



Waggle

**A platform for distributed smart wireless
sensors and in-situ parallel computation**

Pete Beckman

Senior Scientist, Argonne National Laboratory

Co-Director, Northwestern / Argonne Institute for Science and Engineering (NAISE)

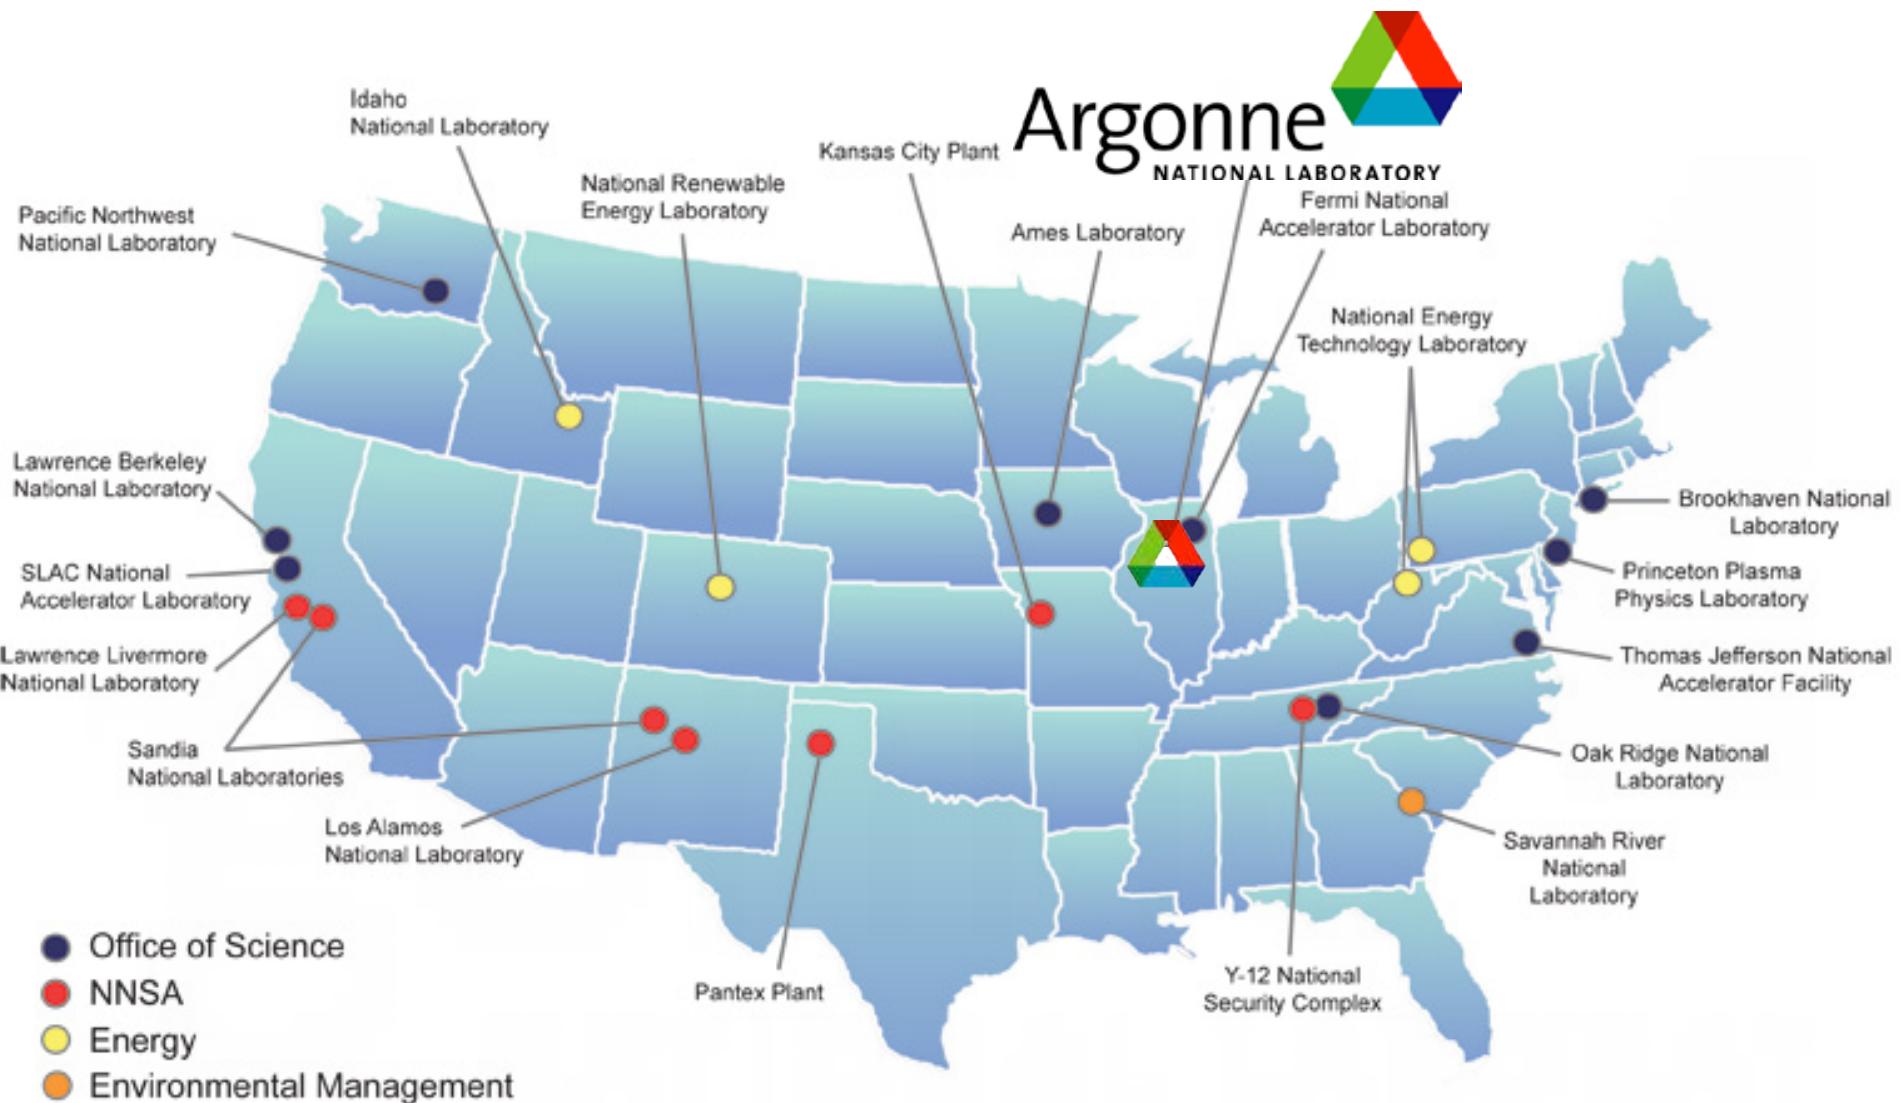
Senior Fellow, University of Chicago Computation Institute

Argonne National Laboratory

- \$675M /yr budget
- 3,200 employees
- 1,450 scientists/eng
- 750 Ph.D.s

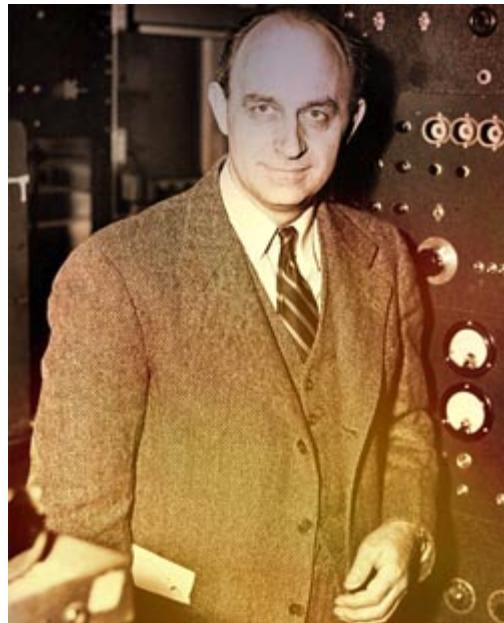


Argonne: Vital part of DOE National Laboratory System



- Office of Science
- NNSA
- Energy
- Environmental Management

Direct descendant of Enrico Fermi's Metallurgical Laboratory



- Opened in Feb 1943 (as new site for Chicago's Metallurgical Laboratory)
- Became Argonne National Laboratory in July 1946 (first national laboratory)



User facilities

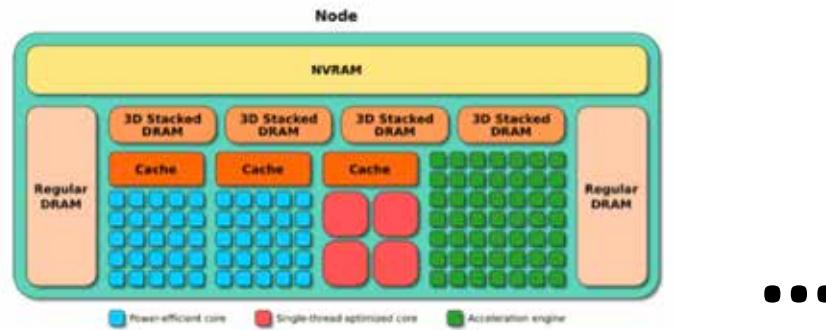
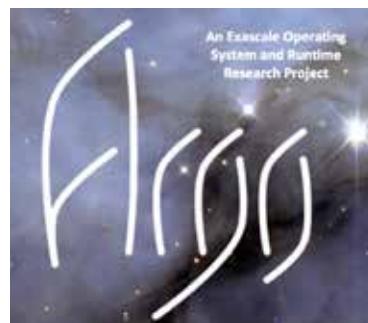


Mira: Argonne's Power-efficient Supercomputer

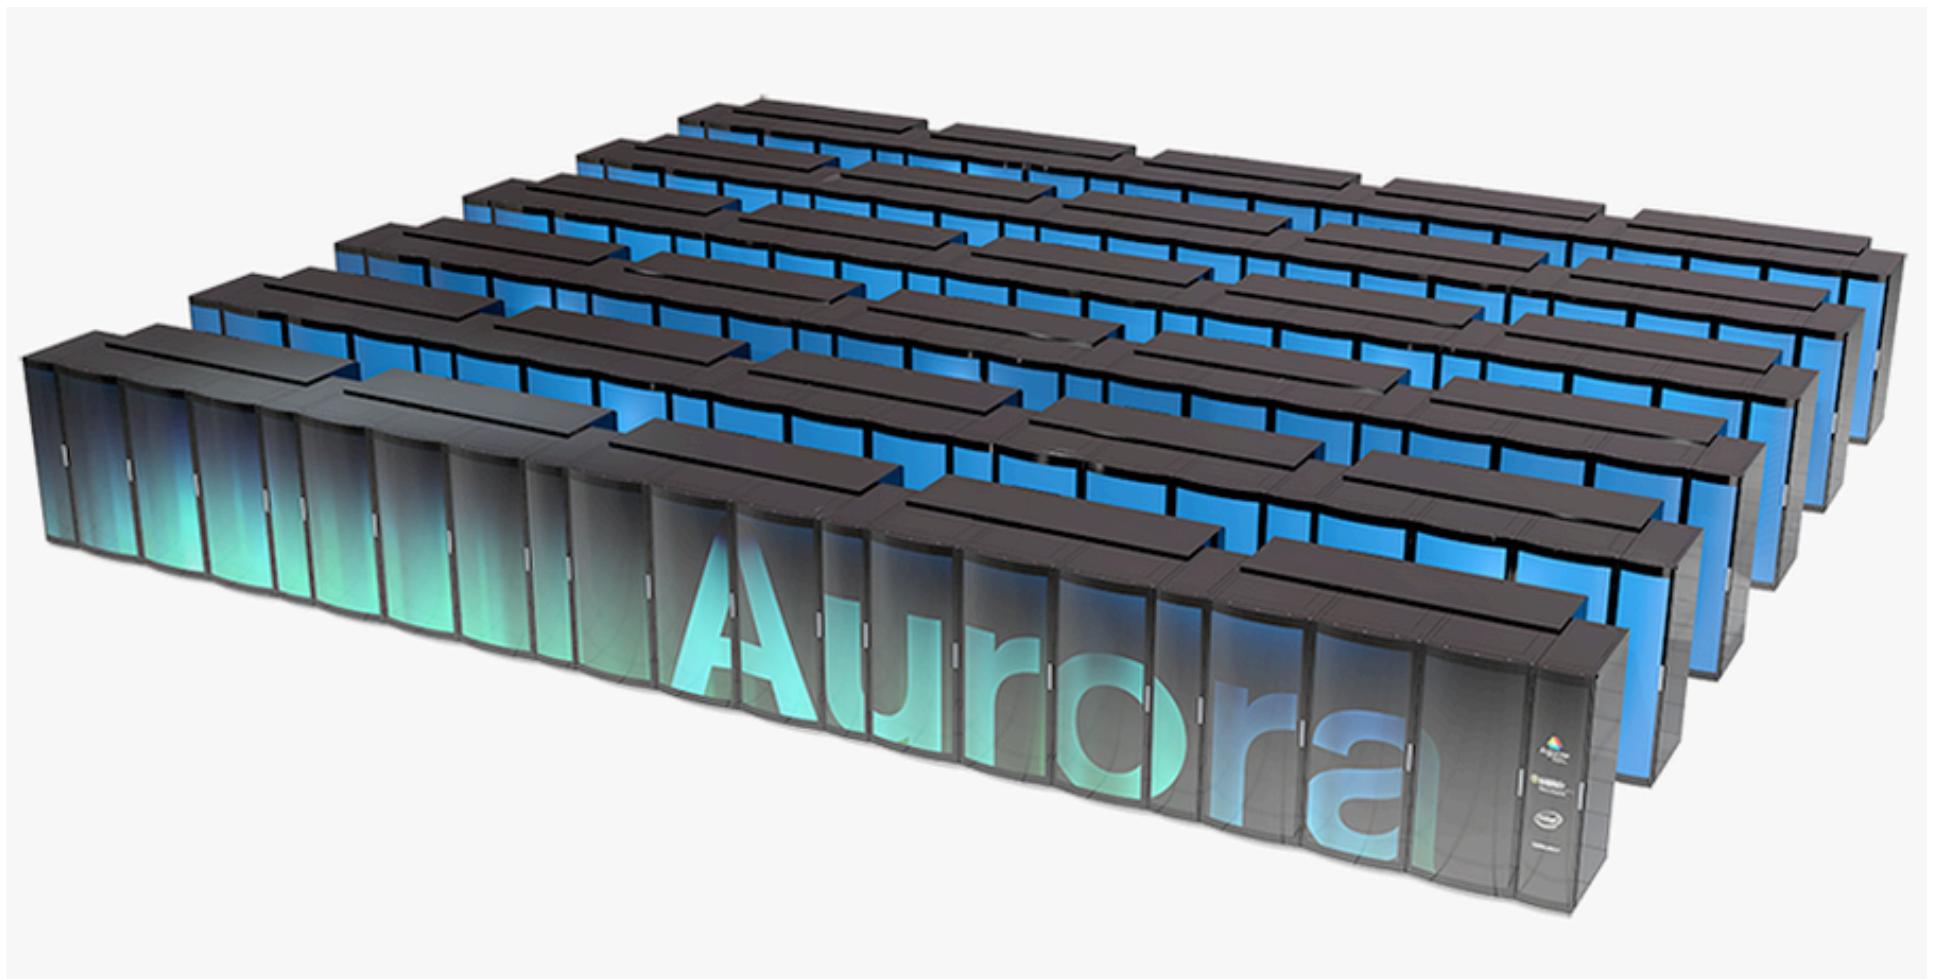
- Blue Gene/Q System
 - 48 racks
 - 786,432 cores
 - 786 TB of memory
 - Peak flop rate: 10 PF
- Storage System
 - ~30 PB capability
 - 240GB/s bandwidth (GPFS)



Argonne: Operating System, File System, Message Layer, etc., System Software Research



Argonne's Next Big Machine: Aurora



Argonne:

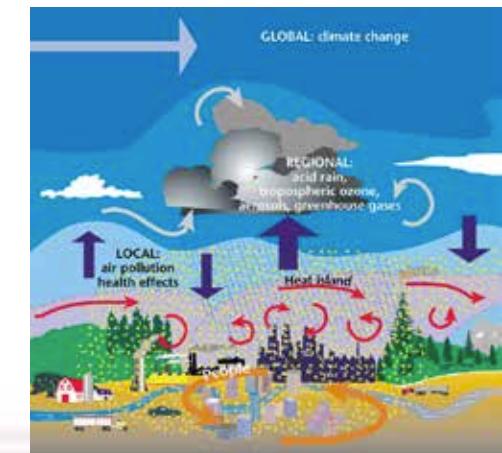
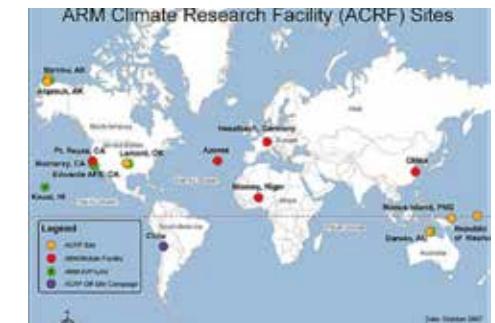
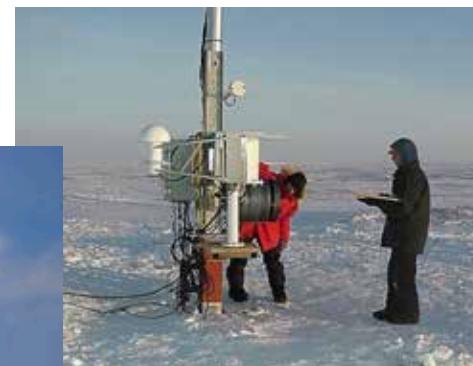
Develops new sensors

Runs large sensor networks

Does climate modeling and simulation

- ... chemical, biological, nuclear and explosive materials

Atmospheric Radiation Measurement Climate Research Facility



Internet of Things... Yawn?



University of Cambridge, 1991



Amazon Dash



IoT: Canonical and GE's FirstBuild Collaborate on Smart Refrigerator

By Canonical on 11 May 2015



"ChillHub is a refrigerator with two USB ports and built-in Wi-Fi connectivity. In addition, ChillHub has an open-source iOS-compatible app [...] Ubuntu is the favored platform for developers of all kinds – particularly those innovating around the Internet of Things."



Example: ESP8266

\$5

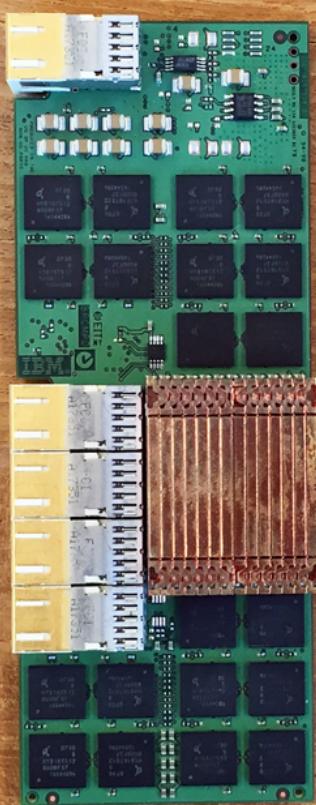


Disruption of Interest: Internet of *SMART* Things

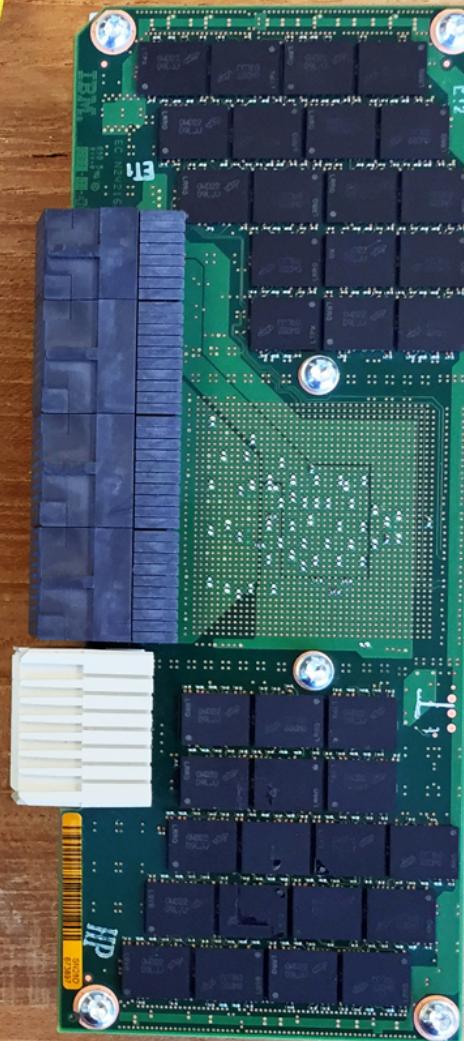
- **Sensors:** Explosion of nano tech (new Moore's-like law)
 - Increasingly smaller, lighter, more accurate, energy efficient
 - Examples: micro GasChrom (NDA), micro air quality sensors, weather, motion, etc.
- **CPUs:** *Capable* low-power CPUs embedded everywhere: cities, people, infrastructure
 - E.g: CPUs we swallow (Freescale K02 prototype), shoes we wear (Nike), to large-scale city water and electrical infrastructure
 - CPUs can be very capable
- **Big Data:** Sensors generate more data than can be stored
 - Sensors+CPUs = new programming model for *in-situ computation*
 - HPC Analysis that can be fused with cloud-based data sets

Opportunity: Move from observing to predicting:
Smart Sensors + Supercomputers = predictions and analysis





IBM BlueGene/P
Supercomputer; 2007
4 compute cores
0.85 GHz, 4 GB RAM
Peak: 13.6 GF/s



IBM BlueGene/Q
Supercomputer; 2011
16 compute cores
1.6 GHz, 16 GB RAM
Peak: 205 GF/s

Introducing Waggle (www.wa8.gl)

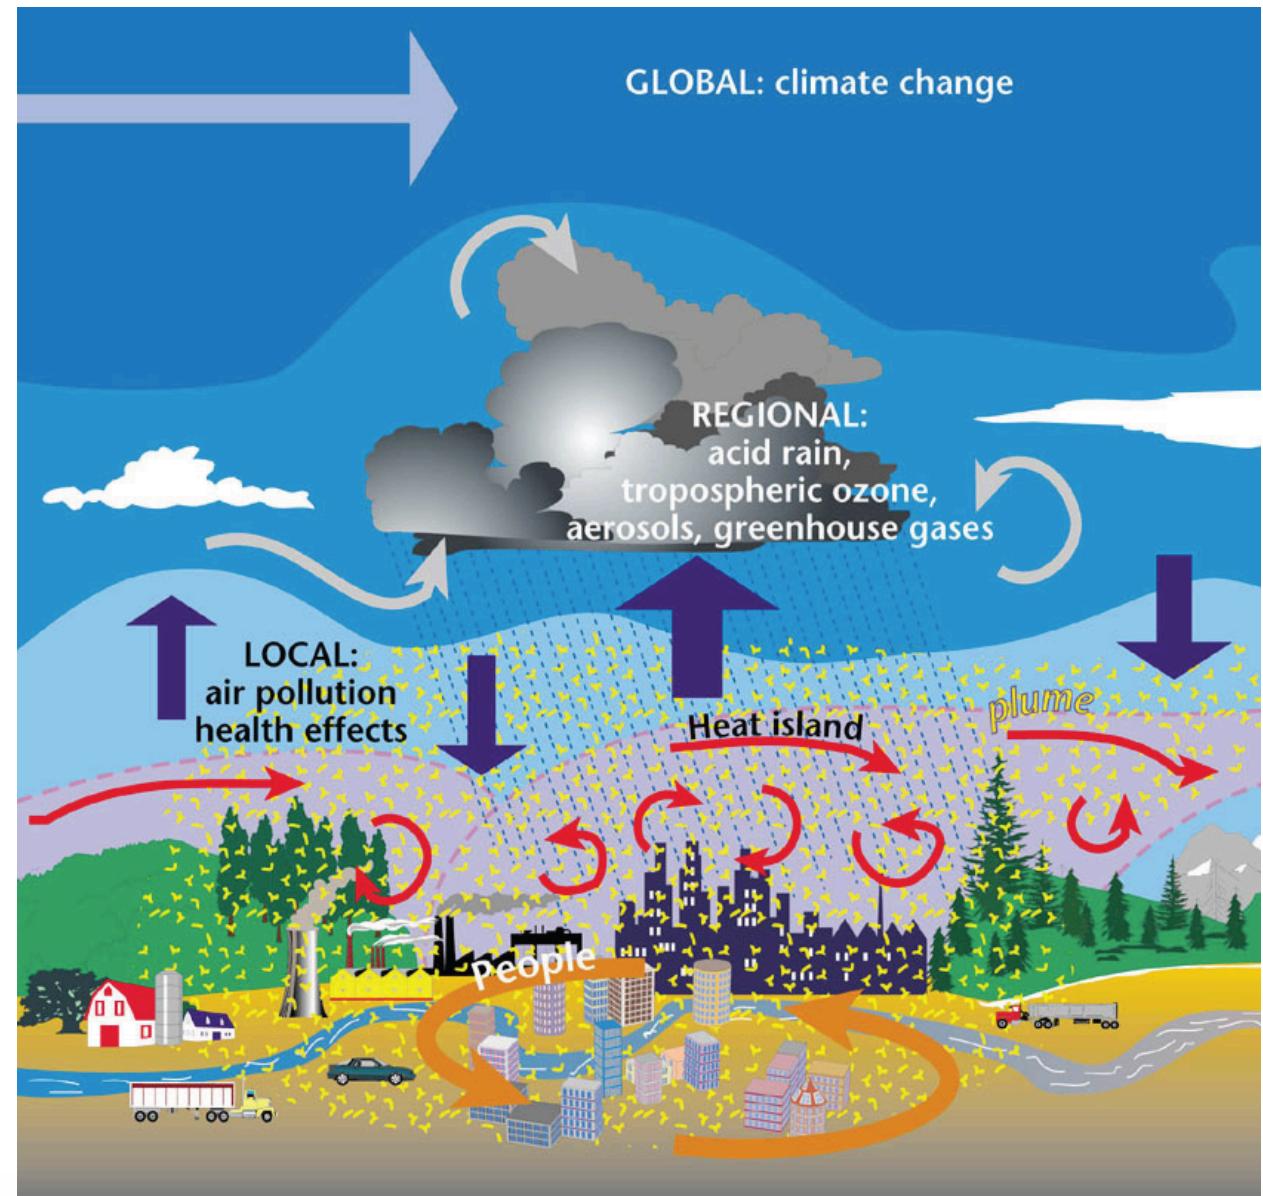


- Open source architecture to leverage disruptive technology
 - A standard building block instead of hack-a-RPi.
- **Powerful CPU**, accurate sensors
- Supports ***In-Situ computation*** for adaptive feature detection, attentive control
- “Deep Space Probe” design for resilience (safe mode, multiple kernels, heartbeats)
- Scalable to 100Ks of nodes; can be linked to supercomputer predictions
- Scalable/hackable design can be adapted for new sensors or control systems, host active education community



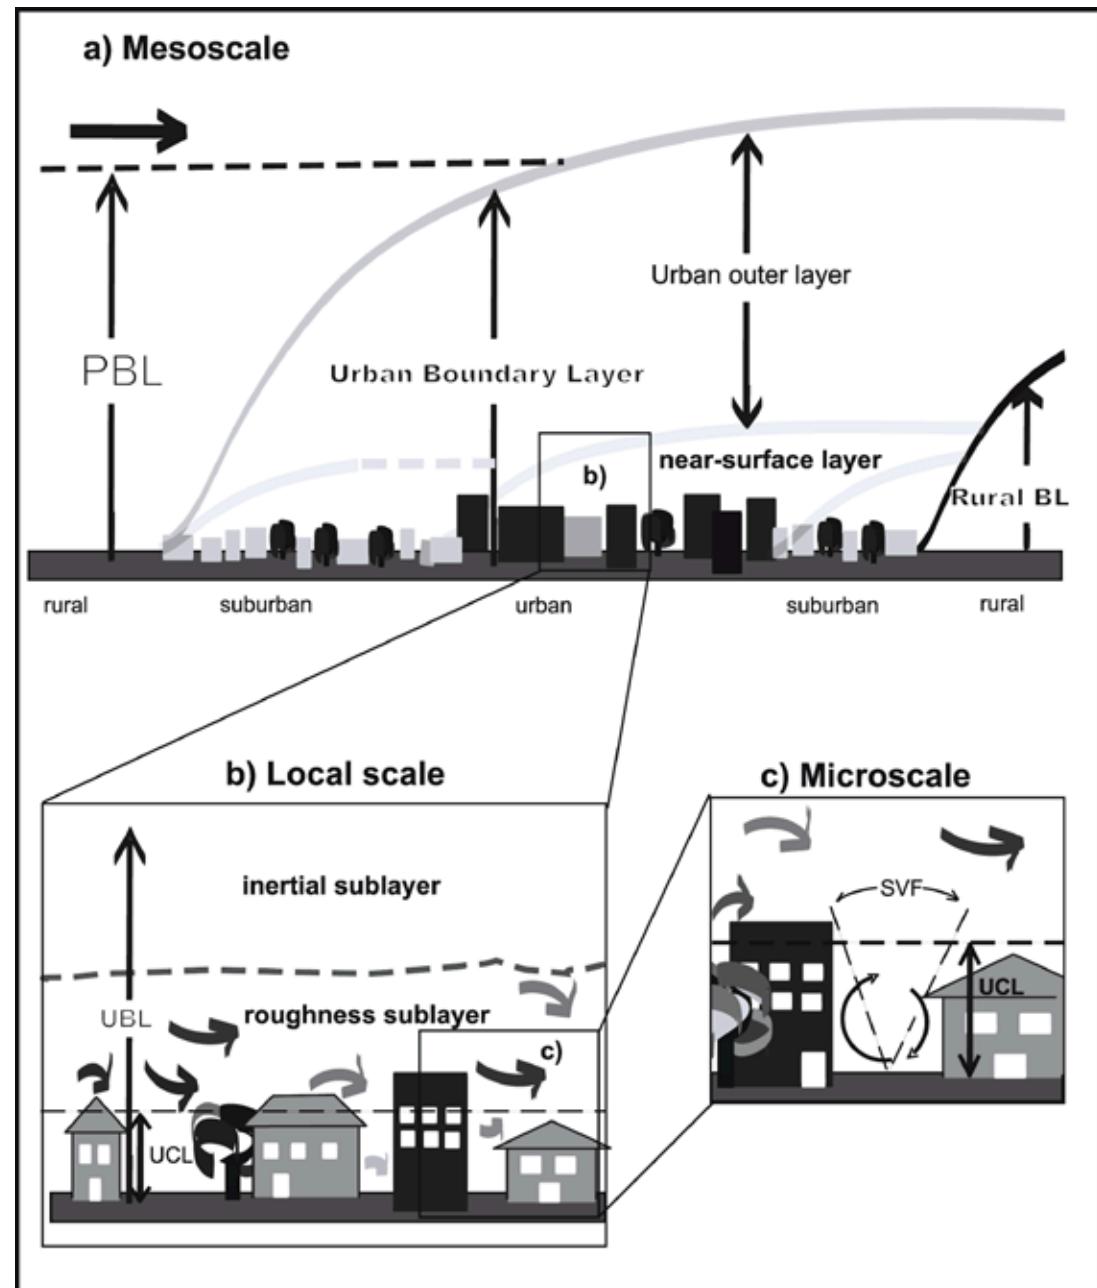
The connection between urban and regional climate

- Cities can alter their local climate through their built environment.
 - Temperature (urban heat island) and precipitation (storm splitting and initiation) are the most widely known examples.
- Cities alter the surrounding regional climate primarily through emissions carried downwind as an “urban plume”.
- Predicting urban climate change requires interactive modeling of regional and urban climate systems.



The Planetary Boundary Layer

- Where we live
- After emissions, primary control on pollution levels is interaction between PBL and free atmosphere.
- In Urban areas
 - Urban Boundary Layer
 - Urban Canopy Layer

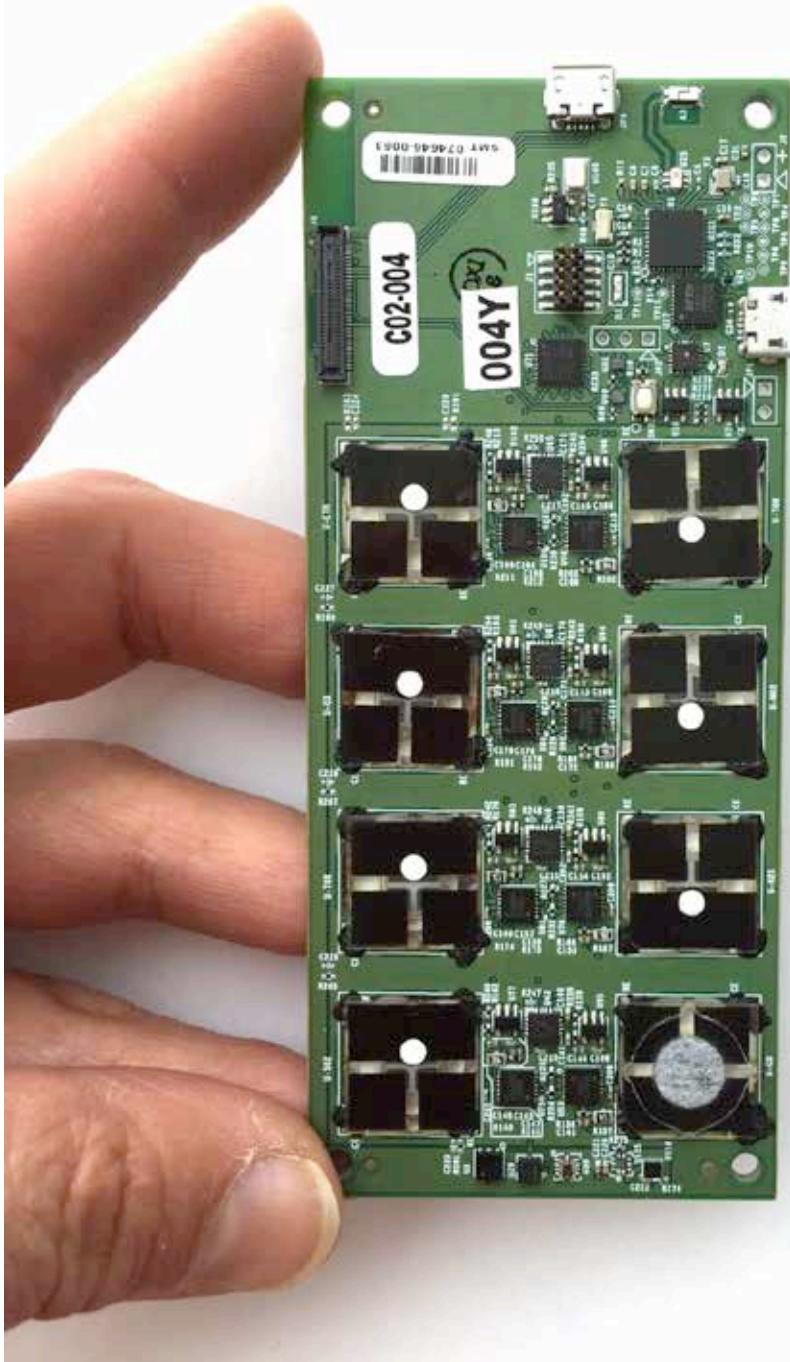


Source: NRC "Urban Meteorology: Forecasting, Monitoring, and Meeting Users' Needs"

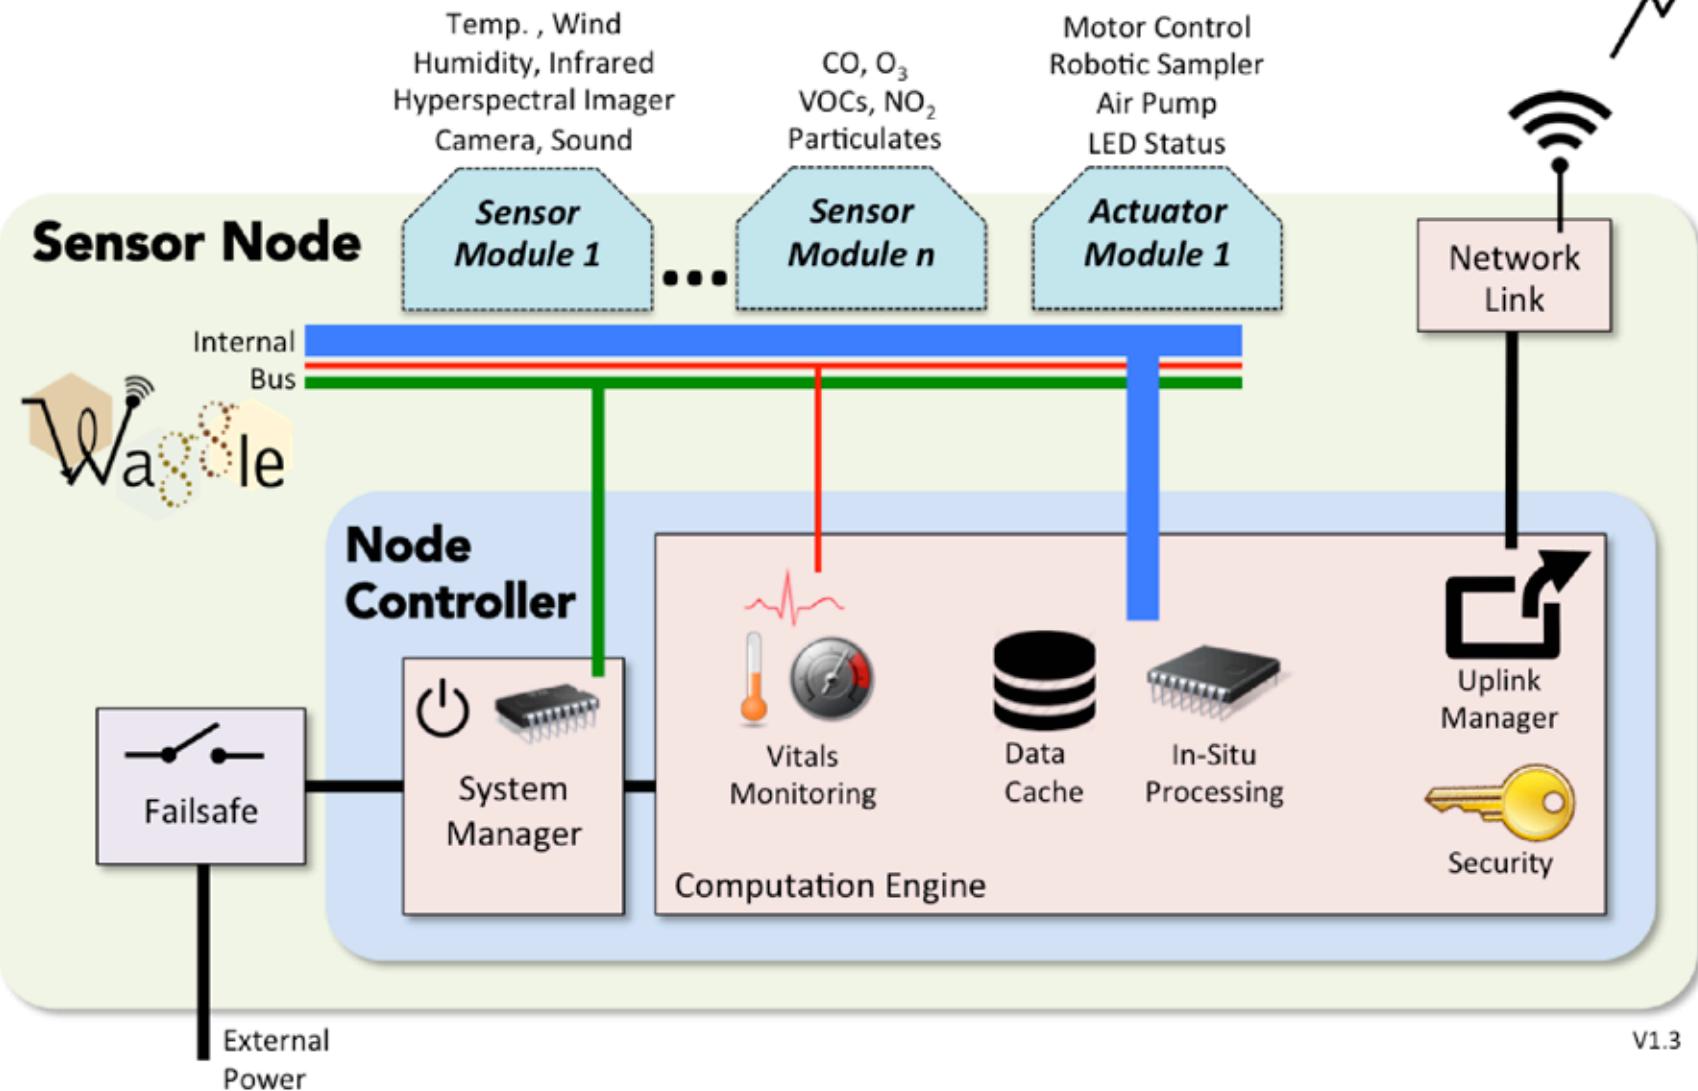
New Advanced Sensors

- NO₂ (Nitrogen Dioxide): <2 ppb
- O₃ (Ozone) < 5 ppb
- CO (Carbon Monoxide) < 1 ppm
- SO₂ (Sulfer Dioxide) < 15 ppb
- H₂S (Hydrogen Sulfide) < 2 ppb
- Soon:
 - HCHO (Formaldehyde)
 - VOC (Volatile Organic Compound)
 - CH₄ (Methane)

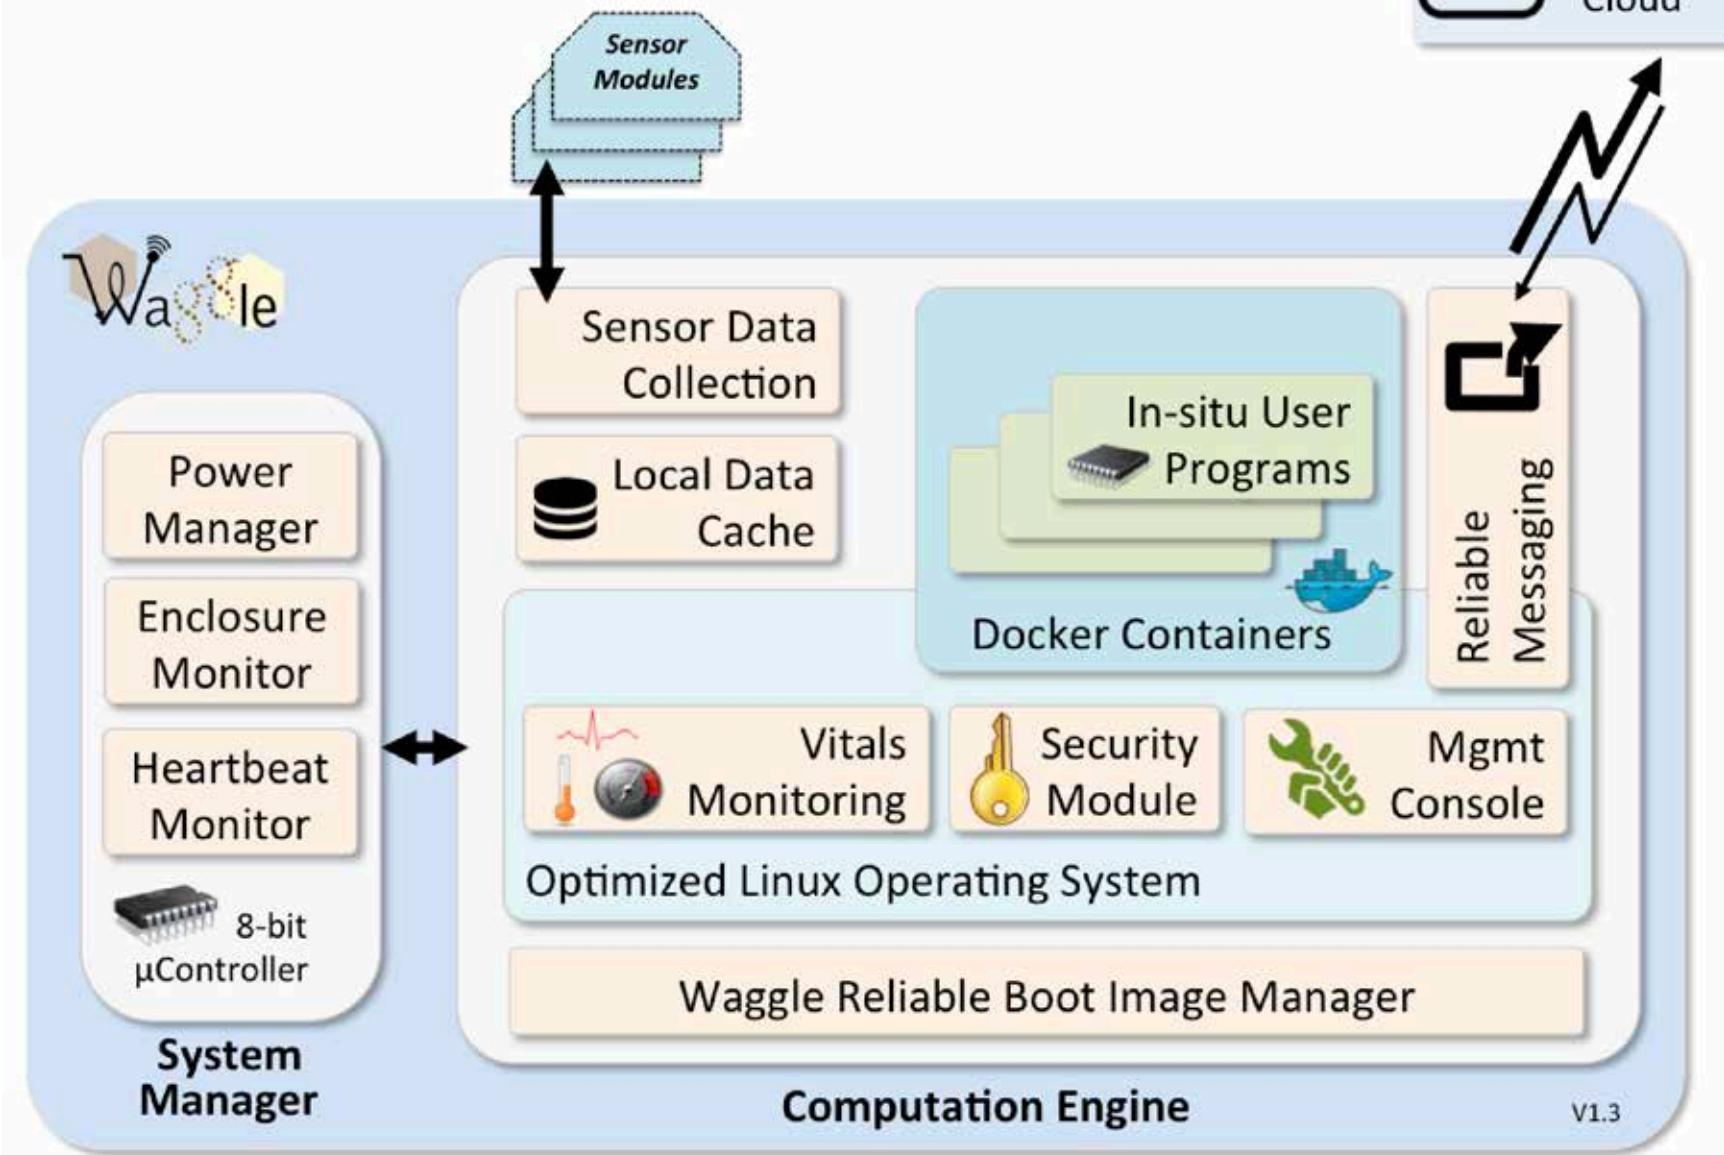
This prototype board would likely sell for < \$75 in production



Waggle Modular Sensor Node Architecture



Waggle Node Controller Software Architecture



V1.3



Many, Many, Research Questions: Are compute-capable sensors interesting?

- # of devices
 - Parallel in-situ processing? Algorithms? Collectives?
- Ultra-resilient, periodically connected operation
- Control systems for actuators (motors, etc.)
- Embedded secret required, and hard to manage (pairing, Apple HomeKit)
- Assume physical access by mischievous hands
 - Cheap Trojan counterfeits, Report false data, Mod actuators
- How does massive, lower-quality data compare to fewer, high quality data?

Sensors

Waggle EnvSense V1.0 Nov2014 board

Waggle Sensor Array Interface Control Document Version 0

1 Overall Packet Structure

The structure of the packet (and the data sub-packet) relies on byte positions and known values, rather than delimiters.

| Field | Value | Byte Position | Length |
|-----------------------------------|-------|---------------|--------|
| Start | 0xAA | 0 | 1 |
| Protocol version | 0x00 | 1 | 1 |
| Length of data (not whole packet) | | 2 | 1 |
| Data | | 3 | 194 |
| CRC of data (not whole packet) | | 197 | 1 |
| End | 0x55 | 198 | 1 |

Table 1: Overall packet structure

Important: If the length of data is 0, the packet immediately ends, meaning the data, CRC, and end fields do not exist.

2 Data Sub-Packet

The data sub-packet consists of 32 "chunks" (30 sensors and 2 MAC addresses). Each "chunk" follows one of seven formats.

2.1 Data Formats

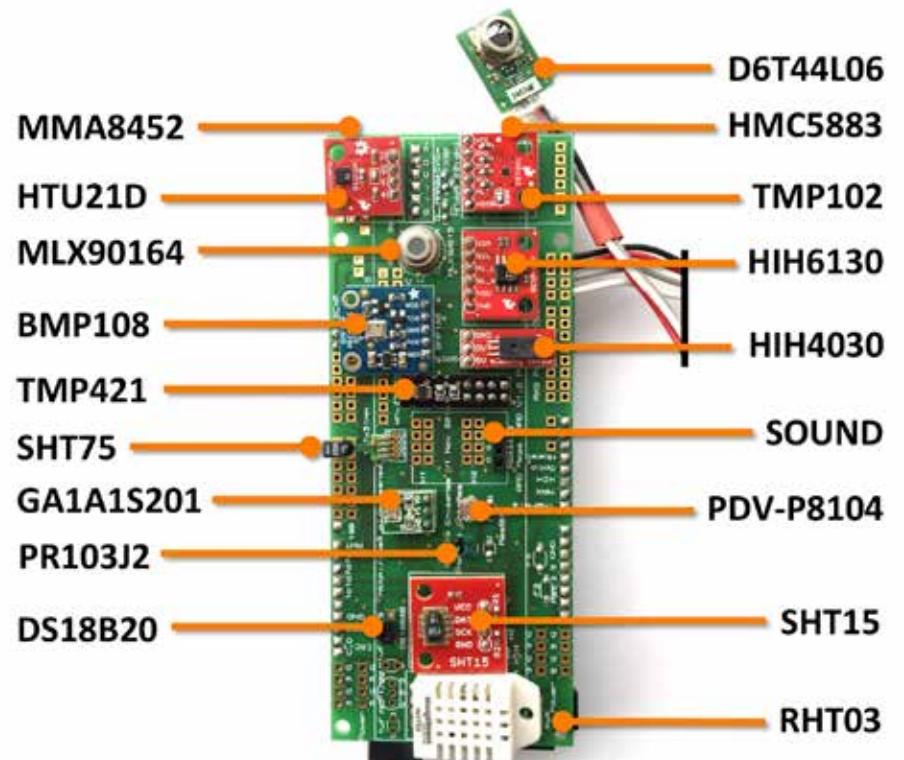
| Format | Byte 0 | Byte 1 | Byte 2 | Byte 3 | Byte 4 | Byte 5 |
|--------|--|------------------------|------------|-----------------|--------|--------|
| 1 | (1 < 7) integer | (neg < 7) fractional | - | - | - | - |
| 2 | 7 MSb | LSB | - | - | - | - |
| 3 | Addr5 | Addr4 | Addr3 | Addr2 | Addr1 | Addr0 |
| 4 | (1 < 7) (neg < 6) (4-bit int < 2) 2 MSb of frac. | 8 LSb of frac. | - | - | - | - |
| 5 | (neg < 6) 6 MSb | 8 LSb | - | - | - | - |
| 6 | (1 < 7) (neg < 6) 6 MSb | Middle 8 bits | 8 LSb | - | - | - |
| 7 | First 8 "chunks" | 8 "chunks" | 8 "chunks" | Last 8 "chunks" | - | - |

Table 2: Data formats

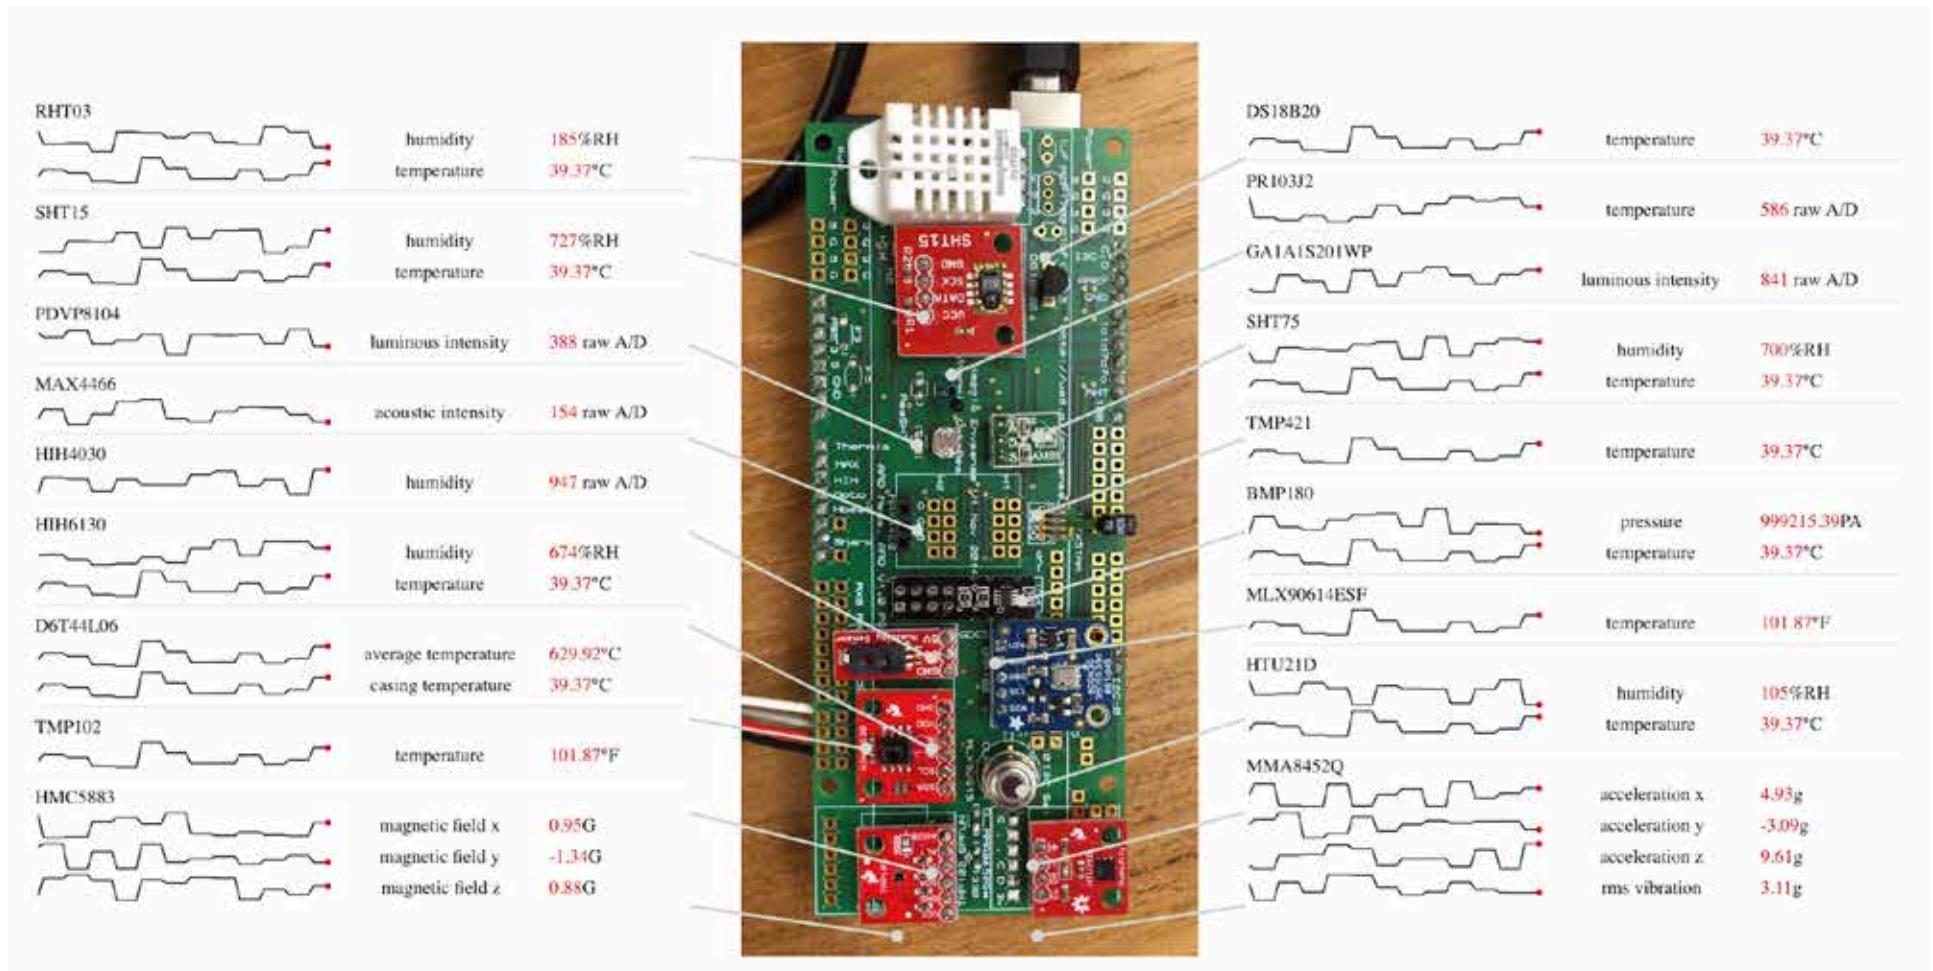
This version of the Waggle protocol does not use standard representations for floating point numbers. Instead, the location of the decimal point is pre-determined (between the integer and fractional components, if applicable).

The most significant bit in byte 0 of formats 1, 4, and 6 means the data is already converted. Formats 2 and 5 contain raw data. Formats 1, 4, 5, and 6 contain a "negative" bit. If this bit is 1, the value is negative.

Air Temperature, Humidity, Infrared
Temperature, Barometric Pressure,
Ambient Light, Sound Levels,
Vibration, Magnetometer, etc.



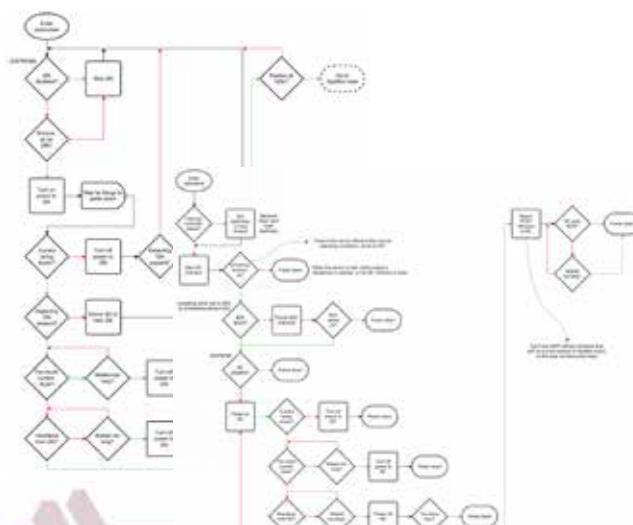
Programmer Interface for Watching



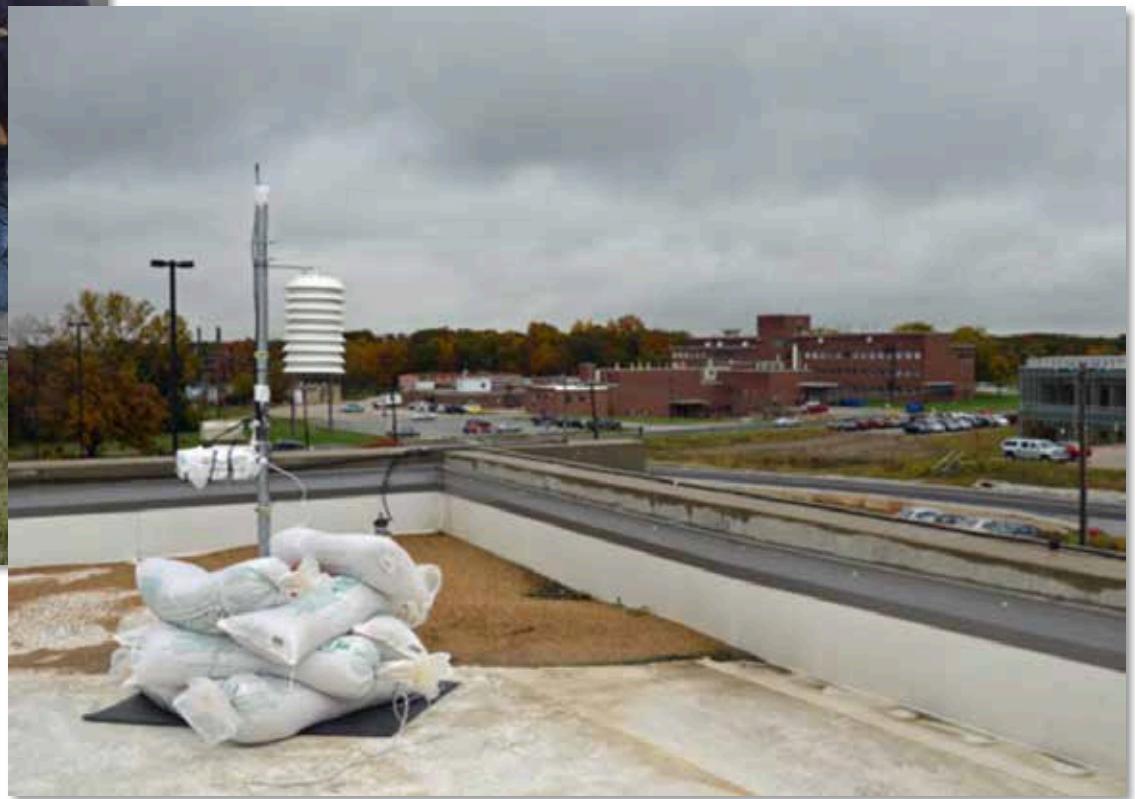
Node Management

Waggle SysMon V1.0 Nov2014 board

- Internal Sensors (intrusion, temp, power (current), etc.)
- Provides “Deep Space” operation with heartbeats, safe mode, hard power control, backup kernels, etc.



(Chicago Winter Survival Test)...



(old) Version 1 Enclosure



Instrumenting Cities

(Preview: Charlie Catlett)

- As cities grow they face increased operational costs and challenges across every area of policy and science.
- There exists no open source, low-cost, general purpose platform to *measure* cities at the spatial and temporal resolution necessary for science and data-driven policy.



Array of Things: Planned Deployment

- Phase 1 (est. 4q2015): 50 Gen3 nodes at CDOT green space “placemaking” project sites.
- Phase 2 (est. 2q2016): 100 Gen3 nodes, Chicago Loop
- Phase 3 (est. 4q2016): 50 Gen4 nodes, Chicago Loop
- Phase 4 (est. 3q2017): 300 Gen4 nodes, mixed density deployment throughout city.

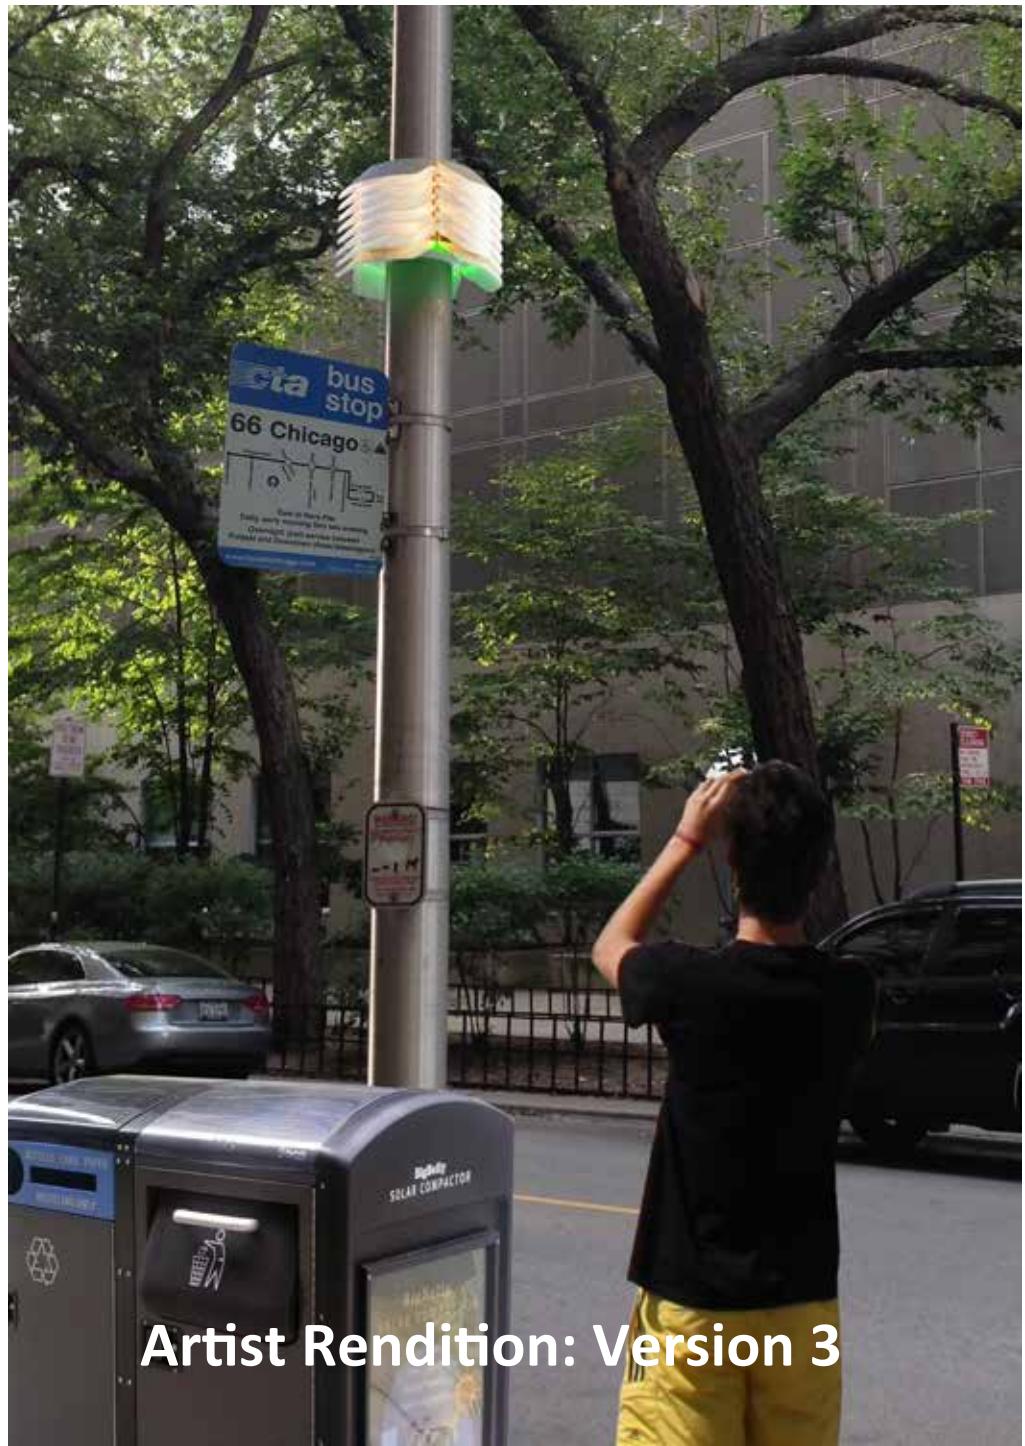


(current) Version 2

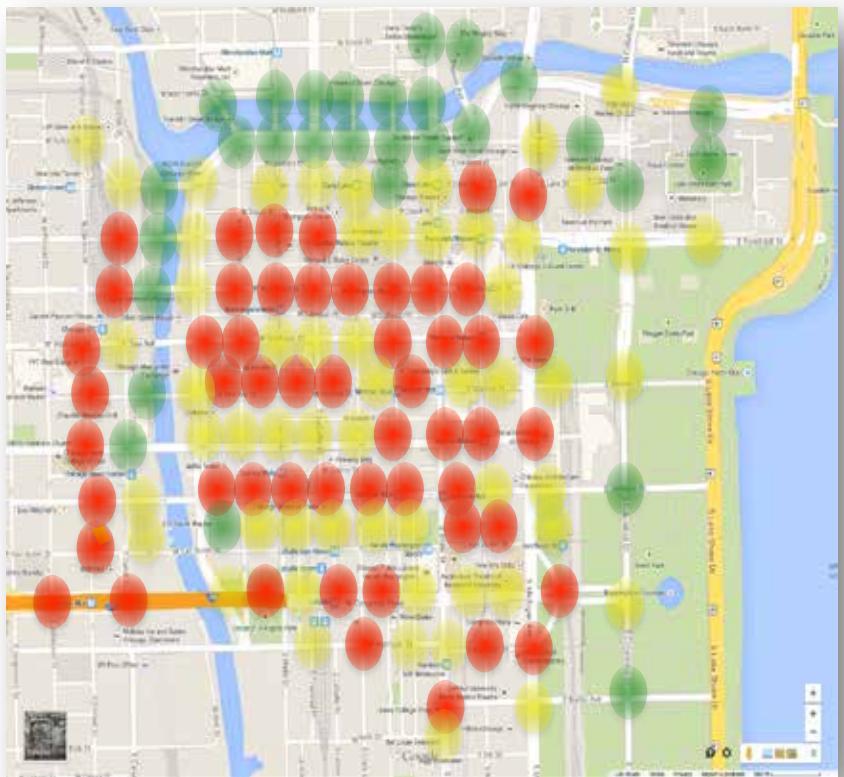
Discussions also under way with Seattle, San Francisco, Los Angeles, St. Louis, Boston, Barcelona, London, Bristol, Glasgow, Umeå, Bolzano.



(current) Version 2



Artist Rendition: Version 3



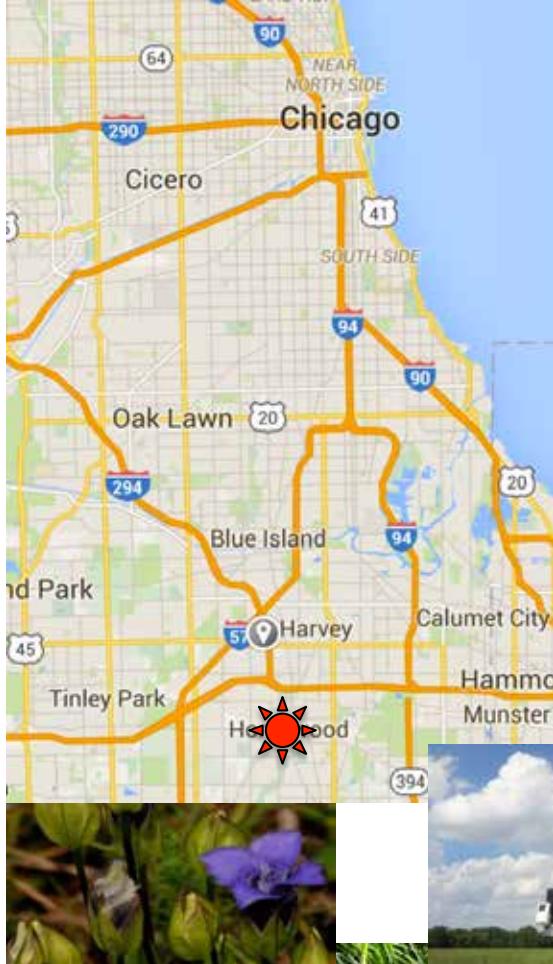
Deployment Map



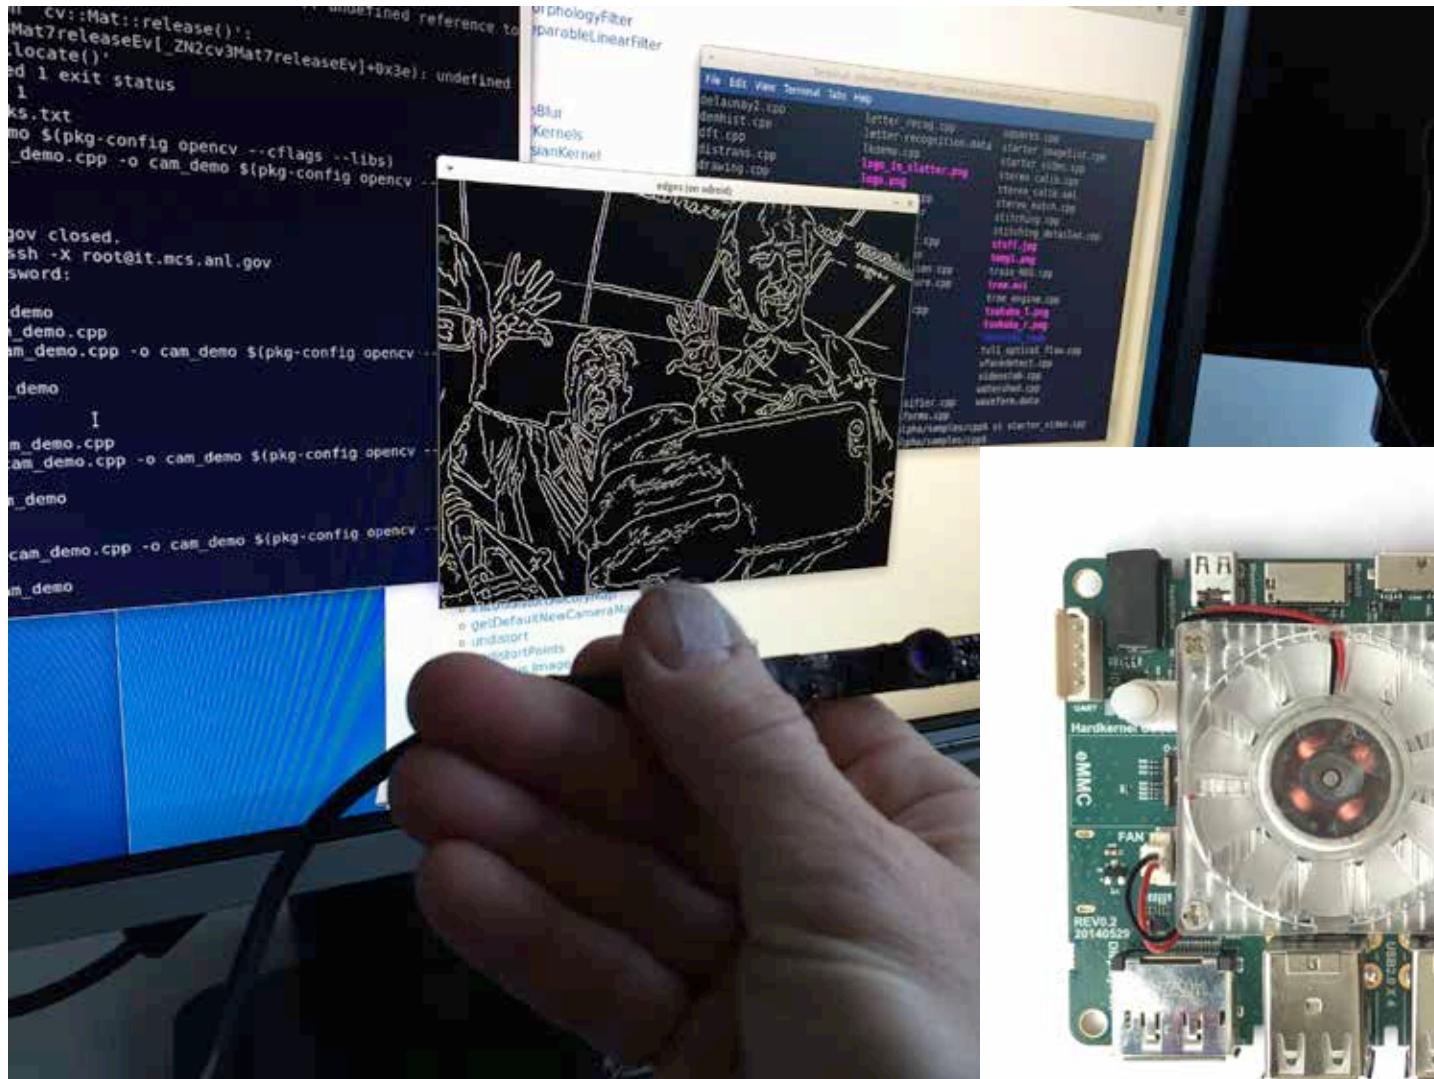
Gensburg-Markham Prairie

370 acres, owned/managed by Nature Conservancy and Northeastern IL Univ.

Registered as National Natural Landmark



In-Situ Analysis and Feature Recognition





Machine Learning: Cat or Dog?









Privacy...

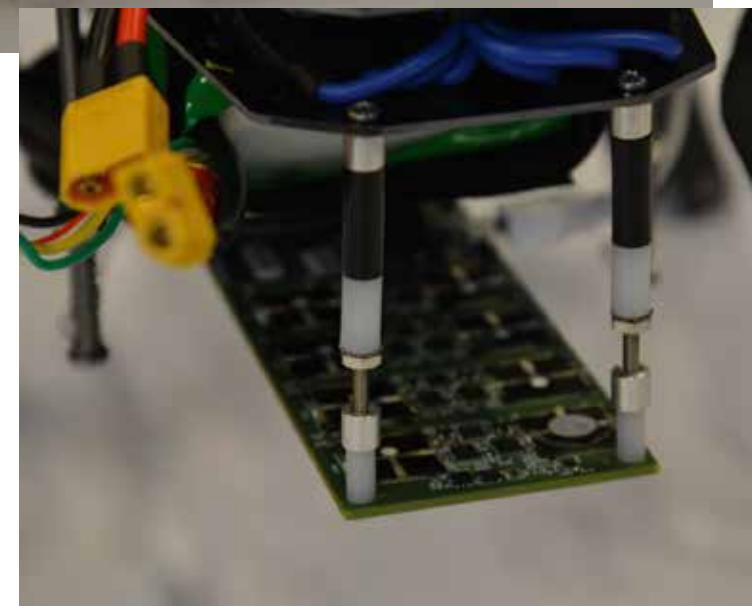
The image displays two screenshots from a mobile application interface. The top screenshot shows the 'BLE Scan Results' screen with a title bar indicating AT&T signal strength and the time 9:09 AM. It lists various Bluetooth devices with their names and signal strengths:

- <no name> Signal strength: -88 dB
- <no name> Signal strength: -96 dB
- Flex Signal strength: -64 dB
- Basis Peak Signal strength: -88 dB
- Charge HR Signal strength: -66 dB
- One Signal strength: -92 dB
- Charge HR Signal strength: -87 dB
- Flex Signal strength: -83 dB
- <no name> Signal strength: -88 dB
- UP24 Signal strength: -85 dB

The bottom screenshot shows the 'Device Data Services' screen with a title bar indicating AT&T signal strength and the time 2:07 PM. It lists device-specific data strings:

- Firmware Revision String Value: 0x302e342e 3200 Ascii: 0.4.2
- Hardware Revision String Value: 0x303800 Ascii: 0B
- Manufacturer Name String Value: 0x41756775 73742048 6f6d652c 20496e63 2e00 Ascii: August Home, Inc.
- BD4AC610-0B45-11E3-8FFD-0800200C9A6 6
- BD4AC611-0B45-11E3-8FFD-... Value: 0xcb10104b be4ee2ed 1fb772a8 30551360 0200 Ascii: EK%Náír'OU'
- BD4AC612-0B45-11E3-8FFD-... Value: 0x411d0098 ddc614e8 1206c06c 55bcc3a 0000 Ascii: AYÆèAIUÉ:
- BD4AC613-0B45-11E3-8FFD-... Value: 0x330f5922 bd2c5e63 2ac27256 50f82d2a 0f01 Ascii: 3Y%"^c"ArVPo-





Looking Forward

- Does community need open platform?
 - Featherweight, middle-weight, CPU-heavy archetypes
- There are good, meaty CS problems in this space...
 - Autonomous decisions, collective behavior, billion-way control, de-centralized mgmt, trusted unprotected hw...
 - Programmable, HPC-style env. with in-situ computation
 - Machine learning, classification, processing audio, etc.
 - Neuromorphic computing
- Selling nodes on SparkFun?
- Privacy? In-situ computation? Air Gap?



*An open platform for
intelligent attentive sensors*

<http://www.wa8.gl>

Home

About

Team

Design ›

News (Blog)

Code & Pubs

Data

Science ›



A customizable, powerful, hackable sensor network platform for science





*An open platform for
intelligent attentive sensors*

| Name | Group | Location | Coordinates | Status | Current Data |
|-----------------|---------|--------------------|-----------------------|-----------|---------------------------|
| UCAOT01 | AoT | Regenstein Library | 41.791848, -87.599566 | Active | View Data |
| UCAOT02 | AoT | Logan Center | 41.785573, -87.604143 | Active | View Data |
| UCAOT03 | AoT | Ratner Center | 41.793373, -87.602294 | Active | View Data |
| UCAOT04 | AoT | Smart Museum | 41.793685, -87.600131 | Active | View Data |
| UCAOT05 | AoT | Mandel Hall | 41.790758, -87.598514 | In-active | View Data |
| UCAOT06 | AoT | Oriental Institute | 41.789009, -87.597298 | Active | View Data |
| UCAOT07 | AoT | Levi Hall | 41.789274, -87.600936 | Active | View Data |
| WXSENETDEPAUL01 | WxSeNet | McGowan Roof | 41.923693, -87.656892 | Active | View Data |
| WXSENETANL01 | WxSeNet | TCS 240 Roof | 41.717602, -87.982415 | In-active | View Data |
| WXSENETANL02 | WxSeNet | Area 460 | 41.701558, -87.994766 | In-active | View Data |

DePaul University | McGowan Roof | Data Archive

Lat: 41.923693 Lon: -87.656892

Installed January 20th, 2015



Jul 30, 2015 3:59:14 pm

RHT03

Temperature NA °C

Humidity NA %RH

HTU21D

Temperature NA °C

Humidity NA %RH

HIH6130

Temperature 31.38 °C

Humidity 31.5 %RH

MLX90014ESF

Temperature 69.49 °F

MMA6452Q

Acceleration x NA g

Acceleration y NA g

Acceleration z NA g

SHT15

Temperature 31.43 °C

Humidity 28.57 %RH

DST44L06

Avg. Temperature NA °C

Casing Temperature NA °C

HIH4030

Humidity 352 raw A/D

TMP102

Temperature NA °C

PR103J2

Temperature 920 raw A/D

HMC5883

Magnetic Field x NA µT

Magnetic Field y NA µT

Magnetic Field z NA µT

SHT75

Temperature 32.2 °C

Humidity 27.38 %RH

BMP180

Temperature 31.89 °C

Pressure 99383 Pa

DS18B20

Temperature 31.62 °C

TMP421

Temperature 31.37 °C

GA1A1S201WP

Luminous Intensity 710 raw A/D

PDV98104

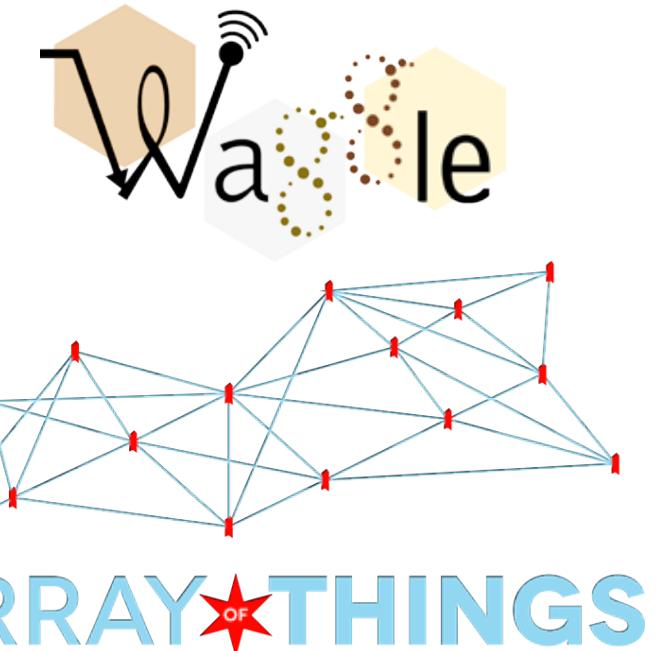
Luminous Intensity 1015 raw A/D

MAX4460

Acoustic Intensity 212 raw A/D

More Information

- Waggle
 - <http://wa8.gl/>
 - Contact: Pete Beckman beckman@anl.gov
- Array of Things
 - <http://arrayofthings.github.io>
 - Contact: Charlie Catlett c@anl.gov



Funding: Argonne's Laboratory Directed Research and Development program

