Data Analytics and Distributed Systems Technology for Autonomous Energy Grids

David E. Culler

University of California, Berkeley

Autonomous Energy Grids workshop

Sept 13, 2017

BECI

BERKELEY ENERGY & CLIMATE INSTITUTE

LBNL & UC BERKELEY

Goal of this talk …

- Encourage you to question assumptions of technological advance …
  - Google, Facebook, Netflix, social nets, ... Are not the only drivers
- The energy challenge can drive "Big Data" network systems technology too
  - Cloud, Premises
  - Data Management, Security, … Trust
  - Heterogeneity
  - Solution synthesis
  - Algorithms
- "Deploy complex distributed software systems that outlive any of the vendors …"
- The Real World is Really messy ...
  - "Autonomous" mean "Not Manual" on the messy parts too

Inter-Infrastructure Network Systems

Sustainable Water

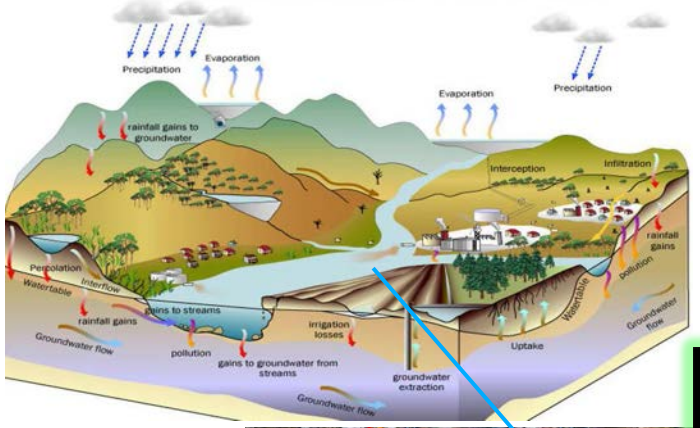

Sustainable Electricity

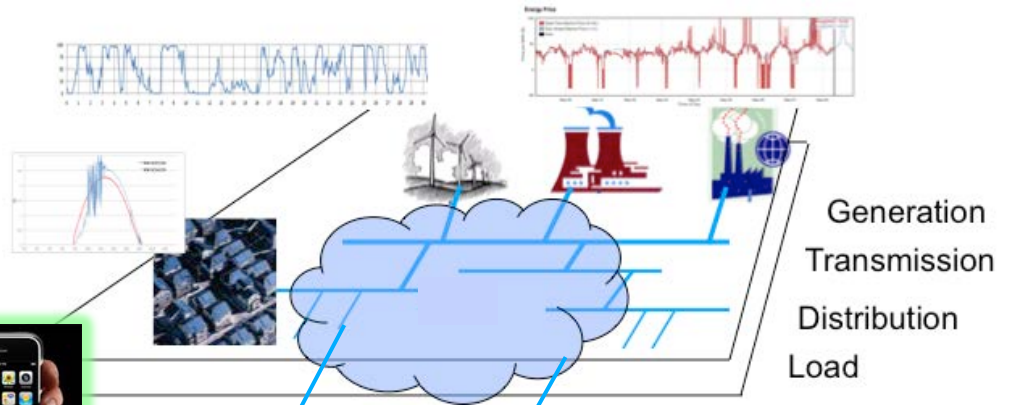

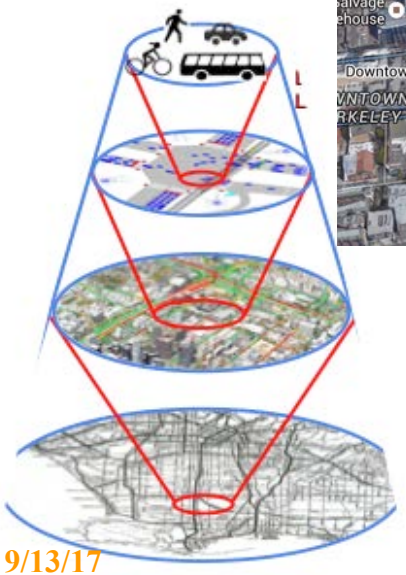

Sustainable Transportation

Sustainable Buildings

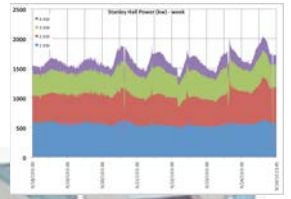

CA2050: GHG 90% below 1990

The short answer: **Yes, we can**

- We can achieve 80% cuts in emissions and still meet our energy needs.
- We can get ~60% of the cuts with technology we largely know about.
  - We basically know how to
  - A lot of this technology is i
    - Deployment will depend on
    - Note: We excluded extreme
- We can get the rest of the this will require new techn development.

CHAIR'S LECTURE: CALIFORNIA ENERGY FUTURES STUDY RESULTS

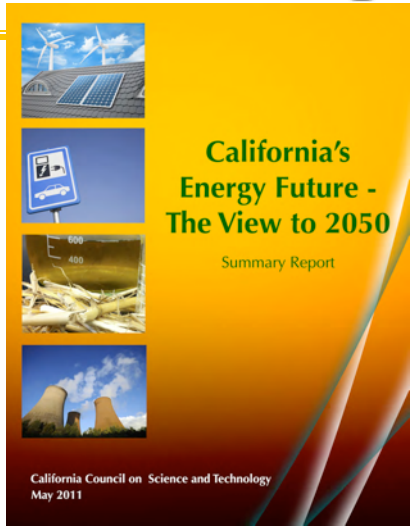

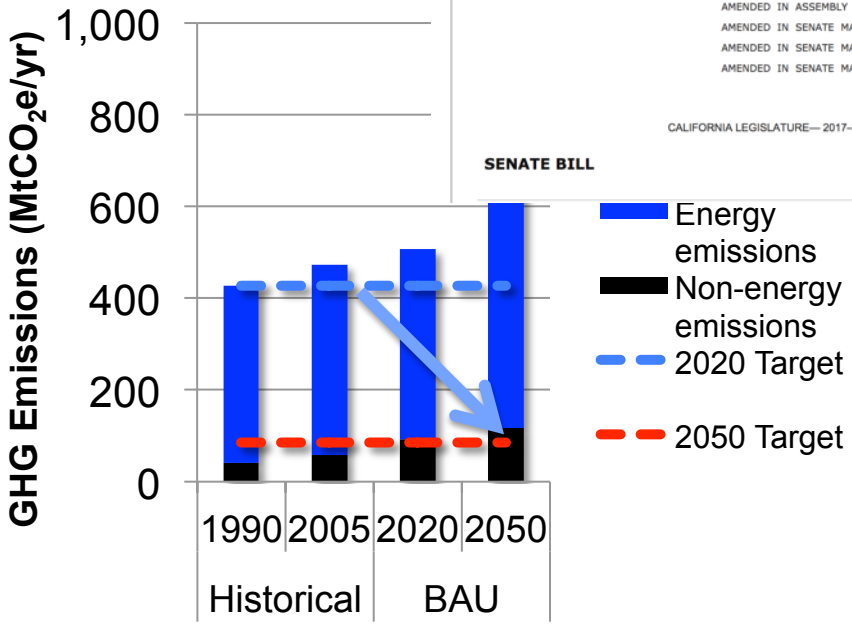

...chnology limitations will ...o exceed the target:

- ...ufficient technology for load balancing without emissions
  - This is an especially big deal if we don't have baseload power
- We don't have enough technology choices "in the pipeline" for de-carbonizing fuel.
  - Need advanced biofuels, but it likely won't be enough
  - CCS may play a larger role in fuels than in electricity

The Problem: Supply-Demand Match

Baseline + Dispatchable Tiers

Oblivious Loads

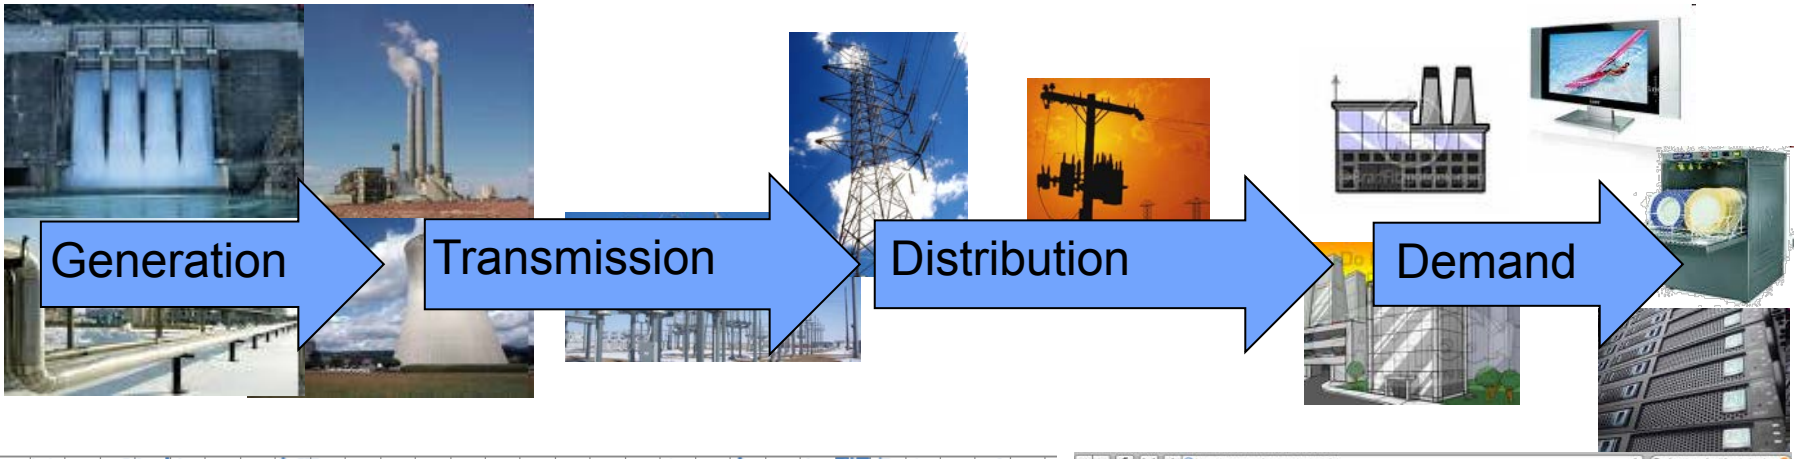

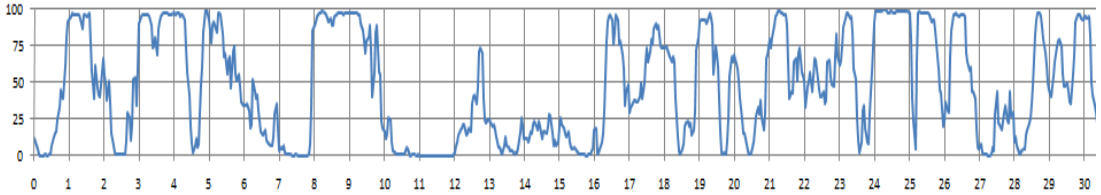

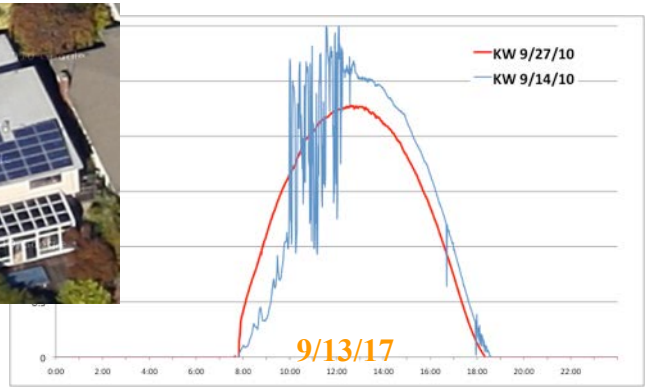

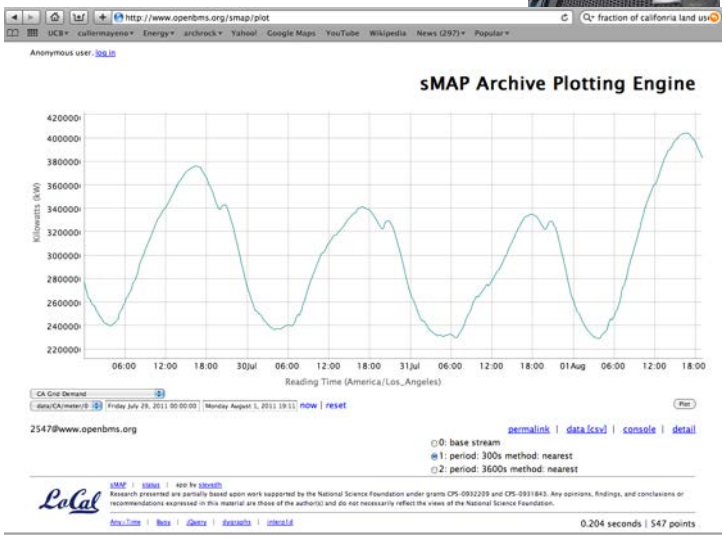

NREL BETS

Towards an 'Aware' Energy Network

Baseline + Dispatchable Tiers

Oblivious Loads

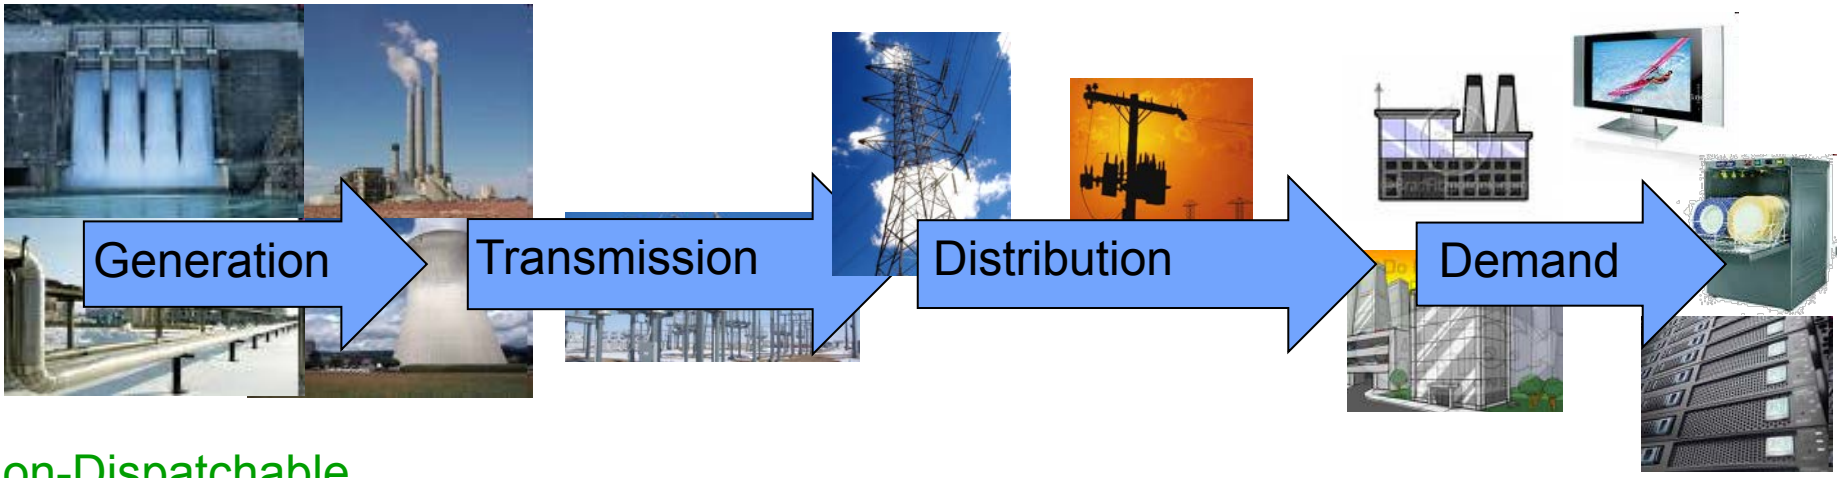

Non-Dispatchable Sources

Aware Interactive Loads

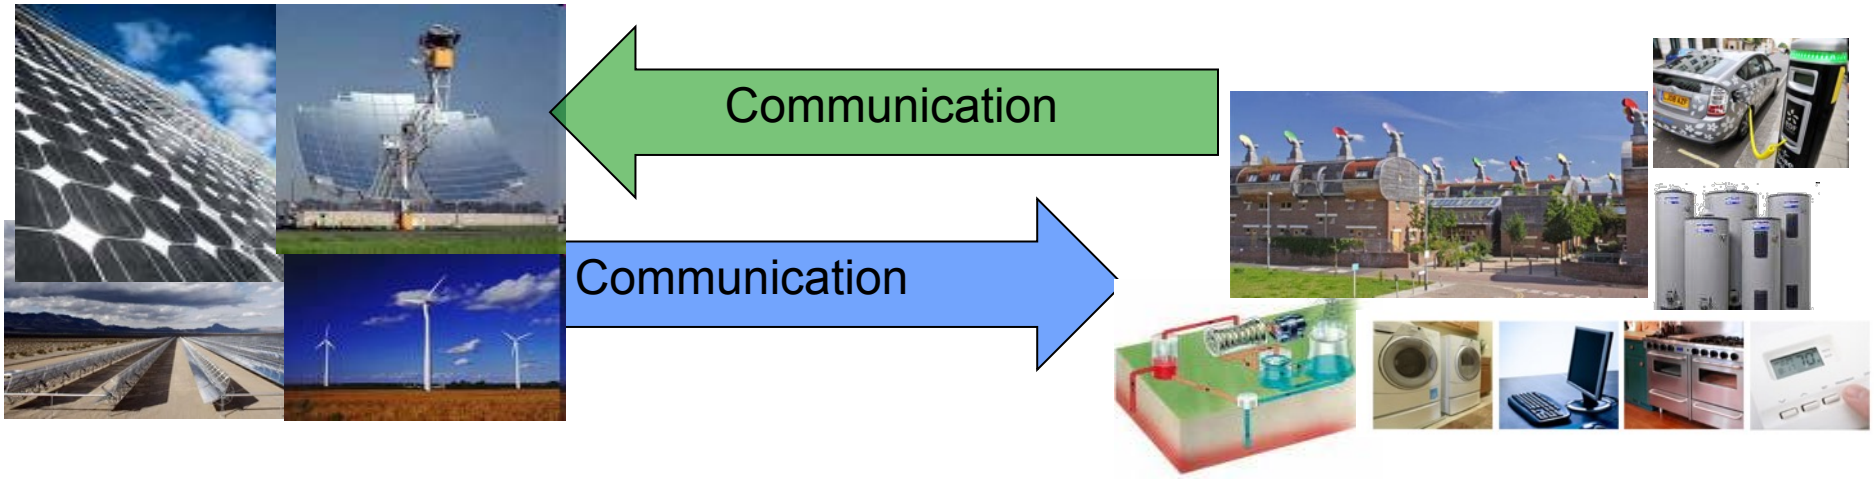

News › World › Americas

California breaks energy record with 80% of state's power generated using renewable methods

Golden State generated 67% of its energy from renewables in one day

Rachael Revesz | @RachaelRevesz | Monday 22 May 2017 10:35 BST |

9K shares

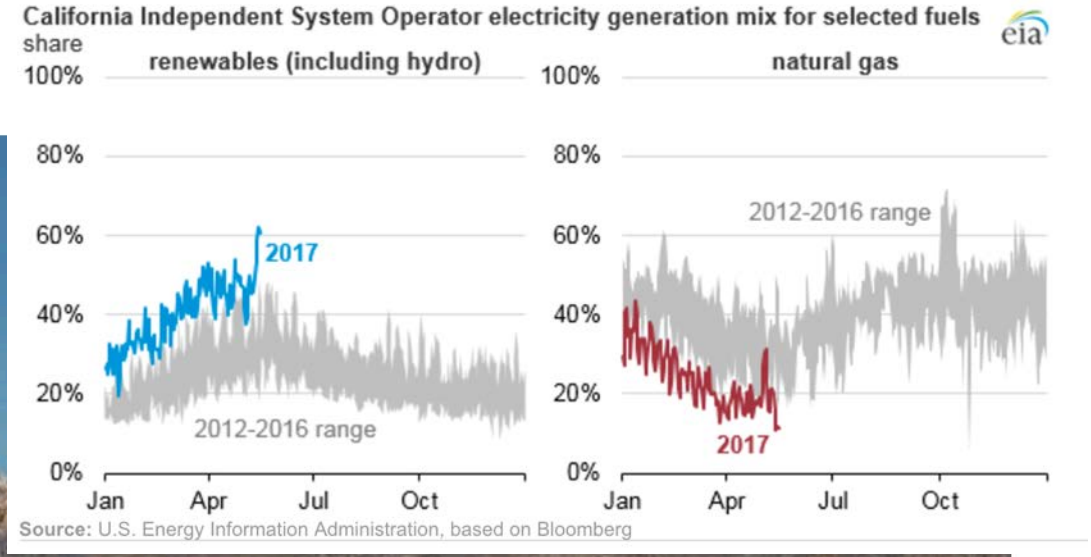

Classical view of the Energy Challenge

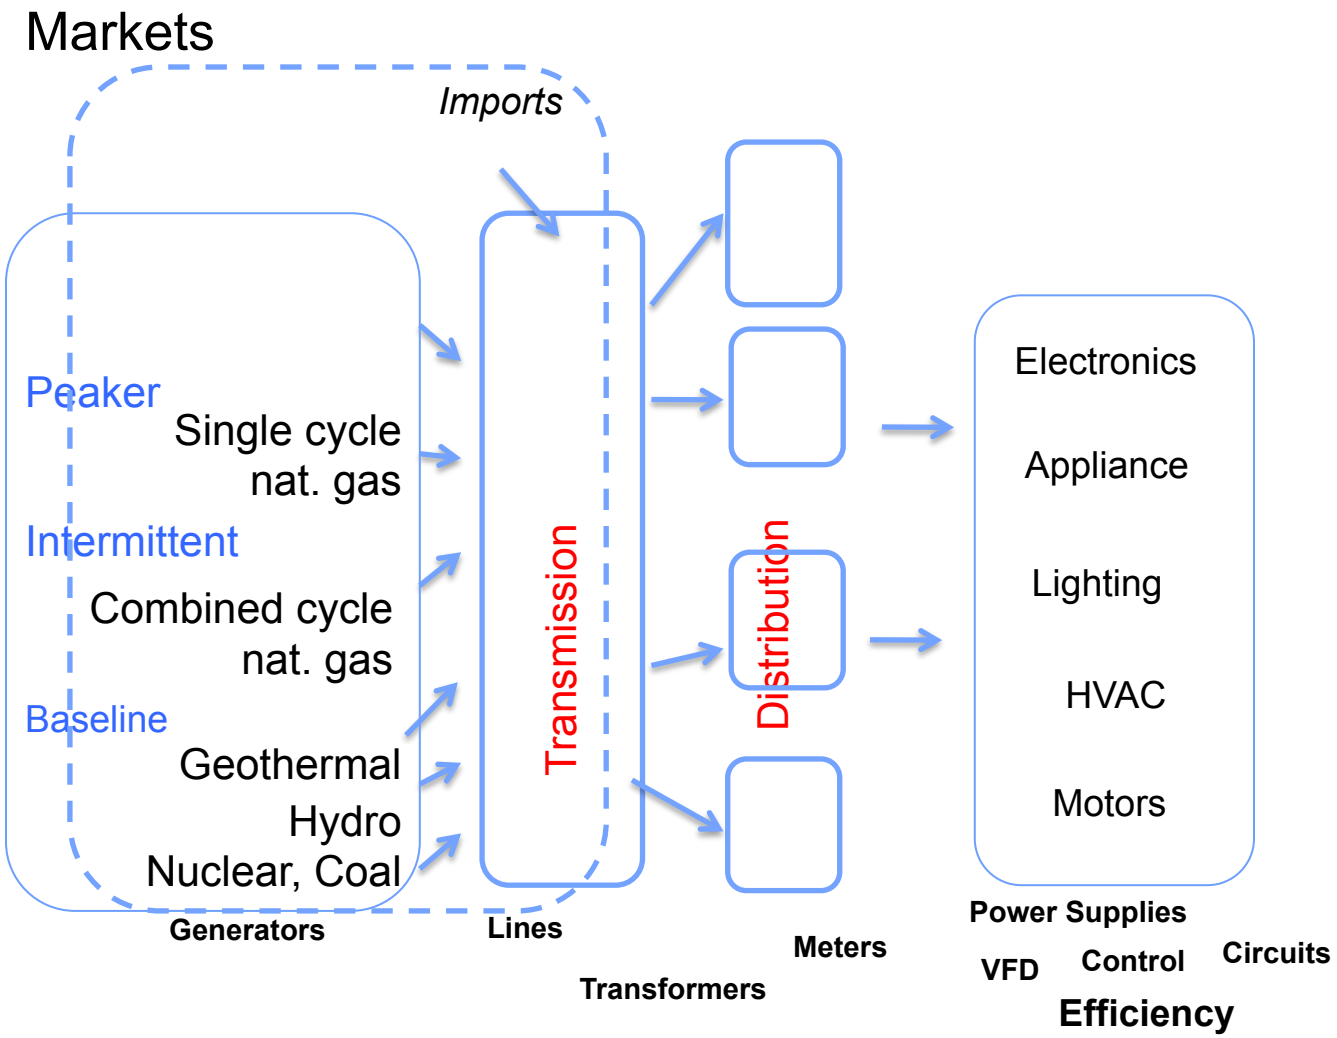

Energy Network as a System

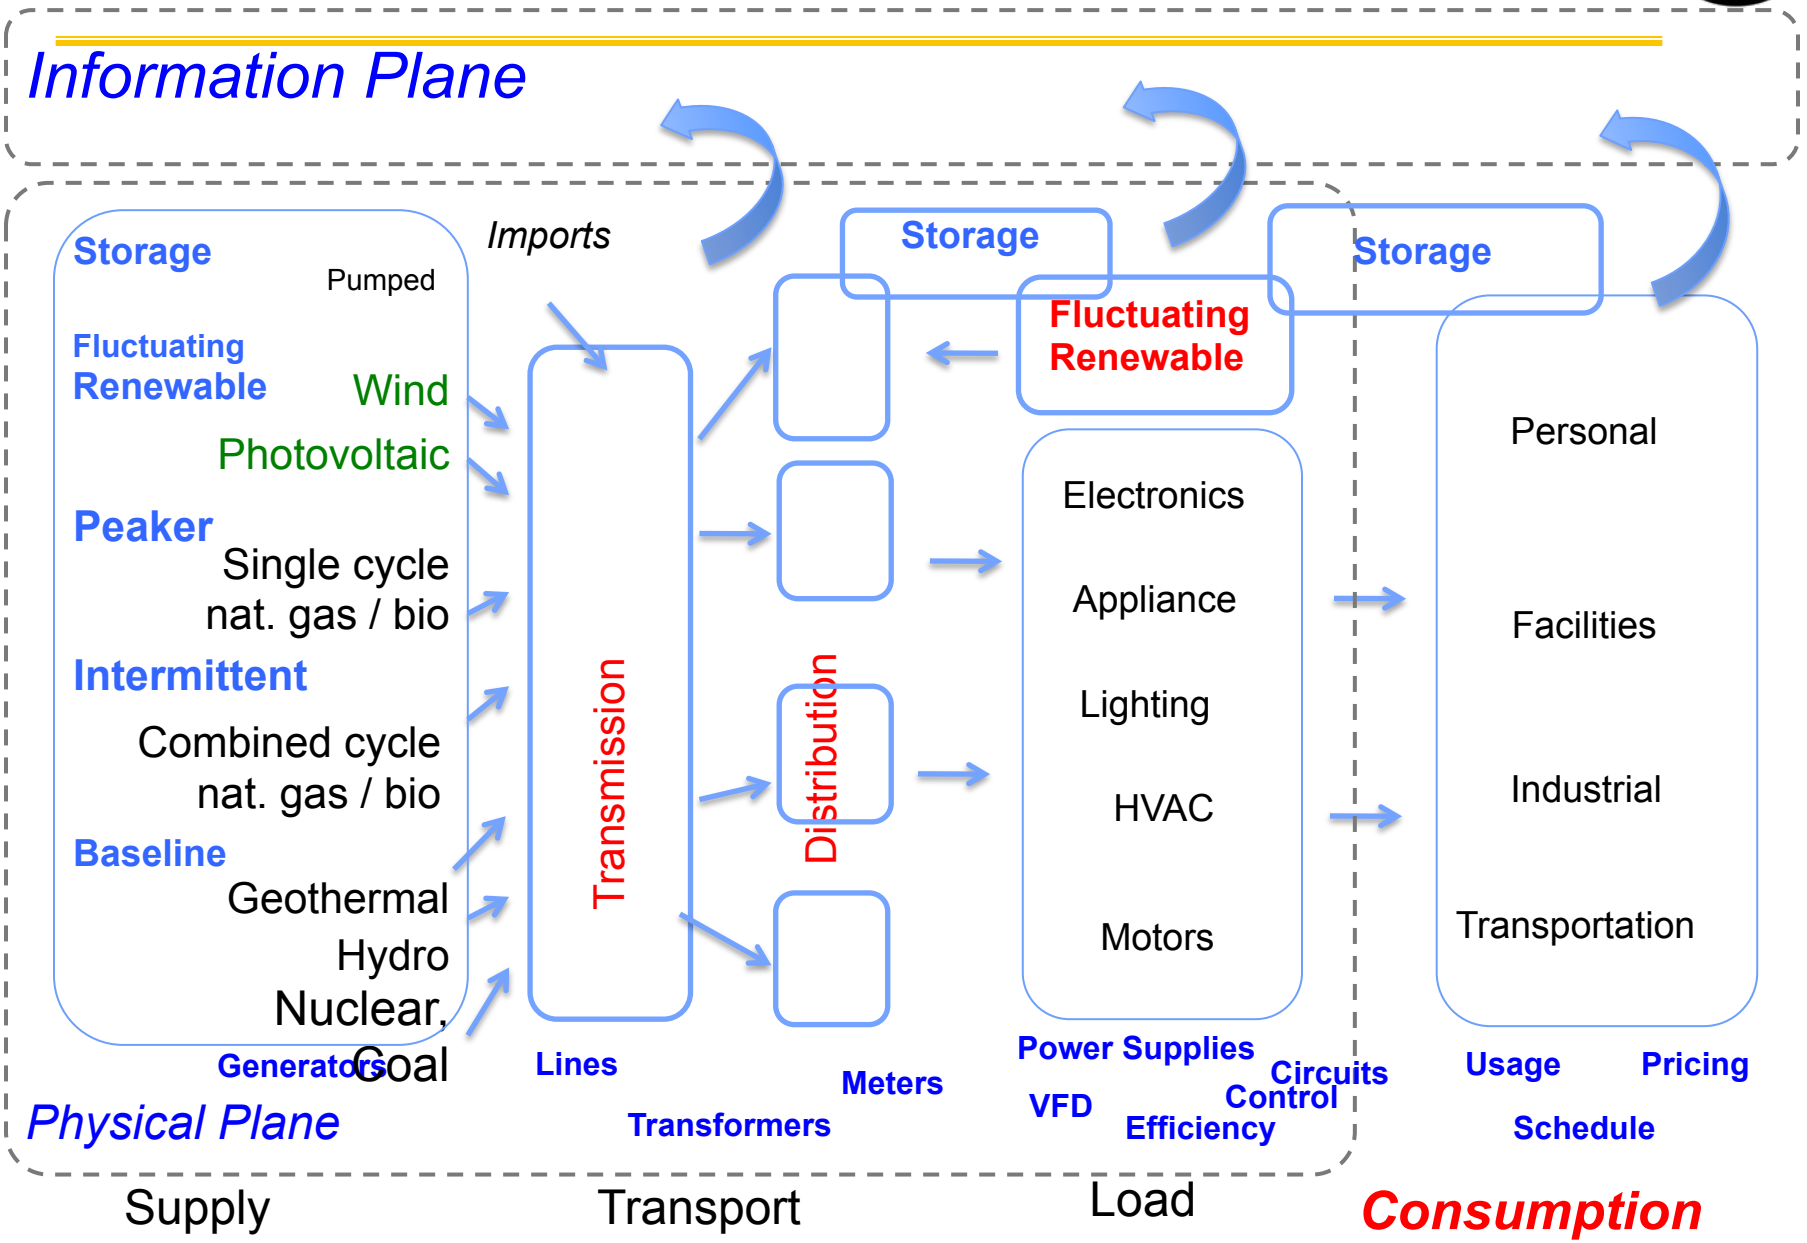

Uniform Access to Diverse Physical Information

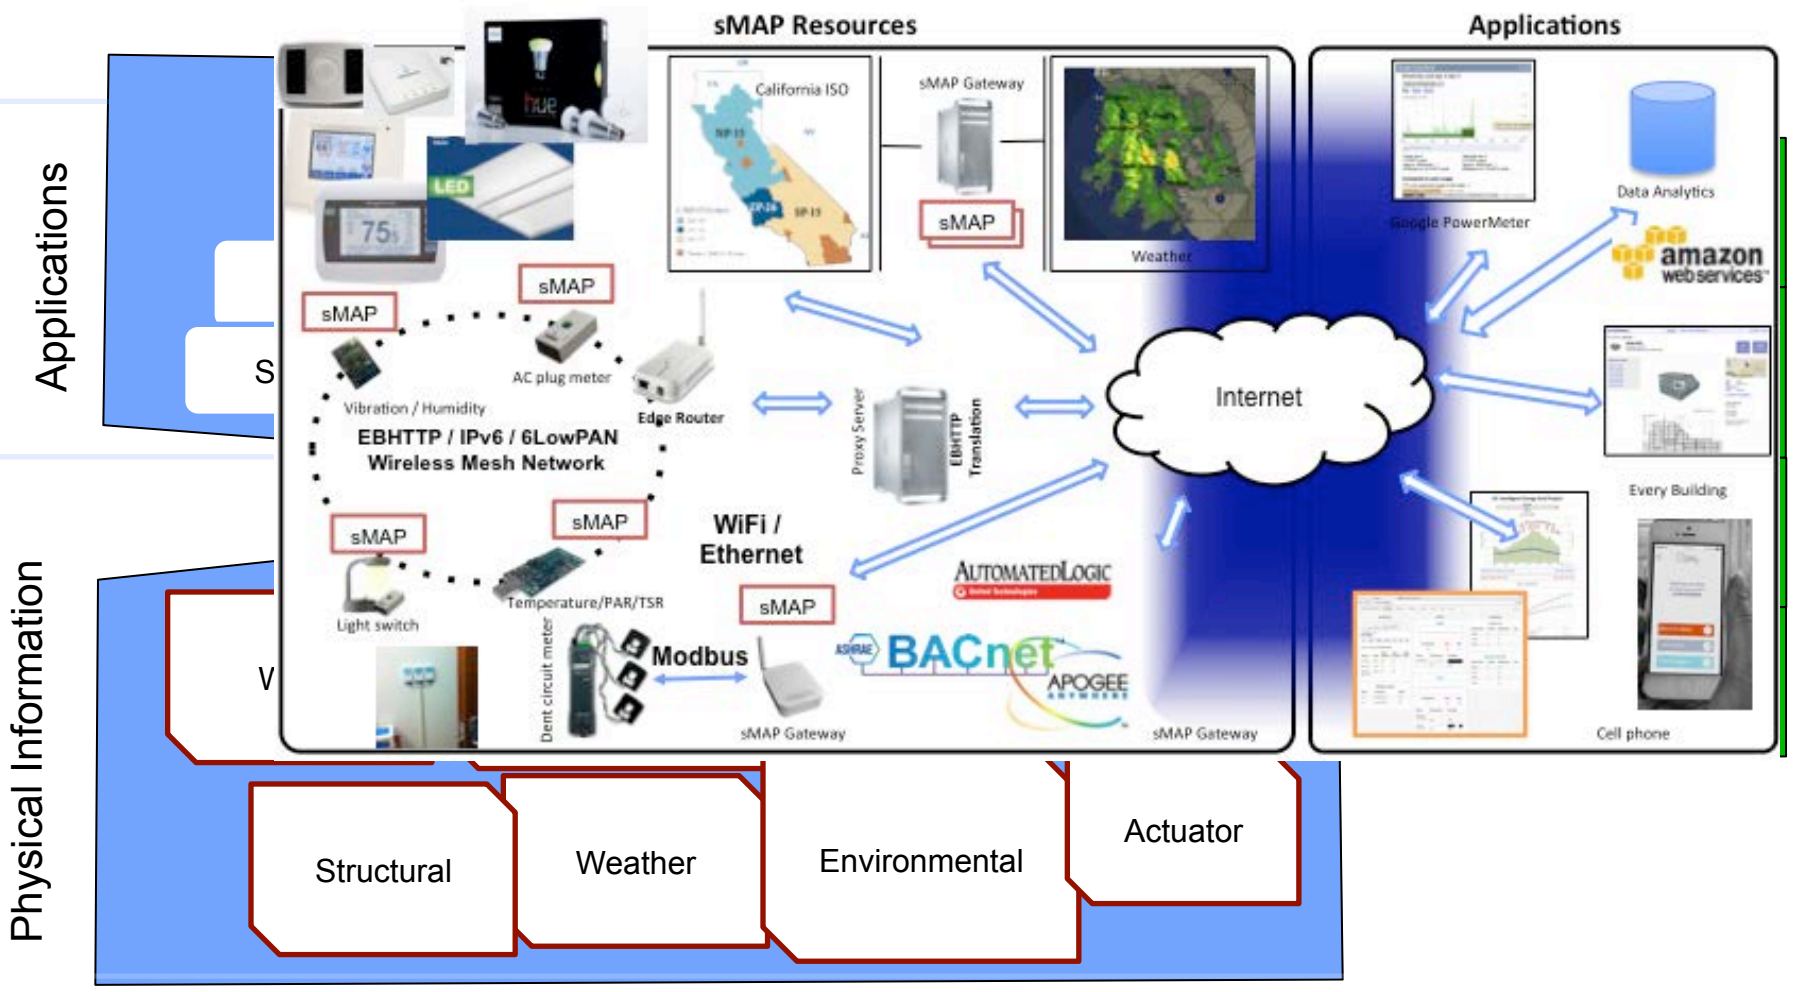

A year in a grid

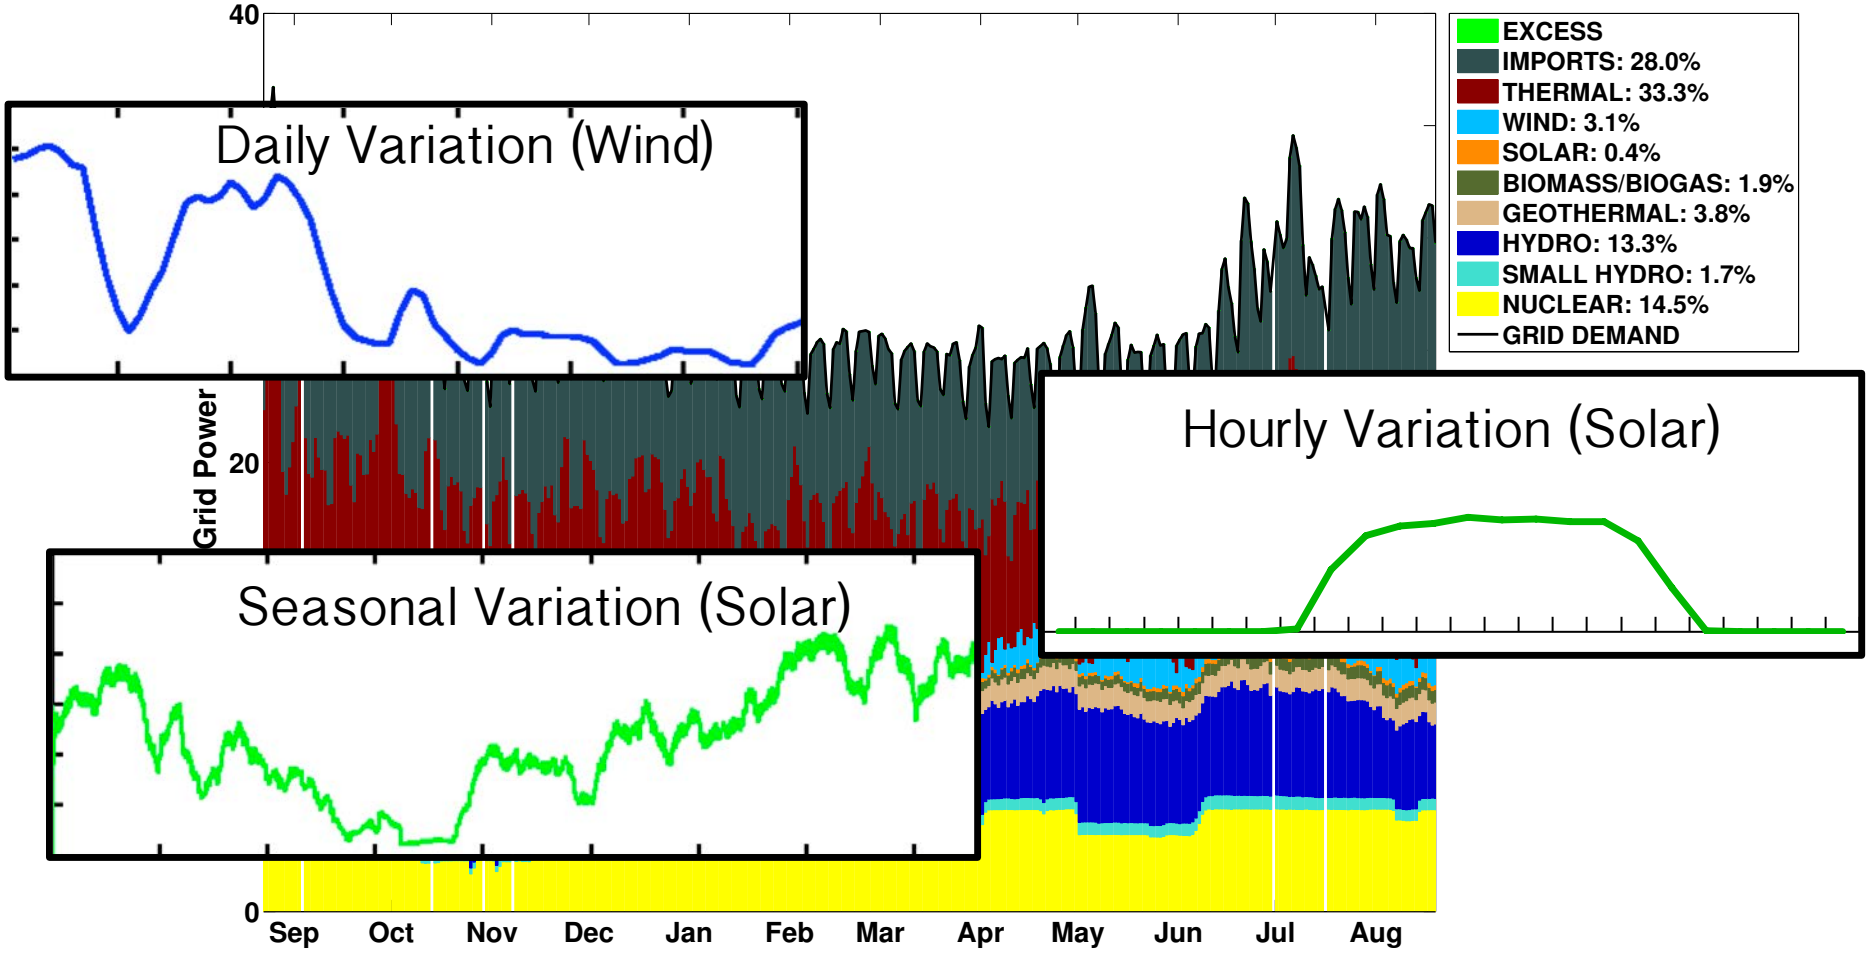

Deep penetration / Current Load

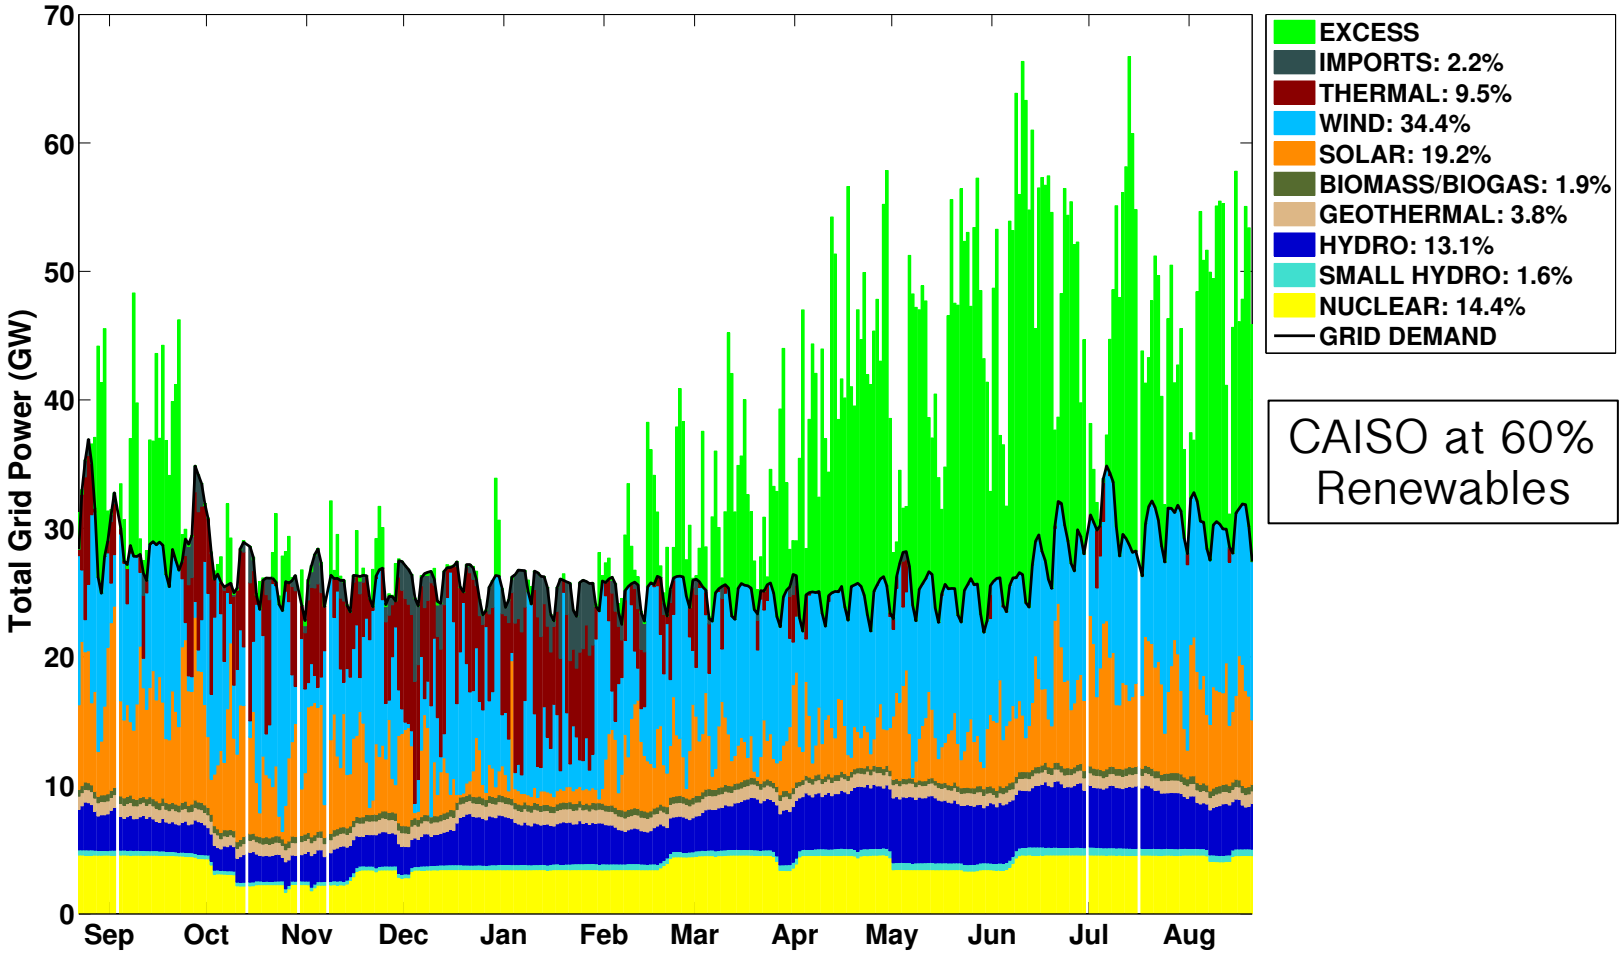

CAISO at 60% Renewables

Every Grid / Load is Different

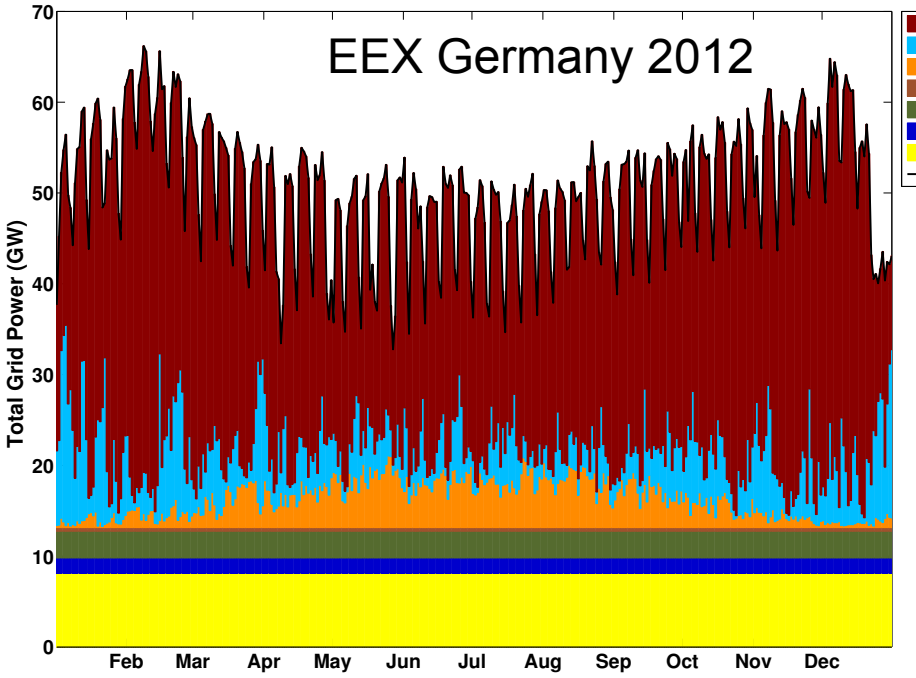

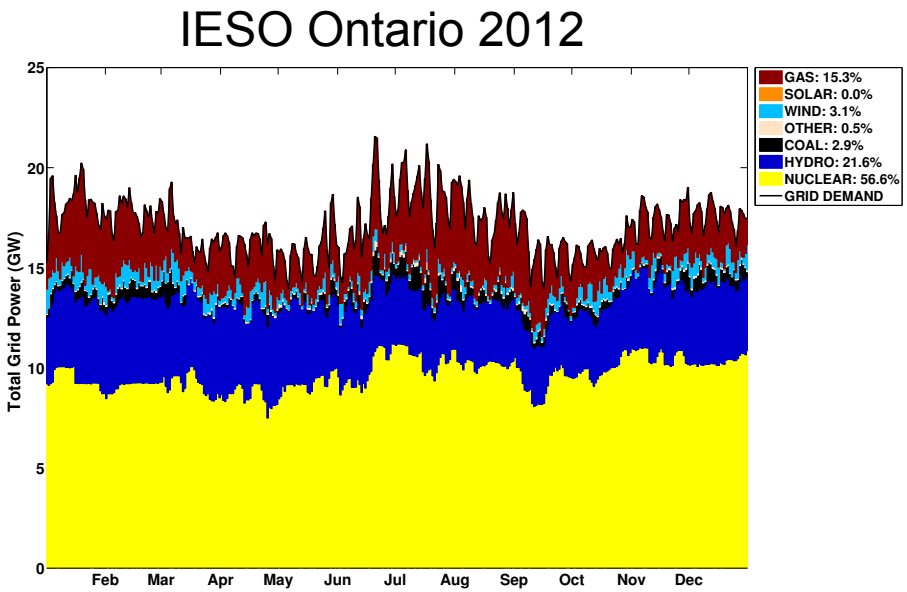

Living Lab Approach: Innovate in a "Virtual Grid"

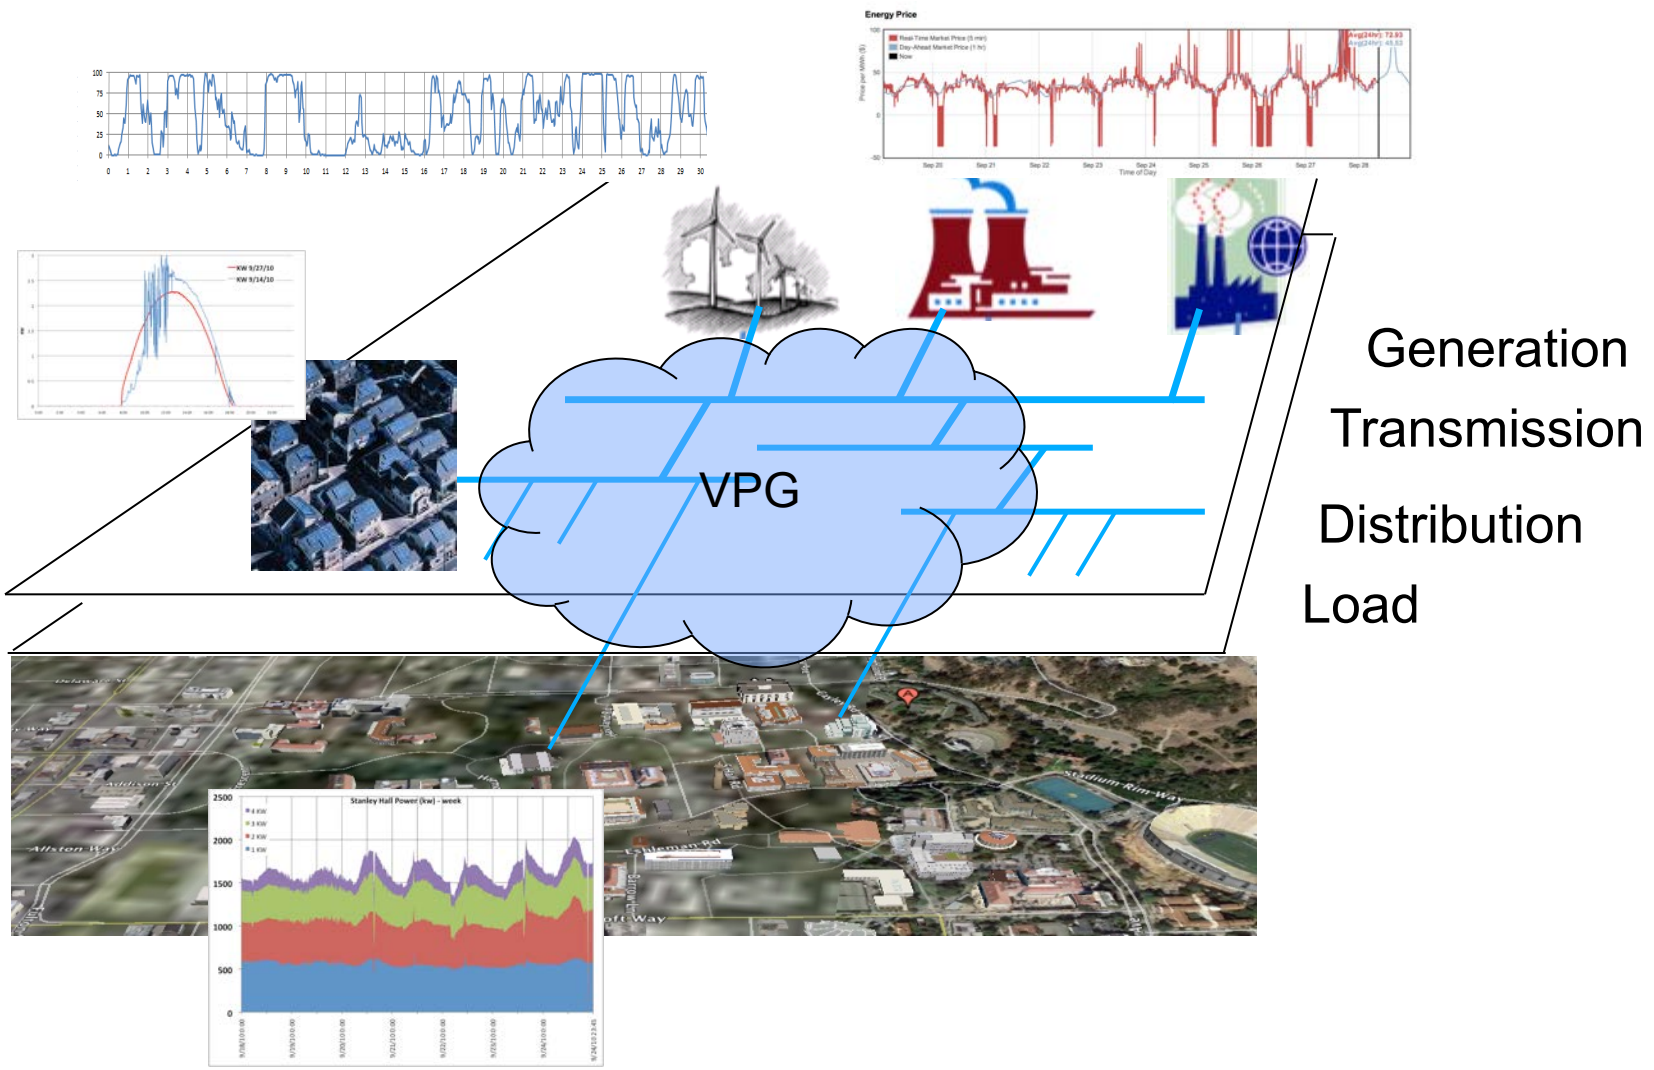

Demand Side: Living Lab Testbeds

The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

The "Other" Energy Usage

Wireless plug meters on 611 of 1200 loads

Bldg 90

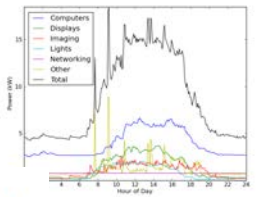

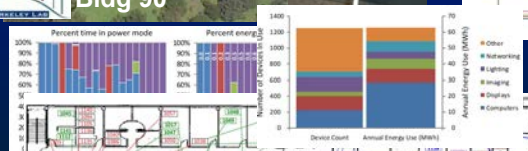

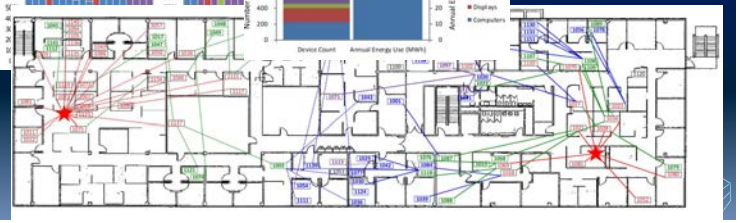

Sutardja Dai Hall

Software controllable!

4 years old

7 floors, 140 k sq. feet

Collaboratories, offices, classrooms, auditorium & nanofab

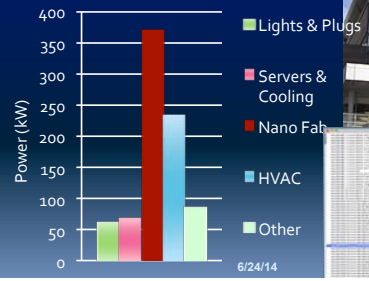

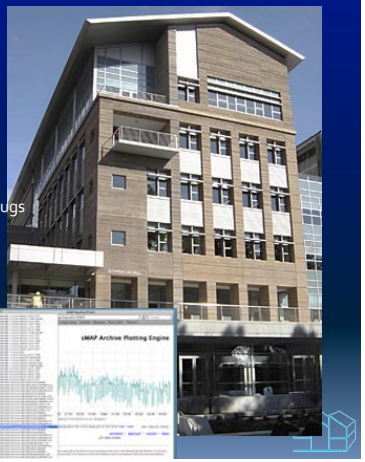

22 6/24/14

Full of Interesting Sensor Networks

151 Temperature Sensors

312 Light Relays

> 6,000 Sense and Control Points

50 Electrical Sub-meters

ble Speed Pumps

12 Variable Speed Fans

138 Air Dampers

21 Controllable Valves

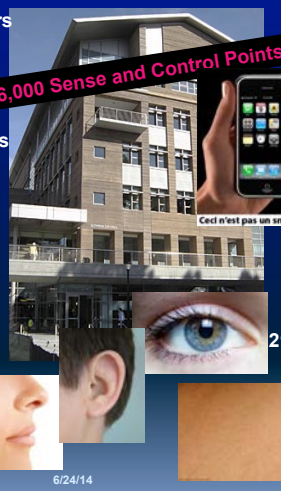

23 6/24/14

Load? - Buildings …

- Where we spend 70% of our electricity (USA)
- Where we spend 40% of our energy
- Where we spend 40% of CO2 emissions
- *Natural counterbalance to fluctuating renewables*

- Where we spend a lot of our $'s
- Where we spend 90% of our lives
- 2/3rds of occupants are uncomfortable

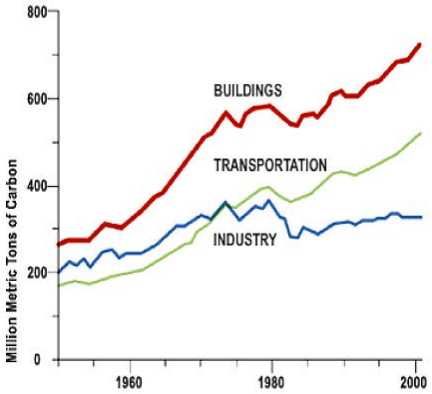

- And once they are built all we can do is …
  - use their hard-wired capabilities,
  - decorate, or
  - "retrofit"

Renewable energy consumption

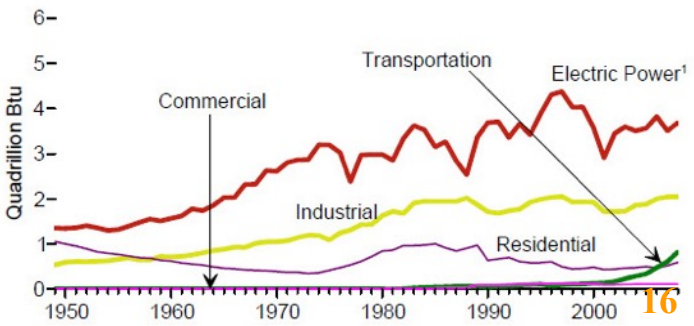

A Complex CyberPhysical System

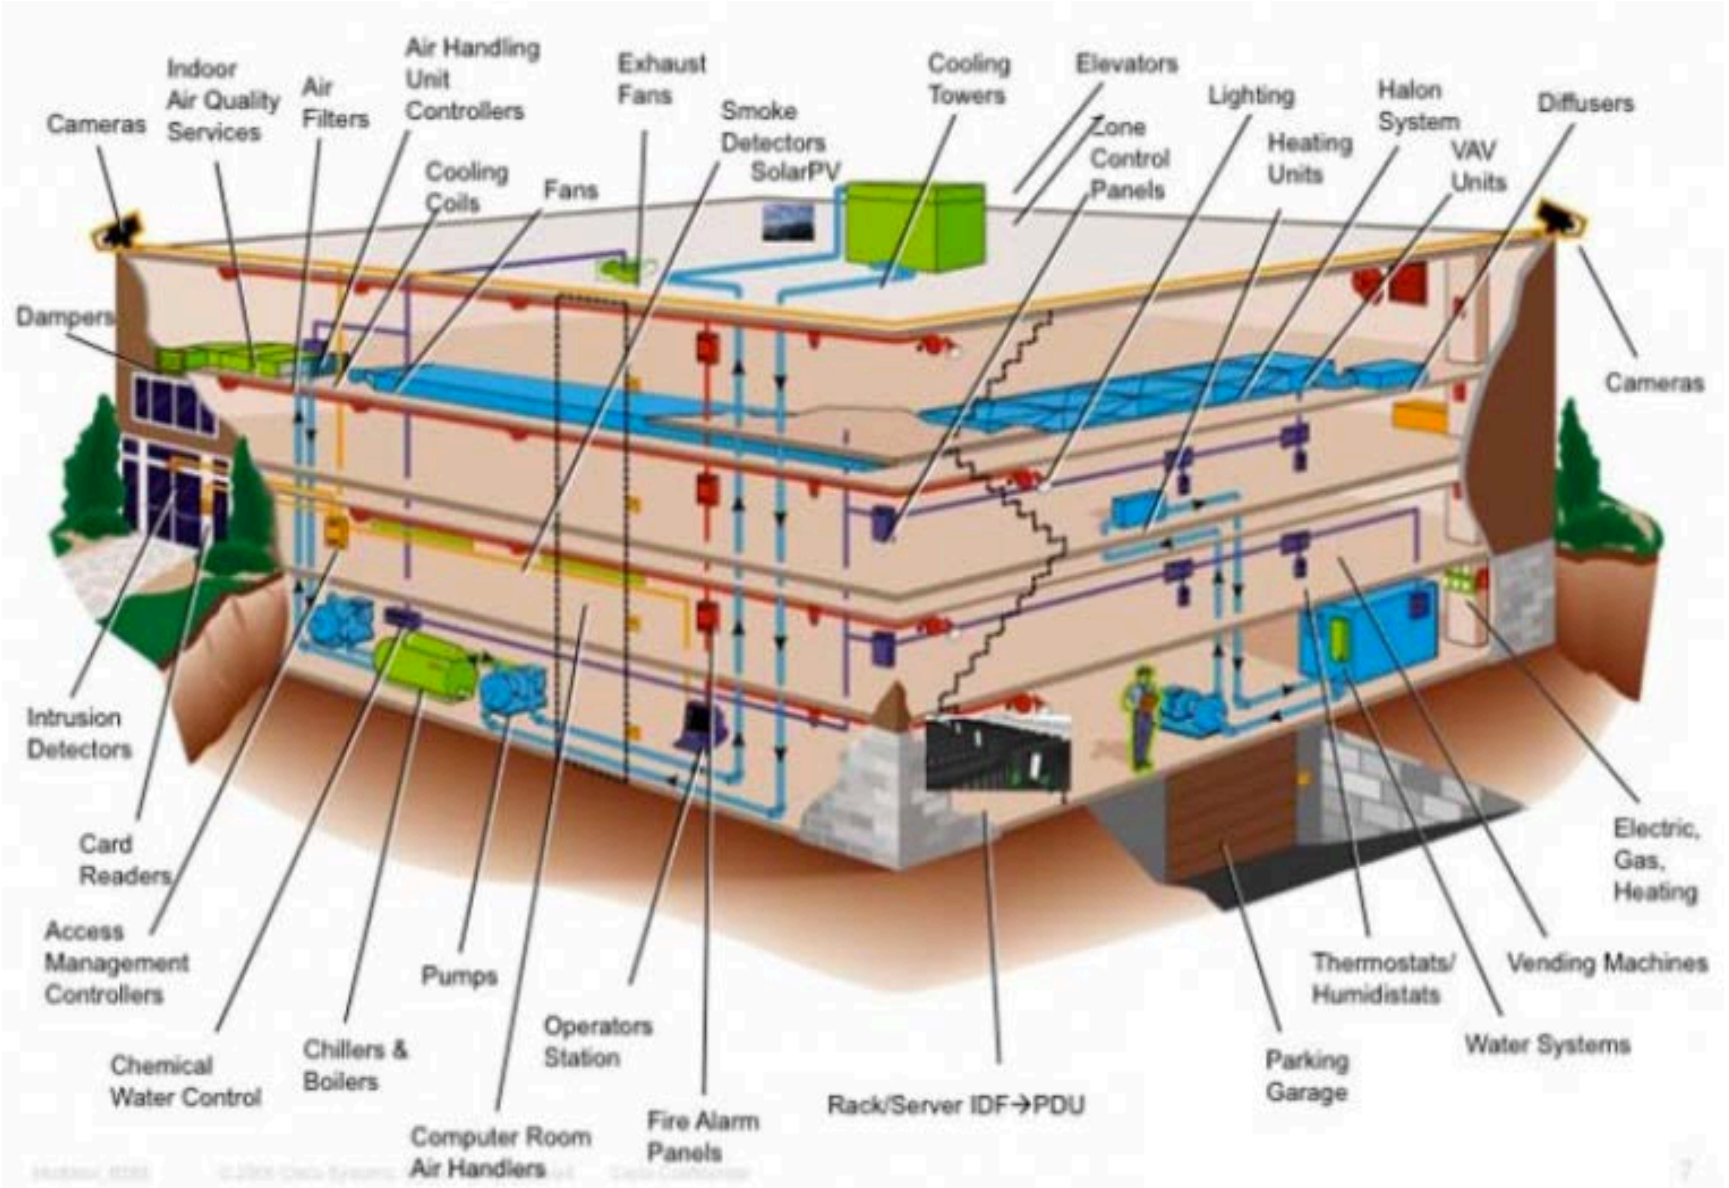

Our Buildings

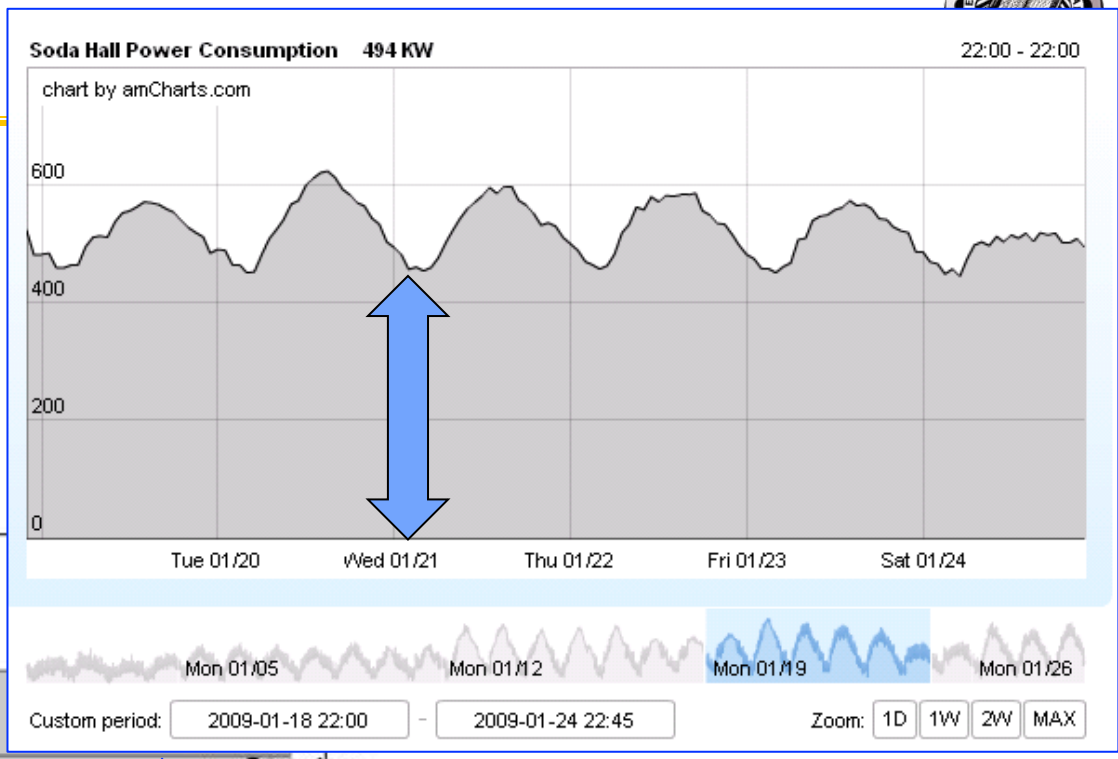

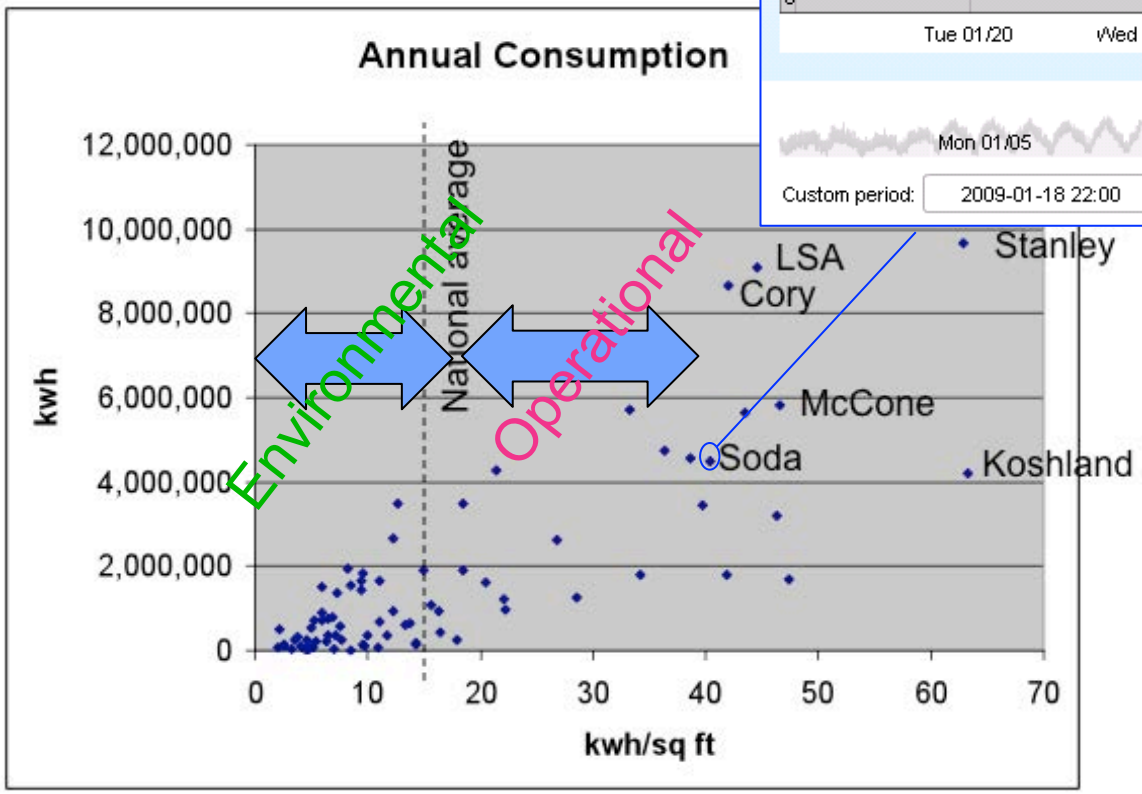

How can we transform buildings into fundamentally more agile machines?

- Programmable
- Separation of the hardware capabilities (primitives)
- from the universe of potential behaviors (applications)
- allow them to be tailored to our desires
  - To the full extent of the underlying capabilities
- And become good citizens of the grid

Elements of a Software Defined Building

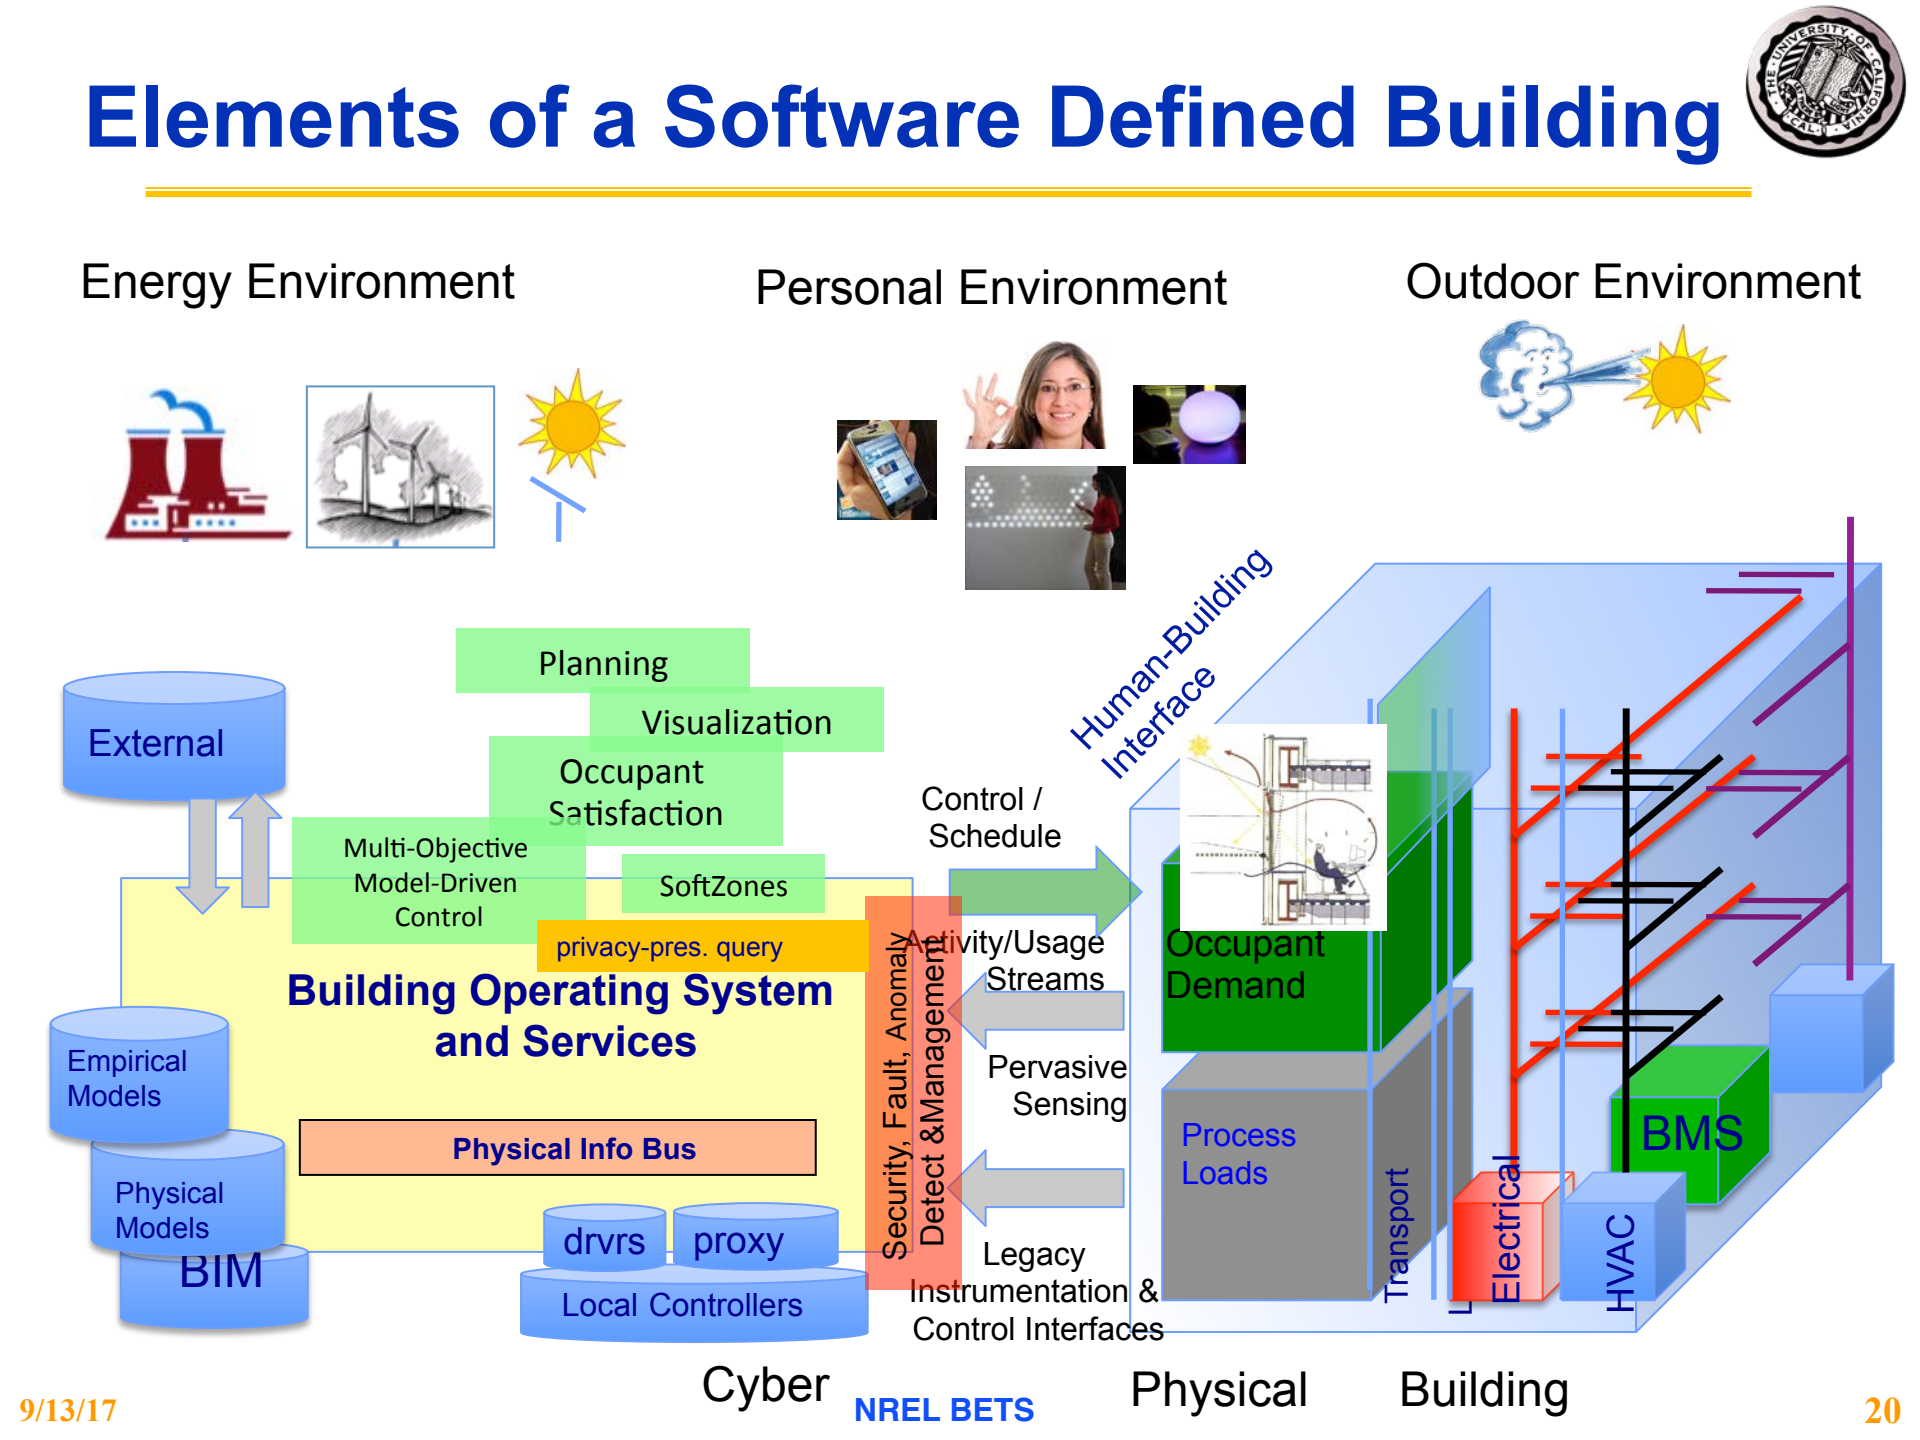

The Building Operating System …

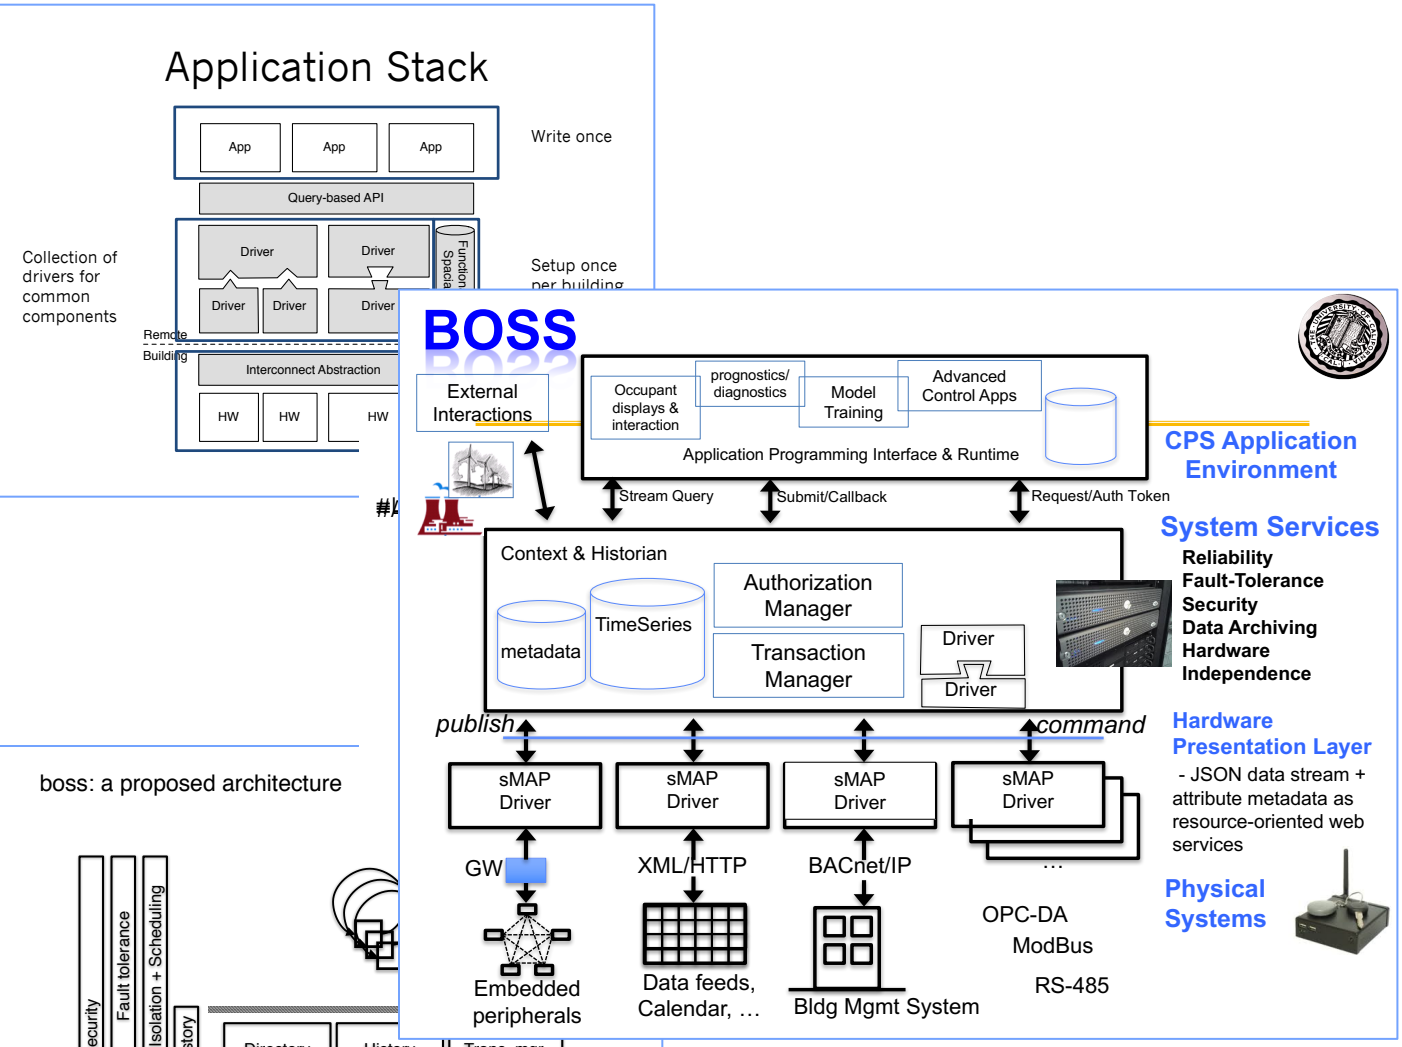

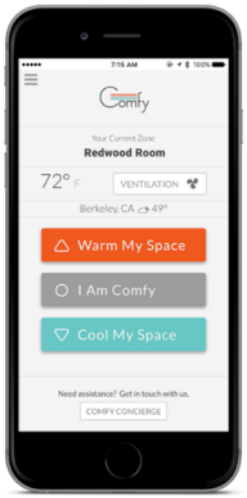

A building app for that …

Energy Environment

Personal Environment

Outdoor Environment

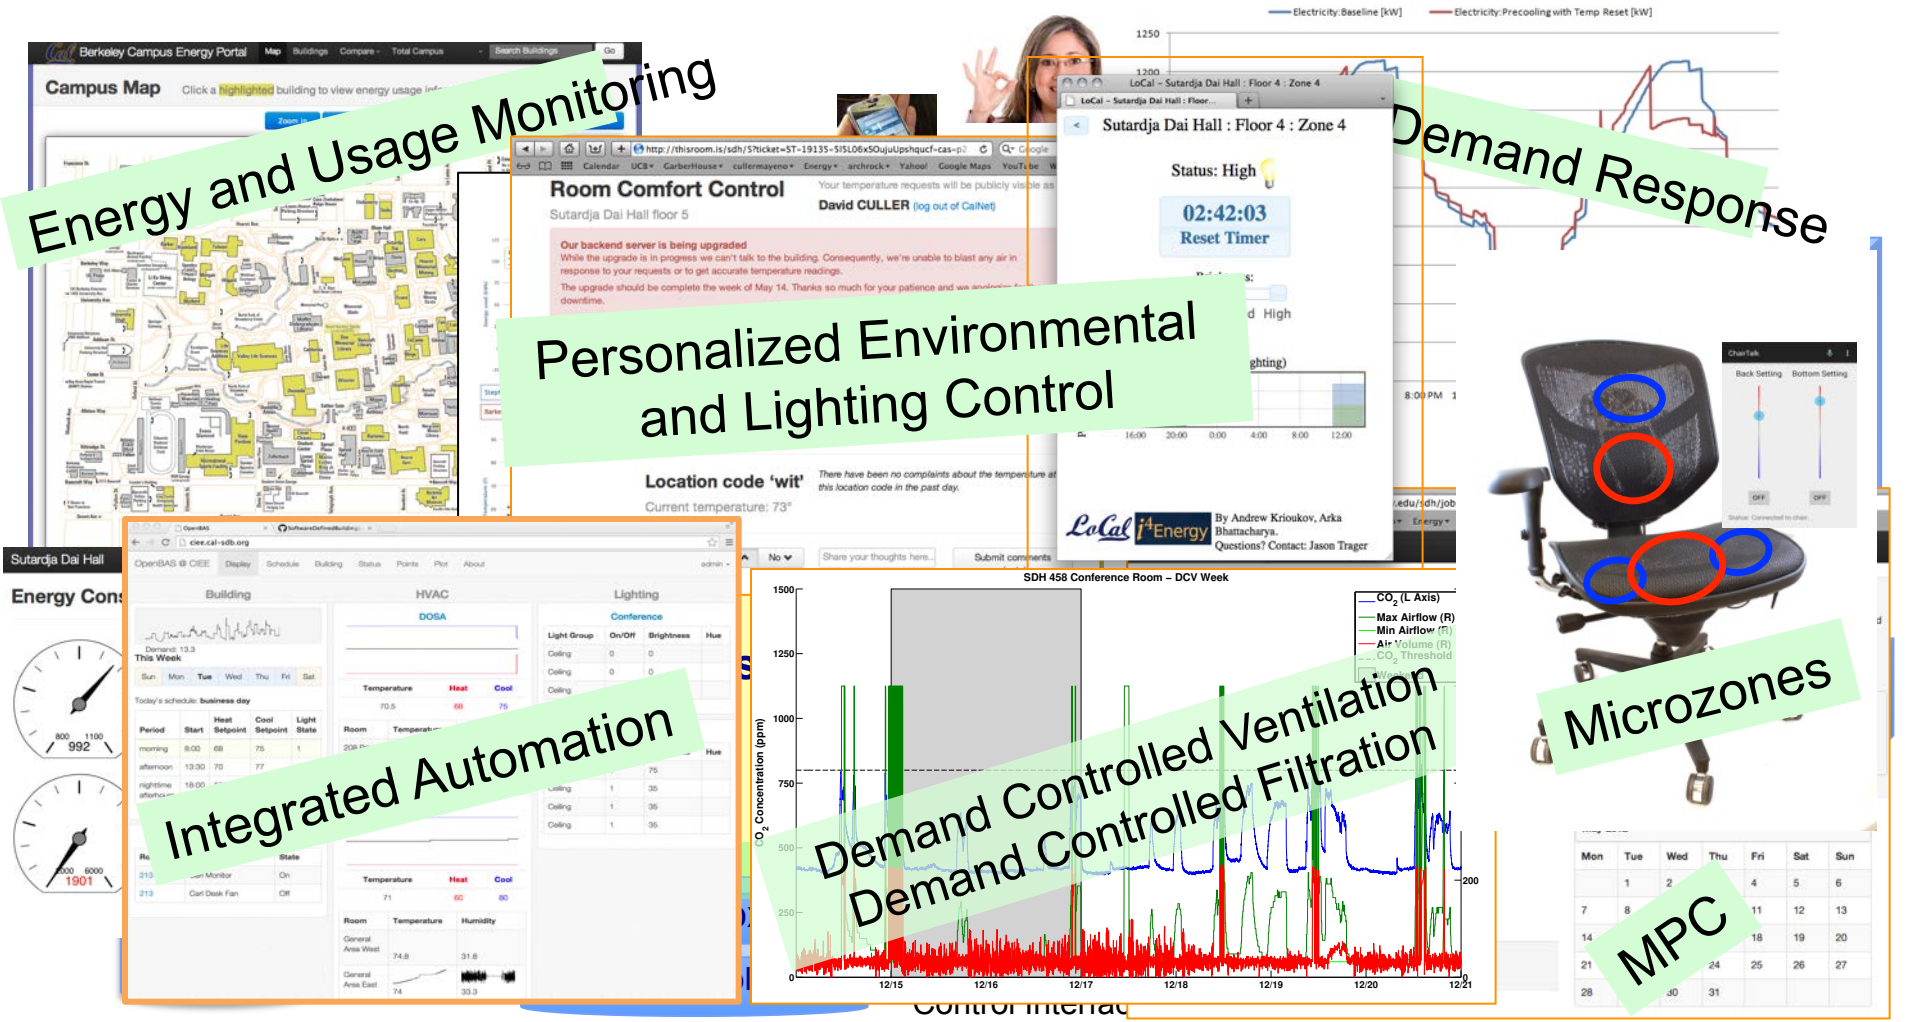

Cyber

Physical

Building

Well-Connected Microzones for Increased Building Efficency and Occupant Comfort

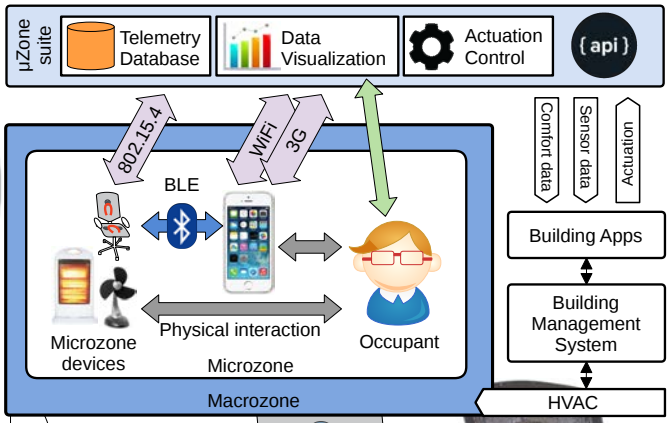

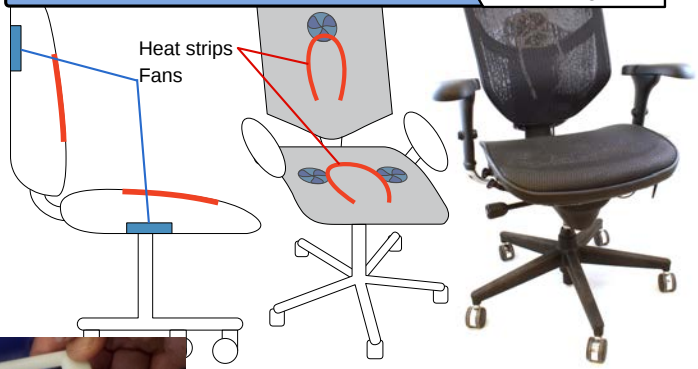

CBE

Supply-following Fridge w/ Ice Battery

Thermal performance =>

Measured power =>

Real Time CAISO signal =>

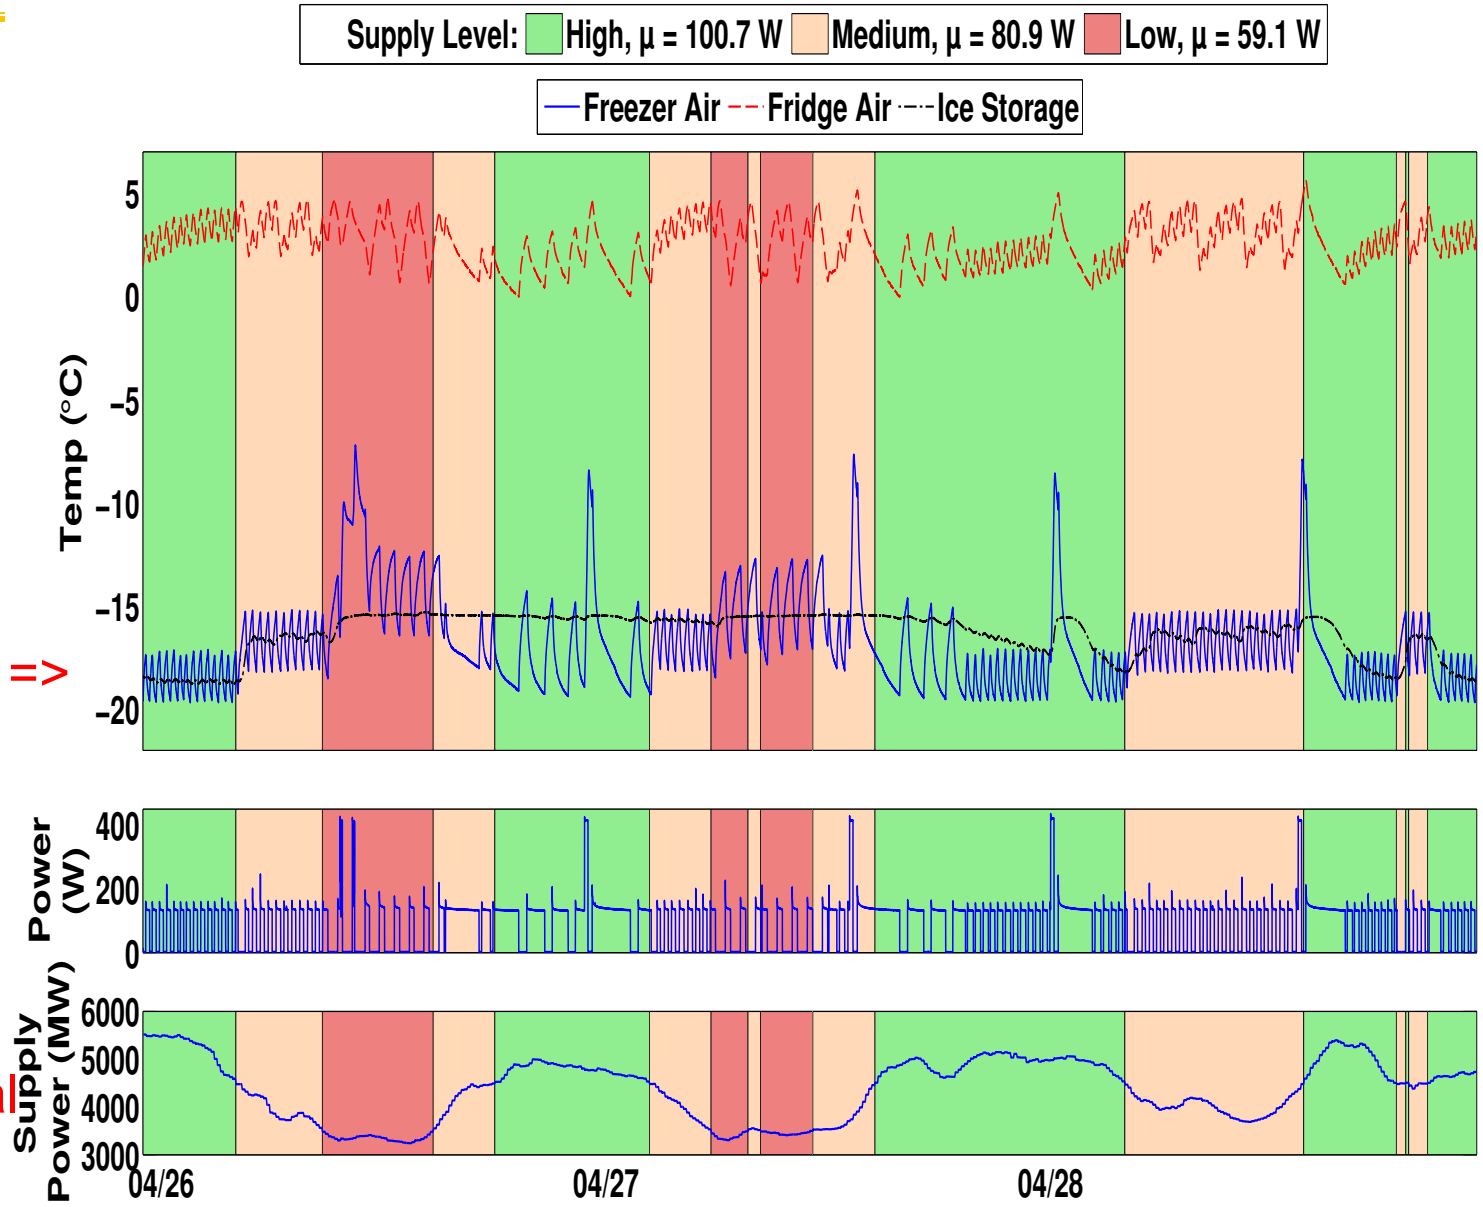

Its about metadata …

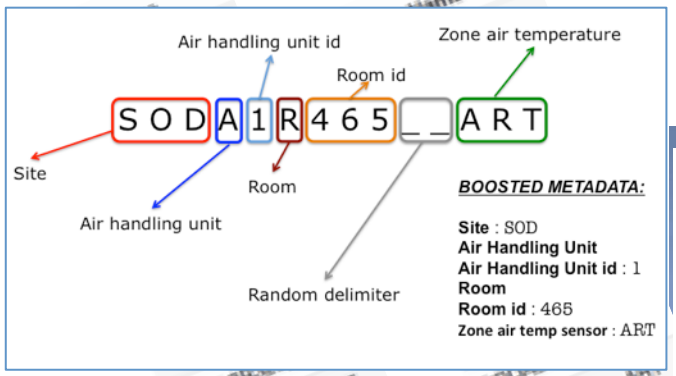

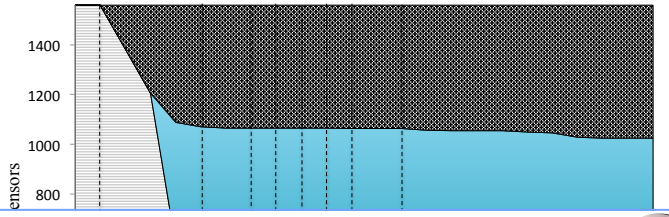

Primitive

Syntactic Exa
Selectior

Can we design one that would pass this test ?

- **Collaboration Universities (UCB, UCSD, UCLA, CMU, UVA, USD) and industry (IBM, Intel)**
- **Each brought a ground-truthed large building tag set and iterated on developing the proposal**

9/26/16 NSD>Sus 63

Coverage	Haystack	IFC	Semantic
Feedback / MPC / FDD	100%	100%	40%
Web Displays	75%	100%	100%
NILM / DR	50%	50%	50%
Energy Apportionment	57%	86%	57%
Occupancy Modelling	42%	58%	25%

- To scale advanced control, etc. to millions of buildings, acquisition of semantic content from tags, readings, CAD, etc. must be automated.

Critical "Building Load" is its vehicles

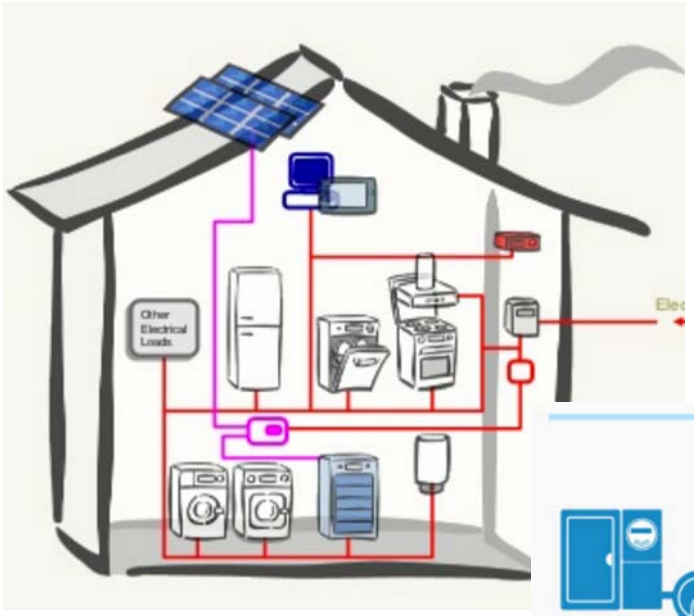

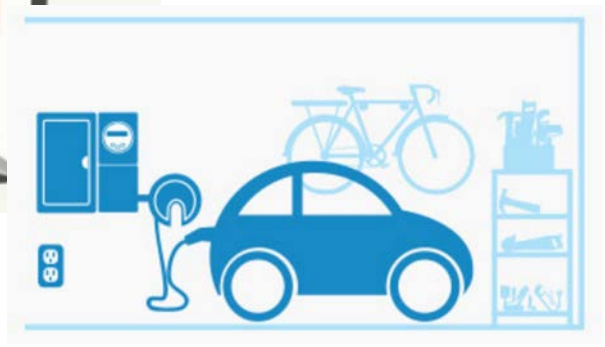

http://www.pevcollaborative.org/

- **Plug-In Electric Vehicles (PEV) in CA: 1.5 million clean vehicles by 2025**
- **70-85% of charging occurs at home**

CALIFORNIA ENERGY COMMISSION

Visibility into the Distribution Tier

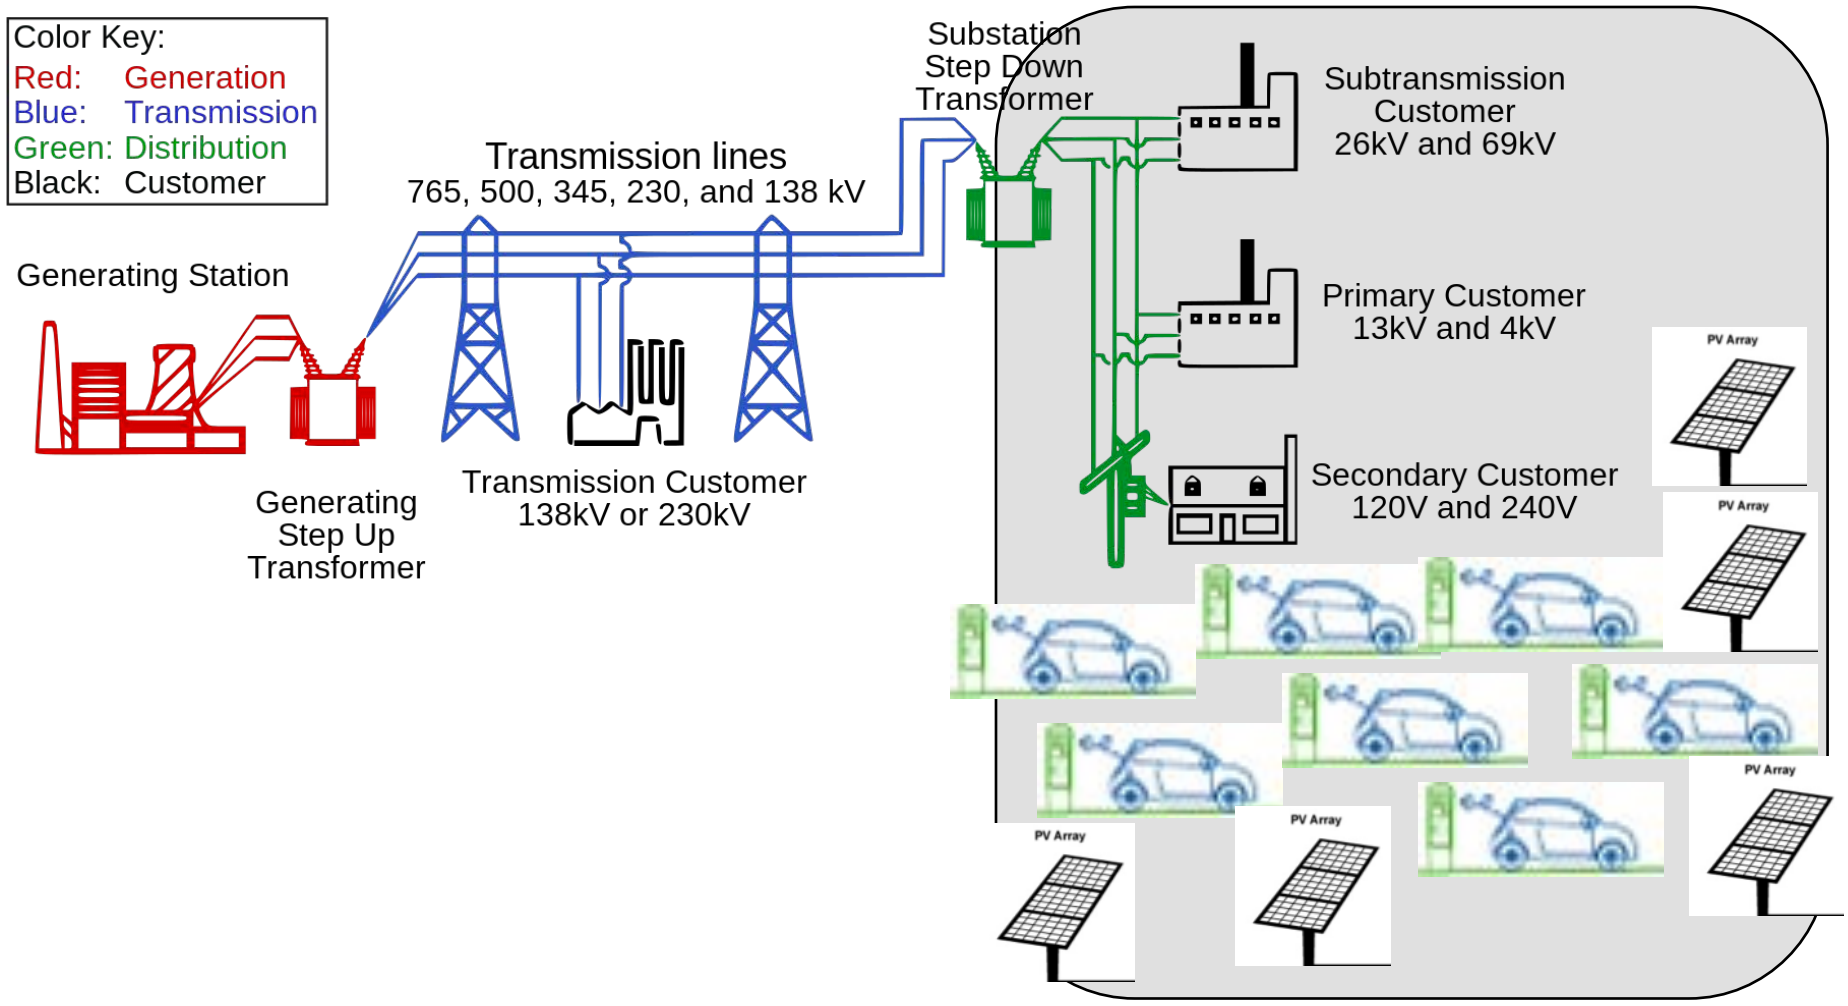

Impacts of multiple electric vehicles charging, renewable energy, storage on the distribution grid?

uPMU grid situational awareness

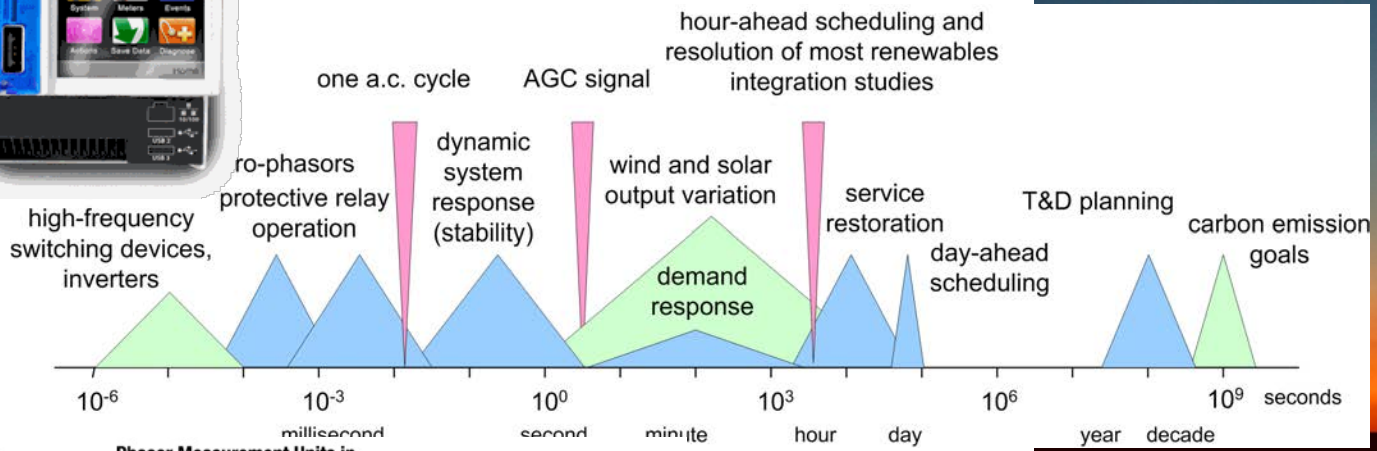

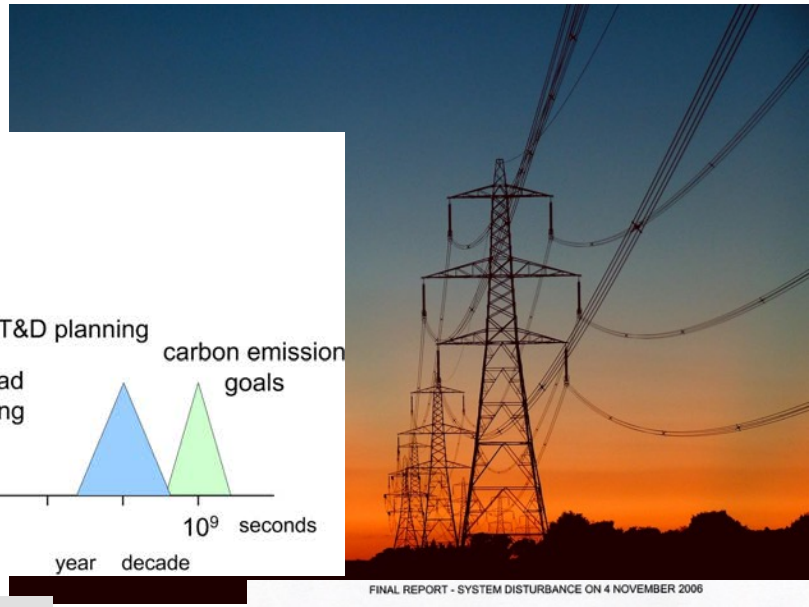

FINAL REPORT - SYSTEM DISTURBANCE ON 4 NOVEMBER 2006

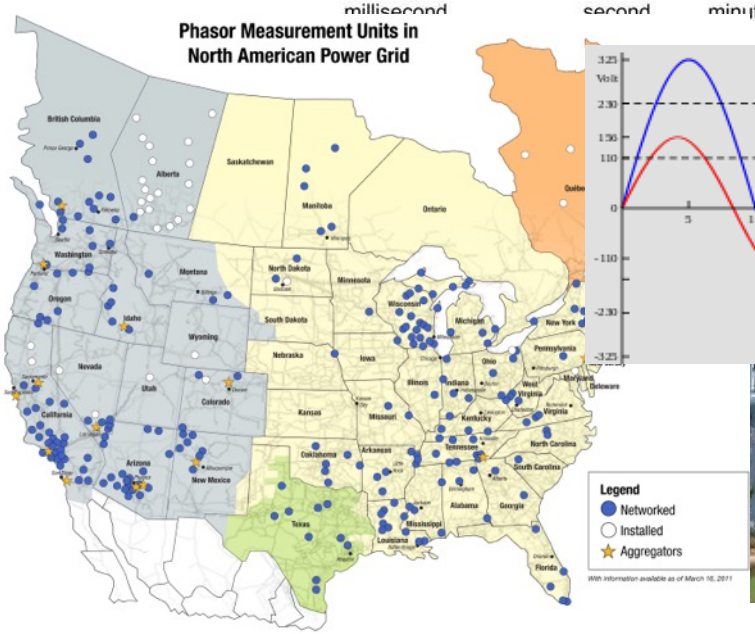

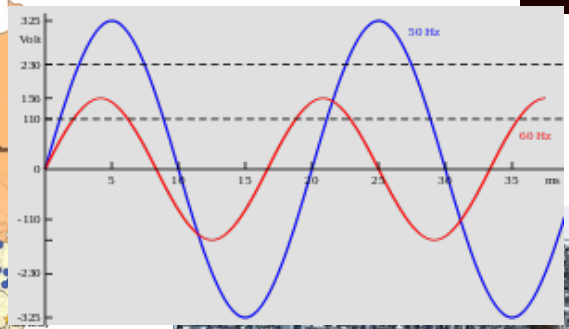

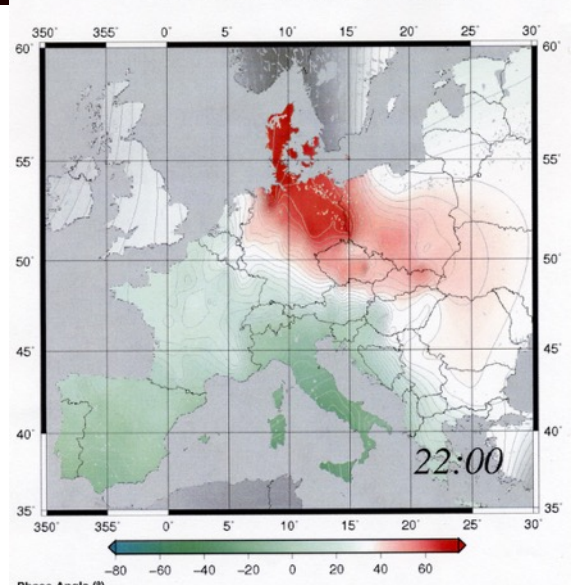

Fig. D1a: Voltage phase angle differences in the UCTE system at 22:00 /ELES/

ARPA-E µPMU Project

Field installations:

UC Berkeley/LBNL

Southern California Edison

Riverside Public Utilities

Alabama Power (Southern Co.)

Tennessee Valley Authority

Interplay of Awareness and Action

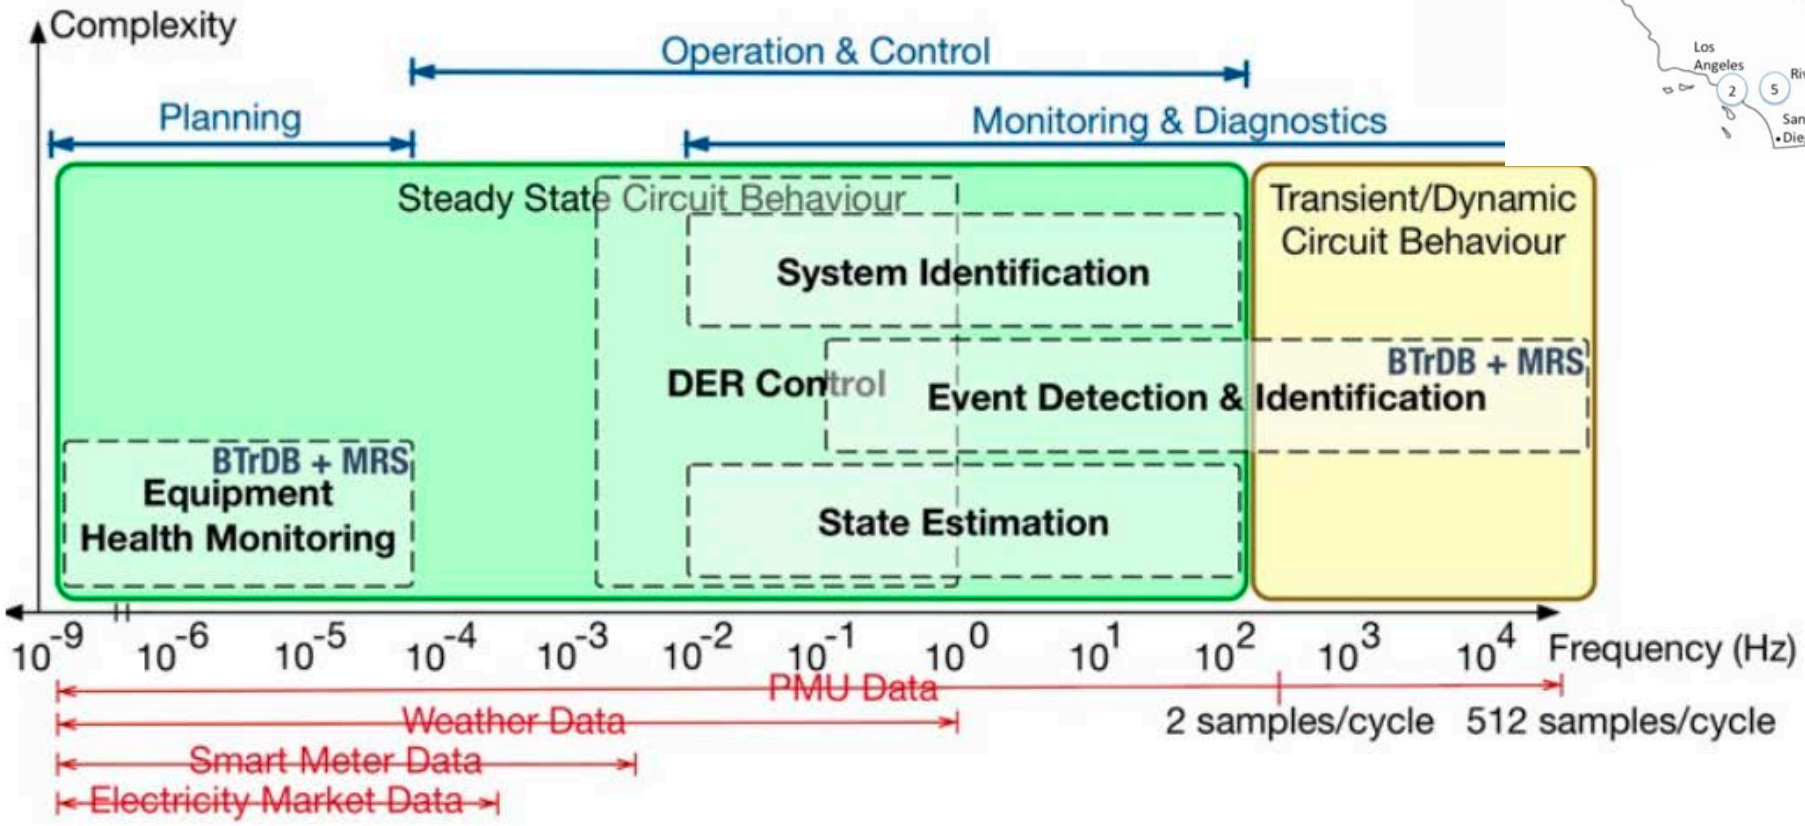

Aware-Grid System Architecture

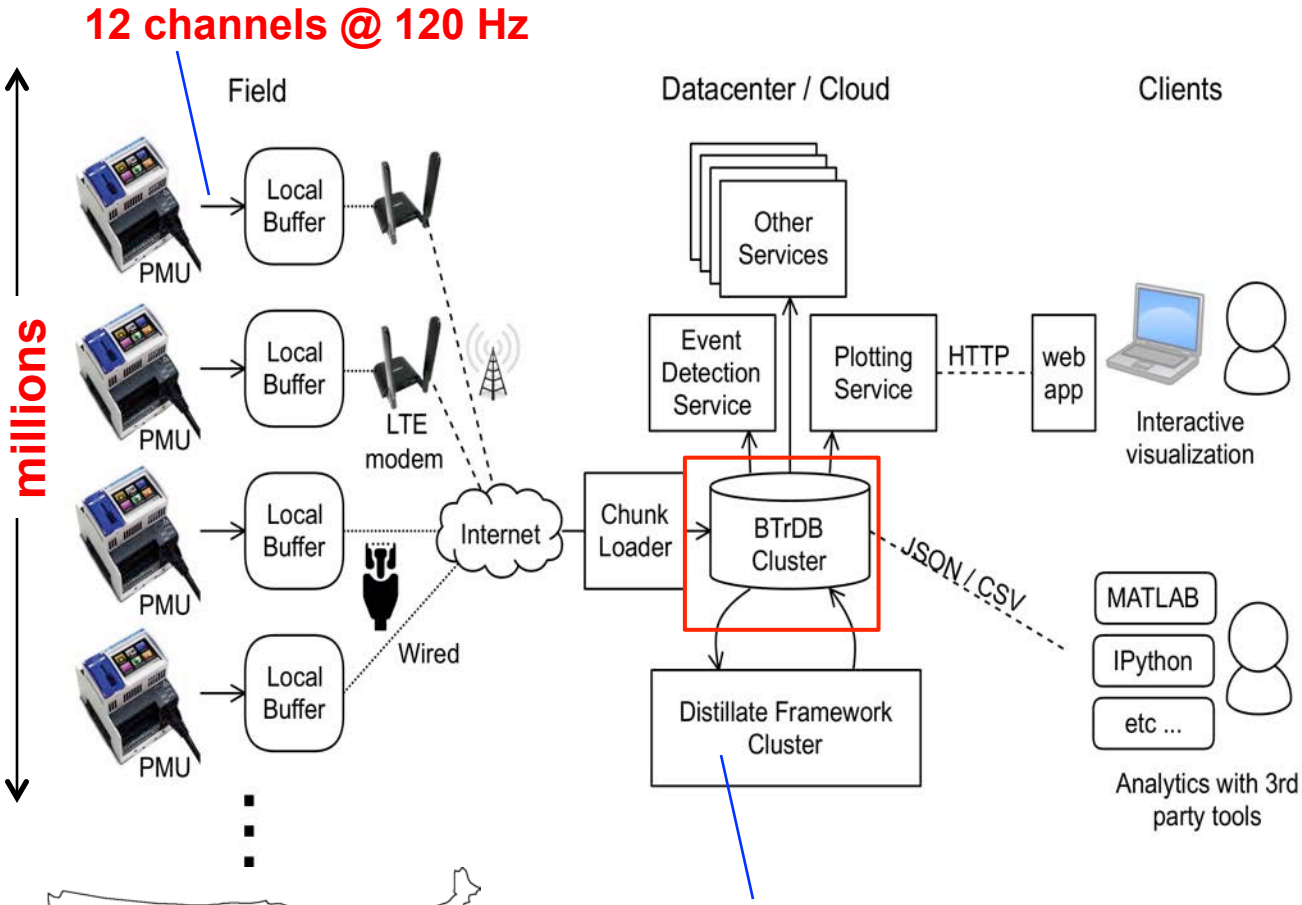

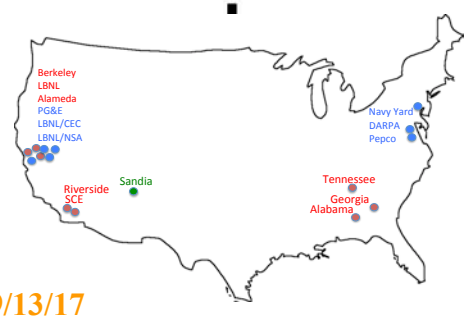

Eventually consistent derivative streams
- out-of-order arrival
- holes, corrections
- multi-stream flows

- Archiver / Database
- Stores (T, V) pairs
- High Density
- Nanosecond precision
- Varying Lag & OOO
- Fault tolerant
- Highly scalable
- Aggregated queries
- **Unique abstraction**
  - query range (ver)
  - insert values => ver
  - delete range => ver
  - query statistical (ver)
  - compute diff(v1, v2)

Example Use Case

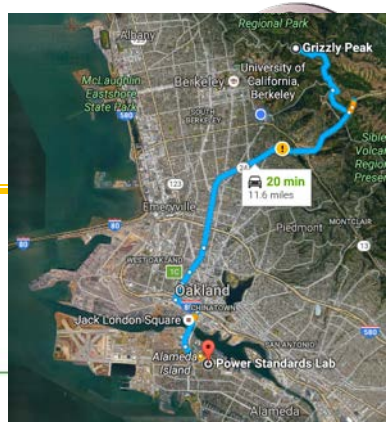

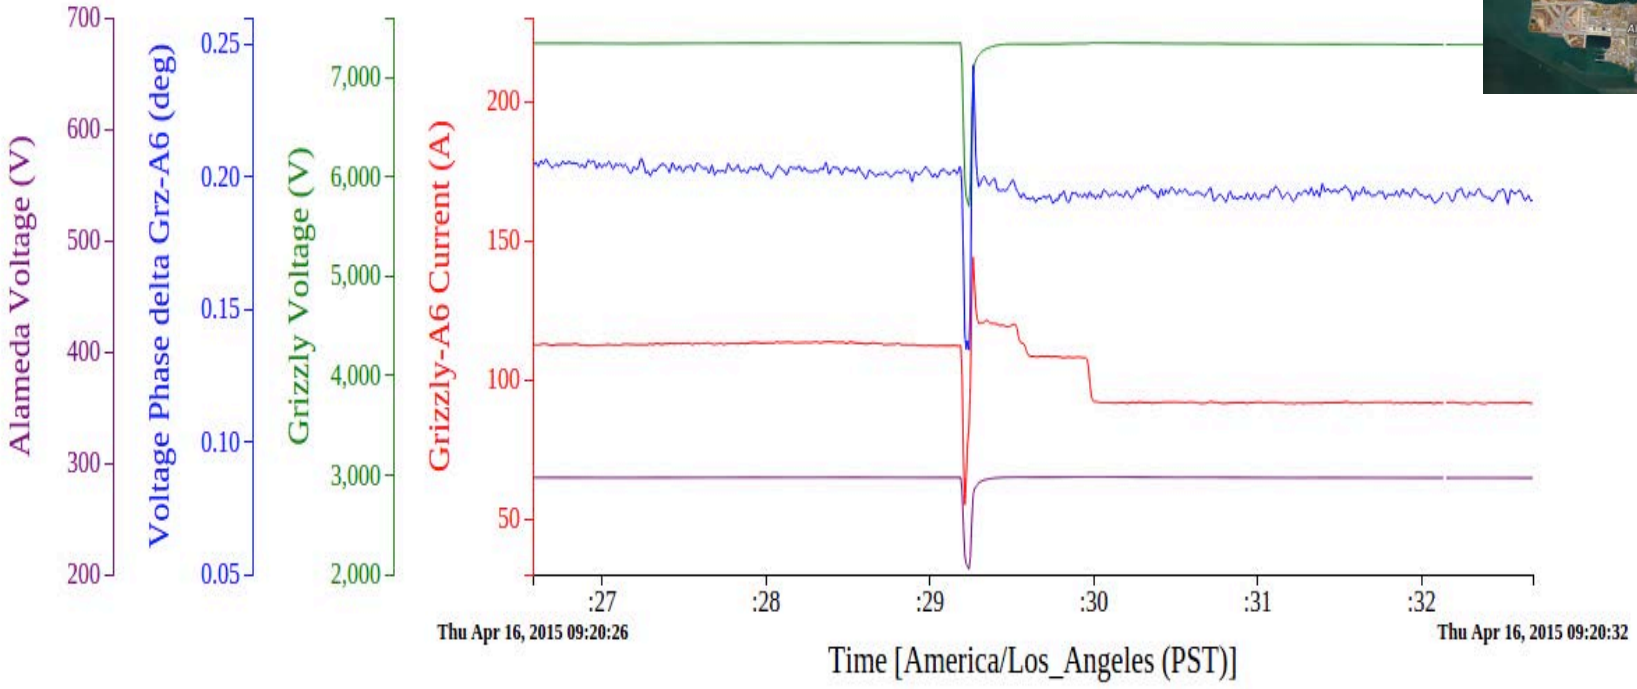

Voltage sag on the transmission system results in a current transient and loss of load. The voltage phase difference between locations on the same primary distribution feeder show the disturbance and typical variations too small to observe with transmission level PMUs

Fast stream telemetry

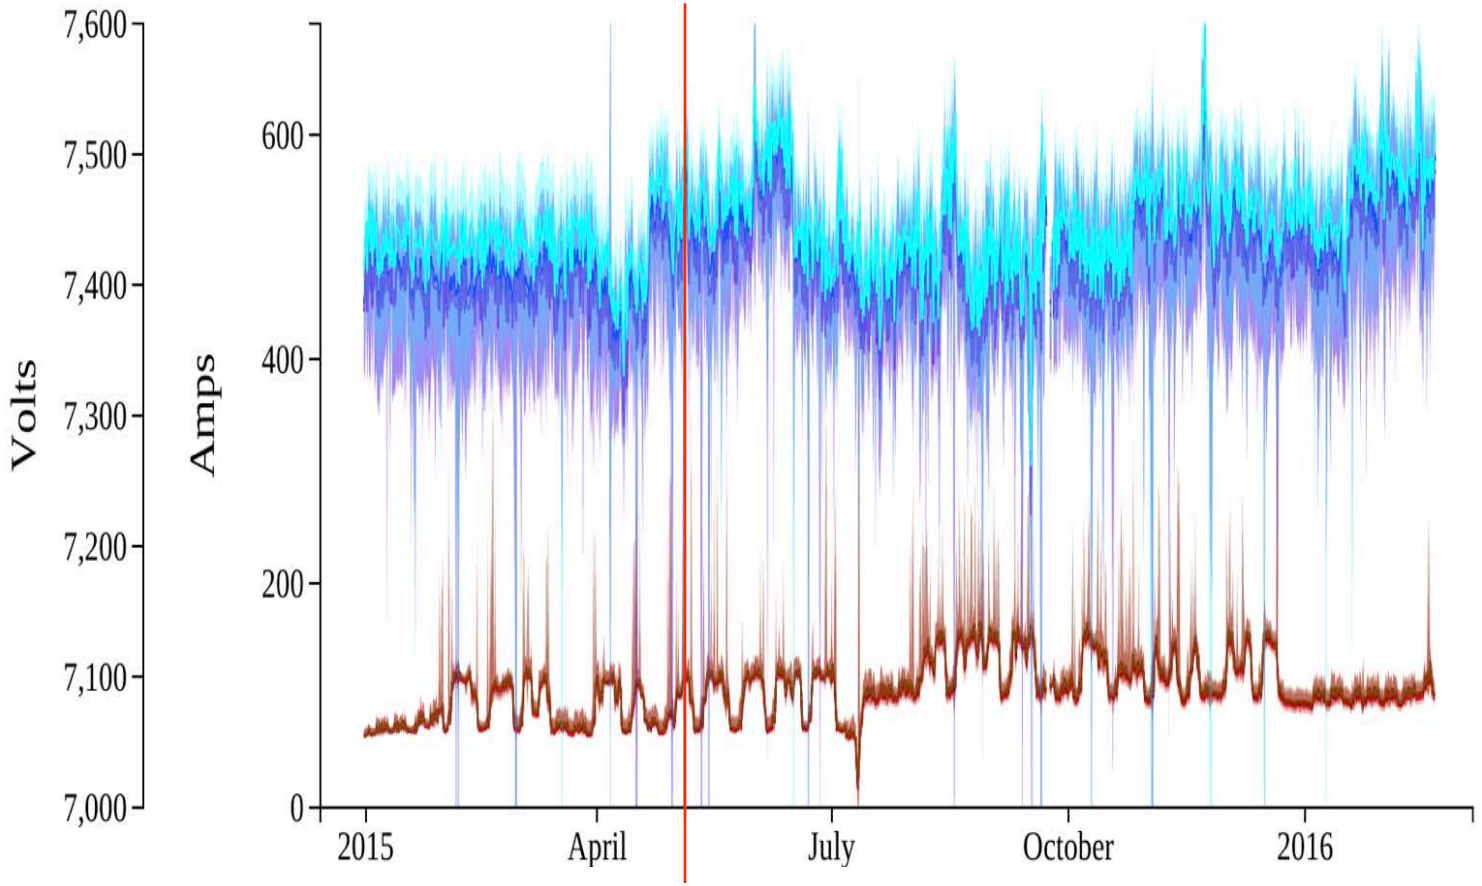

- uPMU-year about 20 billion points
- Each pixel is 4.2 million data points
  - aggregation window 10^6 x sample interval

Multi-Resolution Statistical Store

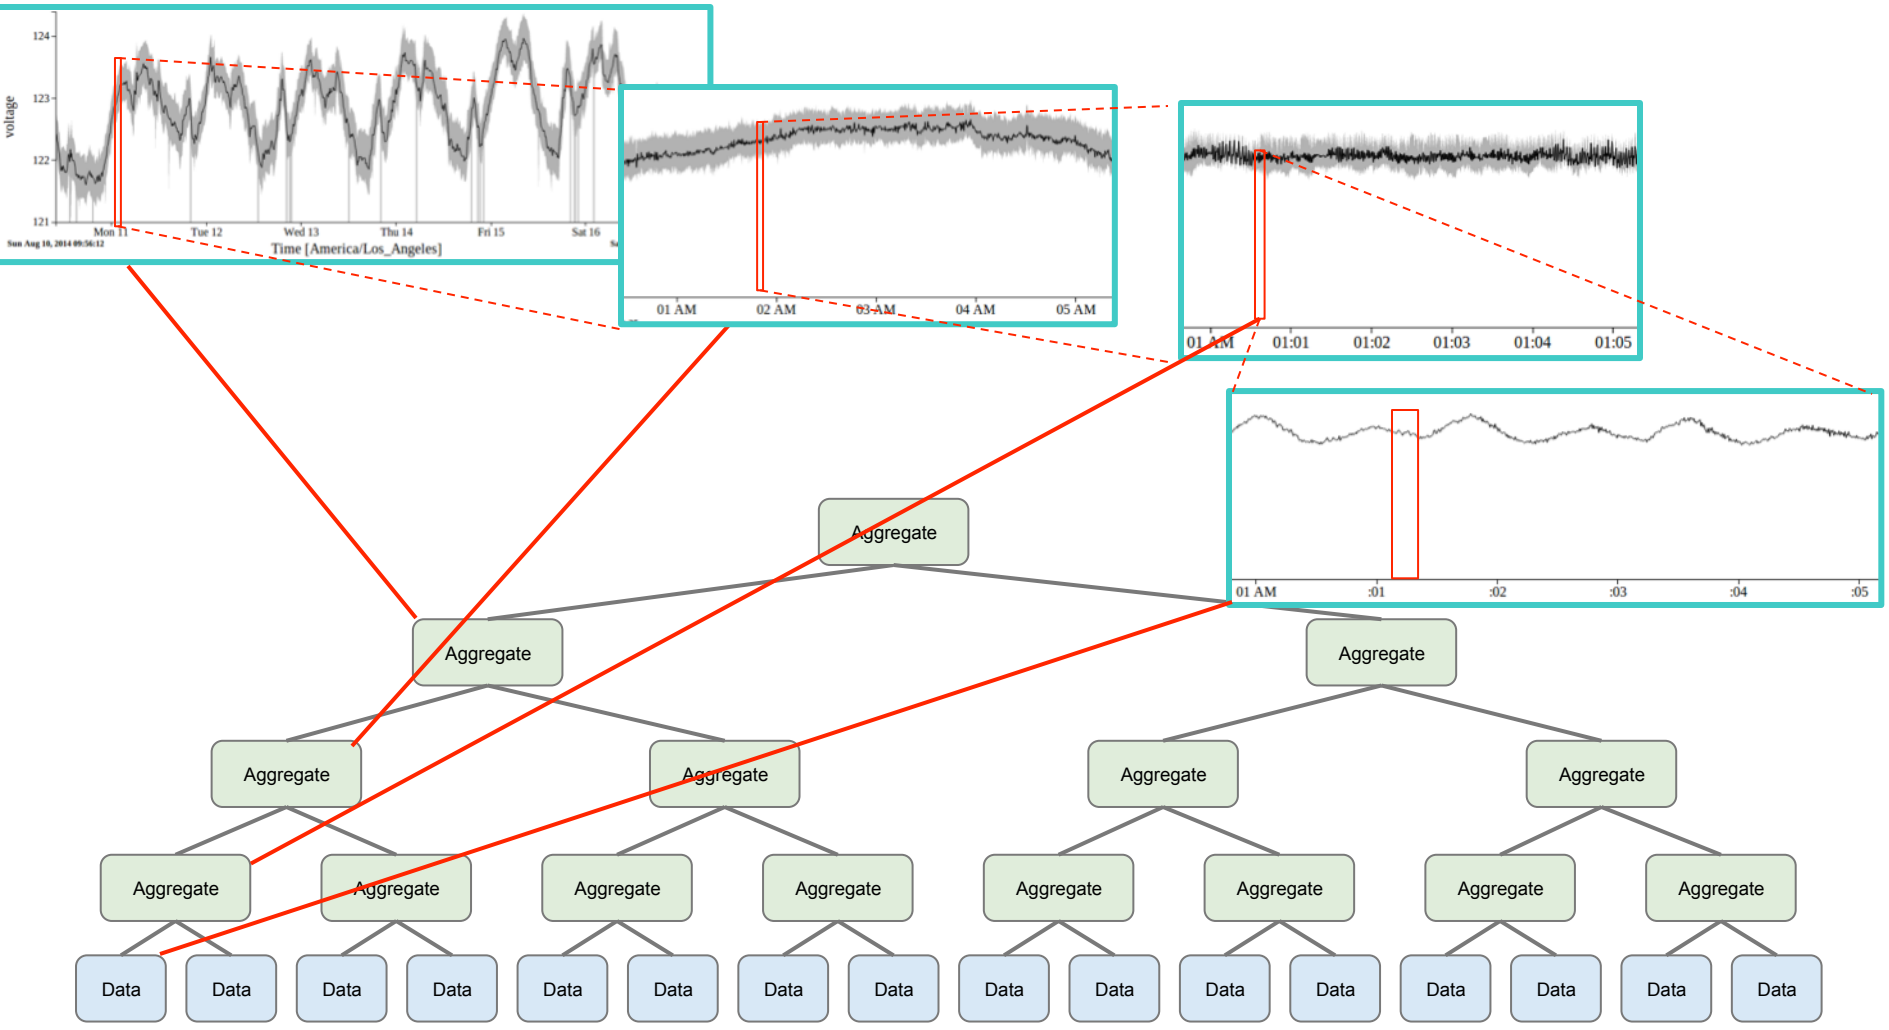

Continuous Distillation Framework

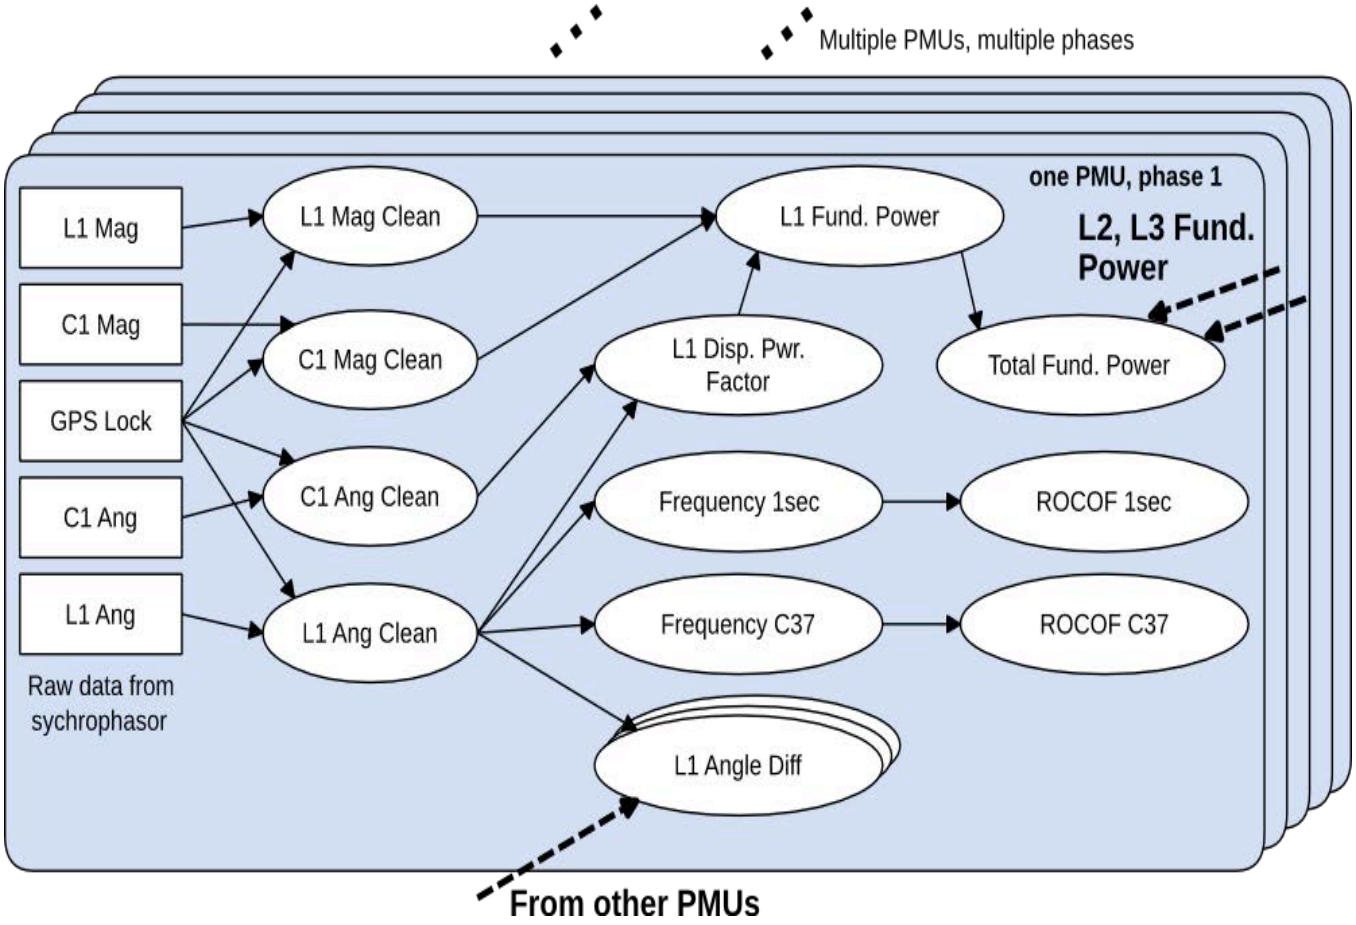

- **Extremely fast change-set determination**
- **Versioned streams**

BTrDB: Partition all time & fill versions

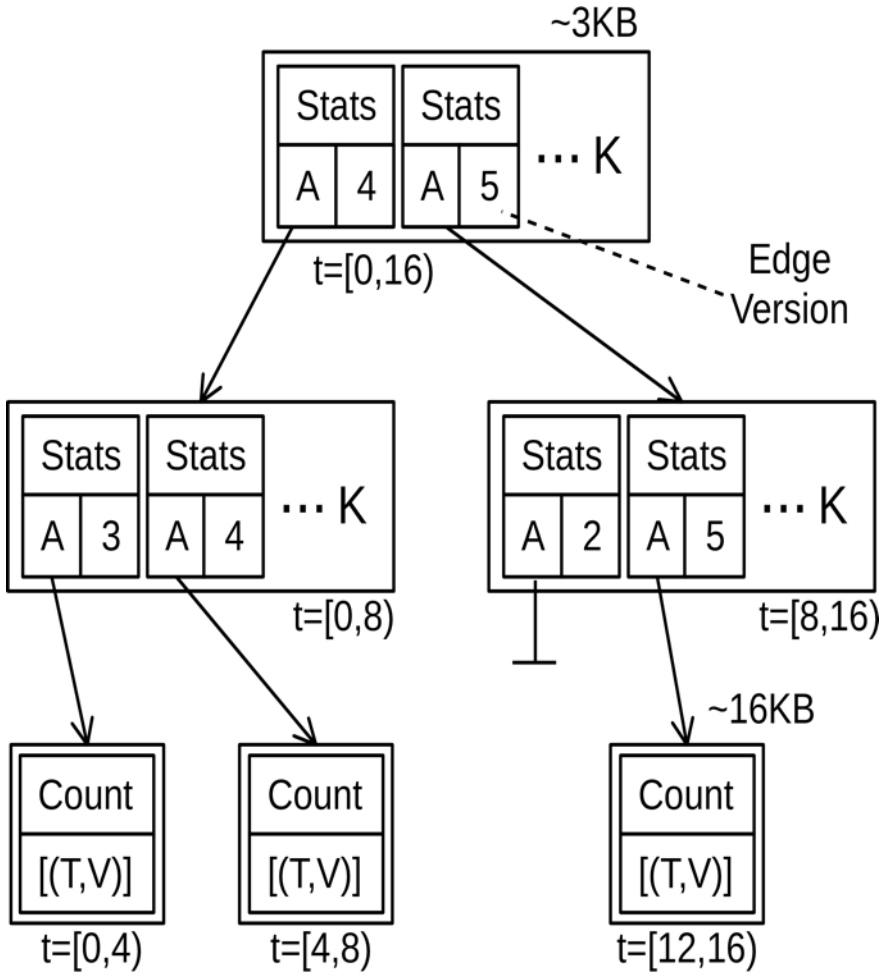

Copy on write K-ary Tree
Partitioning static time (1933 to 2079)

Leaf nodes
- **Time, value pairs + length**

Internal nodes
- **Edges to children**
- **Version annotations for edges**
- **Aggregates for children**
  - **Min, Mean, Max, Count**
  - **Any associative operator**

Tree-Centered Streaming Architecture

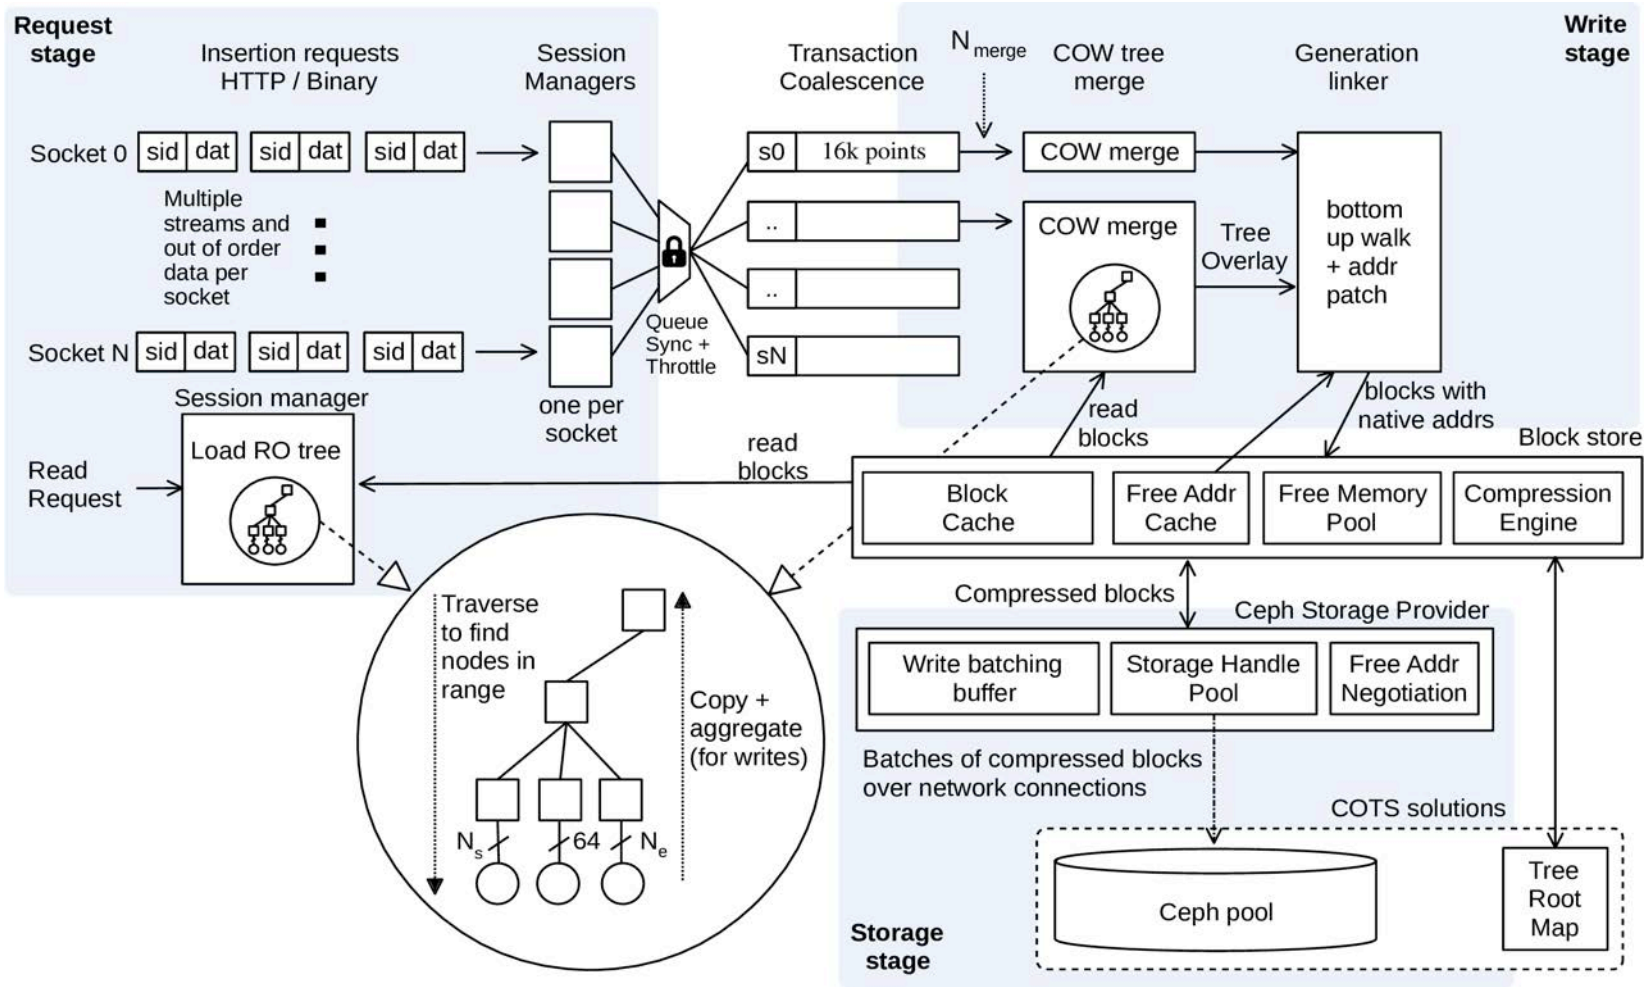

BTrDB Raw throughput

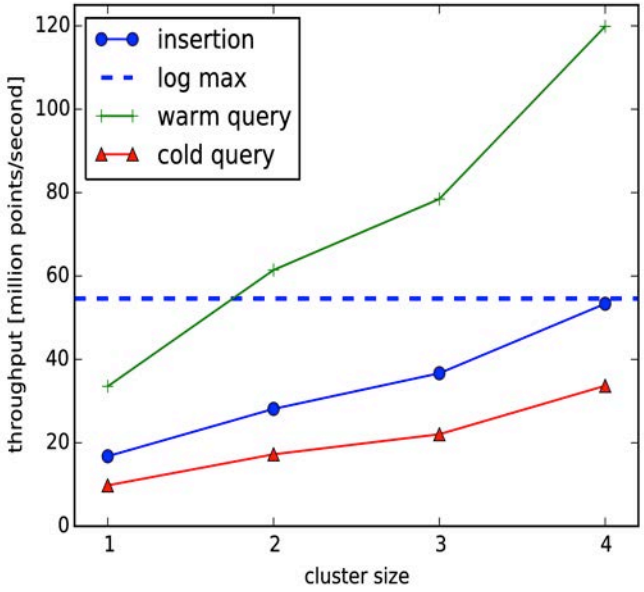

x 1000 = 14 M/s =
x10 analytics

x 1

Roughly $1300
per month on EC2,
excluding redundancy

Scaling studies on commodity Amazon EC2 infrastructure:

- Per node performance: ~1400x faster than Cassandra
- Performance is insensitive to order of data arrival / query

#BTrDB	Streams	Total points	#Conn	Insert [mil/s]	Cold Query [mil/s]	Warm Query [mil/s]
1	50	500 mil	30	16.77	9.79	33.54
2	100	1000 mil	60	28.13	17.23	61.44
3	150	1500 mil	90	36.68	22.05	78.47
4	200	2000 mil	120	53.35	33.67	119.87

	When insertion was	
Throughput [million pt/s] for	Chrono.	Random
Insert	28.12	27.73
Cold query in chrono. order	31.41	31.67
Cold query in same order	-	32.61
Cold query in random order	29.67	28.26
Warm query in chrono. order	114.1	116.2
Warm query in same order	-	119.0
Warm query in random order	113.7	117.2

Statistic queries: visualisation

4.3 M points per pixel column

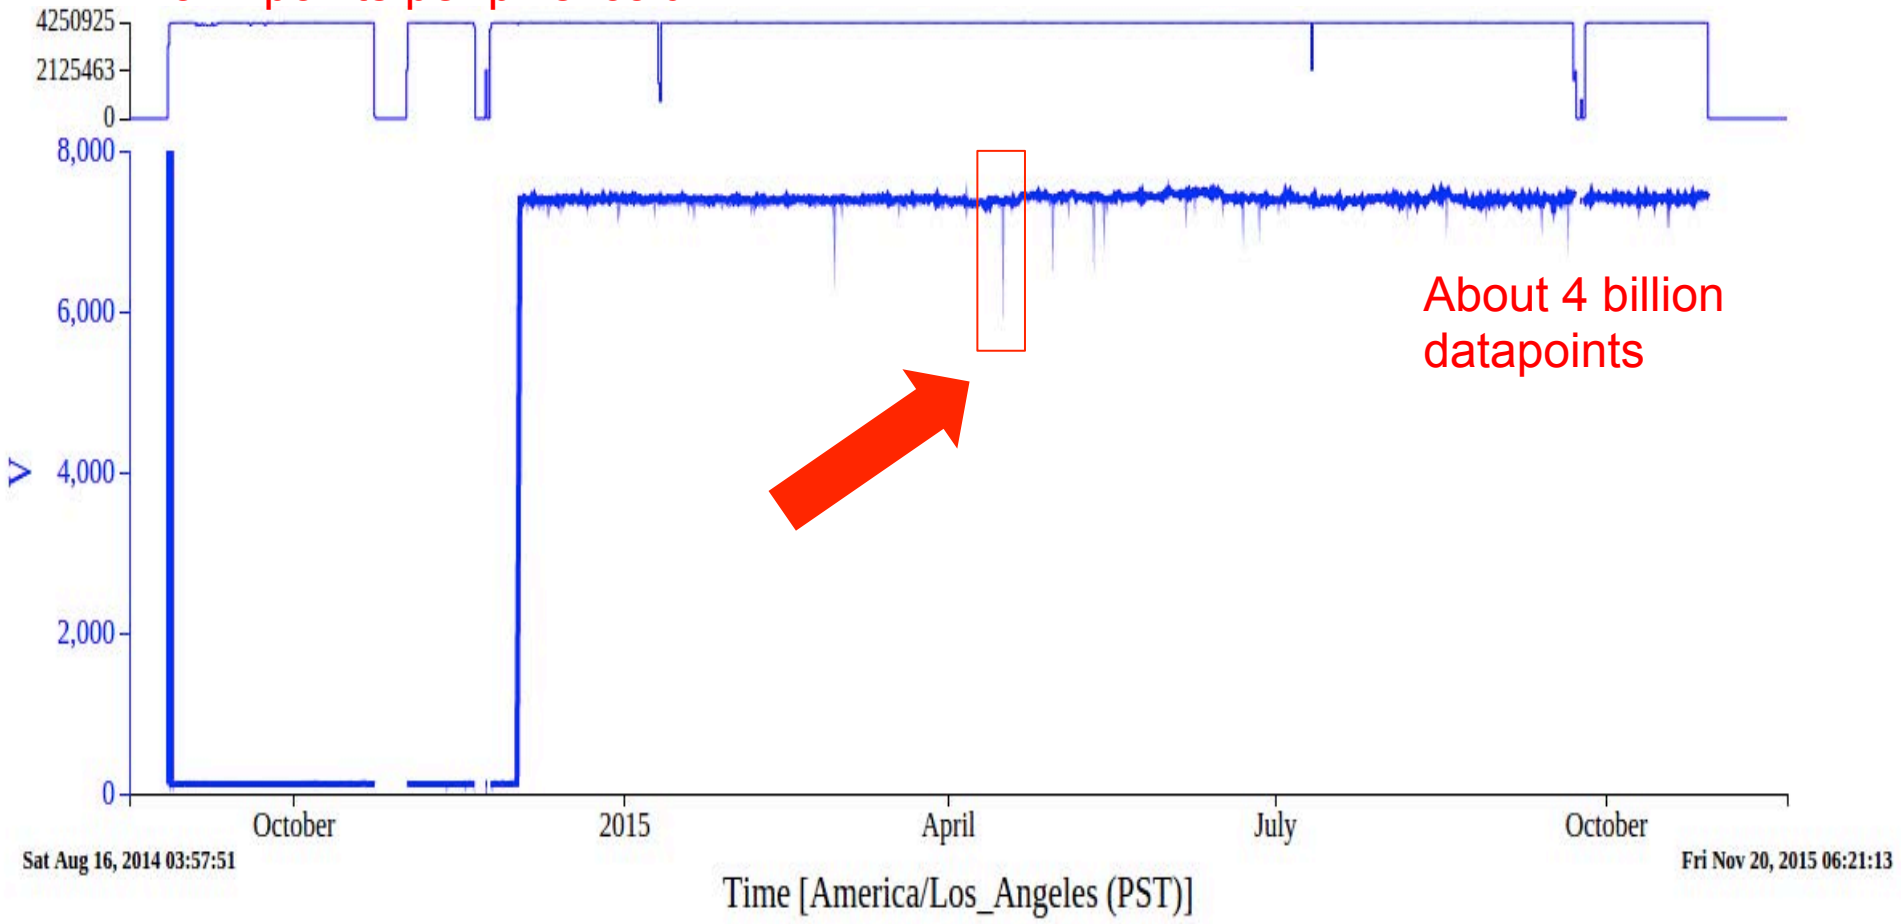

Year

Statistic queries : visualisation

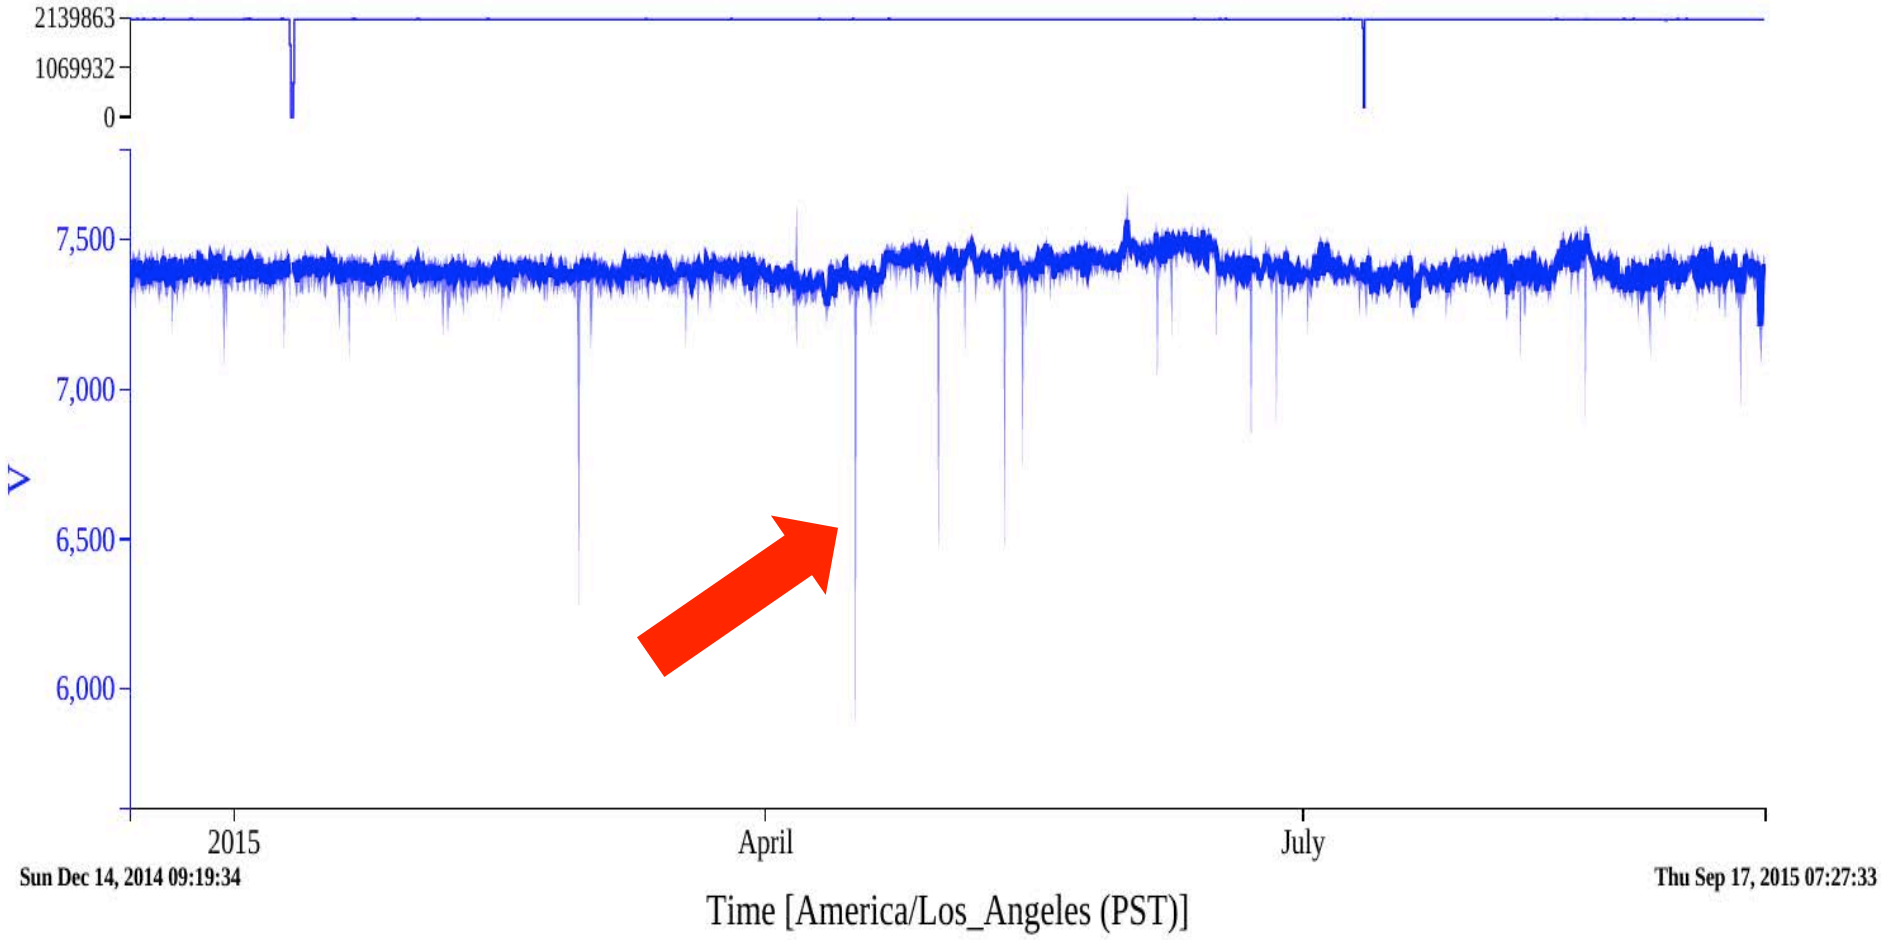

Months

Statistic queries : visualisation

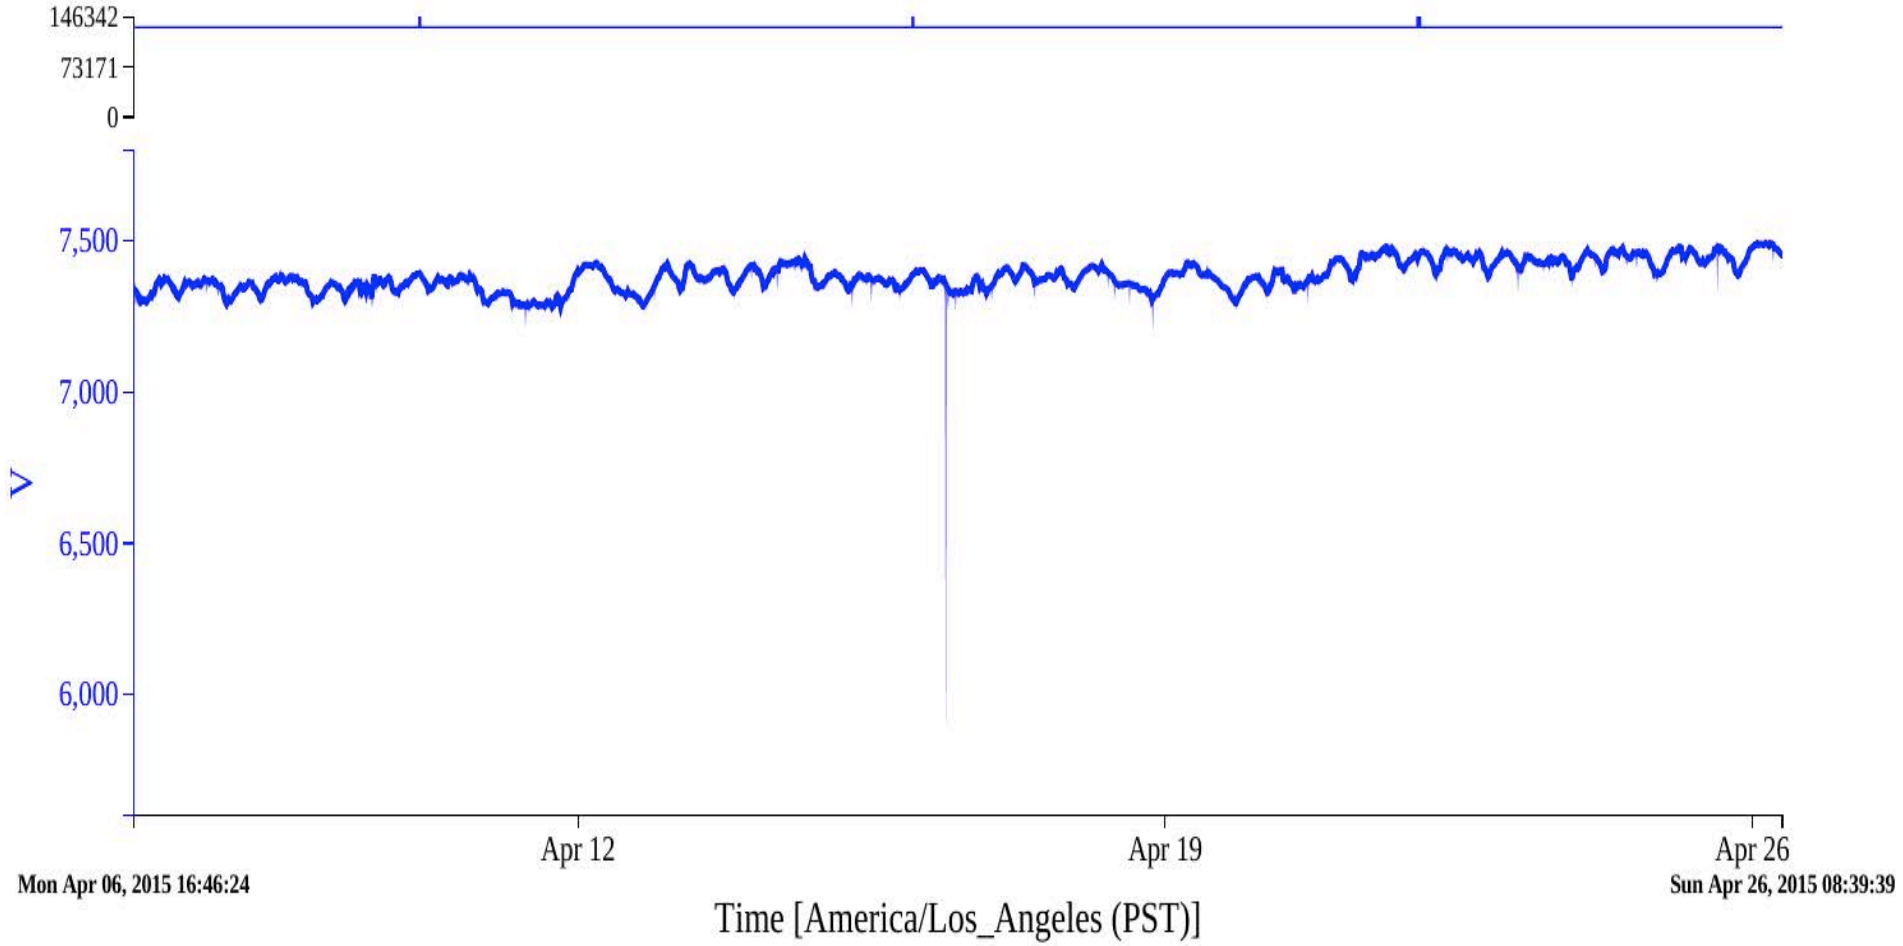

Days

Statistic queries : visualisation

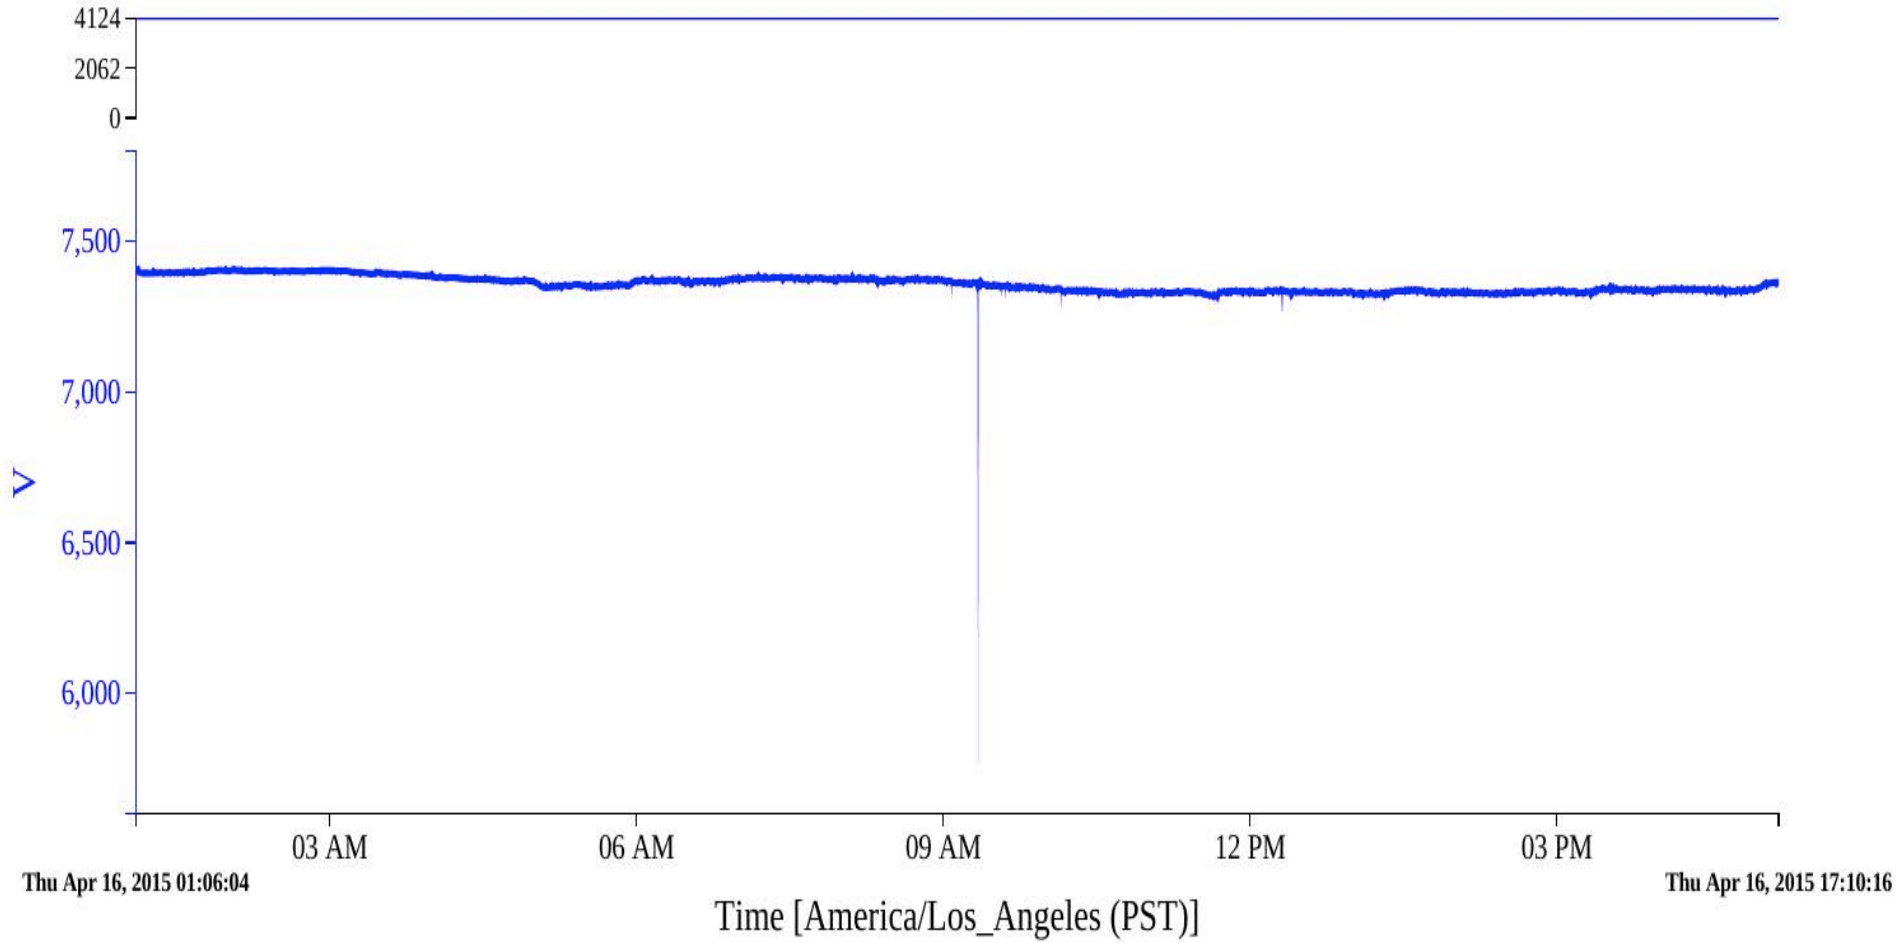

Hours

Statistic queries : visualisation

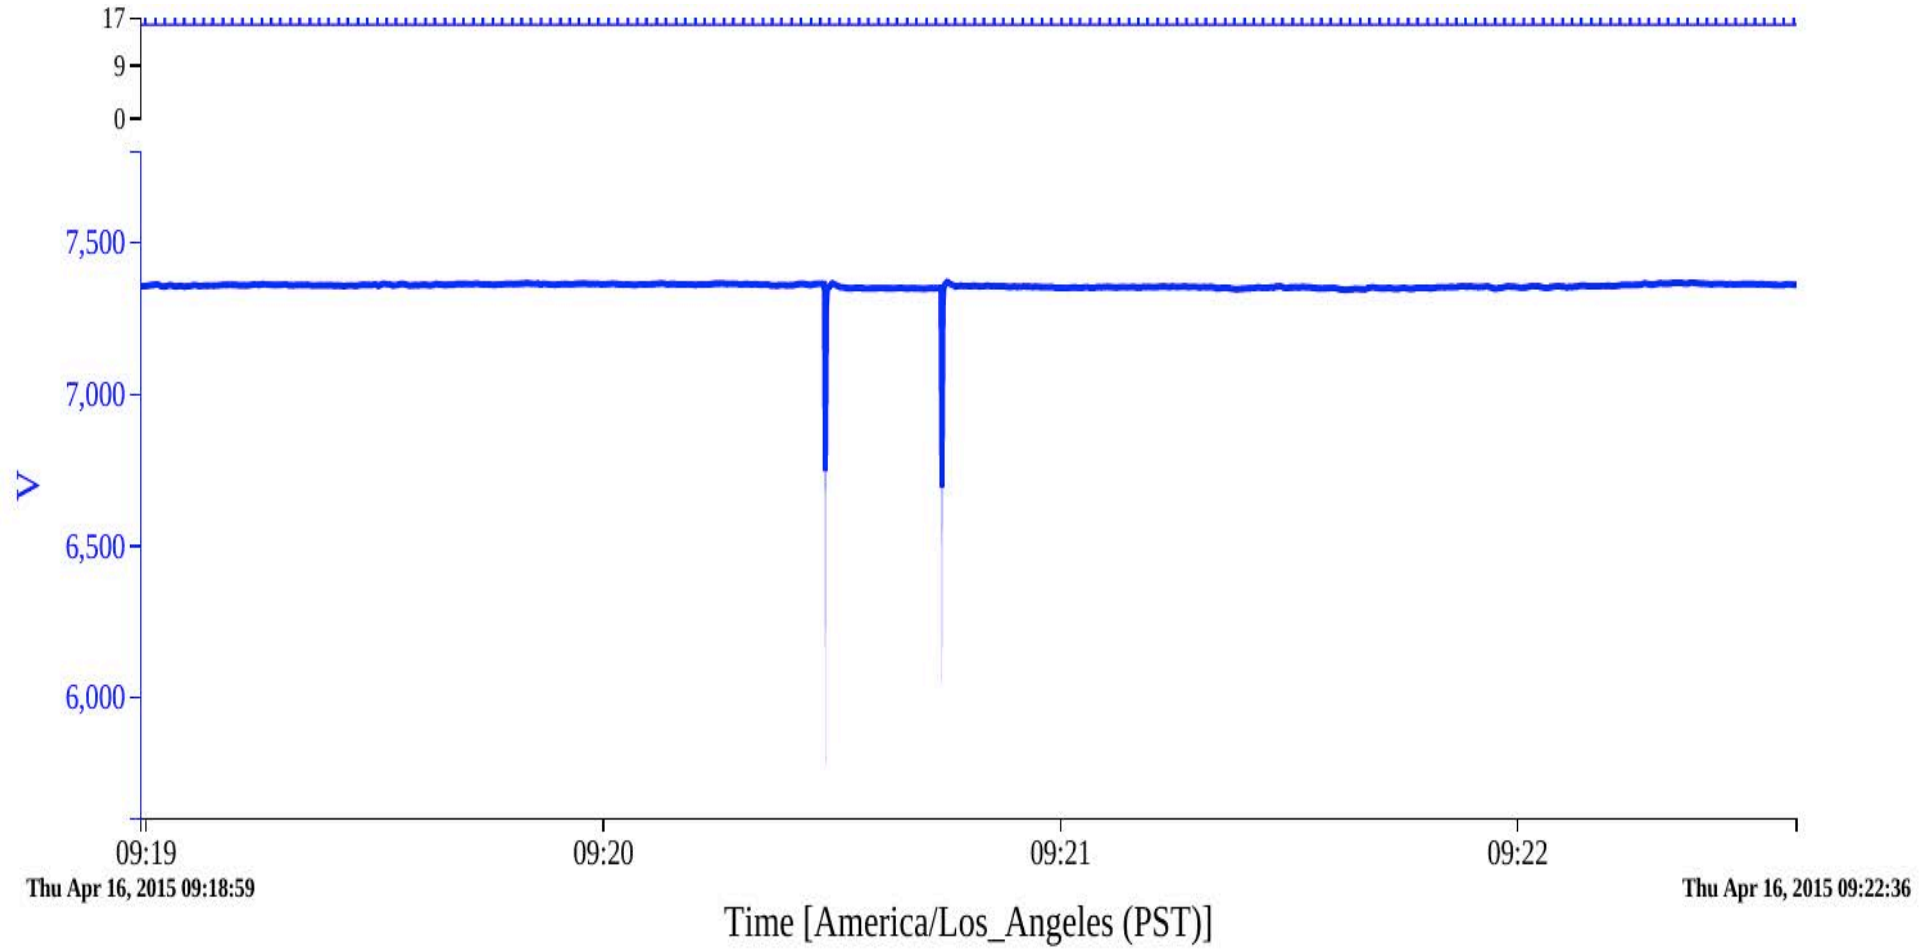

Minutes

Statistic queries : visualisation

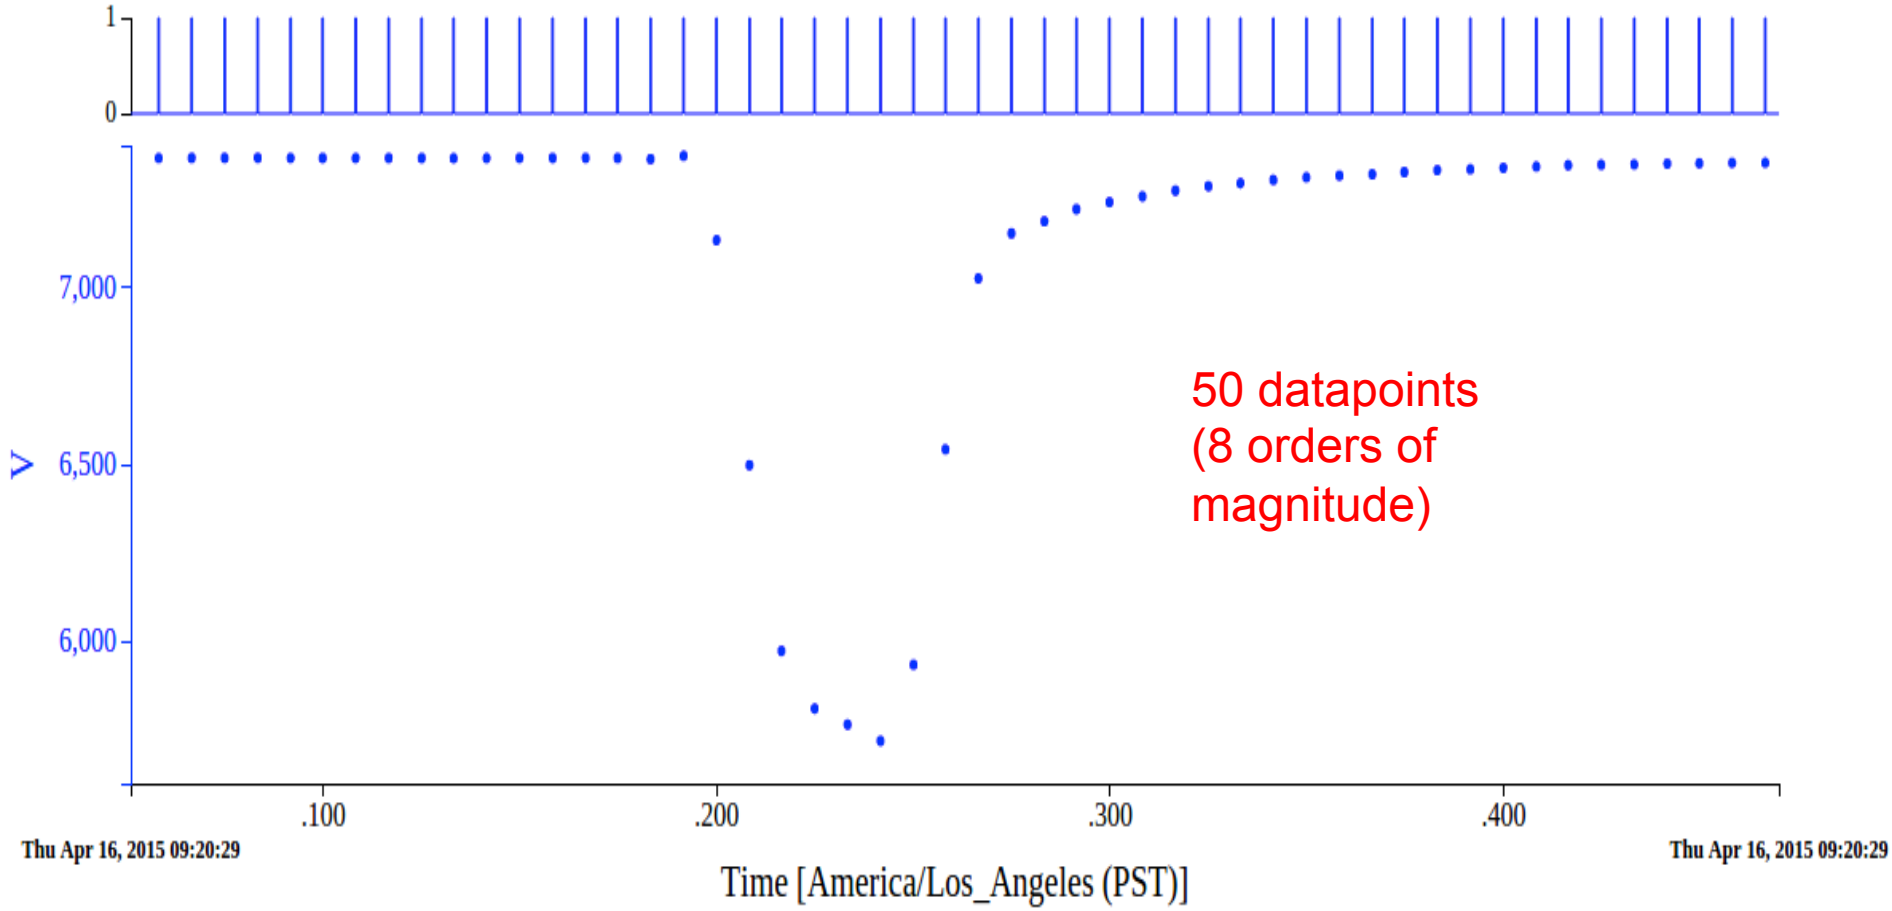

50 datapoints
(8 orders of
magnitude)

Milliseconds

Multi-Resolution Search – V sag

- **Locate and characterize all the voltage sags over a year**

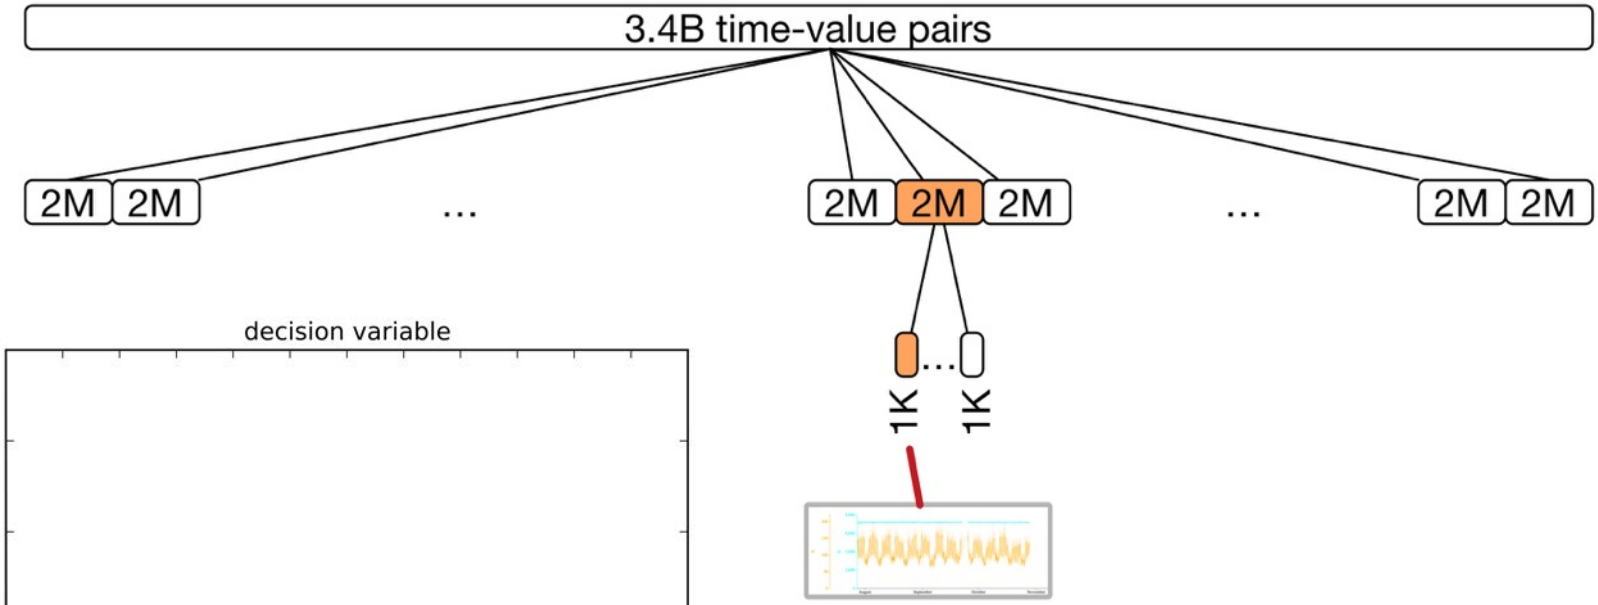

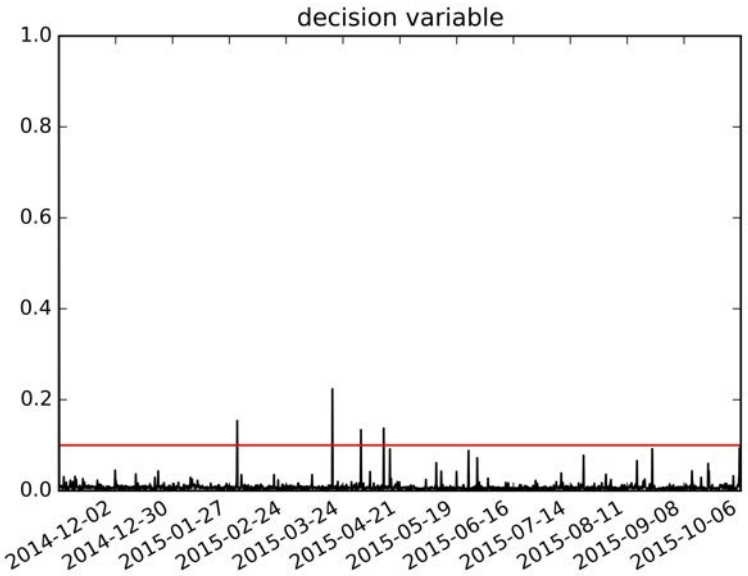

decision variable:
(mean-min)/mean

Multi-Resolution Search – V sag

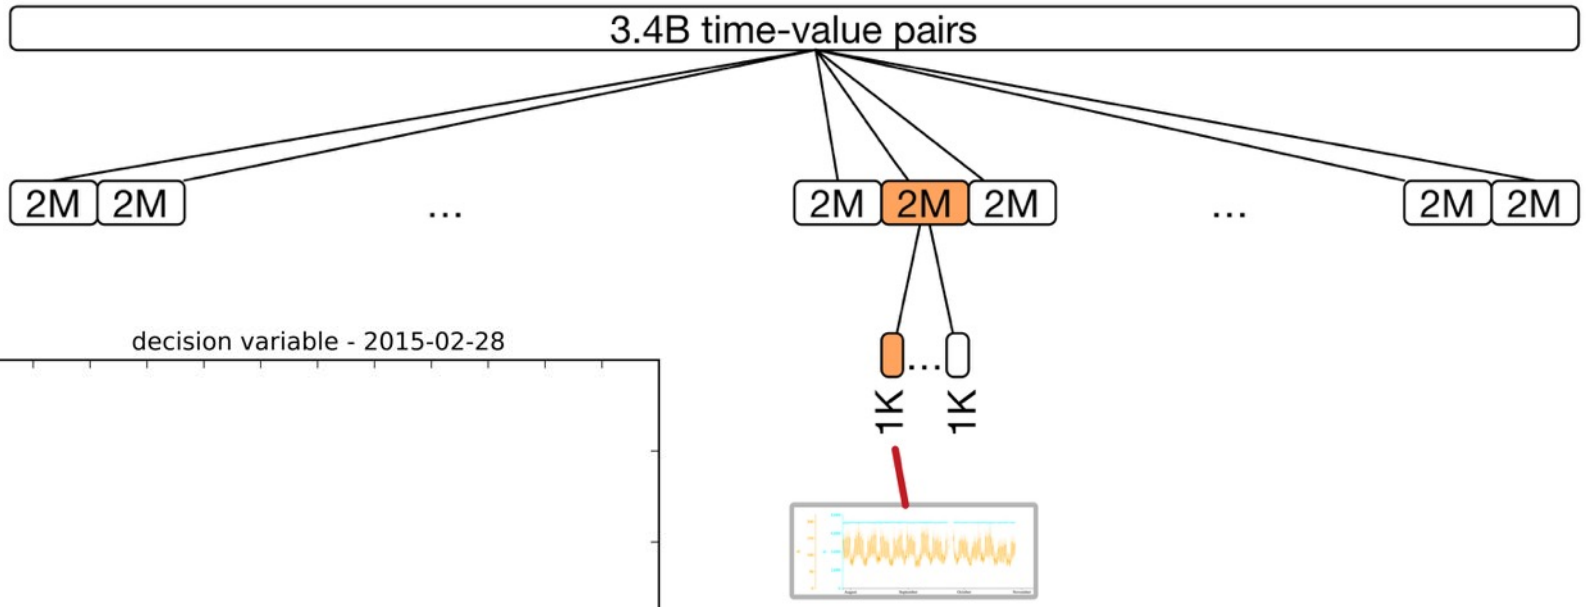

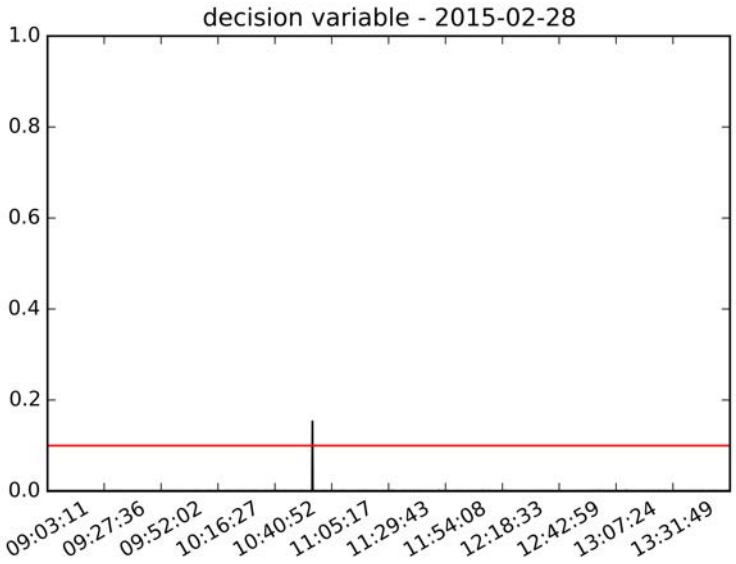

Multi-Resolution Search – V sag

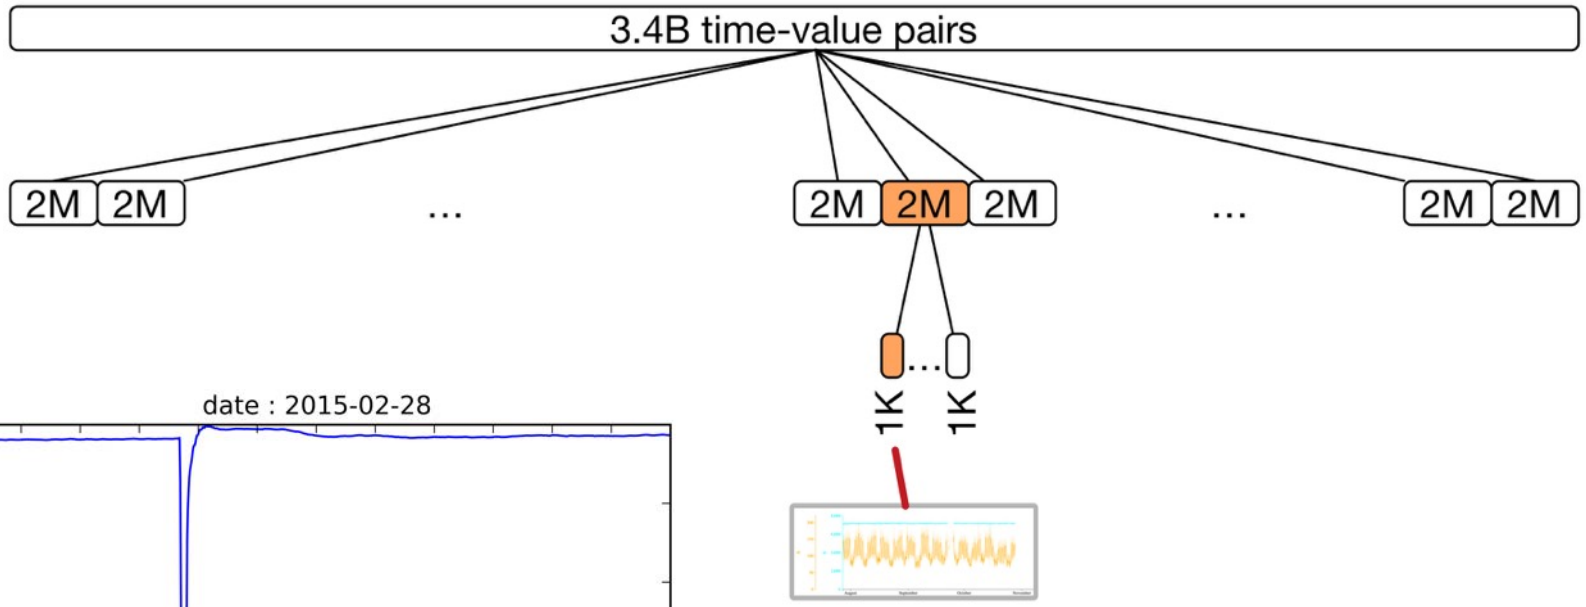

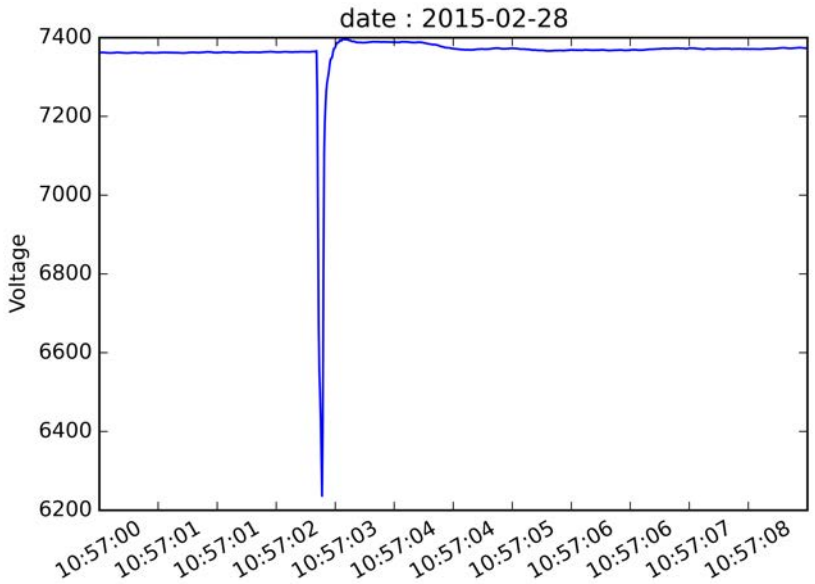

Needles in exponential haystacks

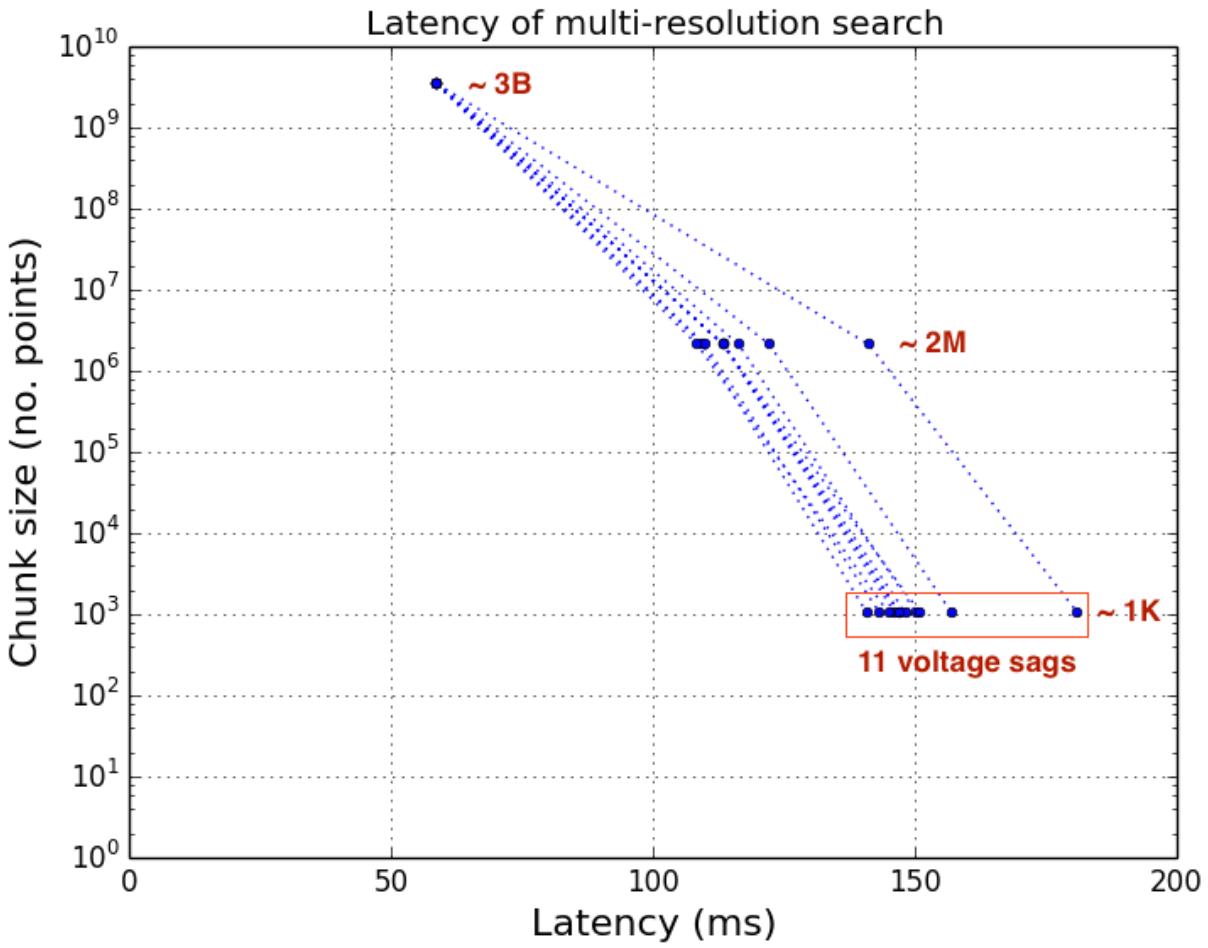

Scaling Potential

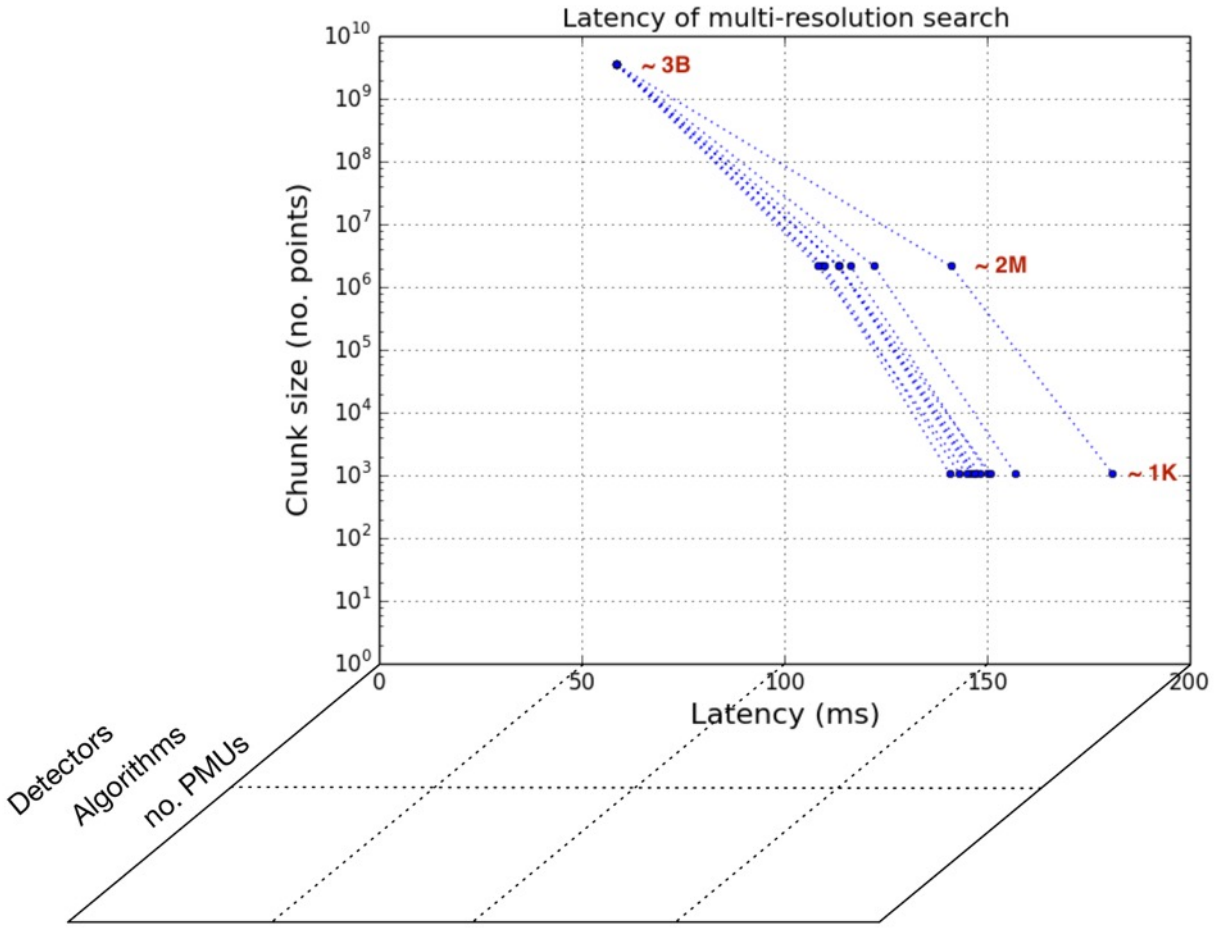

Use case: High-impedance fault detection

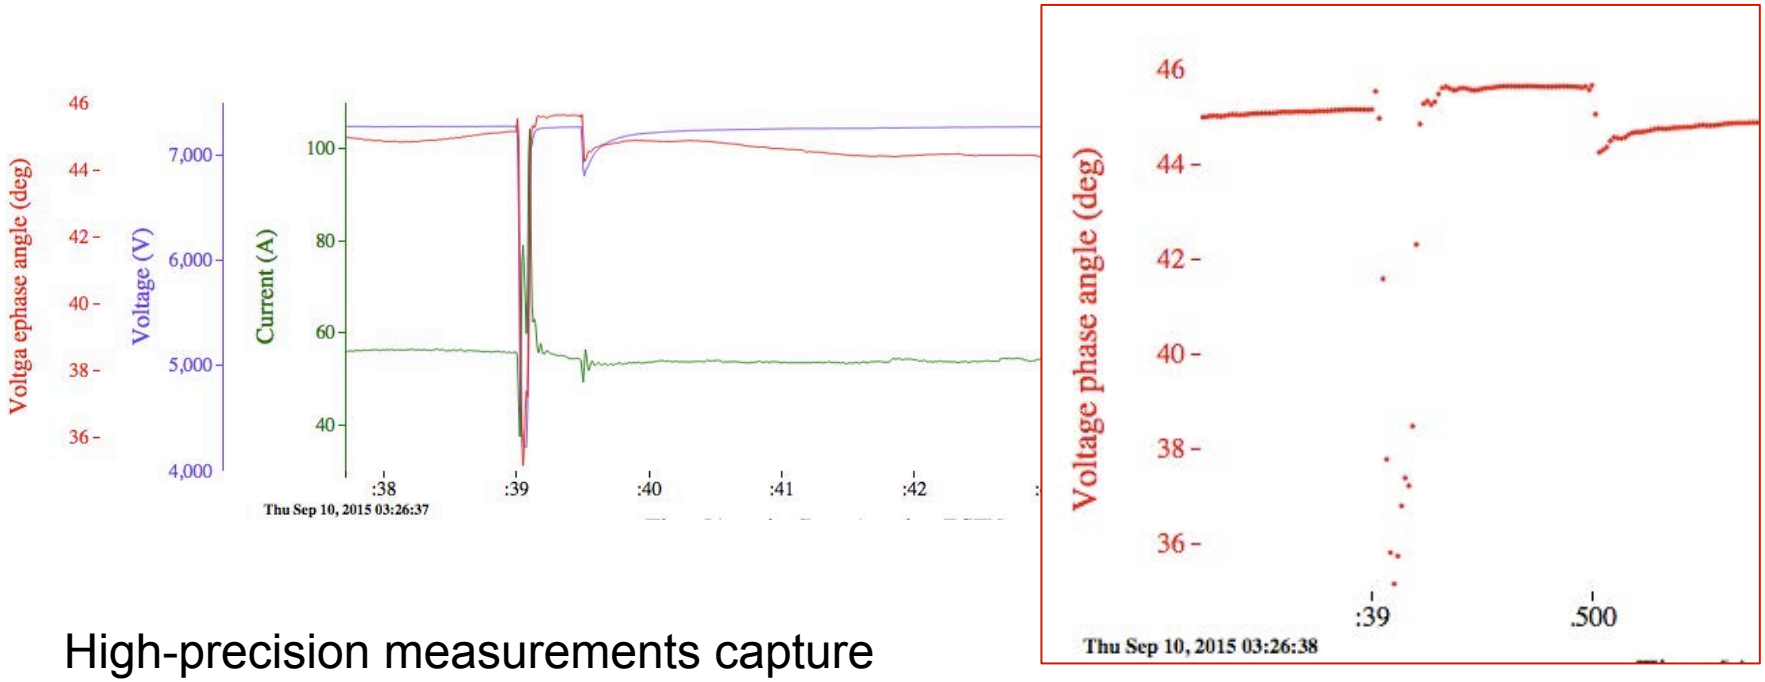

High-precision measurements capture events that do not trip protection, but may impact safety and power quality

Cross-referencing time-aligned data streams supports diagnostics to

- locate disturbance origin
- ascertain proper operation by DG and protection coordination

Use case example: Diagnose cause of PV unit trip

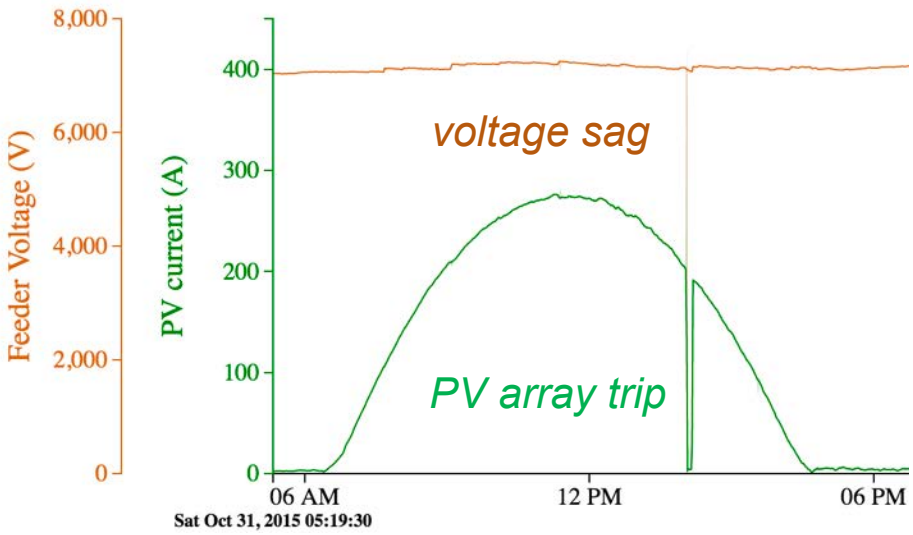

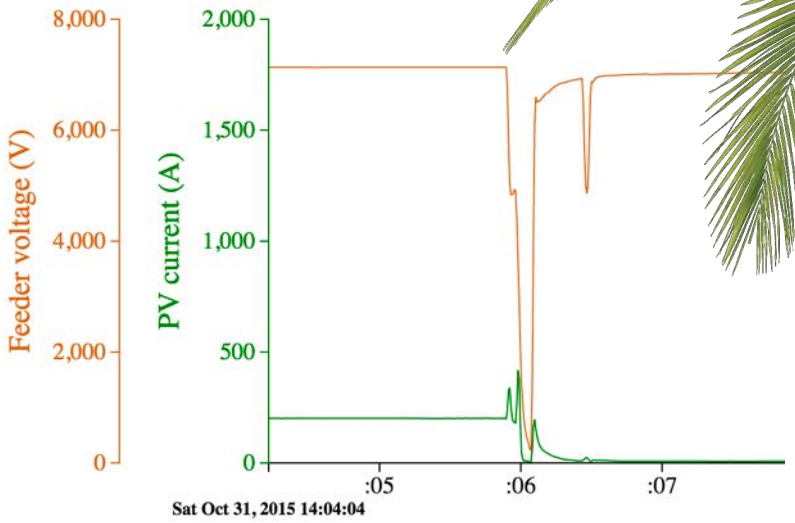

caused by phase B-C fault (palm frond contact) down the feeder

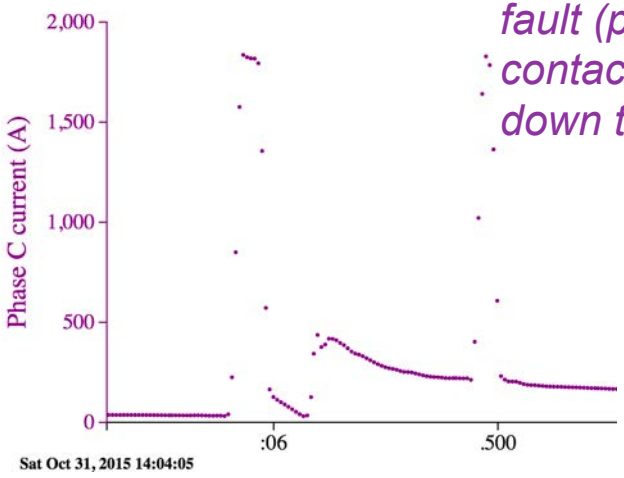

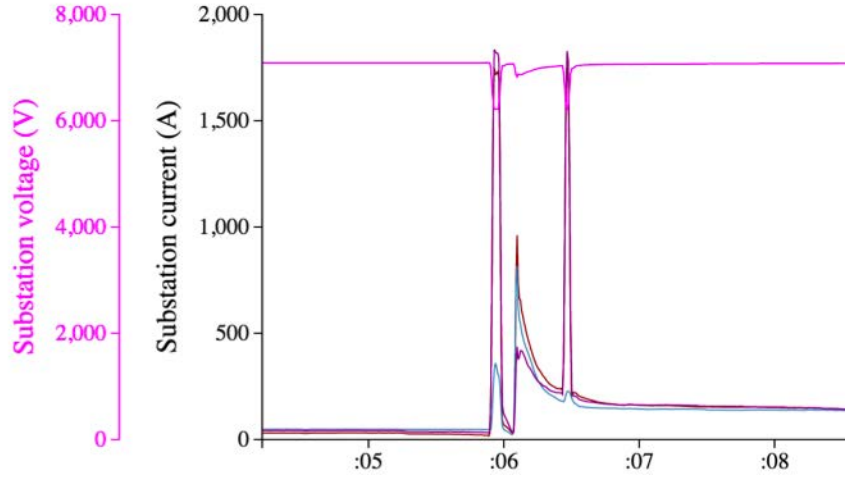

Use case: Detect normal and mis-operation of equipment

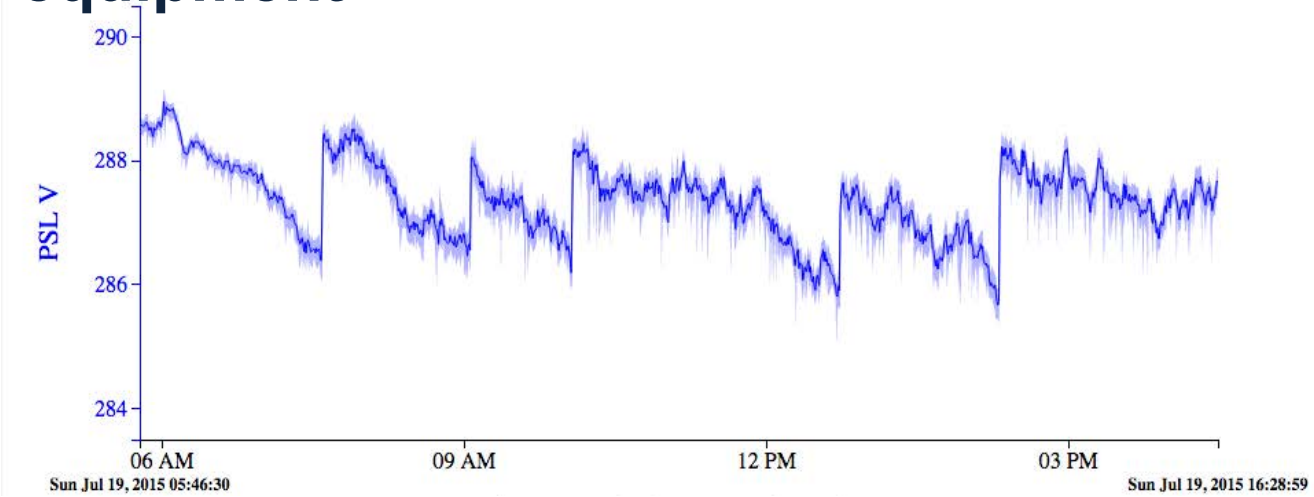

Tap changer at substation transformer steps voltage up and down as load changes over the course of the day

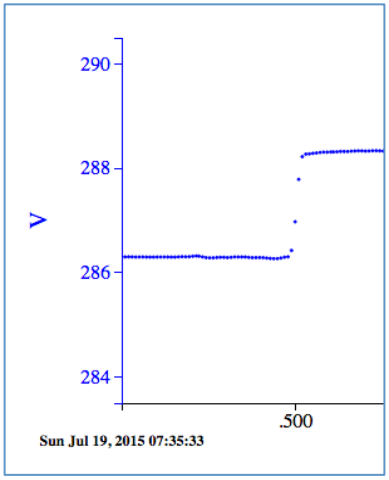

Tap change occurs over ~2 cycles
Graph shows individual 120-Hz samples

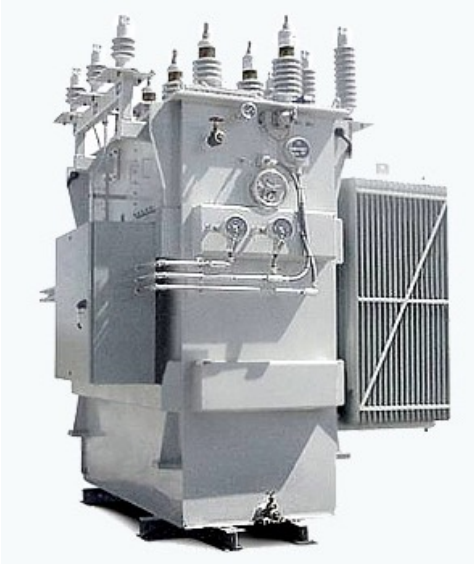

Use case: Detect normal and mis-operation of equipment

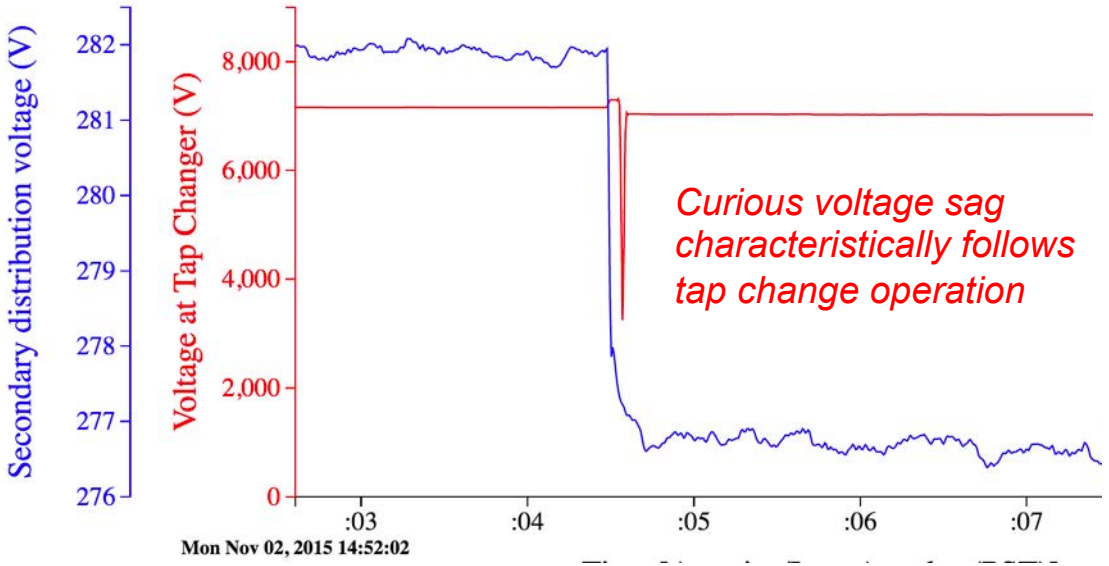

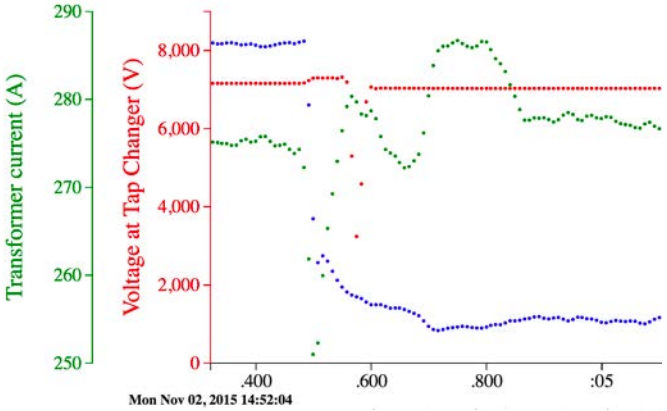

Example:
Anomaly in tap change signature gives early warning of transformer aging or incipient failure

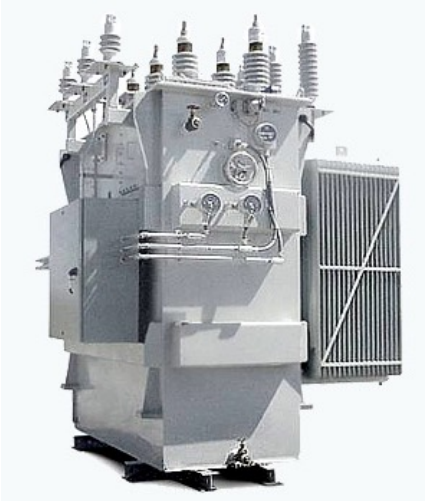

XBOS

```
SUBSCRIBE where Metadata/Sensor/Measure = "Temperature"
           and Metadata/Sensor/Type = "Setpoint"
           and Metadata/Sensor/Setpoint = "Heating"
           and Metadata/HVAC/Zone = "Zone 4";
```

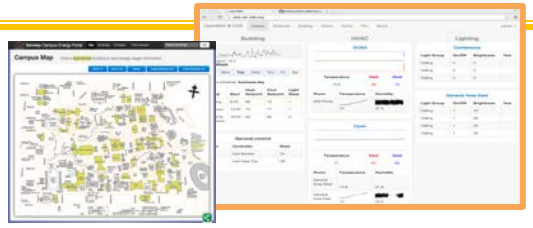

Clients

Control Proc's | Schedulers | Analytics | App Servers

Physical Information Bus
Metadata Query-Based Pub/Sub

JSON/HTTP sMAP Restful web services

JSON/ Websocket

MsgPack

CapnProto

drvr

drvr

Device

Device

Device Sensor Resource

Device Actuator Real-time

APOGEE ANYWHERE | ASHRAE BACnet™ | METER

Plugin	JSON	MsgPack
Message Broker + Kernel API		
Query Processor		

Timeseries Store | Metadata Store | Building Profile

XBOS Kernel

Device Interfaces

- **Set of properties**
  - **Definition**
  - **Data type**
  - **Units**
  - **Required**
- **0+ slots (“write” topics)**
  - **Subset of properties**
- **0+ signals (“read” topics)**
  - **Subset of properties**

Standard XBOS Thermostat Interface

- `info` :
  - `temperature`
  - `relative_humidity`
  - `heating_setpoint`
  - `cooling_setpoint`
  - `override`
  - `fan`
  - `mode`
  - `state`
  - `time`

“info” signal definition

```
type signal struct {
    Temperature       float64 `msgpack:"temperature"`
    Relative_humidity float64 `msgpack:"relative_humidity"`
    Heating_setpoint  float64 `msgpack:"heating_setpoint"`
    Cooling_setpoint  float64 `msgpack:"cooling_setpoint"`
    Override          bool    `msgpack:"override"`
    Fan               bool    `msgpack:"fan"`
    Mode              int64   `msgpack:"mode"`
    State             int64   `msgpack:"state"`
    Time              int64   `msgpack:"time"`
}
```

autogenerated Go code (msgpack on the wire)

Model, model, model …

Model Predictive Control / Learning

- Average energy consumption reduction of 60-85% over DDC mode levels.

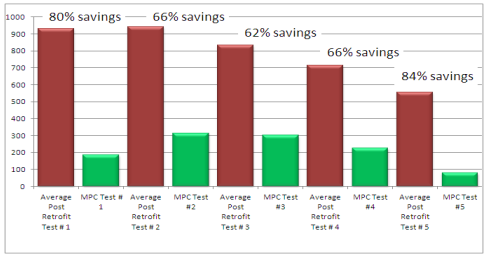

Source: "Model Predictive Control for Mid-Size Commercial Building
done by Dr. Borrelli group in conjunction with UTC Research Center
Published February 2012. US Army Corp of Engineers, Champaign, I

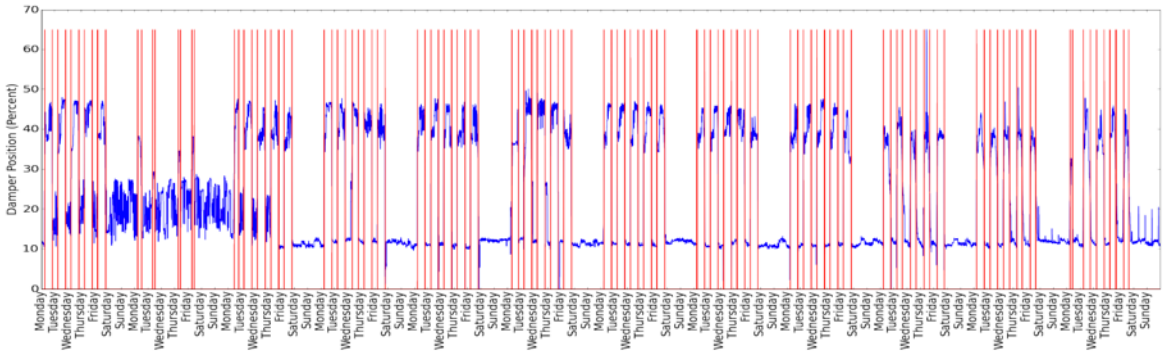

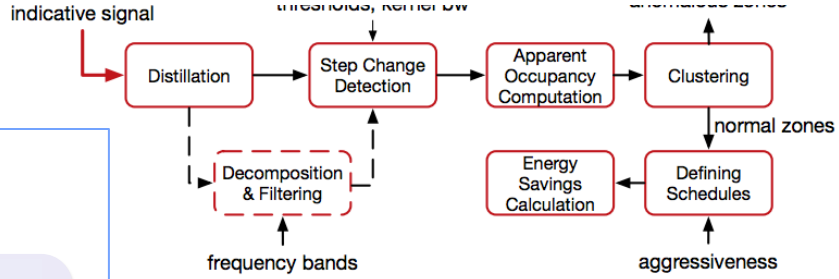

Optimization Formulation

- Predicted states x_k, inputs u_k, disturbances w_k
- At each time step, solve:

$$\min_{\mathbf{U},\mathbf{X}} \sum_{k=0}^{N-1} J(x_{k+1}, x_k, u_k, w_k) \qquad \text{cost function}$$

$$\text{subj. to, } \forall\, k \in \{0,\ldots,N-1\},$$

$$f(x_{k+1}, x_k, u_k, w_k) = 0 \qquad \text{system dynamics}$$

$$g(x_{k+1}, x_k, u_k, w_k) \le 0 \qquad \text{constraints}$$

$$x_0 = x(t) \qquad \text{initial conditions}$$

- Constrained finite time optimal control problem
- Optimization much faster if explicit structure of *J*, *f*, *g* (and derivatives) can be provided

Tony Kelman | MPC in Buildings | January 11th – Slide 5

Basic Idea

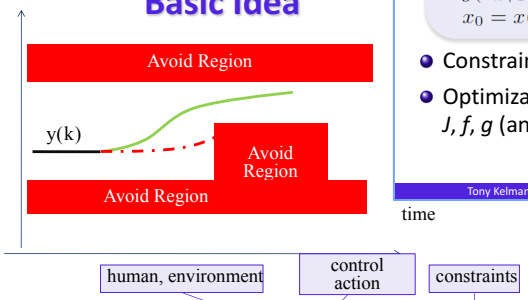

$$y(k+1) = y(k) + b_w w(k) + b_u u(k); \quad y(k)\ 2\ Y(k)$$

At step t decide on ***u(t)*** based on prediction on ***w(t),…, w(t+N), Y(t),…,Y(t+N)***

Two Combined Effects : Anticipation and Coordination

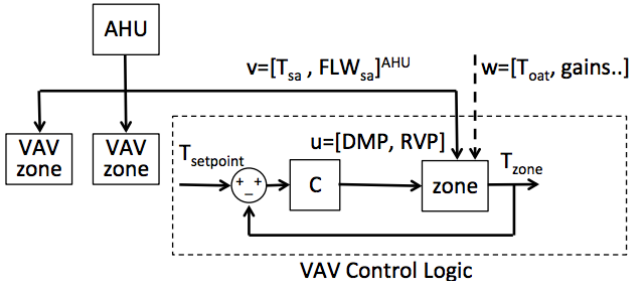

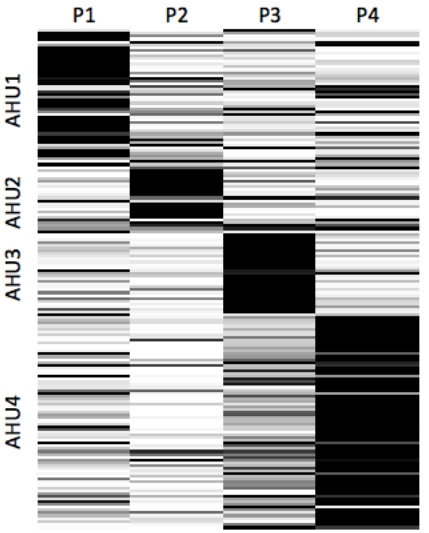

Rogue Zone Analysis

Building Id	Year of Construction	BMS Vendor	Num. of Sense Points	Num. of Thermal Zones	Num. of Hot Rogue Zones	Num. of over-cooled zones
Soda	1994	1	1586	201	5	17
SDH	2009	2	2522	78	2	0
McCone	1961	1	367	42	28	1
Law	1968	1	132	12	1	0
Minor	1941	1	417	48	8	4
Stanley	2007	1	6169	368	35	5
Carleton	NA	1	164	8	3	0
Cory	1950	1	421	20	0	2
DOE	1982	1	277	9	2	0
Tan	1996	1	730	57	10	0
Total			12813	843	94	29

Over-cooled zones, Stuck Dampers, Night time setbacks

Soda Hall, UC Berkeley

Rogue Zones

S.No	Room	% Time (temp > stpt)	avg (Temp - stpt)	AirHandler id	Avg stpt	Avg room temp	Avg Airflow	NightTime Setback ?	Stuck Damper?
1	330B	100.0	10.6	1	72.0	82.6	-1.0	False	NA
2	333	100.0	5.5	4	72.0	77.5	22.1	False	No
3	288	100.0	5.0	1	70.0	75.0	20.5	False	No
4	627	98.7	3.8	3	70.0	73.8	18.1	False	No
5	342	97.7	2.7	1	71.0	73.7	13.7	False	No

Over-Cooled Zones

S.No	Room	% Time (temp > stpt)	avg (Temp - stpt)	AirHandler id	Avg stpt	Avg room temp	Avg Airflow	NightTime SetBack	Stuck Damper?
1	340	0.0	-2.7	1	75.0	72.2	7.9	False	Yes
2	544	0.0	-3.4	1	75.0	71.5	6.0	False	No
3	180_	0.0	-3.5	1	75.0	71.4	5.9	False	No
4	300T	0.0	-3.5	4	75.0	71.4	0.1	False	No
5	420A	0.0	-3.6	1	75.0	71.3	7.9	False	Yes
6	678	0.0	-4.0	1	72.0	67.9	5.4	False	No
7	444	0.0	-4.1	1	72.0	67.8	5.2	False	No
8	384	0.0	-4.3	1	74.0	69.6	5.0	False	No
9	530	0.0	-4.5	1	75.0	70.4	4.7	False	No

Multi-Res. Search on "Load" data

- Find most power hungry of 72 buildings
  - Monthly average power, 3-phase meters
- Extract average weekday/weekend daily profile
- Isolate large spikes in power factor

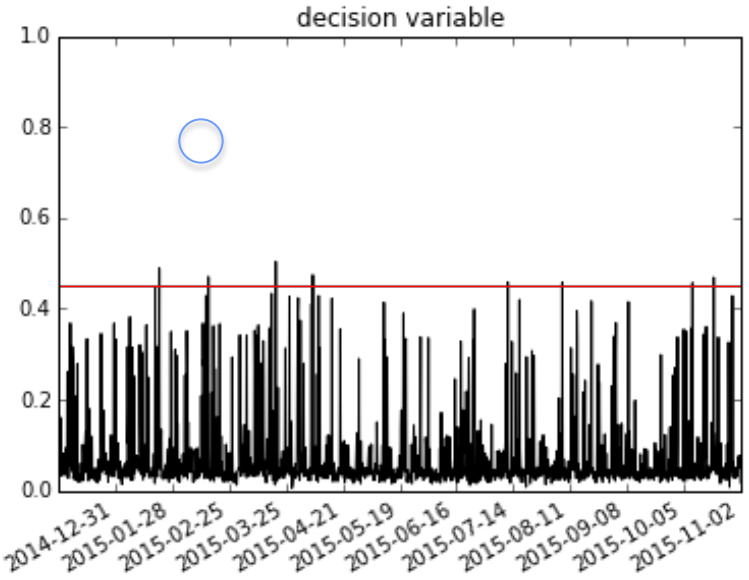

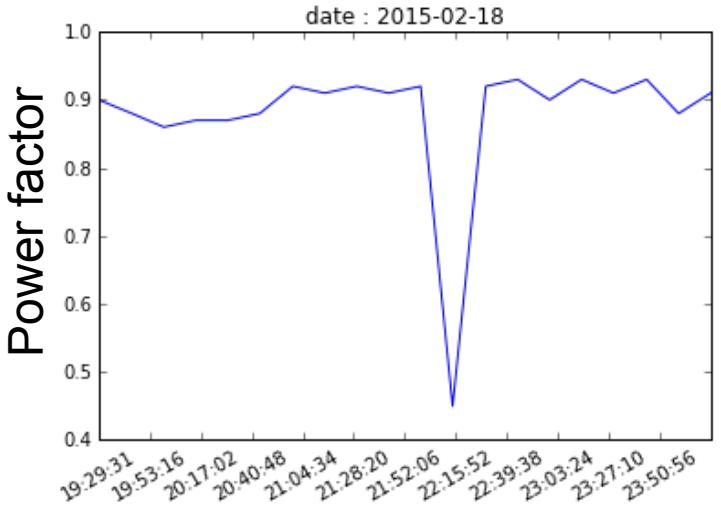

Far beyond the building …

- **Delegation of Authorization**
- **Federation**
- **Protection**
- **Auditability / Privacy**

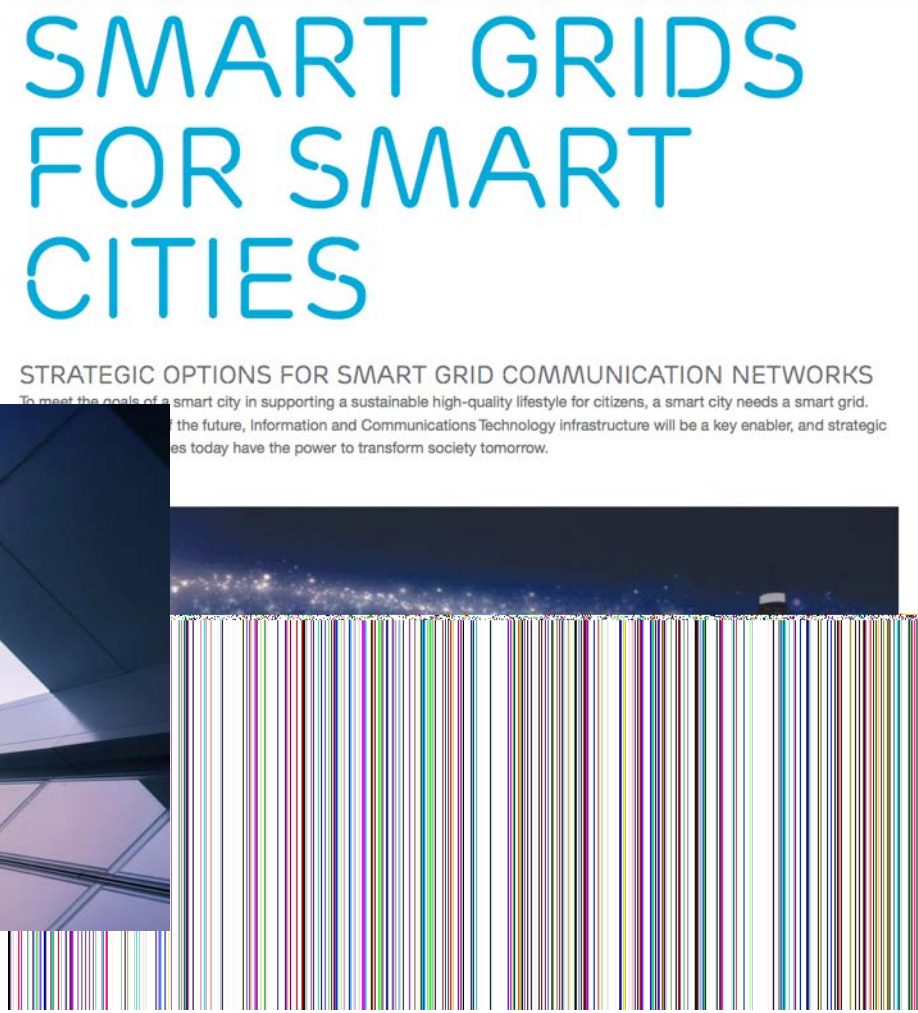

A Broader Foundation … WAVE

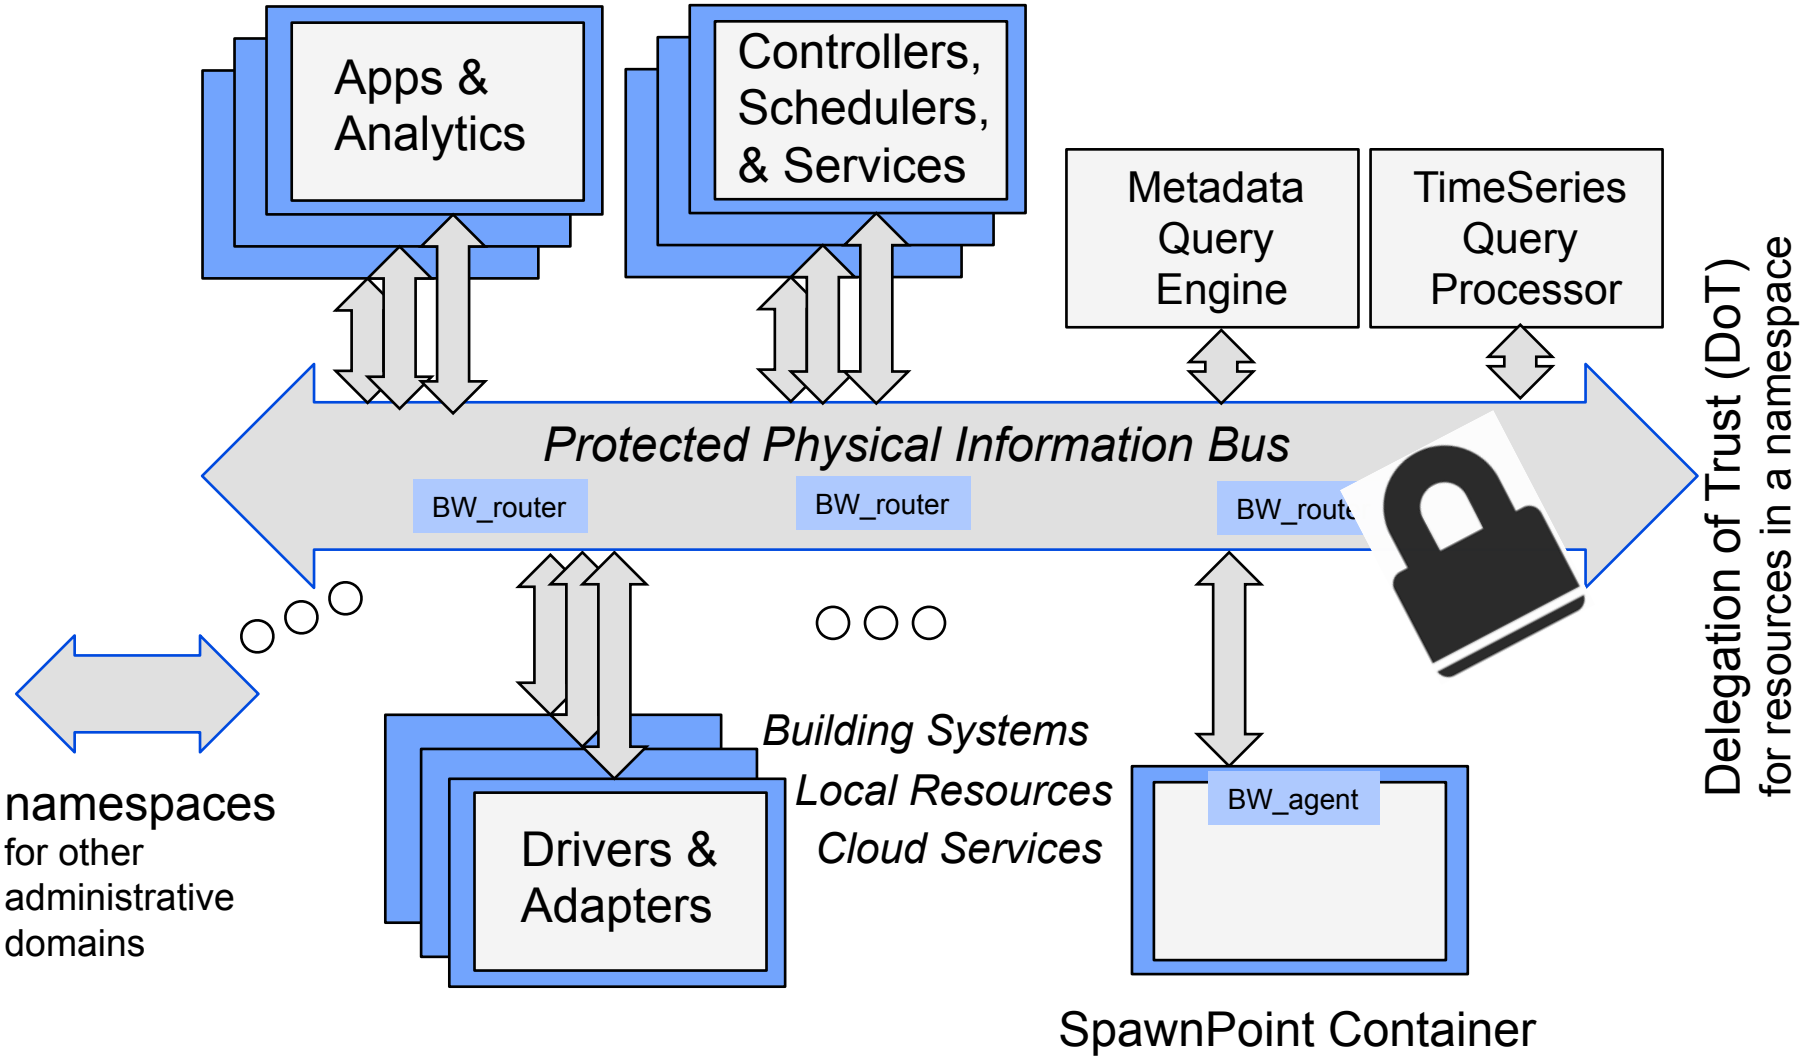

WAVE in a nutshell

- WAVE connects sensors, actuators, controllers, drivers, apps and people
- Backbone of the micro-service architecture forming XBOS
  - Publish/Subscribe syndication (via metadata)
- Rich security model permitting delegation and strong verification
- Every message carries a proof of authorization and authentication
- The whole system operates without any central authorities
- And provides stronger security guarantees than existing systems w.r.t revocation, attack surface, requisite implicit trust etc

Fully Distributed Authentication & Authorization

- Any entity can grant (delegate) permissions
  - Without communicating with grantee or any authority
- Anyone can verify any permission non-interactively
- All actions carry a complete proof of authorization
  - Easily verified by recipient
  - Or by routers – to prevent DOS
- Transparent, Auditable Delegation of Trust (DoT)
  - Or Protected DoT at somewhat higher cost
- Built using micro-contracts over a blockchain (Ethereum)

The Principal: WAVE Entity

- Like: email address / username
- A keypair for signing and verifying: <E_{sk}, E_{vk}>
  - Identified by E_{VK}
  - e.g l0hKkvaVyRDqf_lwt93WJC_a9Zu2F3l61Au6fZtlsCU=
  - Optionally identified by a globally unique, immutable alias e.g mike19
- Represents the holder of the signing key:
  - IoT device
  - Participant
  - Services

Resources

- Like: file paths / URLs
- Within WAVE, interfaces composing services / devices etc are represented by Resource URIs
  - namespace/resource_path

alicehome/hvac/thermostat/setpoint
alicehome/security/door/islocked
caiso/pricing/zone25/electricity

Namespaces

- **Like: Active Directory Domain / Root user**
- **There needs to be "root" from which permissions flow**
- **namespace/resource_path** alicehome/hvac/thermostat/setpoint
  - **All resource URIs begin with the E_{ns} of the namespace entity (or its alias)**
  - **E_{ns} has full permissions on all resources within the namespace**
  - **This namespace prefix is self-proving (as it is a public key)**
  - **The URI alone is enough to identify the namespace authority (no external authority needs to exist)**

To: alicehome/hvac/thermostat/setpoint
Subject: Change value
76F

Delegation of Trust

- **Like: Assuming a role (in RBAC), joining a POSIX group**
- **For other entities to obtain permissions on a resource, they must receive them via a delegation of trust (DoT)**

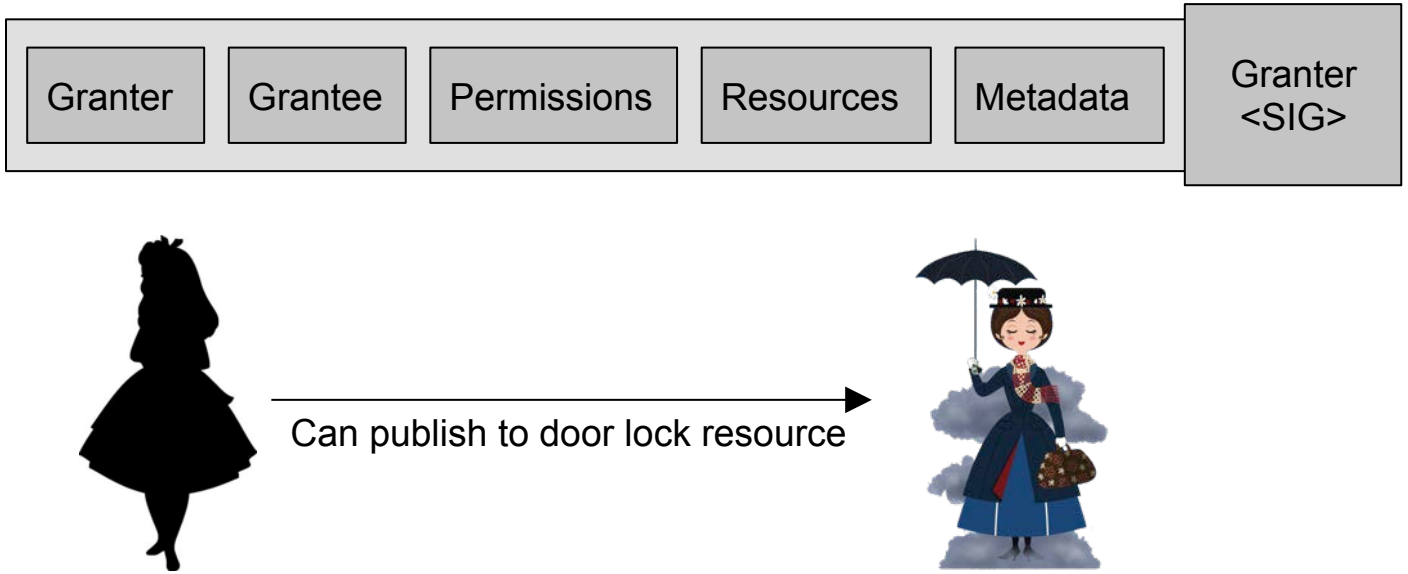

Proof and Verification

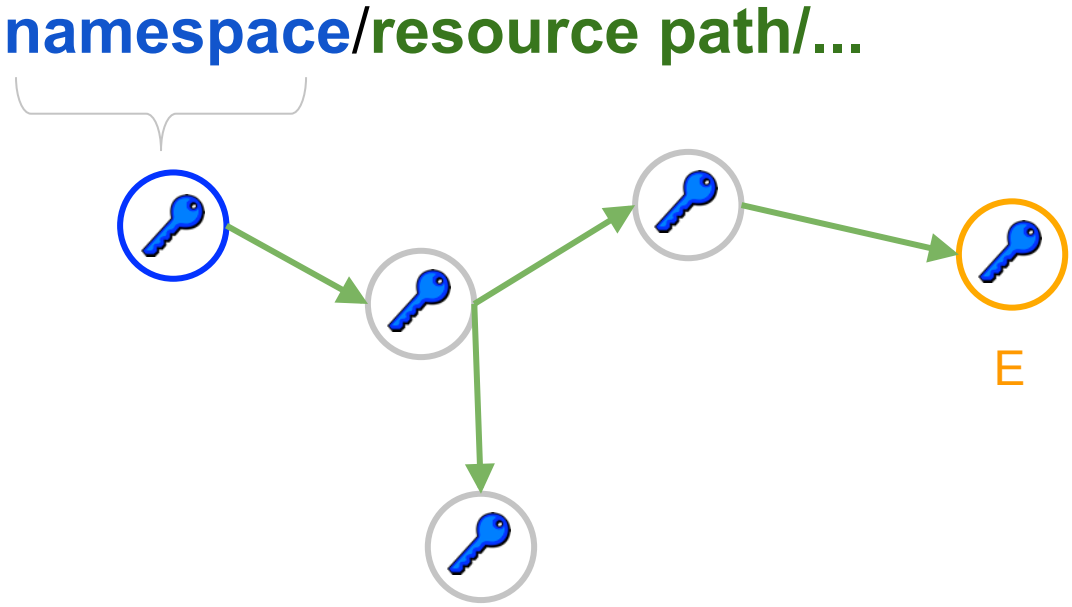

- To prove E has permission P on a URI show there exists a chain of DoTs from E_{NS} to E
- s.t. the intersection of the permissions granted by DoTs on this chain is at least P

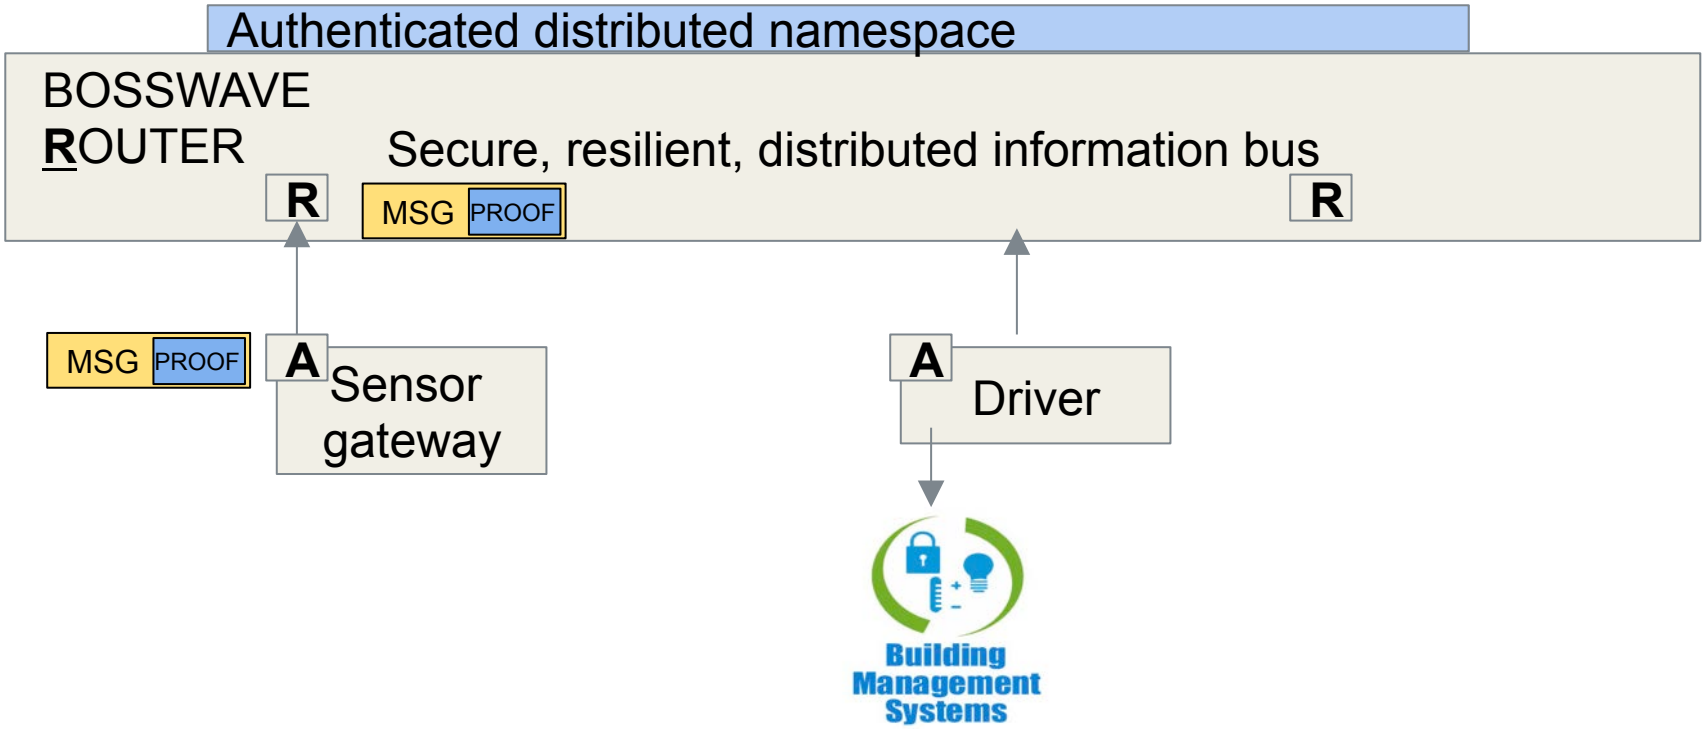
Authenticated distributed namespace
BOSSWAVE
ROUTER
Secure, resilient, distributed information bus
R
MSG
PROOF
R
MSG
PROOF
A
Sensor
gateway
A
Driver
Building
Management
Systems

BOSSWAVE URIs

```
<namespace>/.../<service name>/<instance name>/<interface name>/{signal,slot}/<property>
```

```
ciee/devices/venstar/s.venstar/OpenSpace/i.xbos.thermostat/signal/info
```

XBOS URI Idioms

- **`<service name>`: Identifies instance of driver**
- **`<instance name>`: Identifies device exposed by driver**
- **`<interface name>`: Set of signal/slots and properties to expect**

XBOS Authorization with WAVE

- **Leverage structured URIs, BOSSWAVE URI patterns**
- **Common device permissions**
  - `ciee/*/signal/info`
  - `ciee/*/i.xbos.thermostat/+/+`
  - `ciee/*/i.xbos.thermostat/signal/+`
  - `ciee/*/i.xbos.thermostat/slot/setpoints`
  - `ciee/*/OpenSpace/i.xbos.thermostat/slot/setpoints`
- **Permission granularity limited by messages**

XBOS Components: Spawnpoint

- **Secure, distributed, managed containers**
- **Deployment administration**

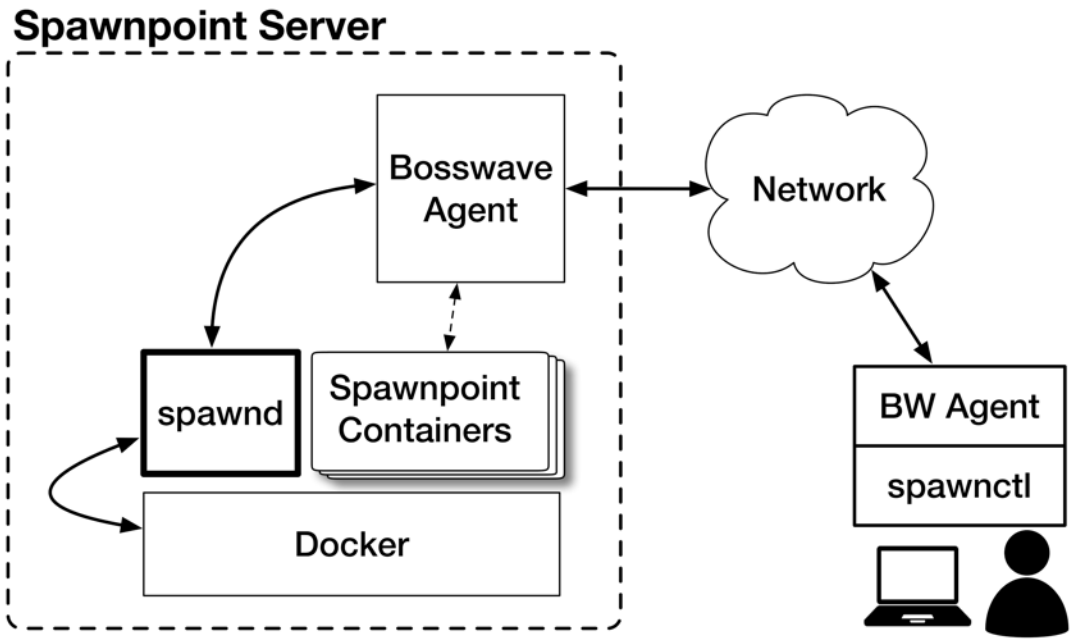

XBOS Components: Archiver

- **Timeseries data storage and retrieval**
- **Applies BOSSWAVE permission model to archival data**
- **BTrDB-backed**

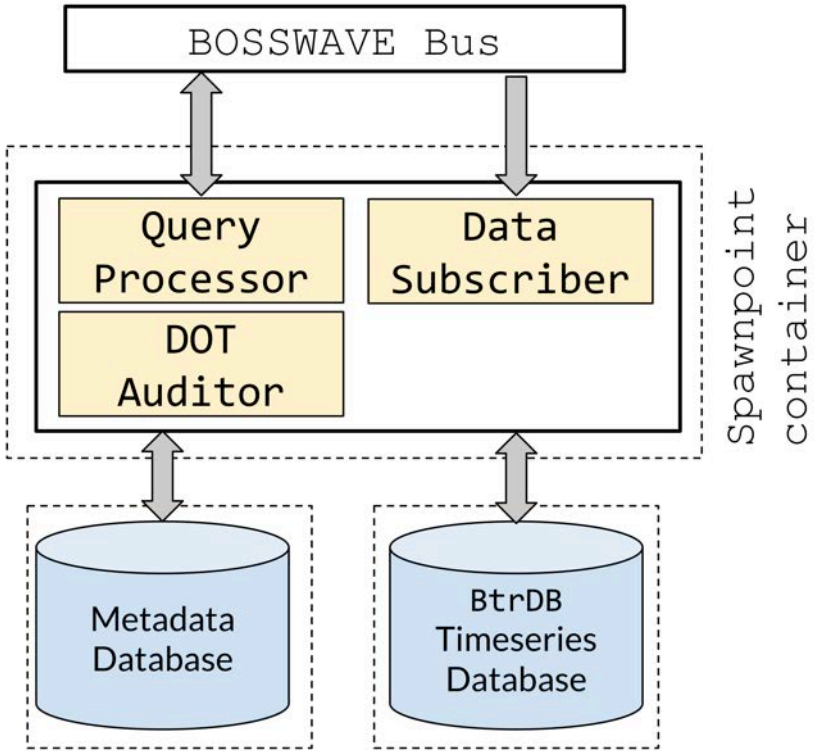

XBOS Components: Building Profile

- **Serves Brick metadata model of building**
- **Describe/relate:**
  - **physical resources**
  - **logical resources**
  - **building subsystems**
  - **equipment hierarchies**

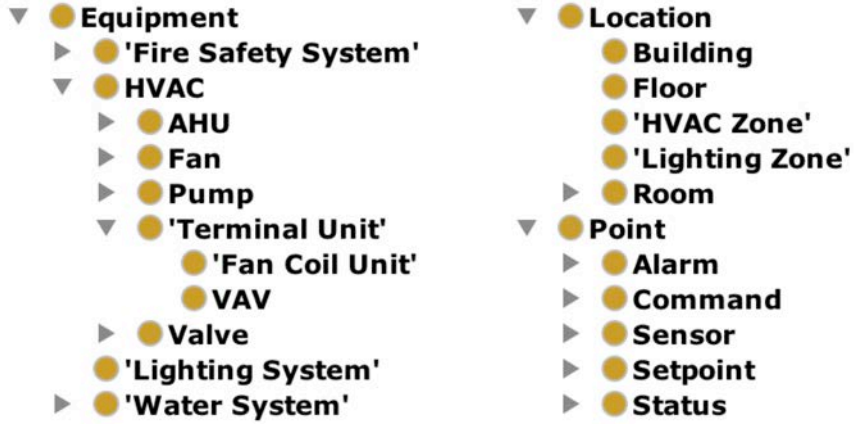

HVAC Monitoring App: CIEE

1. Application: ***"Get me the thermostat for each zone"***

App

HVAC Monitoring App: CIEE

1. Application: *"Get me the thermostat for each zone"*
2. Query HodDB using SPARQL

```
SELECT ?tstat ?temp_uuid ?uri ?zone WHERE {
    ?zone rdf:type brick:HVAC_Zone .
    ?tstat rdf:type brick:Thermostat .
    ?tstat bf:uri ?uri .

    ?tstat bf:hasPoint ?ts .
    ?ts rdf:type brick:Temperature_Sensor .
    ?ts bf:uuid ?temp_uuid .

    ?tstat bf:controls/bf:feeds? ?zone .
};
```

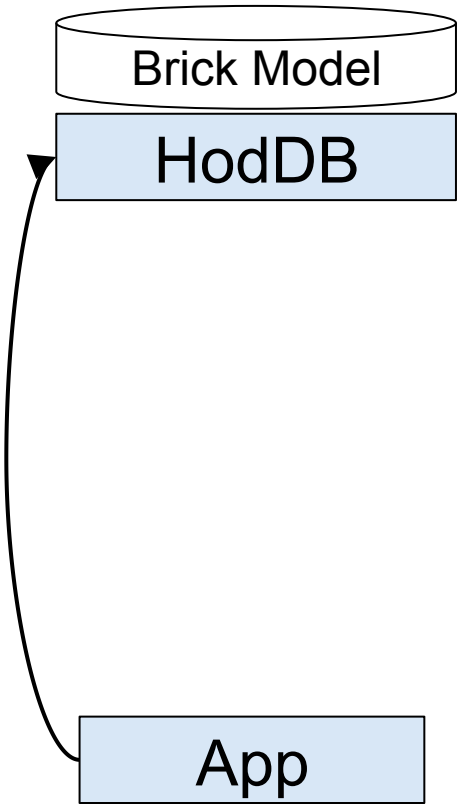

HVAC Monitoring App: CIEE

1. Application: *"Get me the thermostat for each zone"*
2. Query HodDB using SPARQL
3. Configure app using results

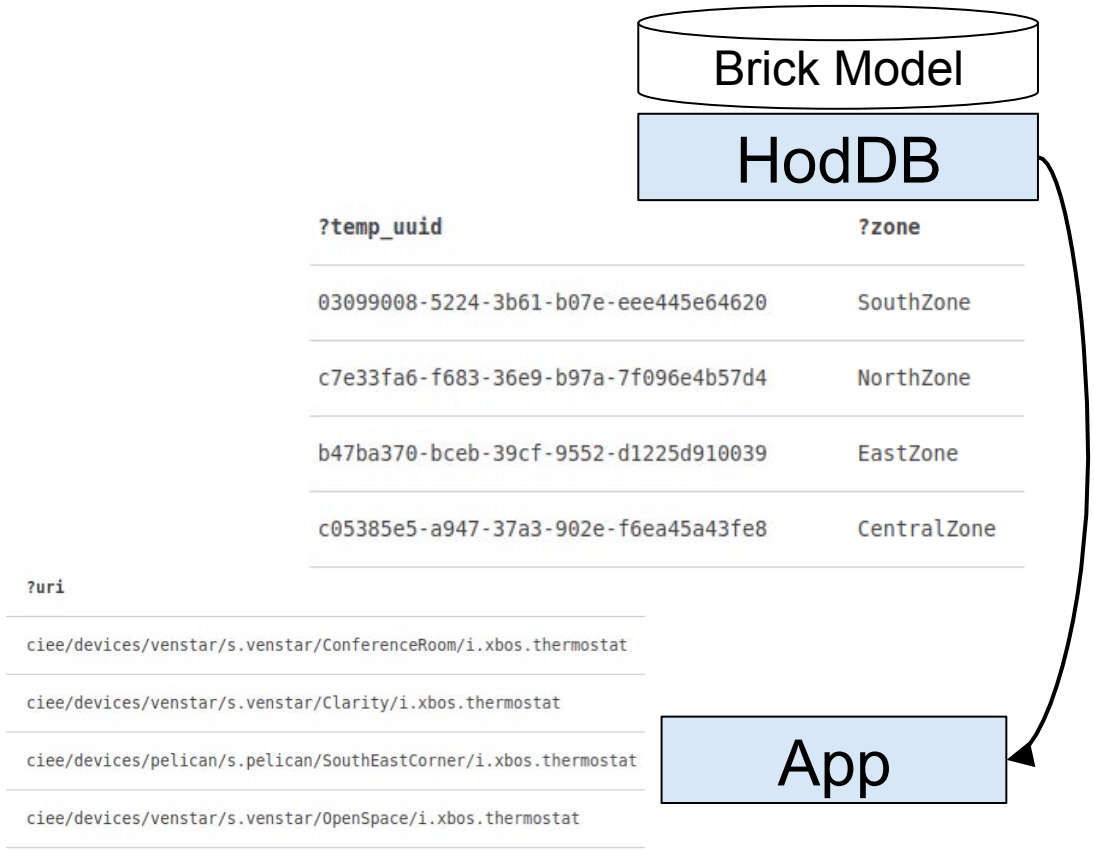

HVAC Monitoring App: CIEE

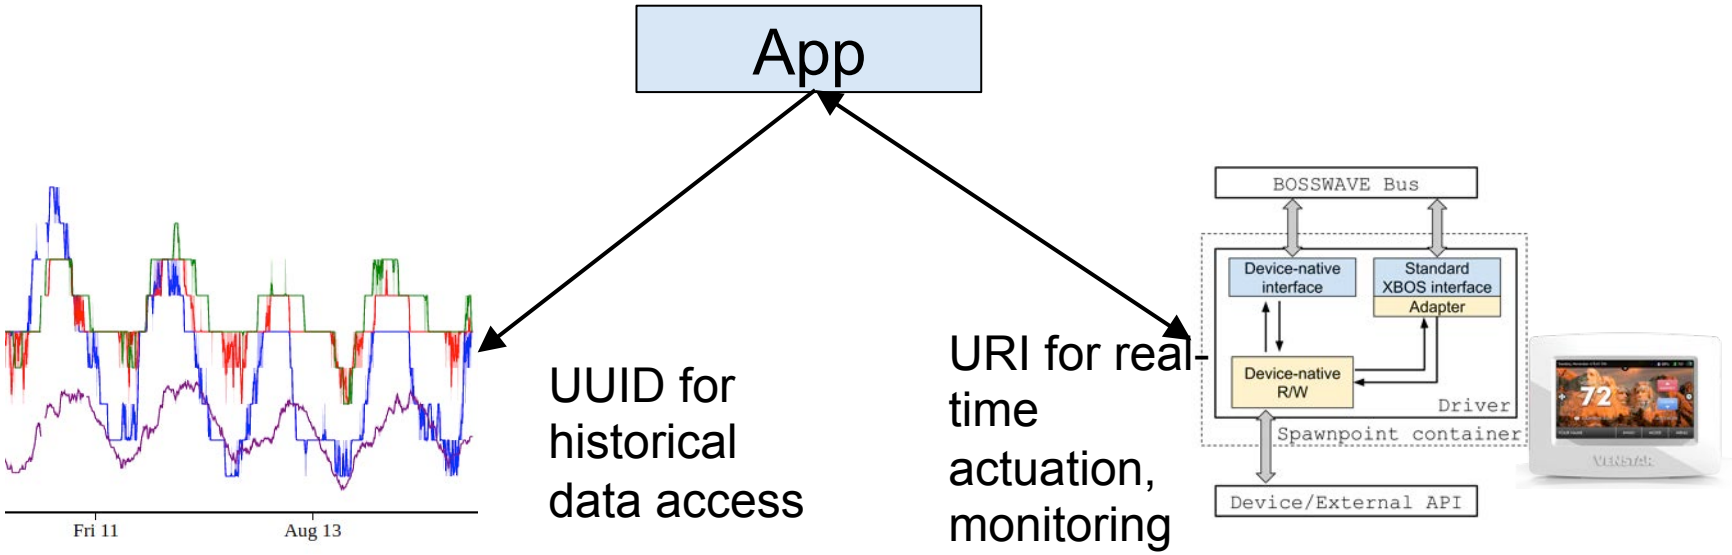

- **Brick model yields logical resource names**
- **How we achieve portable applications**

HVAC Monitoring App: CIEE

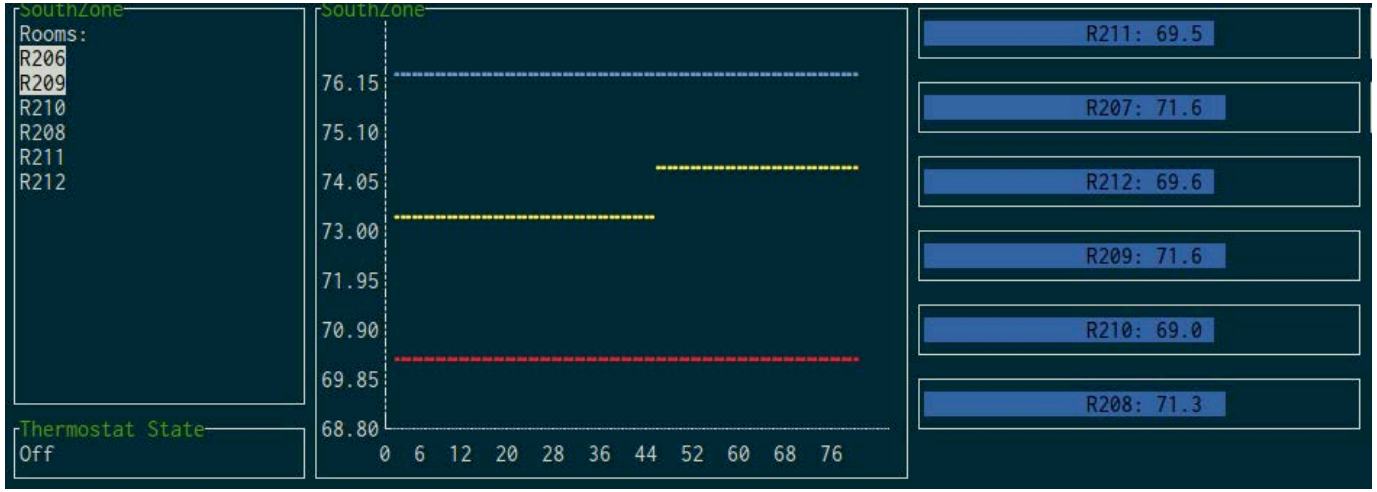

- **Fully auto-generated from Brick**
- **Combination of historical queries + live streaming data**
- **Occupancy notification**
- **Color-coded temperature sensors**
- **Thermostat trends**

Demand Response App

- **Bootstraps from Brick model**
  - **Find devices by class**
- **Apply strategy:**
  - **Widen deadband**
  - **Dim lights**
  - **Find occupancy sensors, turn off equipment in empty rooms**
- **~100 lines of Python:**
  - `https://goo.gl/Mdrh2o`

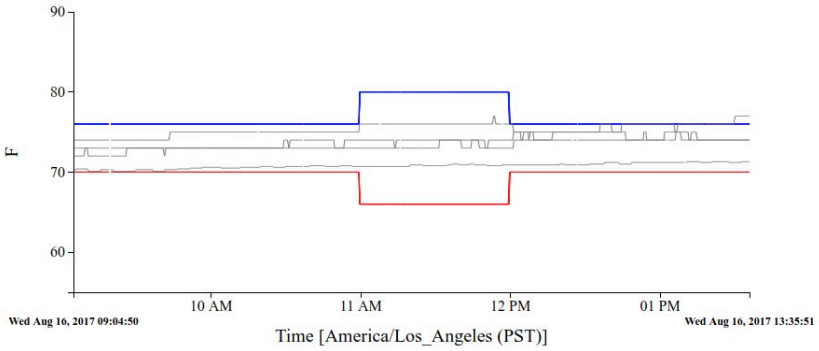

Hamilton: Flexible, Open Source $10 Wireless Sensor System for Energy Efficient Building Operation

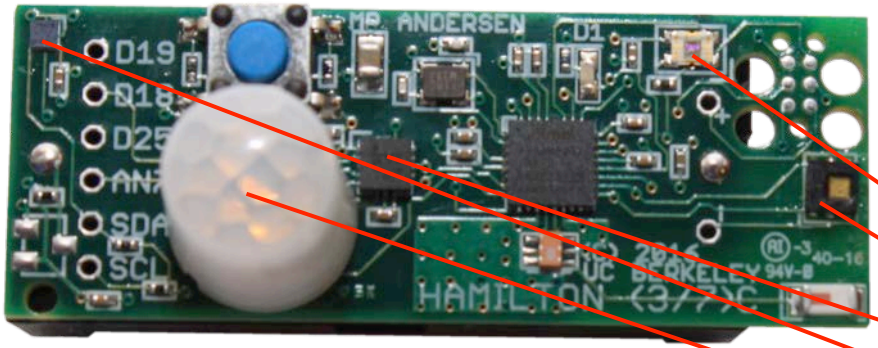

ATSAMR21 - 32-bit 48MHz Cortex M0+
32 KB RAM / 256 KB Flash

- Objective is to have it cost $10
- Sans PIR it costs ~$20 at the moment, manufactured in Oakland
- Illuminance sensor
- Air temp ±0.2 C RH ±2 % (factory calibrated)
- Orientation (3-ax magnetometer & accel)
- Radiant temperature (90° cone)
- Optional high sensitivity PIR

- We expect a **20 second** reporting rate (18 uA) to last **5 years** on a real battery (9.5 years on an ideal 1500mAh battery)
- Each Reading: 3D Mag, 3D Acc, Air Temp, Rh, Illuminance, Radiant Temp, Aggregated PIR

Networking

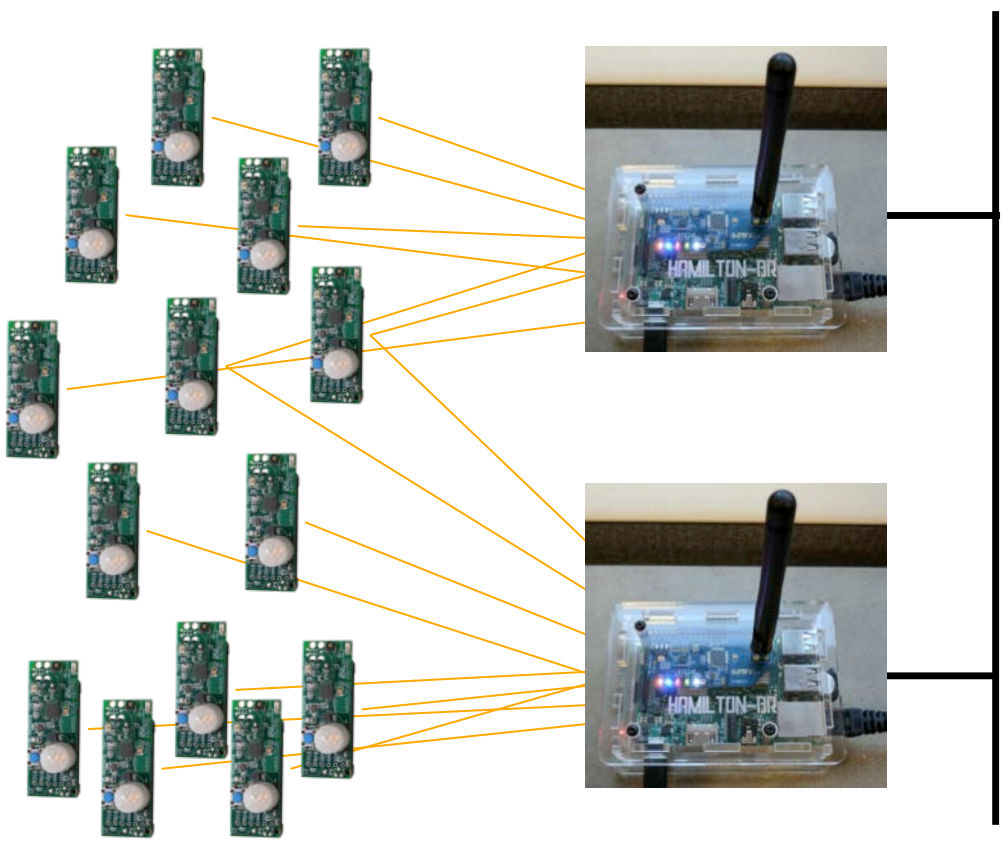

- Archiving
- Real-time analytics
- Building applications
- Third party system adapters

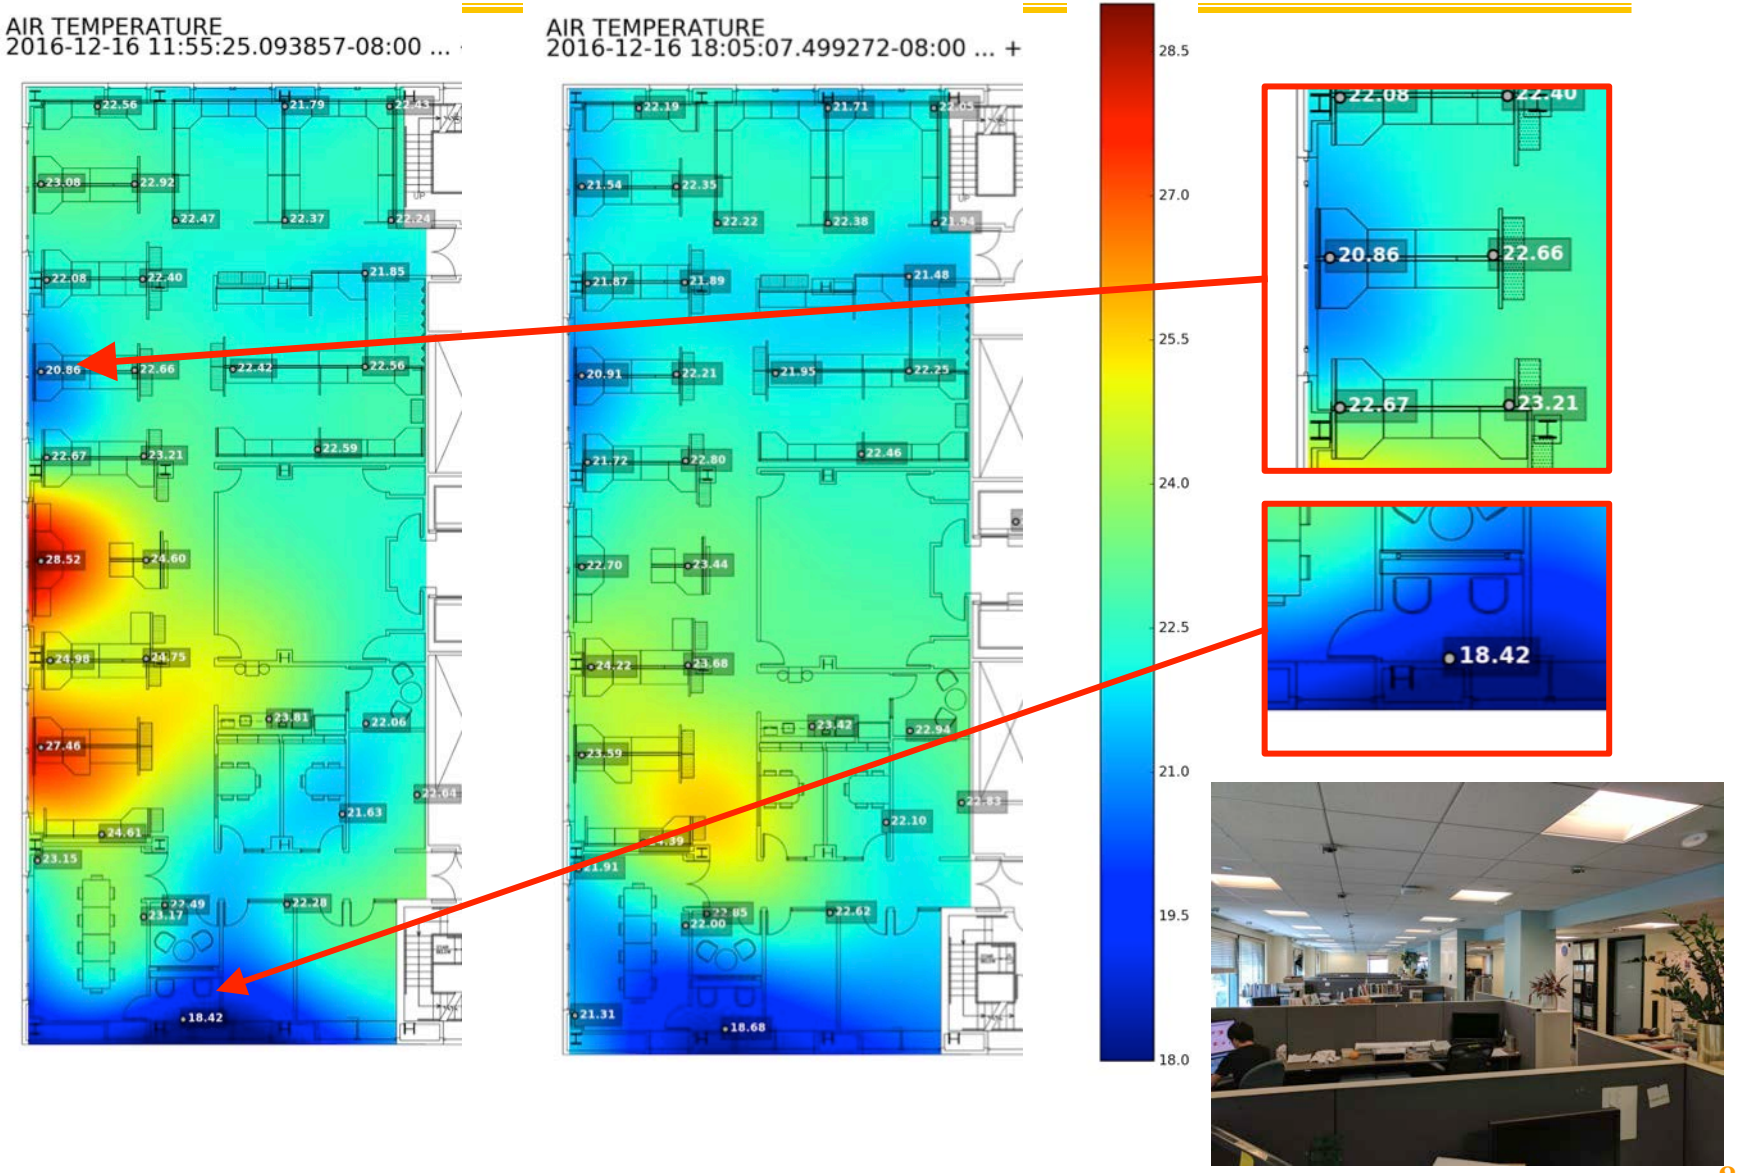
AIR TEMPERATURE
2016-12-16 11:55:25.093857-08:00 ...
AIR TEMPERATURE
2016-12-16 18:05:07.499272-08:00 ... +
28.5
27.0
25.5
24.0
22.5
21.0
19.5
18.0
20.86
22.66
22.67
23.21
18.42

Putting it Together in Demand Response

7.5k sq ft office building

- 5 RTUs
- 1 Building meter
- 2 Plug meters
- ~25 Lights
- ~16 Hamiltons

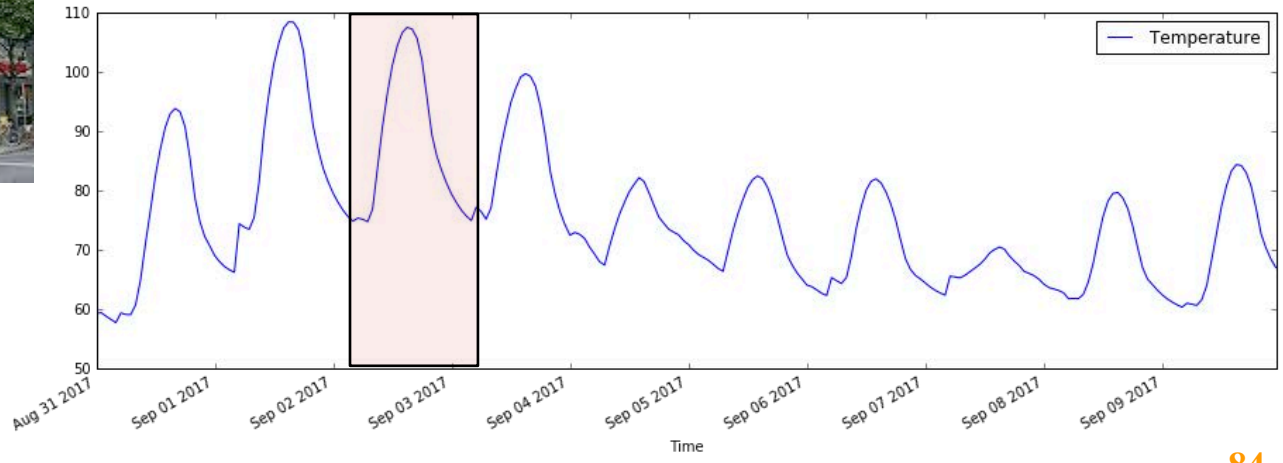

Sept 2 – noDR day

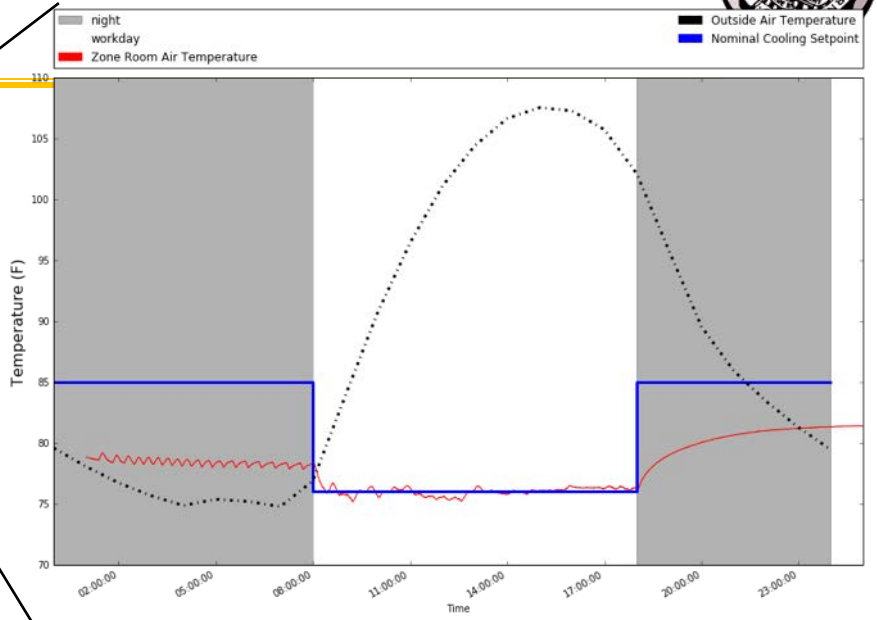

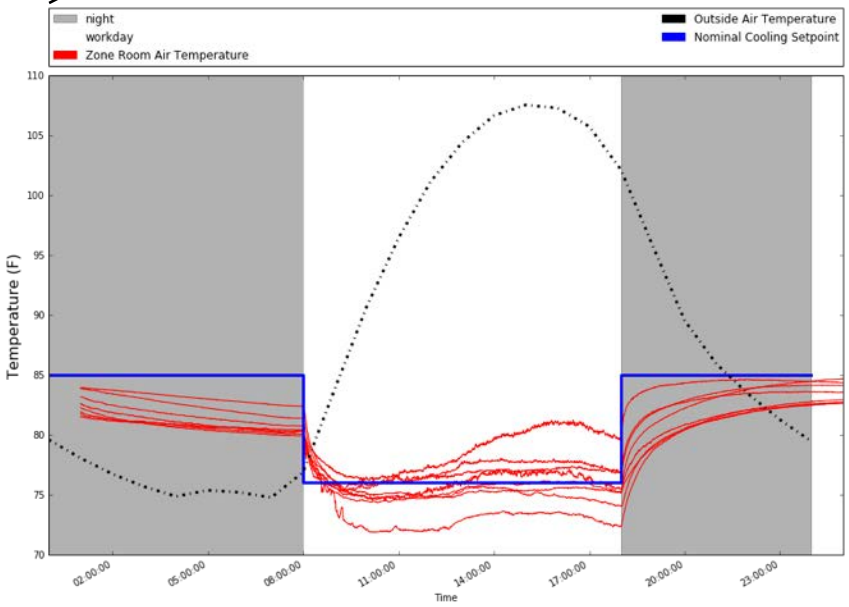

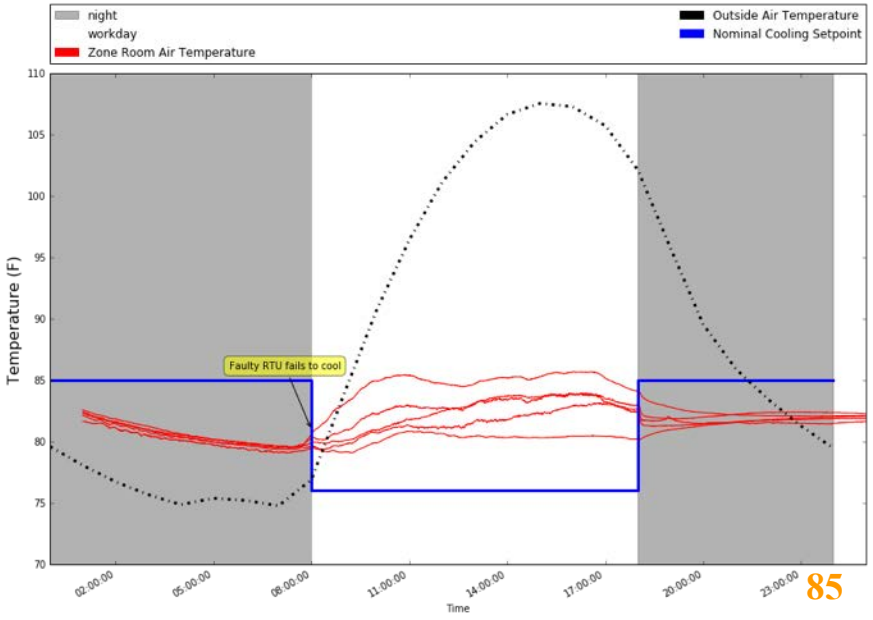

CIEE Building Energy Usage – Sept 2

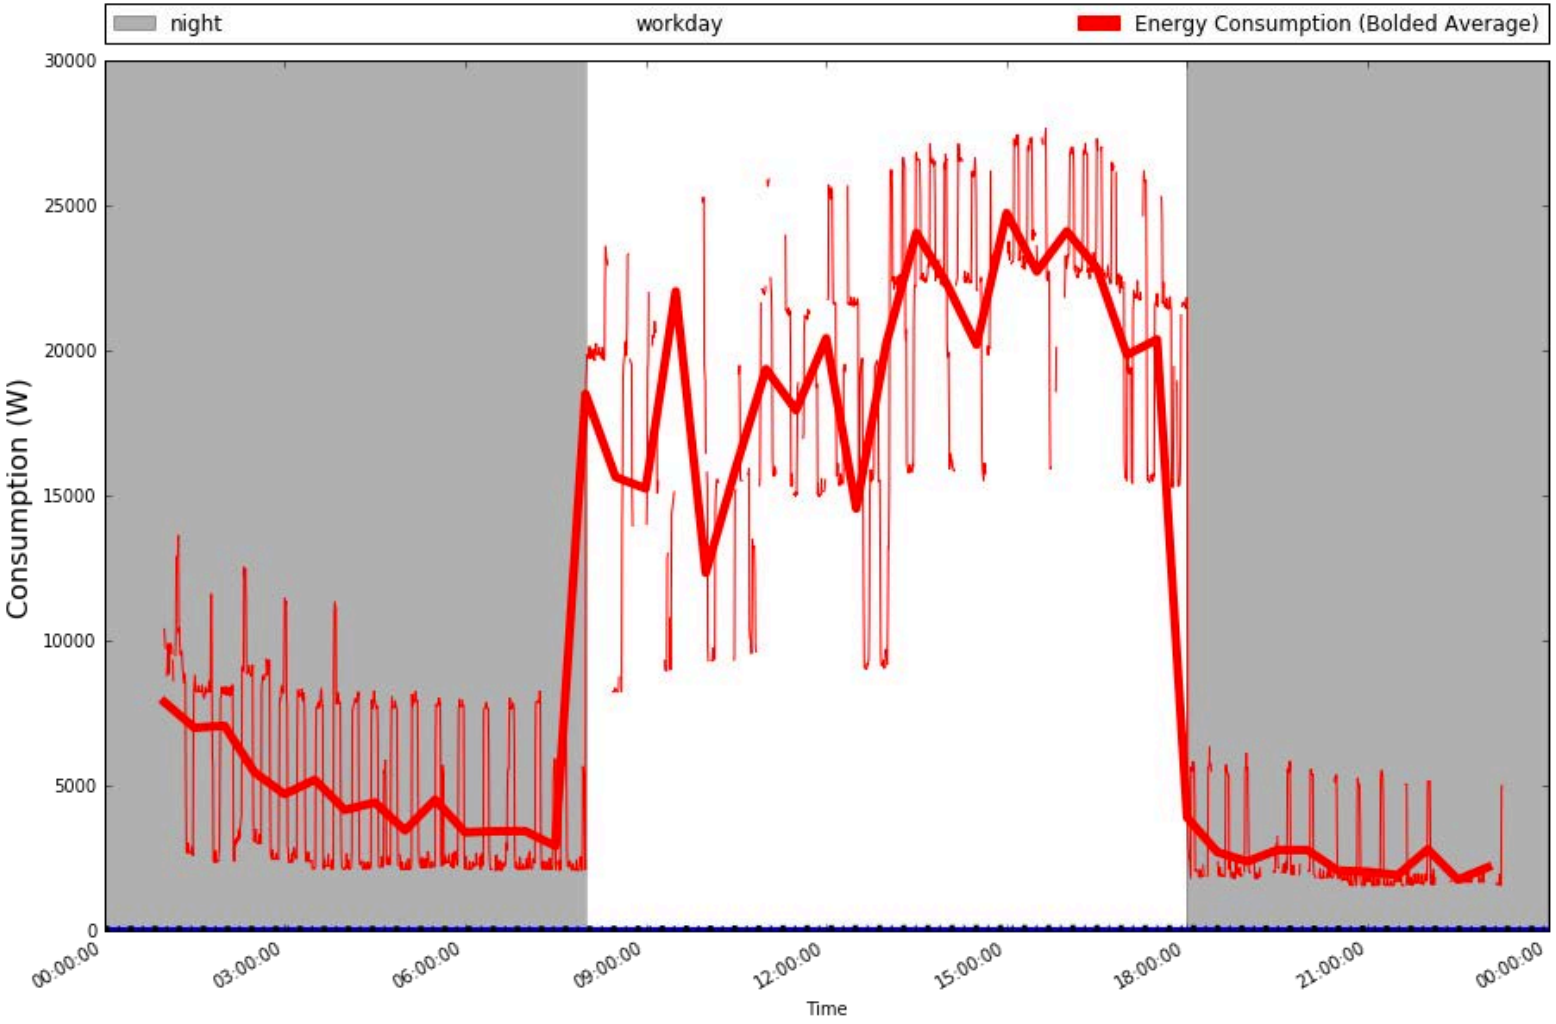

Building Energy under DR

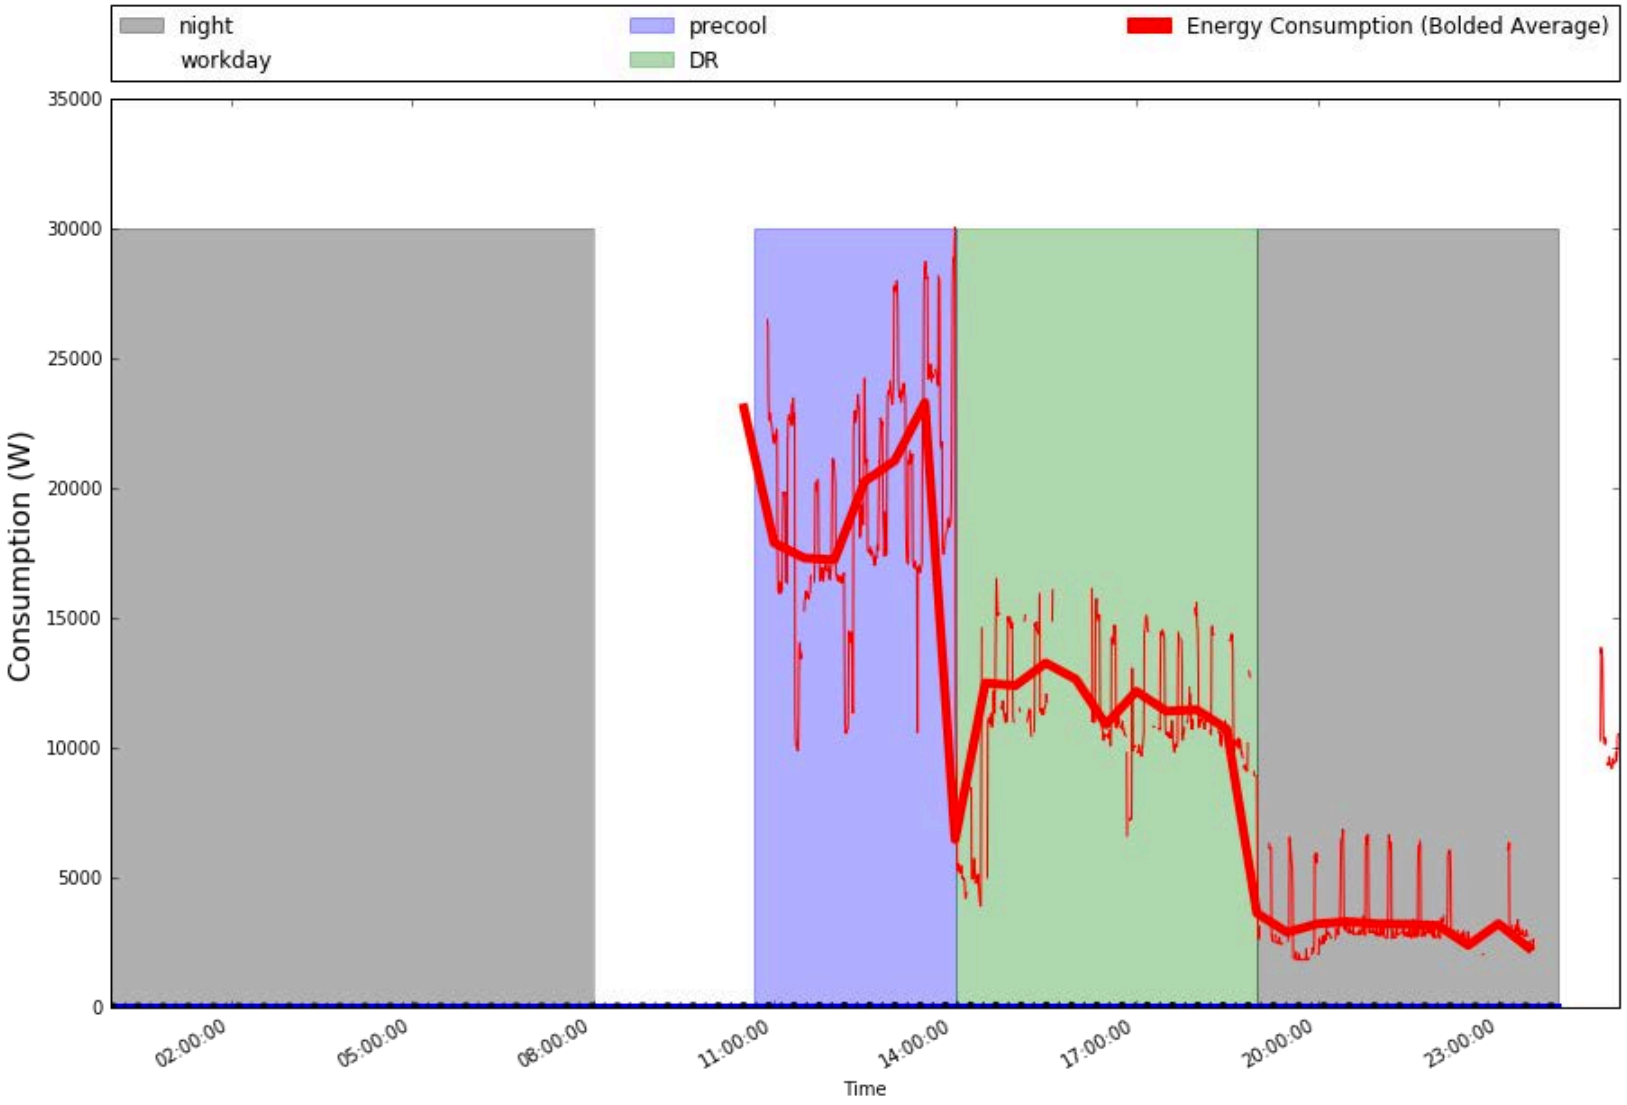

Sept 3 – DR day

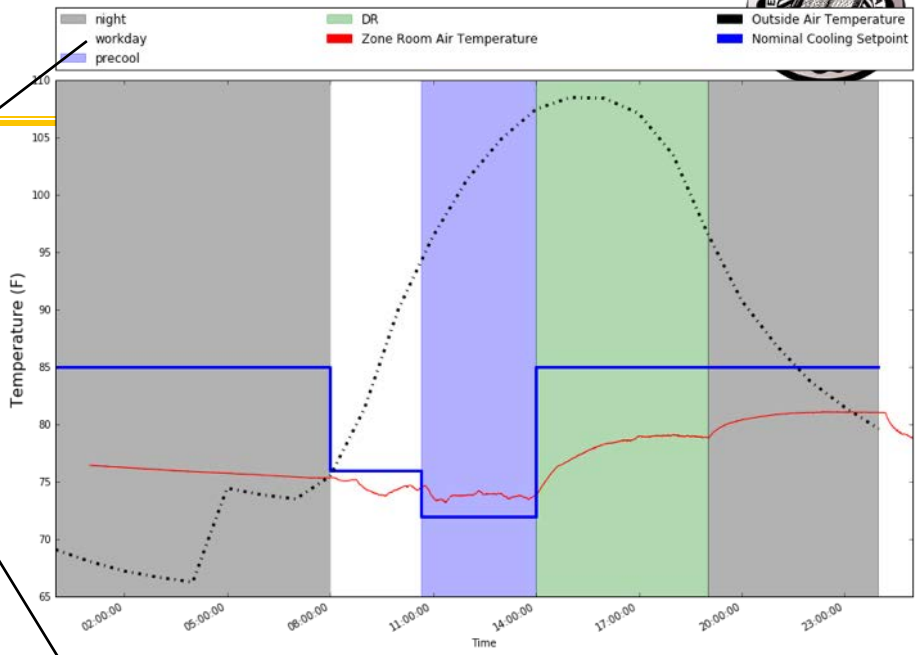
night
workday
precool
DR
Zone Room Air Temperature
Outside Air Temperature
Nominal Cooling Setpoint
Temperature (F)
Time

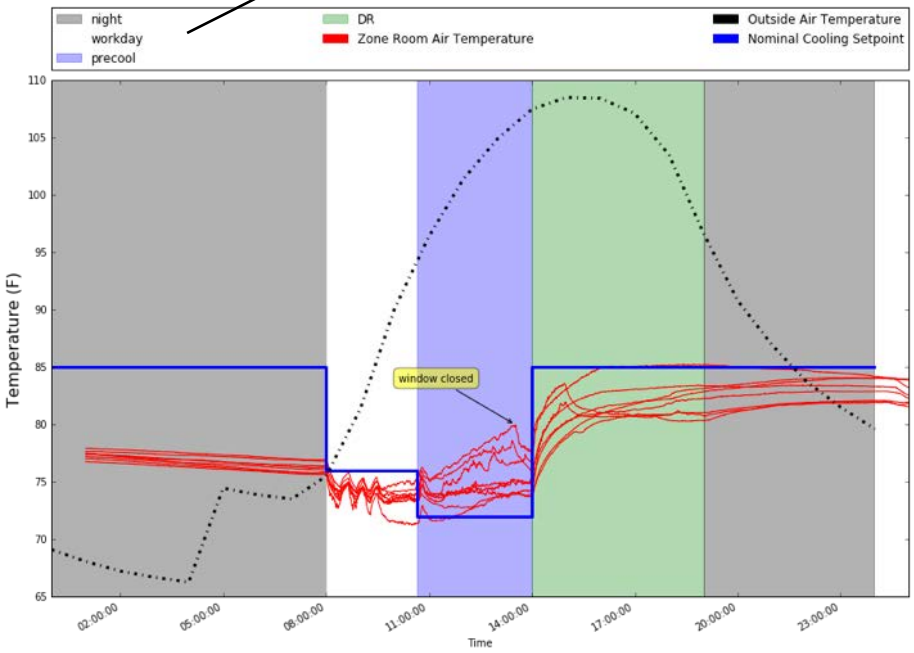
night
workday
precool
DR
Zone Room Air Temperature
Outside Air Temperature
Nominal Cooling Setpoint
window closed
Temperature (F)
Time

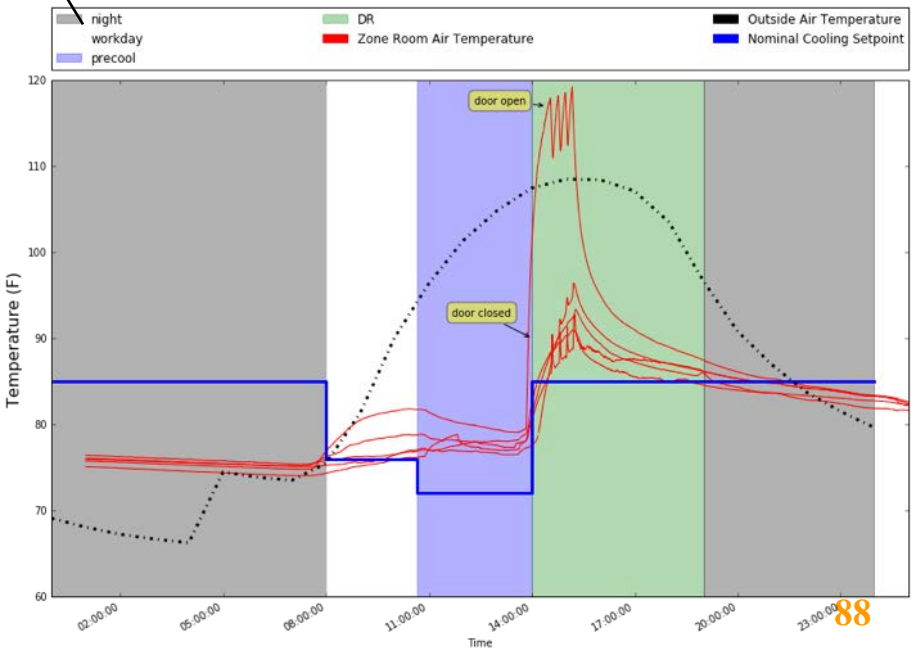
night
workday
precool
DR
Zone Room Air Temperature
Outside Air Temperature
Nominal Cooling Setpoint
door open
door closed
Temperature (F)
Time

Energy Analytics

- Observe building energy and all the control actions
- Easy to identify energy usage per action
  - without NILM

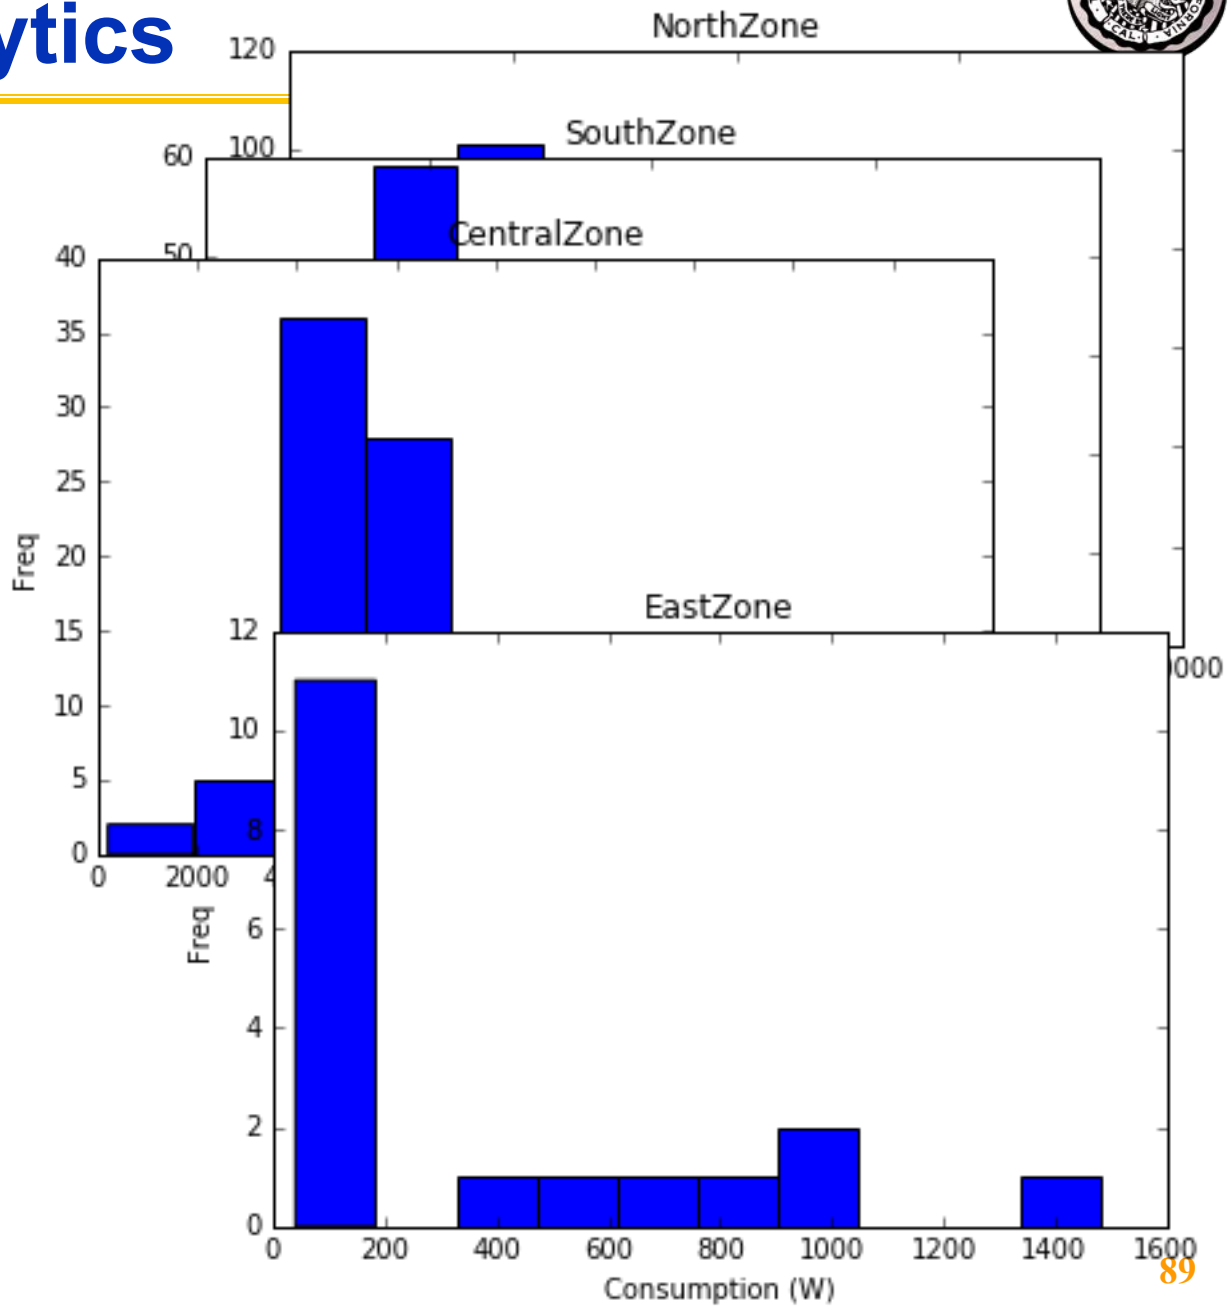

City-Scale Energy Services

- For a city the size of San Francisco (~1 million people) …
  - Property management (Owner => Mgr => Occupant => Guest)
  - Utility / ISO demand side management
  - Energy management services integrate notification => response based on occupant valie
- BOSSWAVE objects driven by housing turnover
- Entities:
  - Occupant, Owner, Utility, Apt lease, House title, Meter, thermostat
- DOTs/Revocations:
  - Begin/end lease
  - Read your electric meter, thermostat
  - Give family/occupant access

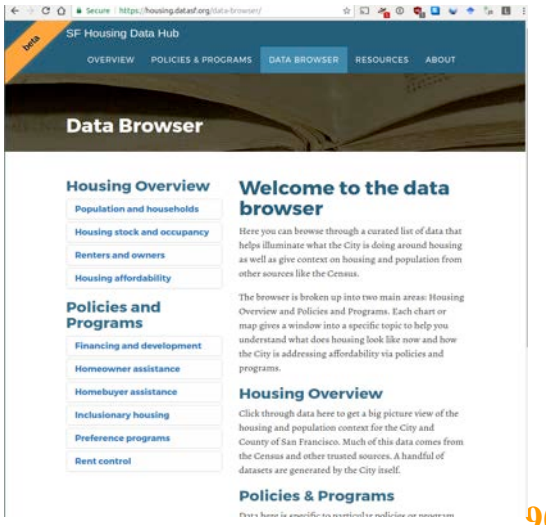

BOSSWave City-Scale Emulation

Housing distribution:
- Land use/parcel data
- 136852 properties; house/apt dist.

Housing turnover:
- Census data (House 8%/year, Apt 17.3%/year)

Occupants per house
- Mean 2.3/house, apply variance of 3

Type	Entities	DoTs granted	Avg Out°
Occupant	951,293	1,312,005	1.38
Apt Owner	15,787	529,562	33.54
Apt Bldg	40,921	40,921	1
Apt Lease	264,781	264,781	1
House Title	95,931	95,931	1
Thermostat	360,712	N/A	N/A
Meter	360,712	N/A	N/A
Utility	603	722,026	1197.39
Total	2,090,740	2,965,226	1.42

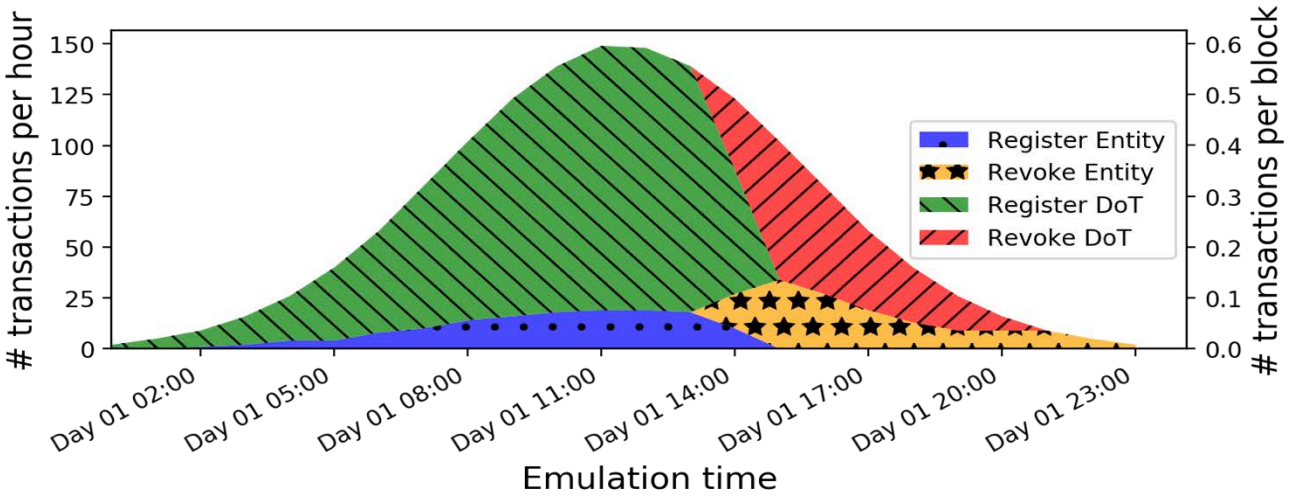

Towards Energy Networked Systems

- Research Methodology: Instrumented Virtual Grid of Living Laboratories
- Physical Information representation & syndication
- Grid/Building Operating System & Services
- Scaling: Automated Metadata acquisition
- Portable Applications on Physical systems
- Nailing: Evidence-based schema standardization
- Distribution Tier Awareness
  - Extremely fast, multiresolution time series query processing with statistical aggregation
- Fully distributed authentication and authorization for establishment of trust

Thanks

U.S. DEPARTMENT OF ENERGY
Energy Efficiency & Renewable Energy

Where to go for more

- **An Information-Centric Energy Infrastructure: the Berkeley View,** Sustainable Computing: Informatics and Systems, 2011.
- **sMAP - a Simple Measurement and Actuation Profile for Physical Information.** 8th ACM Conference on Embedded Networked Sensor Systems (Sensys 2010).
- **Defining CPS Challenges in a Sustainable Electricity Grid,** ACM/IEEE Third International Conference on Cyber-Physical Systems, 2012.
- **BOSS: Building Operating System Services,** 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI '13), 2013.
- **Enabling Advanced Environmental Conditioning with a Building Application Stack**. 4th International Green Computing Conference (IGCC '13), 2013.
- **Well-connected Microzones for Increased Building Efficiency and Occupant Comfort**, 2016 ACEEE Summer Study on Energy Efficiency in Buildings, 2016.
- **Automated Metadata Construction to Support Portable Building Applications,** BuildSys 2015
- **Analyzing Metadata Schemas for Buildings : The Good, The Bad and The Ugly,** BuildSys 2015.
- **Enabling Synergy in IoT: Platform to Service and Beyond**, 1st IEEE International Conference on Internet-of-Things Design and Implementation
- **Enabling Portable Energy Applications with a Building Operating,** ACEEE Summer Study on Energy Efficiency in Buildings, 2016.
- **Brick: Towards a Unified Metadata Schema for Buildings,** BuildSys 2016
- **DISTIL: Design and implementation of a scalable synchrophasor data processing system**, 2015 IEEE International Conference on Smart Grid Communication (SmartGridComm), 2015
- **BTrDB: Optimizing Storage System Design for Timeseries Processing**, 14th USENIX Conference on File and Storage Technologies (FAST 16), Feb., 2016.