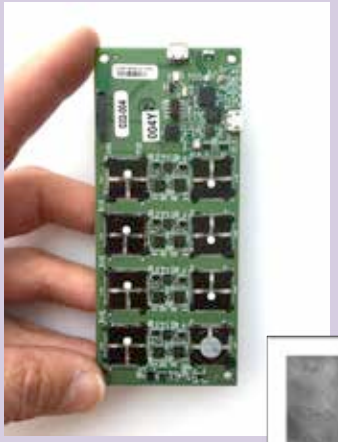

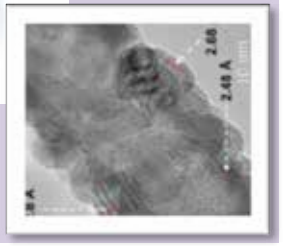

Low Power CPUs

GPU / Smartphones

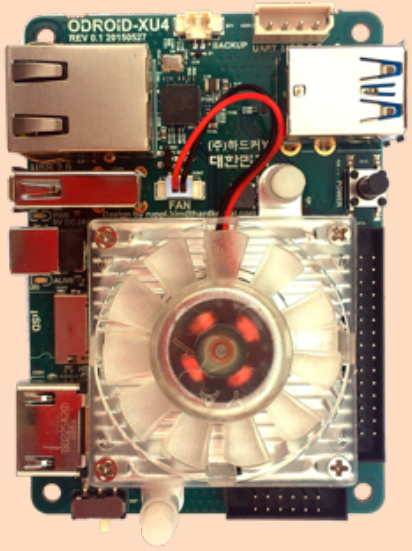

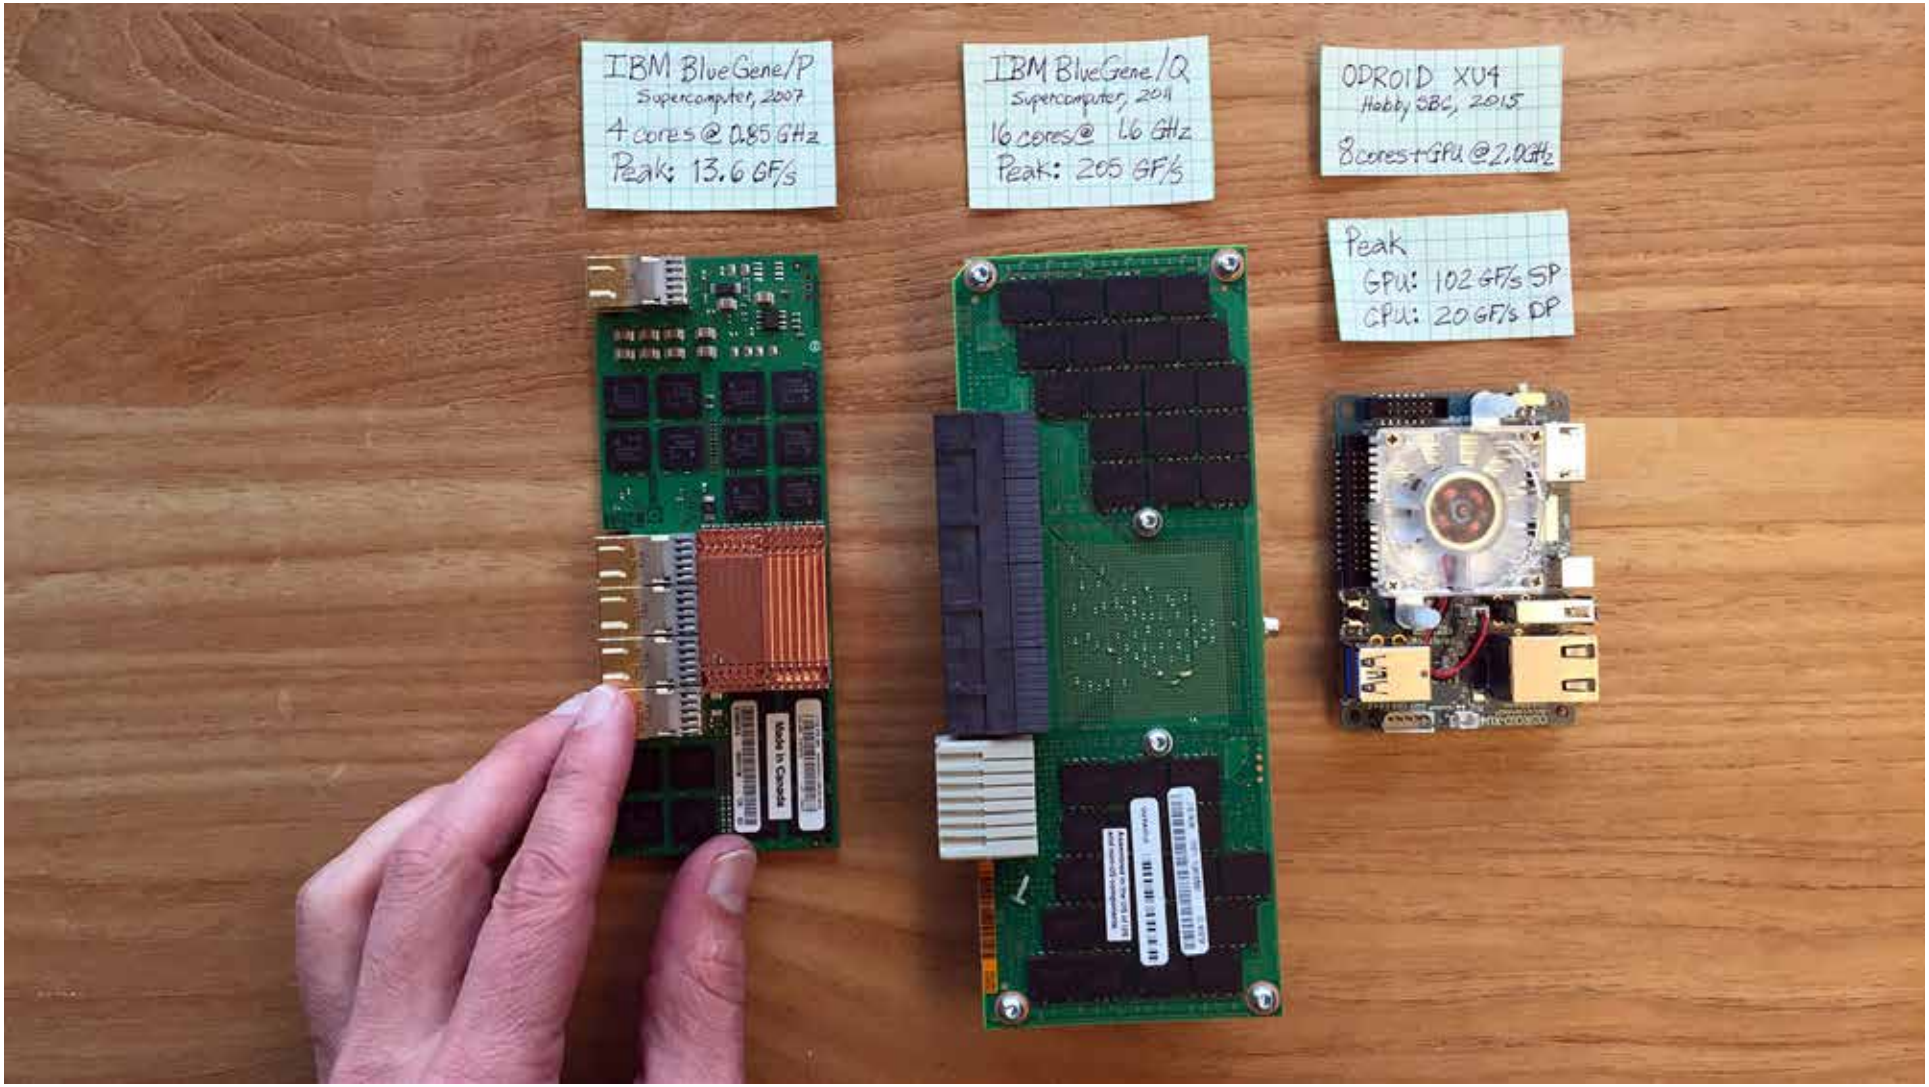
IBM BlueGene/P
Supercomputer, 2007
4 cores @ 0.85 GHz
Peak: 13.6 GF/s
IBM BlueGene/Q
Supercomputer, 2011
16 cores @ 1.6 GHz
Peak: 205 GF/s
ODROID XU4
Hobby SBC, 2015
8 cores + GPU @ 2.0 GHz
Peak
GPU: 102 GF/s SP
CPU: 20 GF/s DP

Powerful, Resilient & Hackable

Multiple boot media (µSD / eMMC)

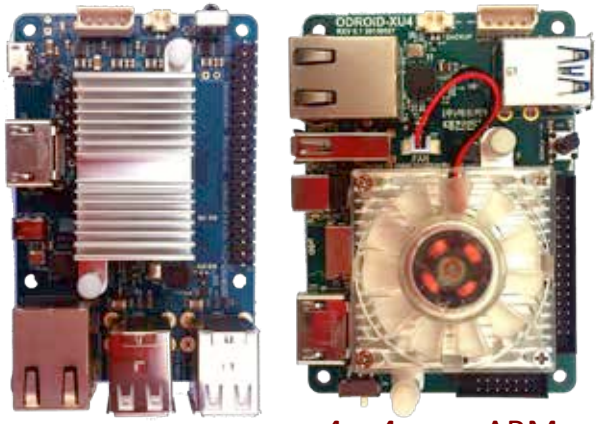

4-core ARM

4 + 4-core ARM
8-core GPU

Node Control &
Communications

In-Situ / Edge
Processing

Linux Development Environment

Relays

Current Sensors

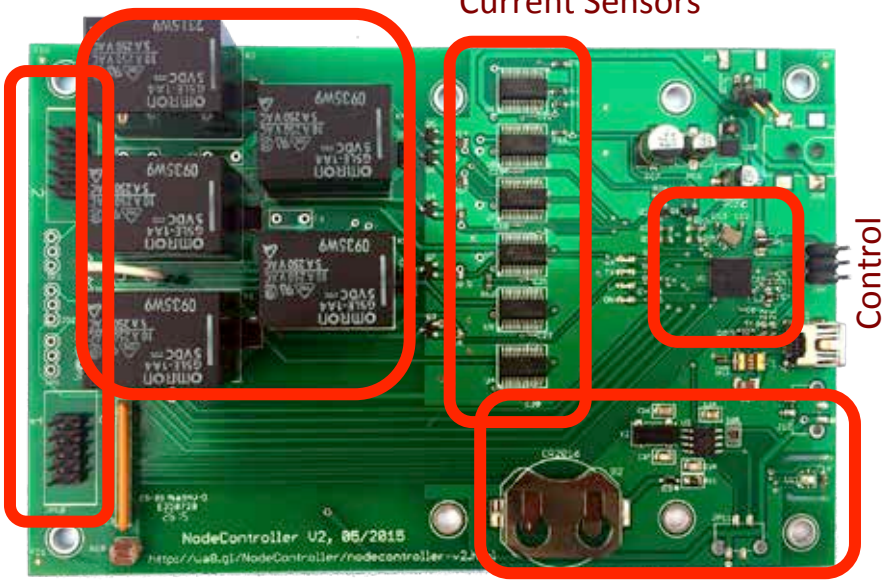

Real time clock & Internal sensors

"Deep Space Probe" Design

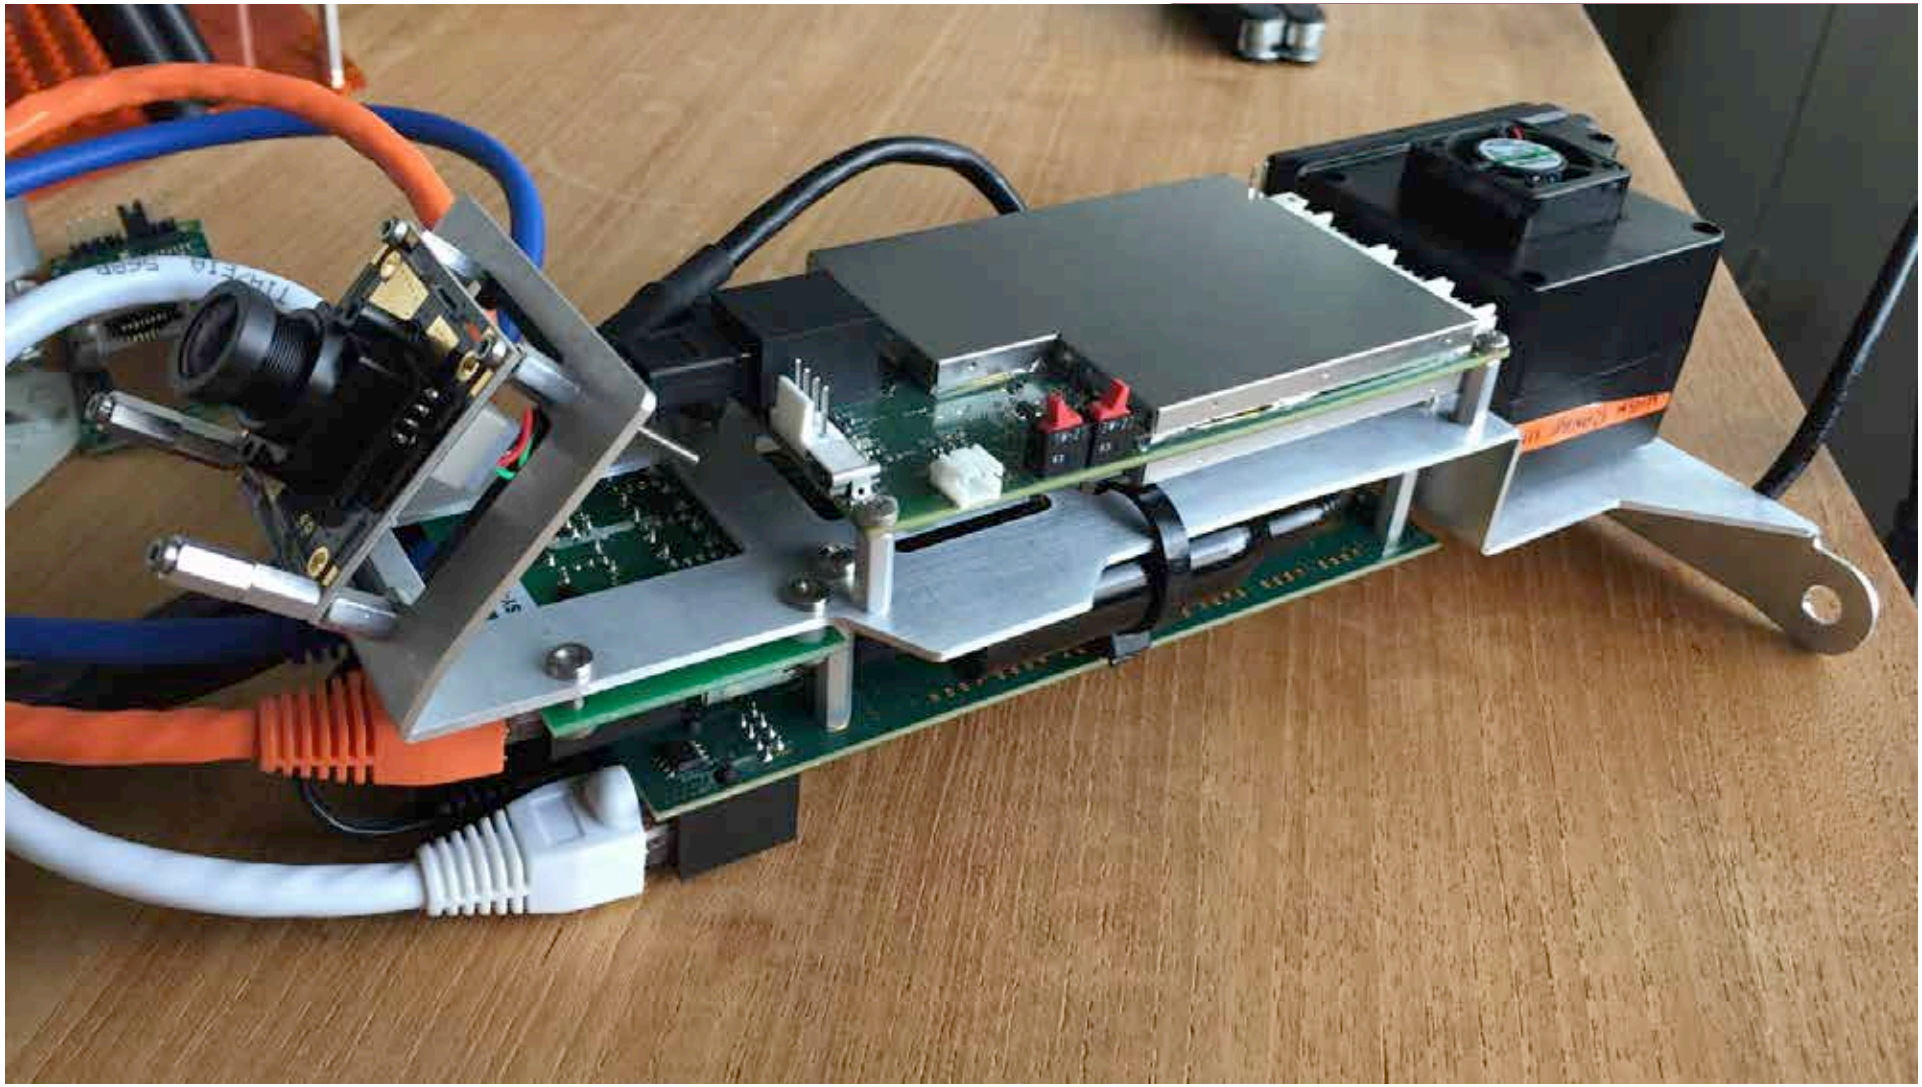

Edge Computing: Analysis and Feature Recognition

- Parallel Computing
- Open Platform
- Deep Learning

Waggle Machine Learning & Edge Computing

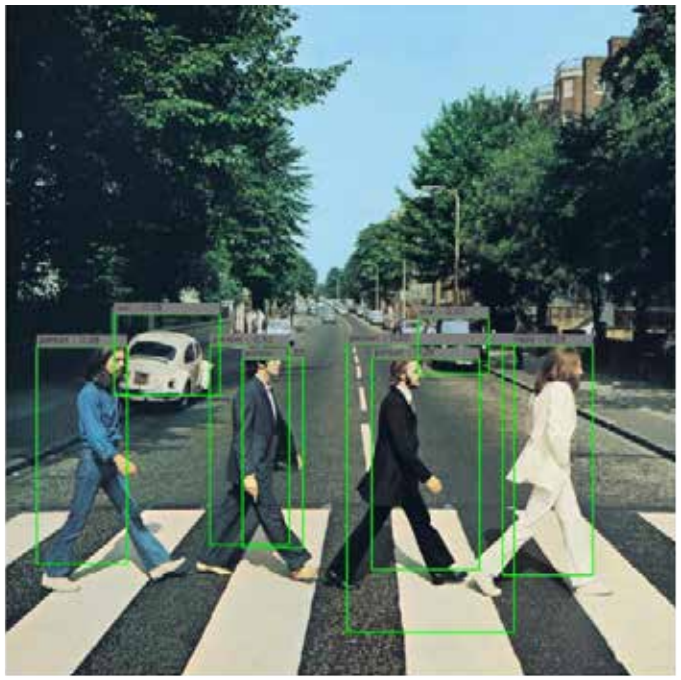

https://waggle-sensor.slack.com/files/noaholsman/F243LQL66/output.jpg

- We are exploring Caffe & OpenCV
  - Convolutional Neural Networks
- Training will be done on systems at Argonne
- Classification on Waggle

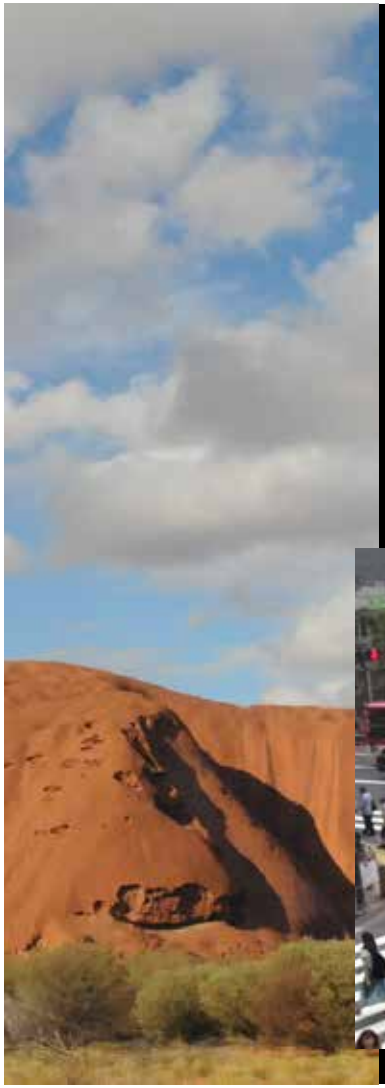


9 FT 6 IN

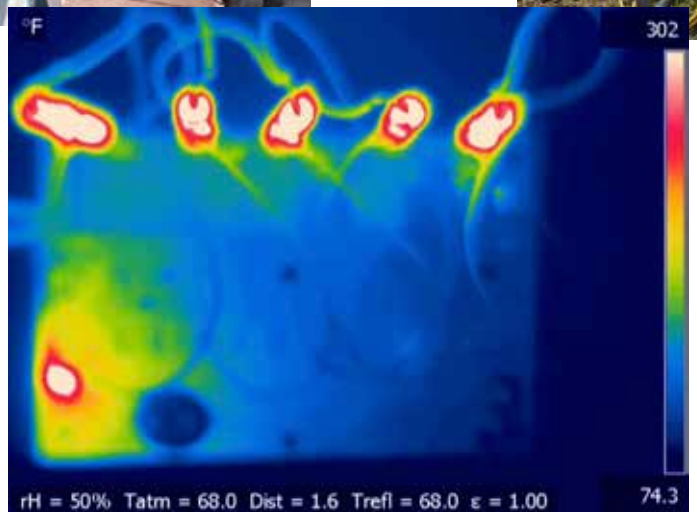

Waggle / AoT Robust Testing

The Data

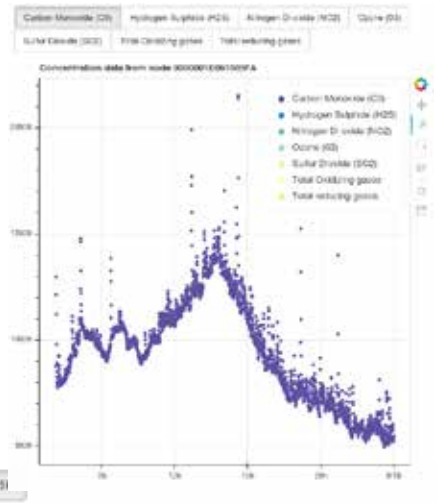

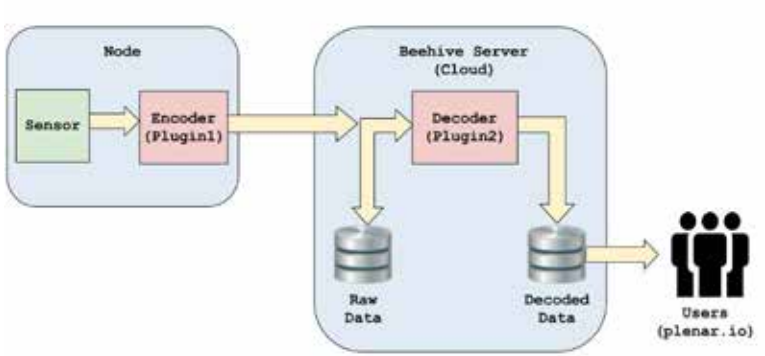
Node
Sensor
Encoder
(Plugin1)
Beehive Server
(Cloud)
Decoder
(Plugin2)
Raw
Data
Decoded
Data
Users
(plenar.io)

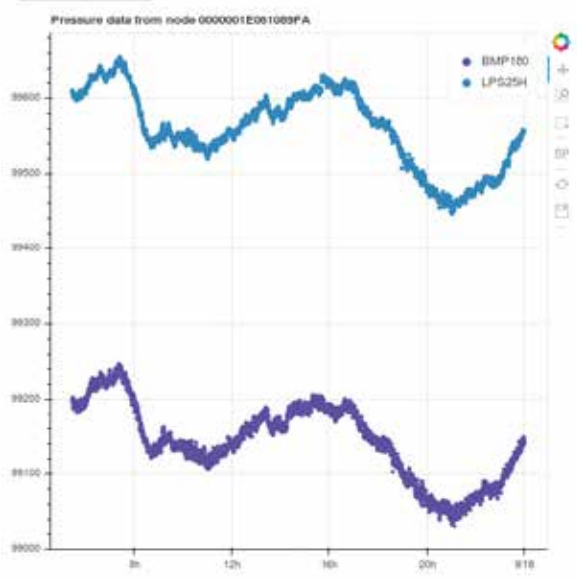
BMP180
LPS25H
Pressure data from node 0000001E061089FA

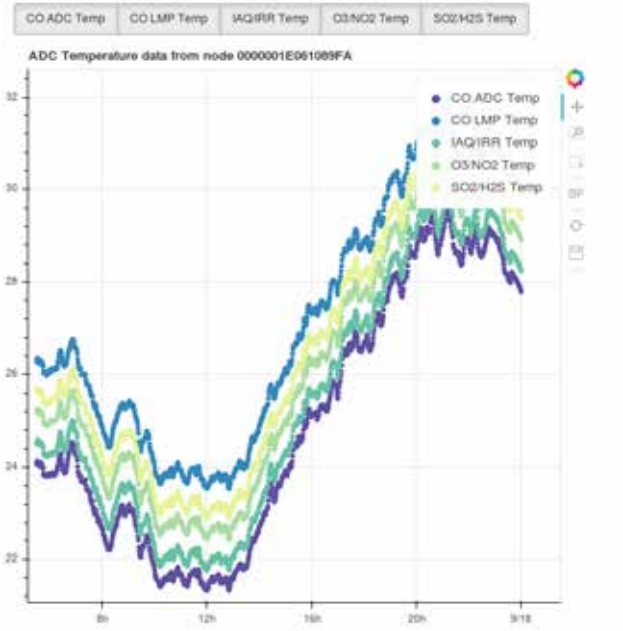
CO ADC Temp
CO LMP Temp
IAQ/IRR Temp
O3/NO2 Temp
SO2/H2S Temp
ADC Temperature data from node 0000001E061089FA

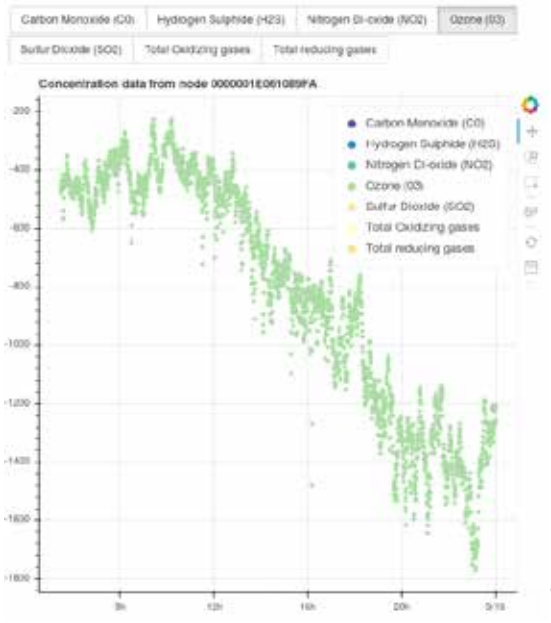
Carbon Monoxide (CO)
Hydrogen Sulphide (H2S)
Nitrogen Di-oxide (NO2)
Ozone (O3)
Sulfur Dioxide (SO2)
Total Oxidizing gases
Total reducing gases
Concentration data from node 0000001E061089FA

Deploying Waggle

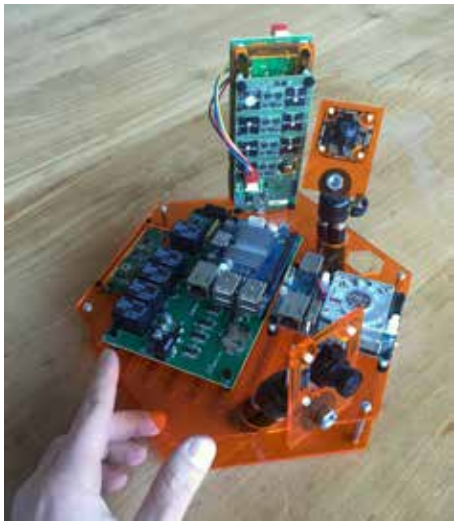

Waggle: A Platform for Research

- **Open Source / Open Platform**
  - Reusable, extensible software communities
- **Machine Learning: Computer Vision**
  - Data must be reduced in-situ
- **Novel Sensors**: Nano / MEMS / µfluidics
  - Explosion of nano/MEMS & imaging tech
- **Low-Power CPUs**: GPU / Smartphones
  - Powerful, low-power, smartphone CPUs

Opportunity: Big Data + Predictive Models

Smart Sensors + Supercomputers/Cloud Computing = predictions and analysis

The Array of Things

THE UNIVERSITY OF CHICAGO

Argonne NATIONAL LABORATORY

URBAN CENTER FOR COMPUTATION AND DATA


THE UNIVERSITY OF
CHICAGO

Why Focus on Cities?

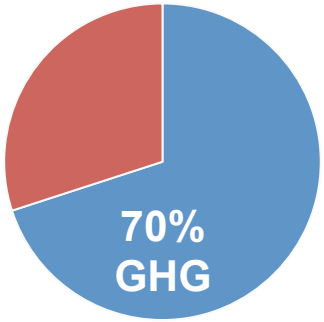

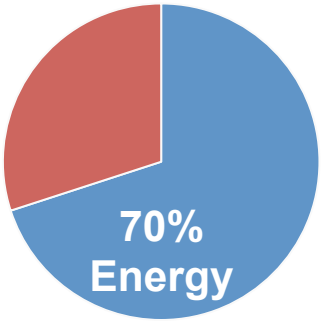

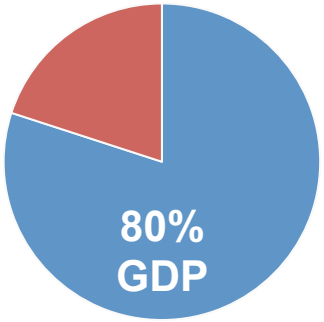

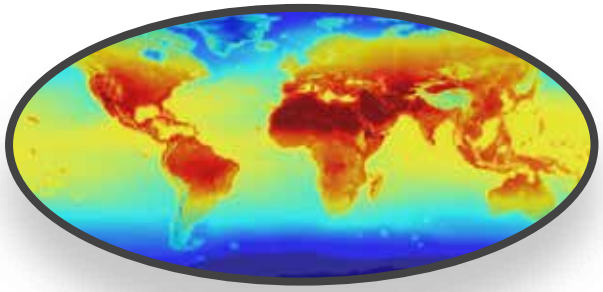

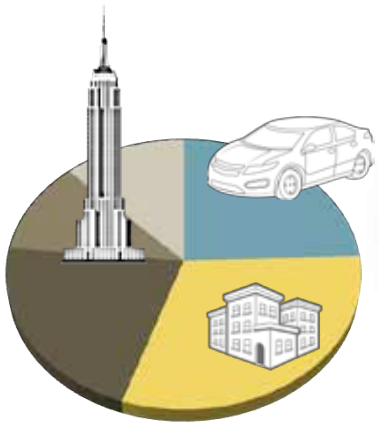

Why measure cities?

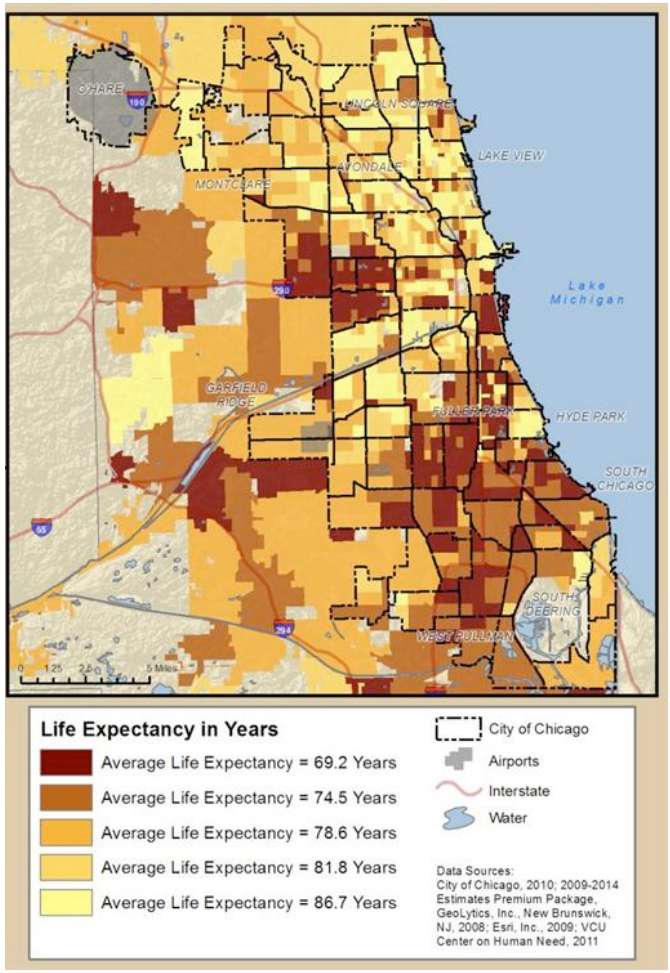

VISION ZERO
FOCUS CITIES
Austin | Boston | Chicago | DC | Fort Lauderdale | LA | NYC | Portland | San Francisco | Seattle

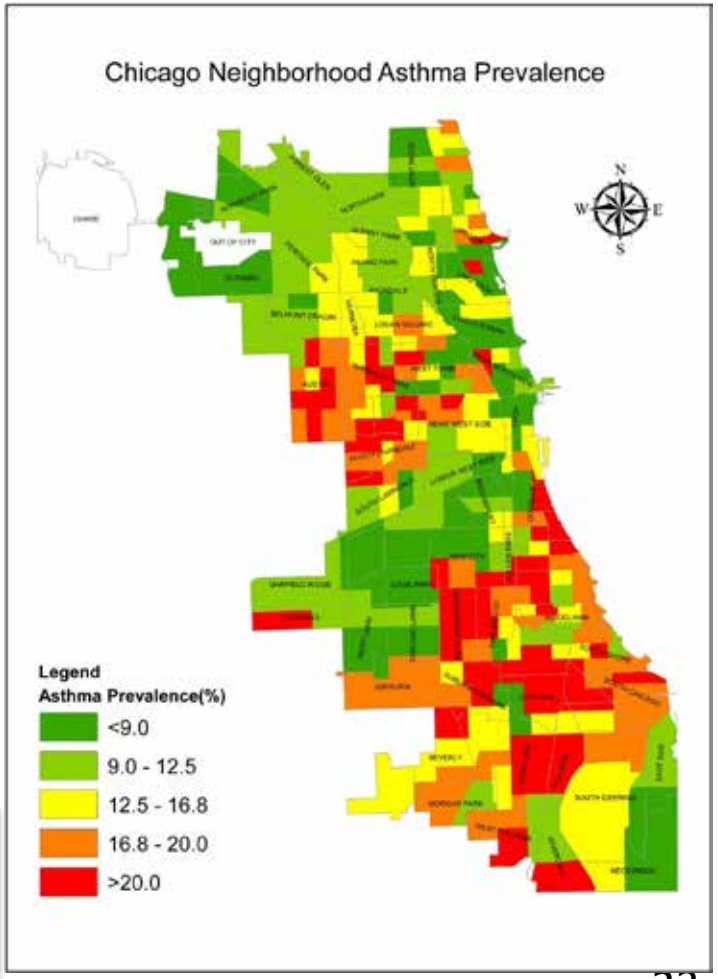

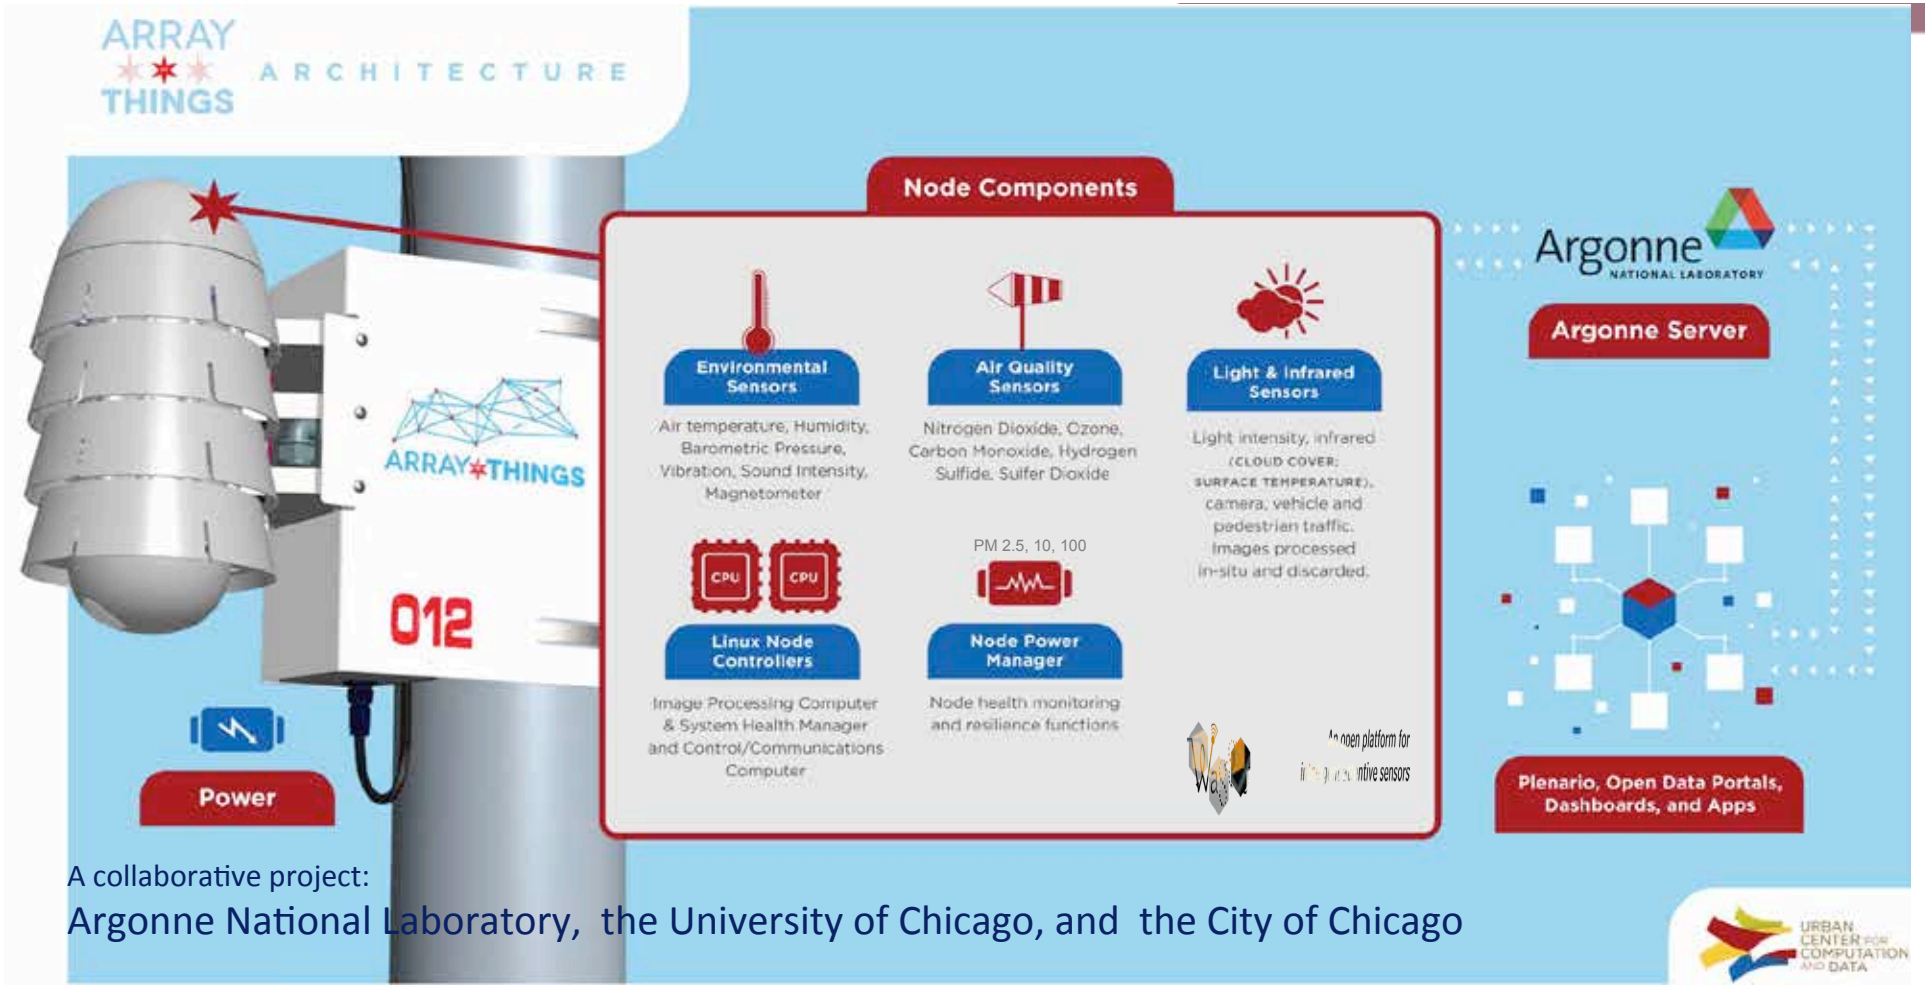

A collaborative project:

Argonne National Laboratory, the University of Chicago, and the City of Chicago

Supported by collaborating institutions and the U.S. National Science Foundation.
Industry In-Kind partners: AT&T, Cisco, Intel, Microsoft, Motorola Solutions, Schneider Electric, Zebra

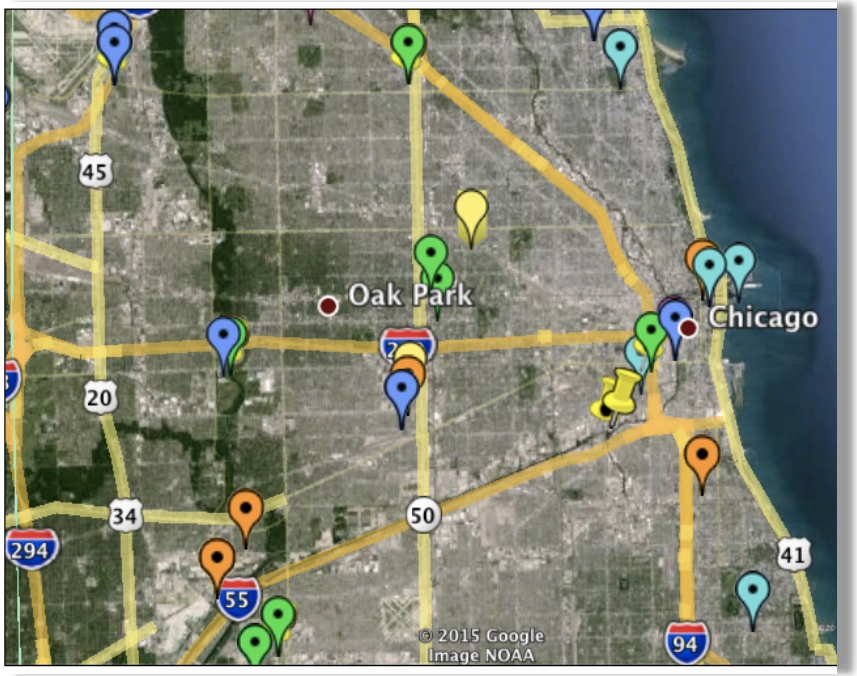
45
Oak Park
Chicago
20
34
50
41
294
55
94
© 2015 Google
Image NOAA

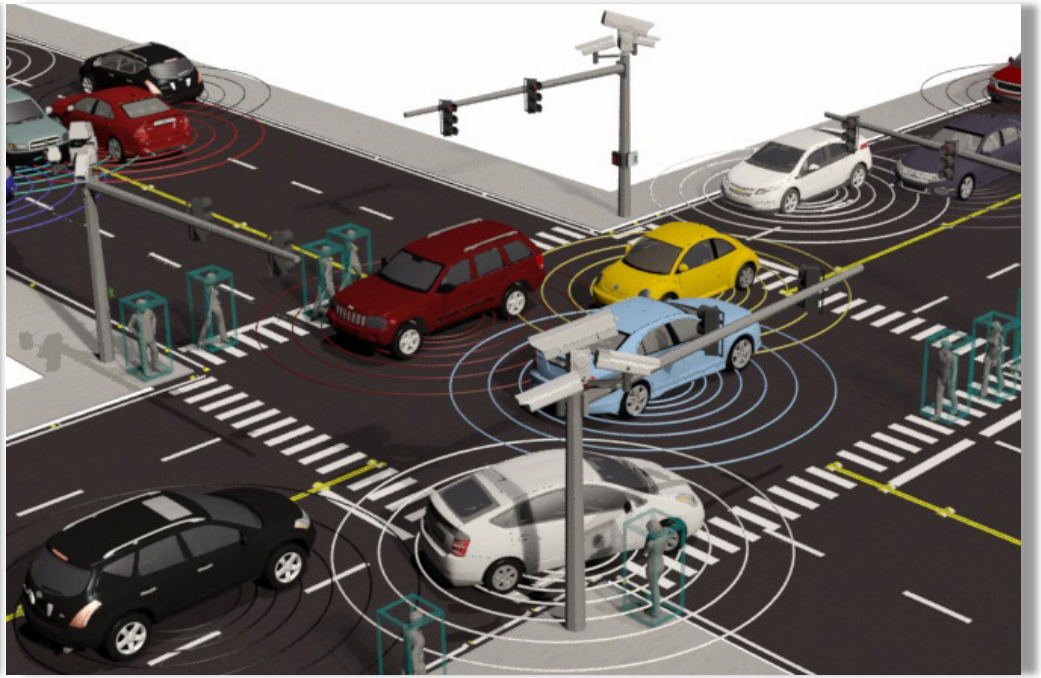


CDOT


N State St


Damen & Ashland
PM 2.5 Alert
Power Outage

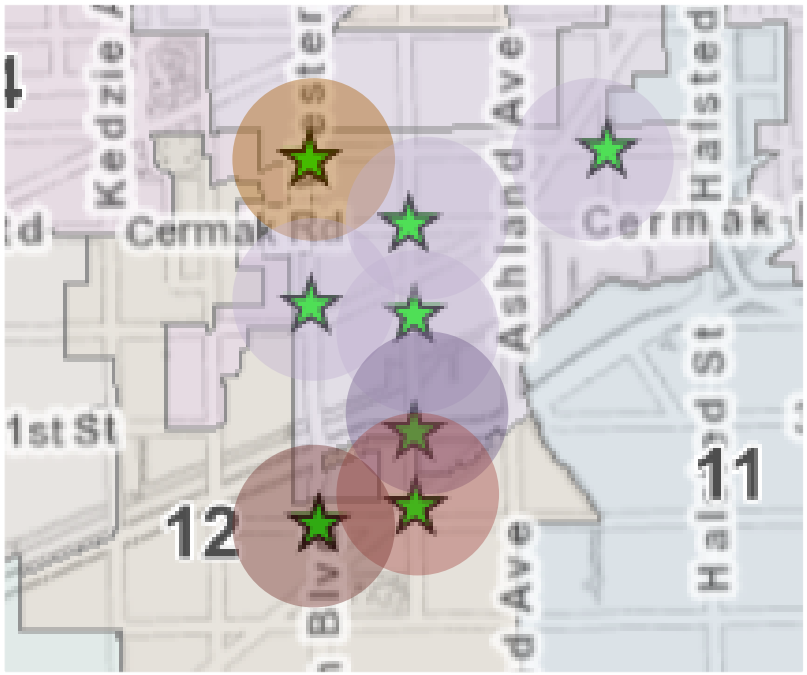
Cermak Rd
Cermak
Ashland Ave
Halsted St
1st St
12
11


LANE TECH COLLEGE PREP
STEM
science • technology • engineering • mathematics
motorola
foundation
SAIC

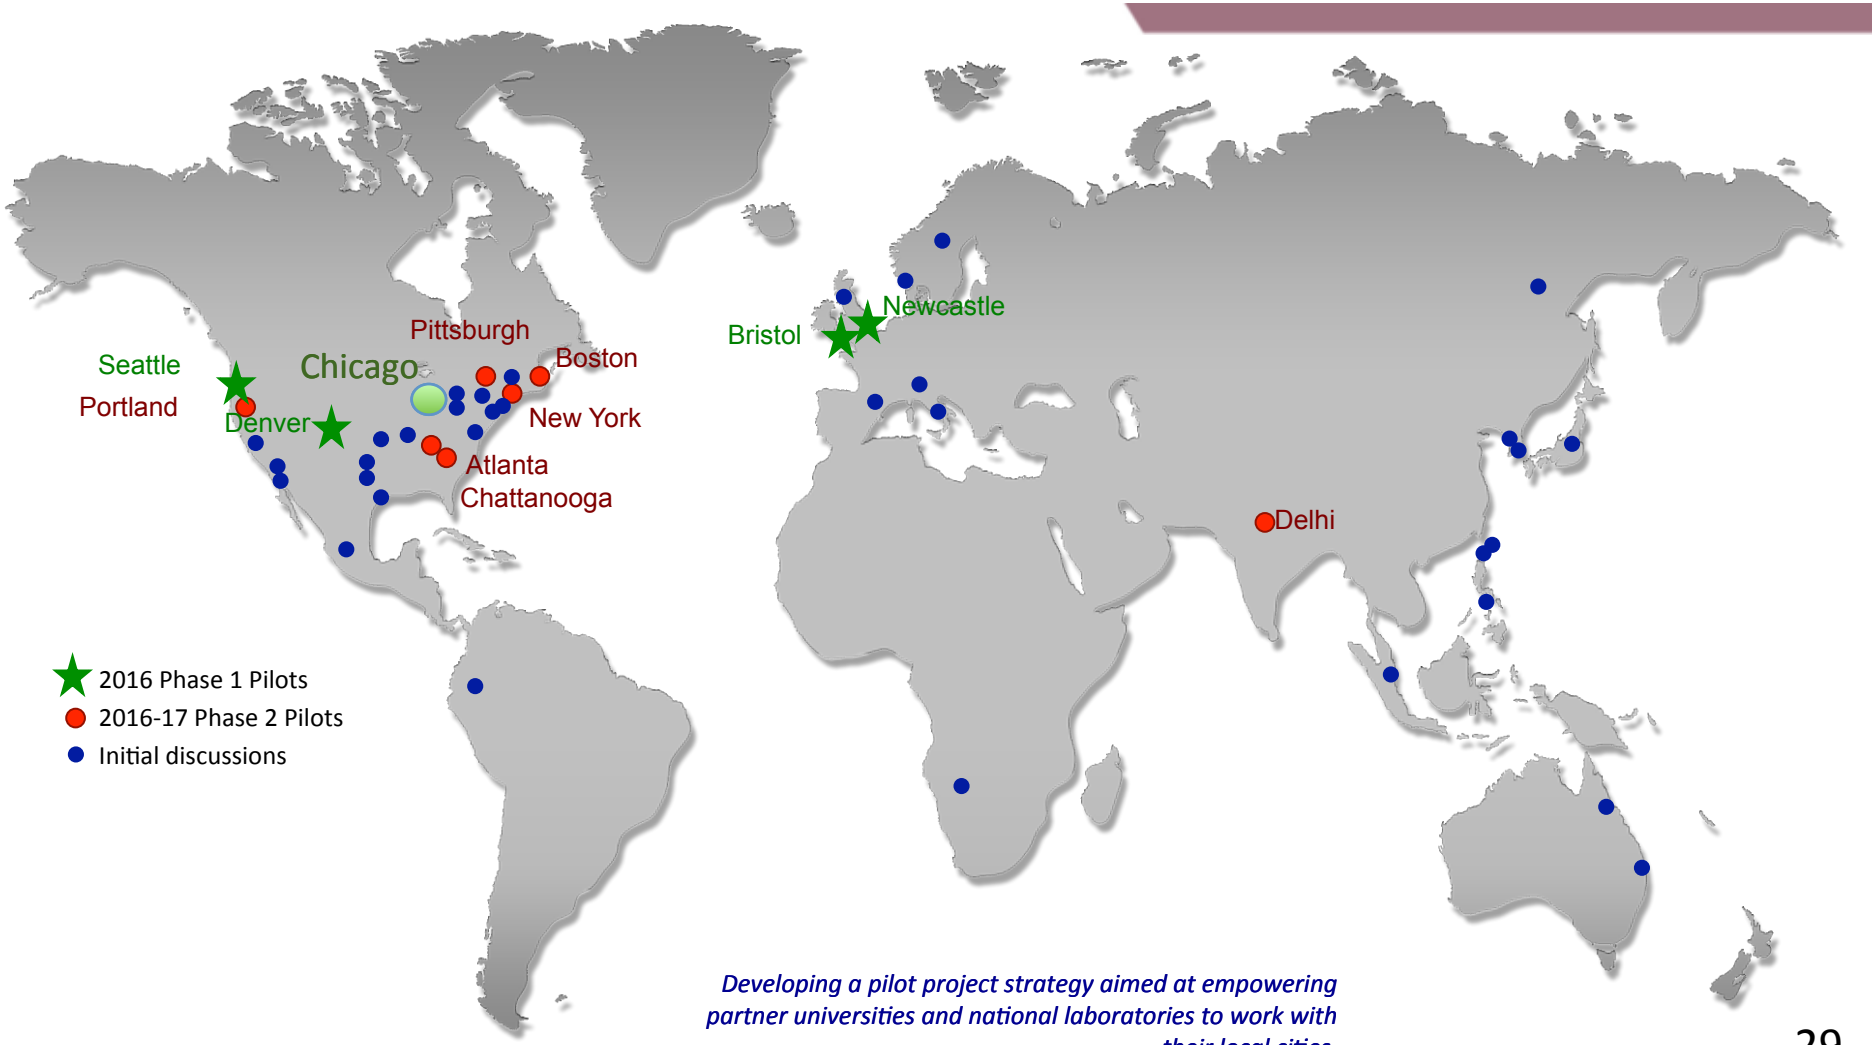

Developing a pilot project strategy aimed at empowering partner universities and national laboratories to work with their local cities.

Why HPC Geeks Should Care

- New sensors are ***programmable* parallel computers**
  - Multicore + GPUs & OpenCL or OpenMP
  - New algorithms for in-situ data analysis, feature detection, compression, deep learning
  - Need new progmod for "stackable" in-situ analysis (for sensors and HPC)
  - Need advanced OS/R resilience, cybersecurity, networking, over-the-air programming
- 1000s of nodes make a **distributed computing "instrument"**
  - New streaming programming model needed
  - New techniques for machine learning for scientific data required
    - Both for within a "node" and collectively across time series
- How will **HPC streaming analytics and simulation** be connected to live data?
  - Can we trigger HPC simulations after first approximations? (weather, energy, transportation)
  - Unstructured database with provenance and metadata for QA/collaboration
- Use novel HPC hardware to solve power issue?
  - Can we use neuromorphic or FPGAs to reduce power for in-situ analysis & compression?
- We are trading precision & cost for greater spatial resolution: What is possible?

Questions?

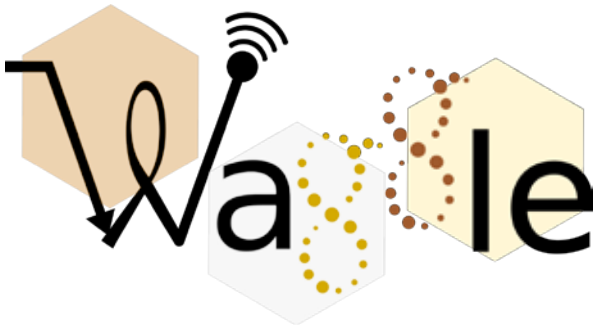

http://www.wa8.gl

http://arrayofthings.github.io