Smart City Services

Tarek Abdelzaher
Dept. of Computer Science
University of Illinois at Urbana Champaign

Recap

- Sensing in social spaces informs data-centric applications in two ways:

Recap

- Sensing in social spaces informs data-centric applications in two ways:
  - First approach:

    Statistics

Recap

- Sensing in social spaces informs data-centric applications in two ways:
  - Second approach:

    Modeling

Recap

- Sensing in social spaces informs applications in two ways:
  - Offers data statistics:
    - Statistics need a lot of data → Can't generate statistics if you did not measure.
    - *Example:* report speed of traffic on different streets by multiple cars and empirically compute the average
  - Allows data modeling:
    - Models allow inferring values of these variables even in places were you did not measure
    - *Example:* generate a model that predicts traffic speed as a function of speed limit, number of lanes, time of day, and day of the week, and weather (dry/rain/etc)

Green GPS

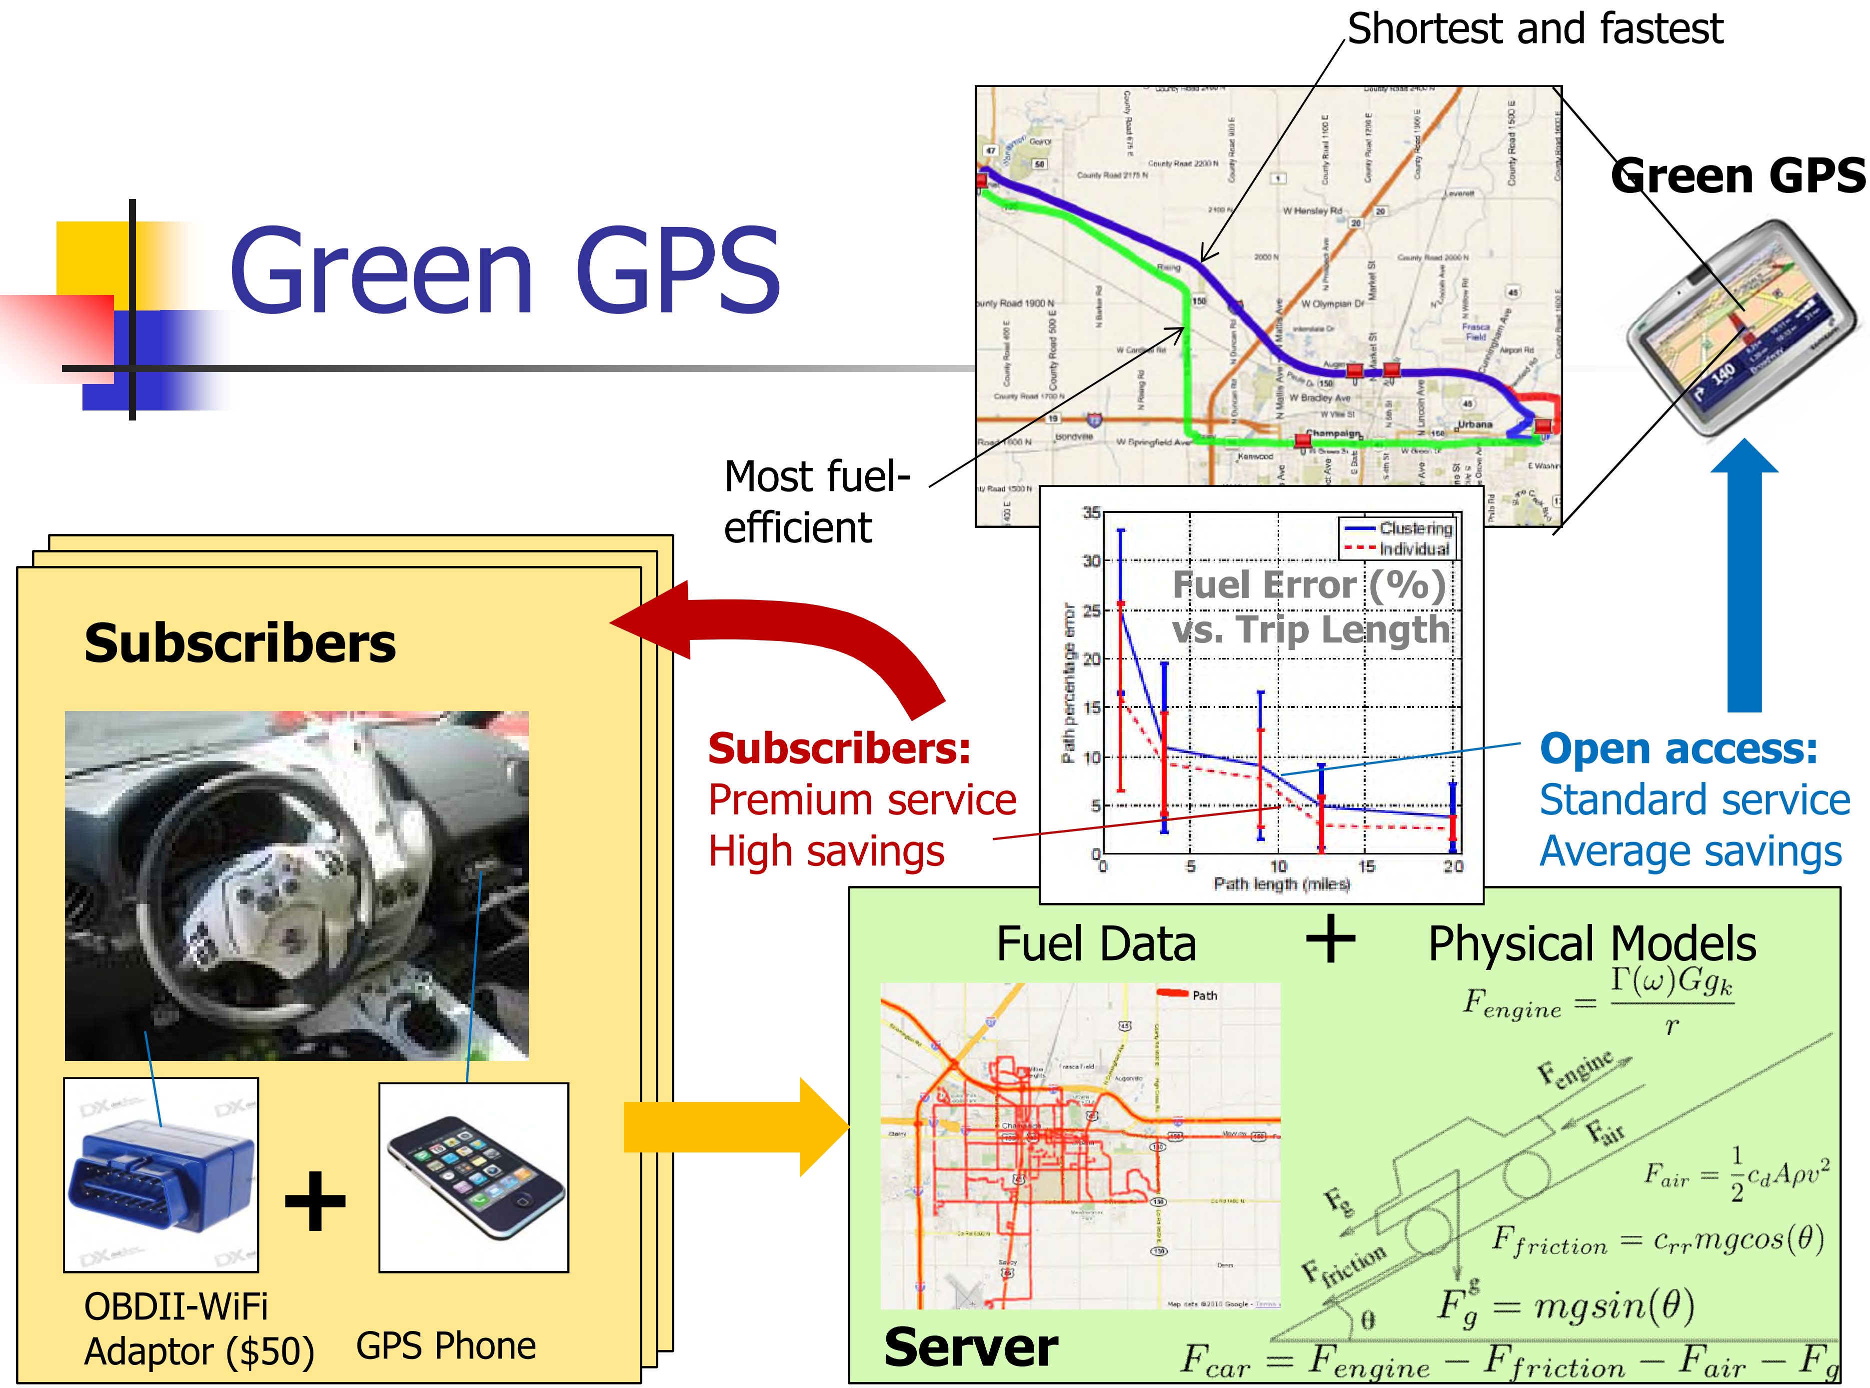

Shortest and fastest
Green GPS
Most fuel-efficient
Clustering
Individual
Fuel Error (%) vs. Trip Length
Path percentage error
Path length (miles)
Subscribers
Subscribers:
Premium service
High savings
Open access:
Standard service
Average savings
OBDII-WiFi Adaptor ($50)
+
GPS Phone
Fuel Data
+
Physical Models
Path
Server
$F_{engine} = \frac{\Gamma(\omega)Gg_k}{r}$
F_{engine}
F_{air}
$F_{air} = \frac{1}{2}c_d A\rho v^2$
F_g
$F_{friction}$
$F_{friction} = c_{rr} mgcos(\theta)$
$F_g^g = mgsin(\theta)$
θ
$F_{car} = F_{engine} - F_{friction} - F_{air} - F_g$

Fuel Consumption Model

- Simple model for fuel consumption derived from physics principles
- Approximate based on easily measurable parameters (e.g. stop signs, traffic lights, speed limits)

$$gpm = k_1 m\bar{v}^2 \frac{ST + \nu TL}{\Delta d} + k_2 m \frac{\bar{v}^2}{\Delta d} + k_3 m cos(\theta) + k_4 A\bar{v}^2 + k_5 m sin(\theta)$$

$$F_{engine} = \frac{\Gamma(\omega) G g_k}{r}$$

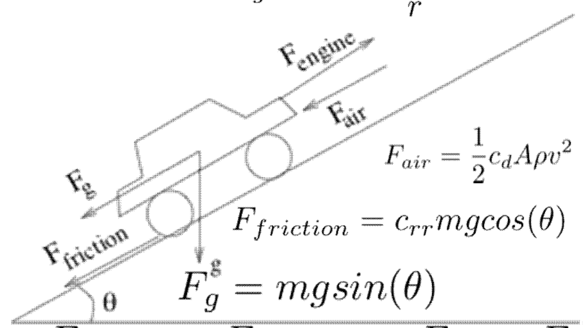

$$F_{air} = \frac{1}{2} c_d A \rho v^2$$

$$F_{friction} = c_{rr} m g cos(\theta)$$

$$F_g = mgsin(\theta)$$

$$F_{car} = F_{engine} - F_{friction} - F_{air} - F_g$$

Sampling Regression Modeling Framework

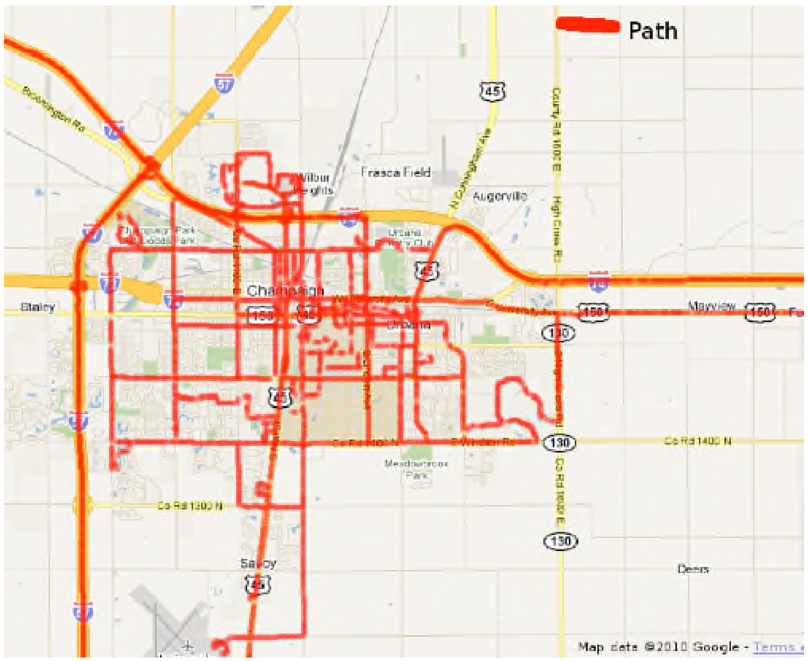

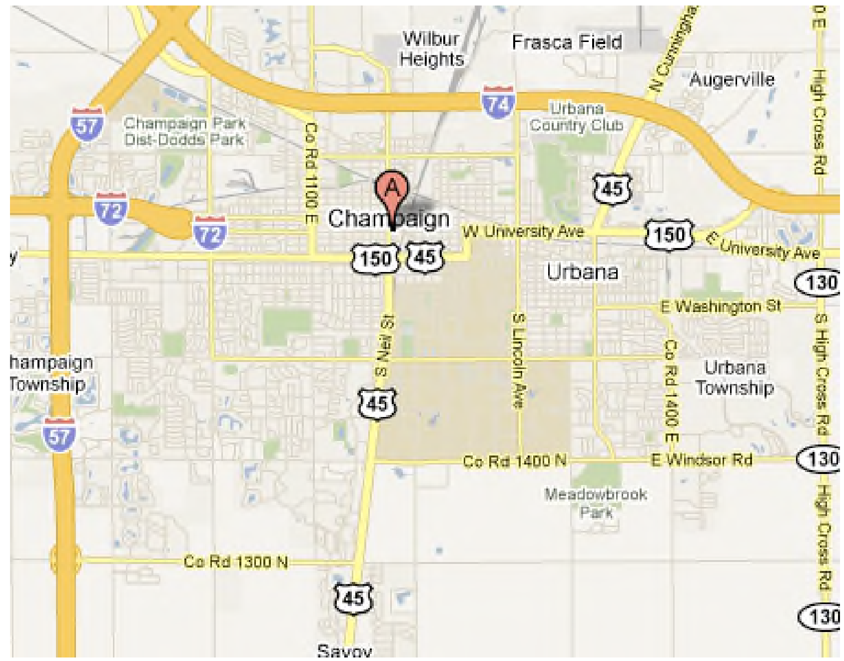

Fuel consumption of
A few cars driven on a
few roads

Predict fuel consumption of
any car on any road in
Urbana-Champaign

Finding Fuel-efficient Routes

- Example

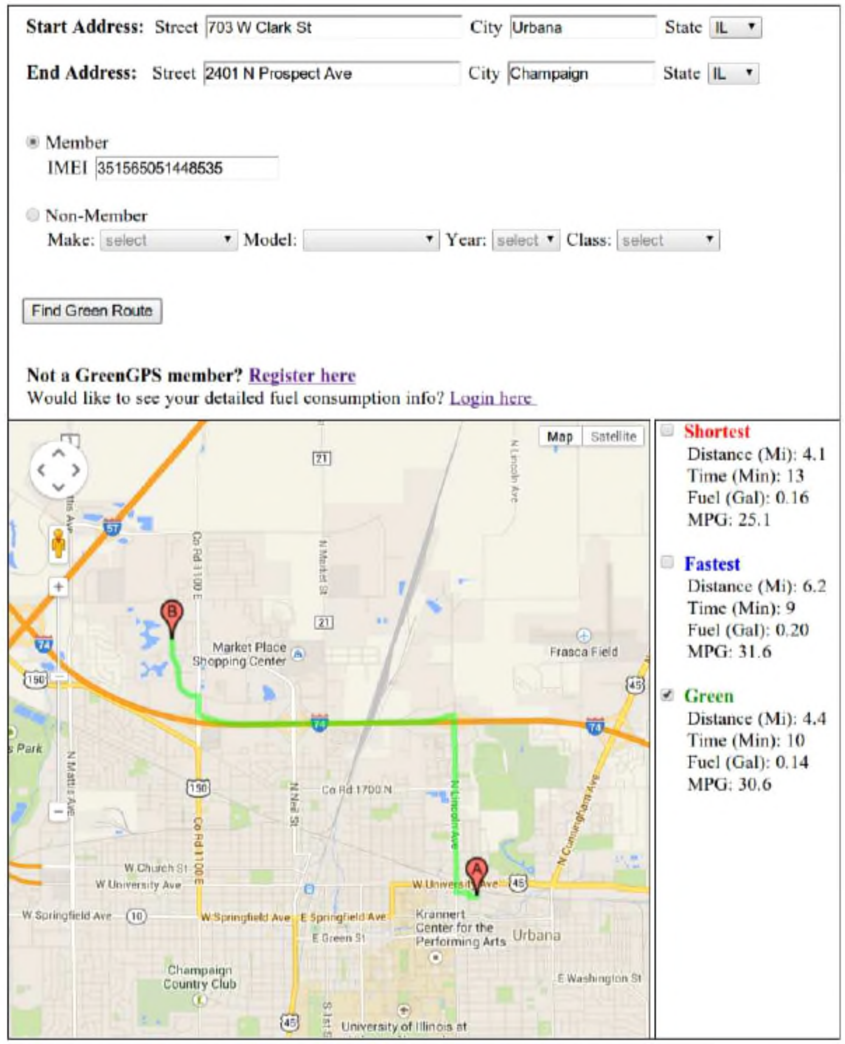

Generalization and Modeling

- Regression modeling:
  - Problem: one size does not fit all. Who says that Fords and Toyotas have the same regression model?
- Regression model per car?
  - Problem: Cannot use data collected by some cars to predict fuel consumption of others.
- Challenge: Must jointly determine both (i) regression models and (ii) their scope of applicability, to cover the whole data space with acceptable modeling error.

Accuracy Results

- The sampling regression cube improves prediction accuracy significantly

- A regression cube without model reduction is even worse than a single model!

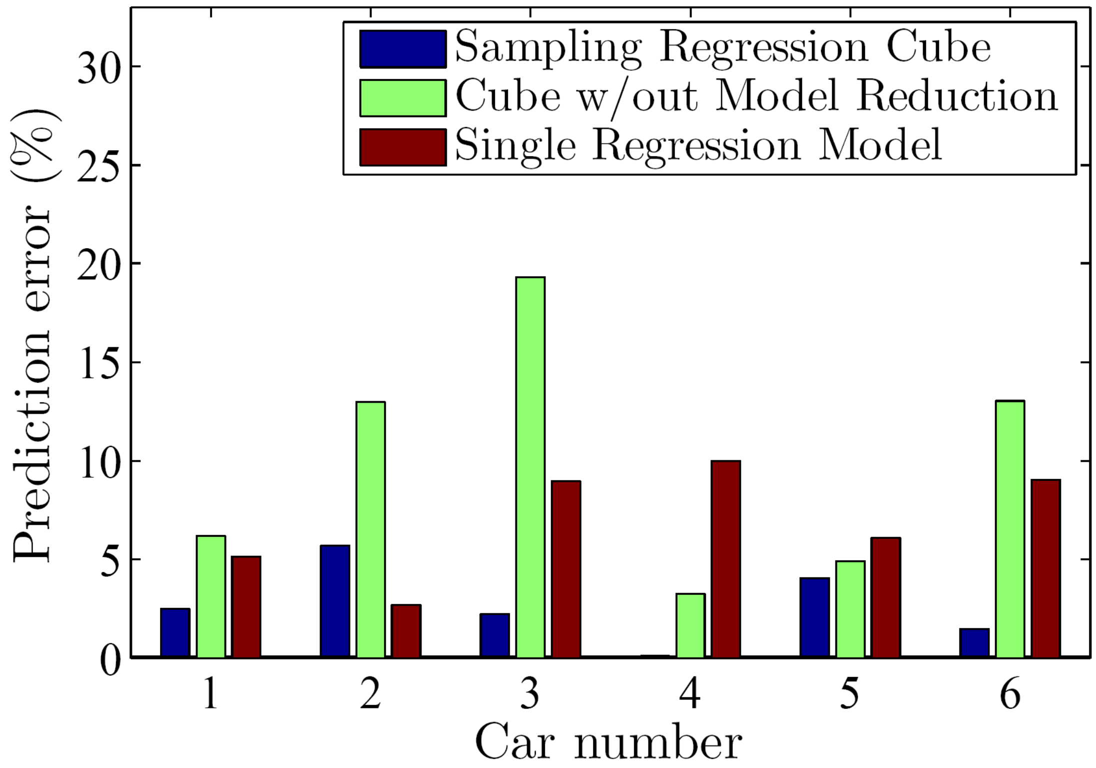

Finding Fuel-efficient Routes

- GreenGPS routes (green) may be different from fastest (blue), shortest (red), and Garmin EcoRoute (purple)

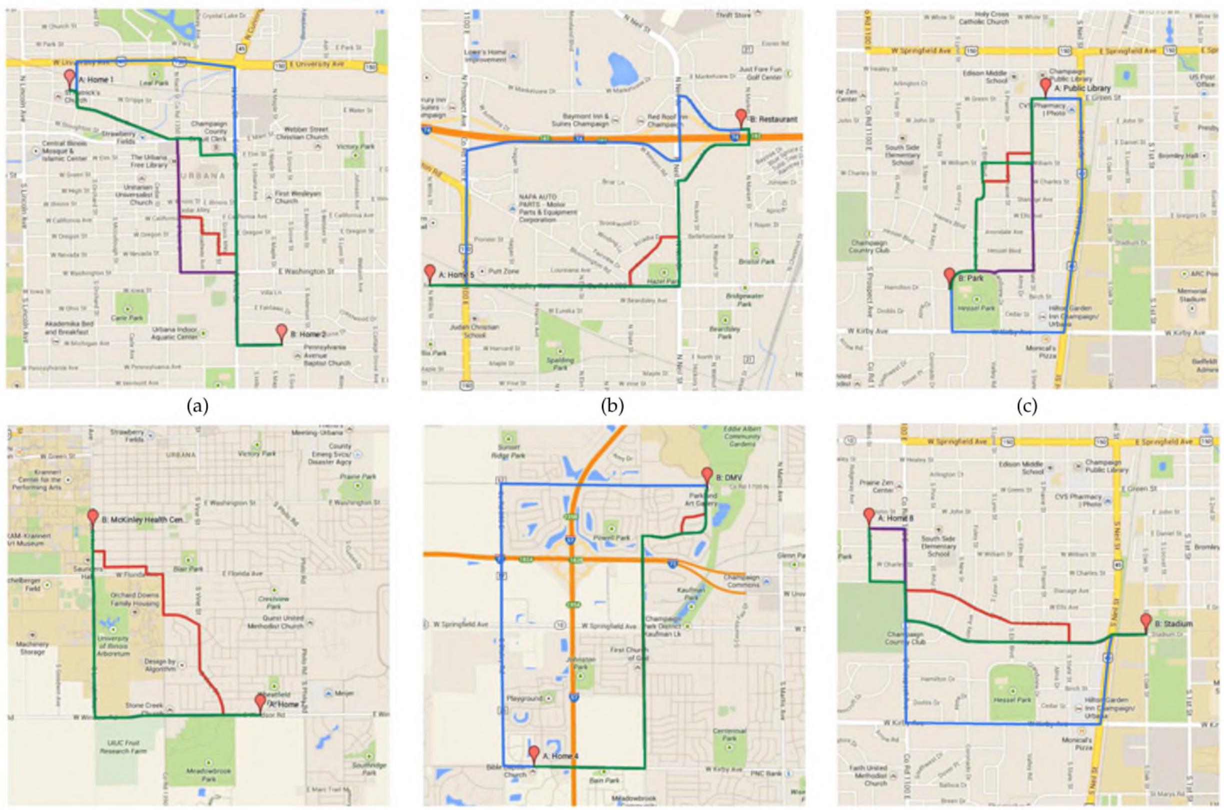

(a) (b) (c) (d) (e) (f)

Fuel Savings Evaluation

- GreenGPS savings compared to fastest route, shortest route, and Garmin EcoRoute

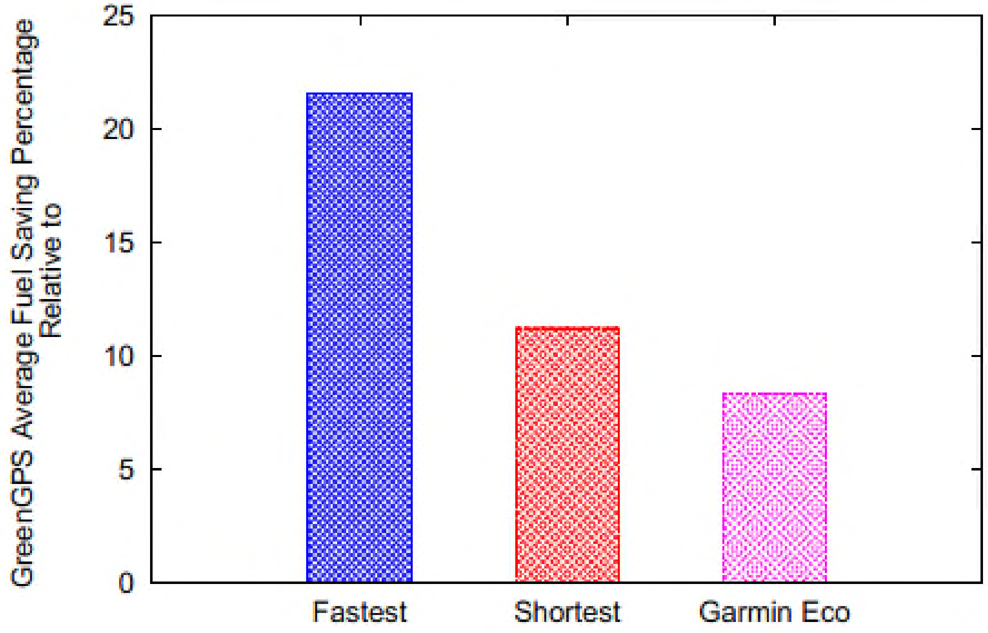

Choosing a Model to Fit the Data

- Must use simpler models when data is sparse

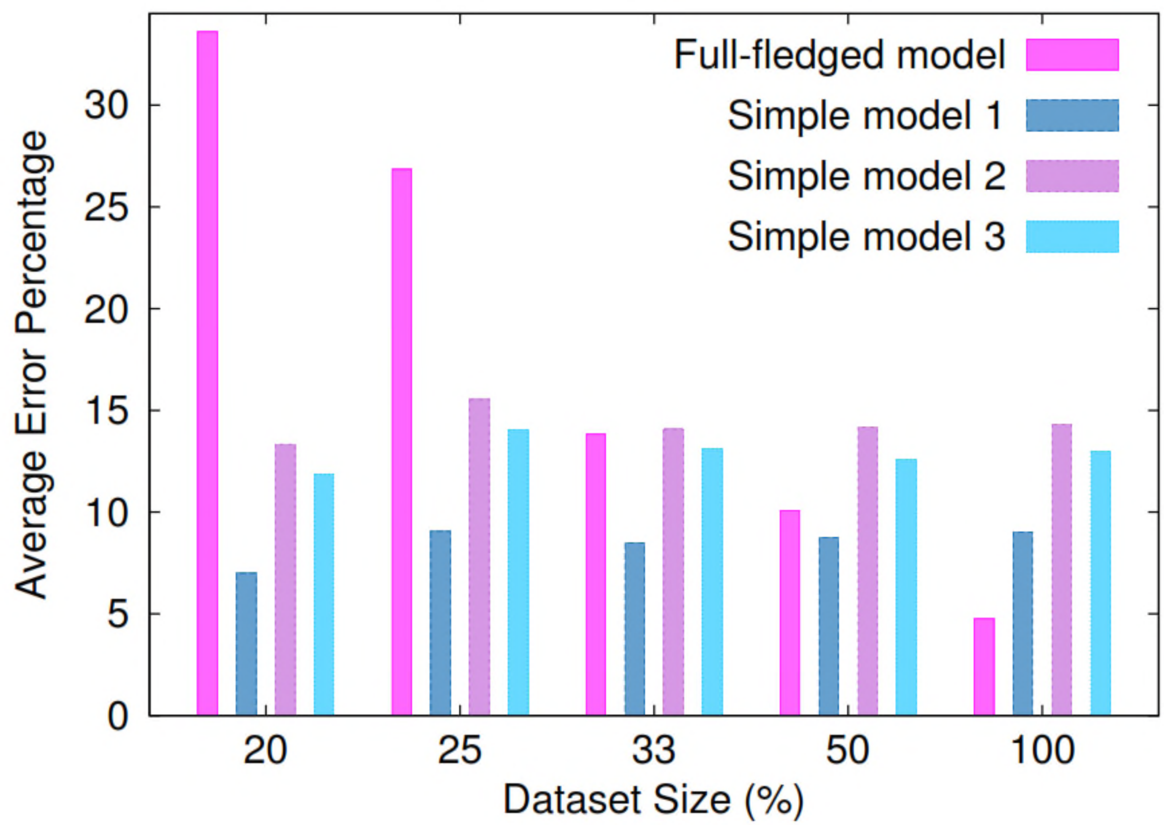

Challenge: The Transition Point

- How much data is enough to use a more detailed model?

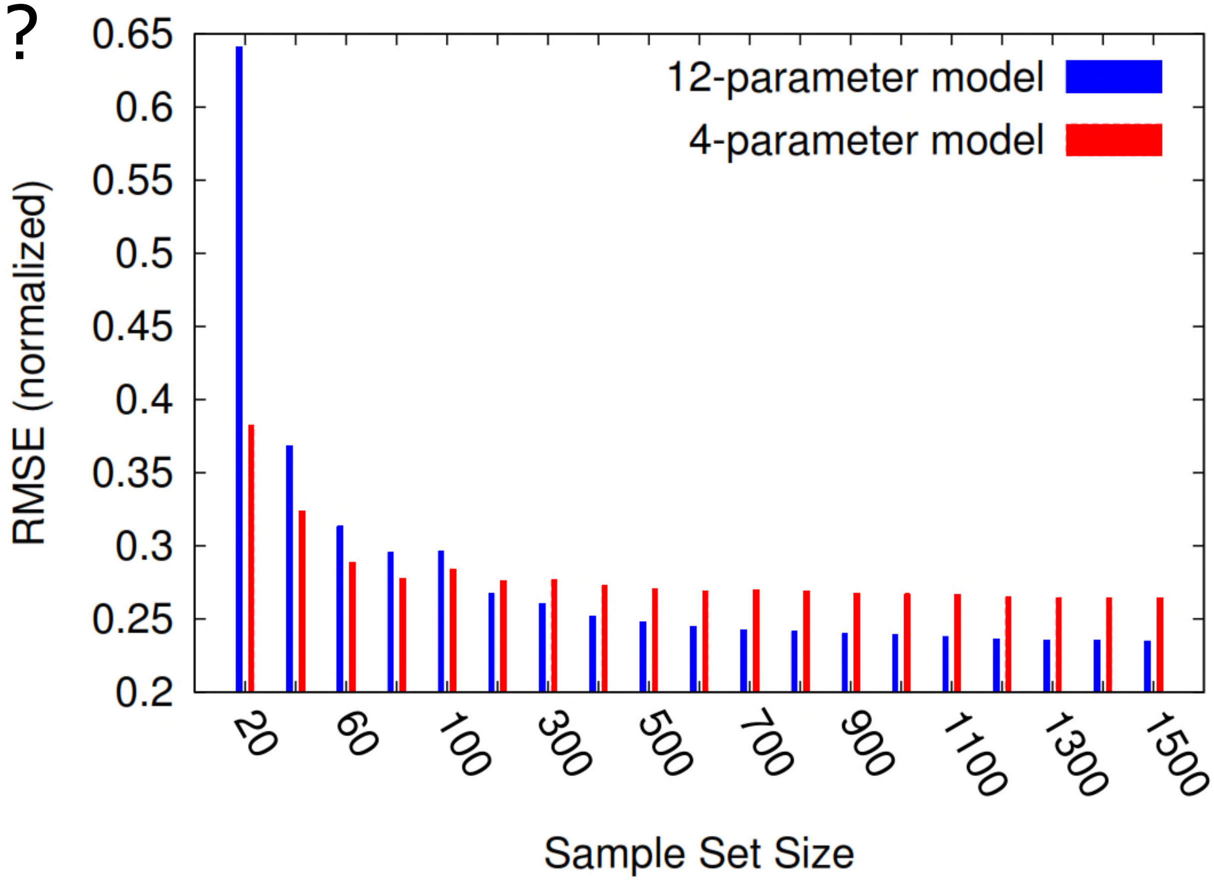

CityDrive: Predicting Traffic Light Schedules

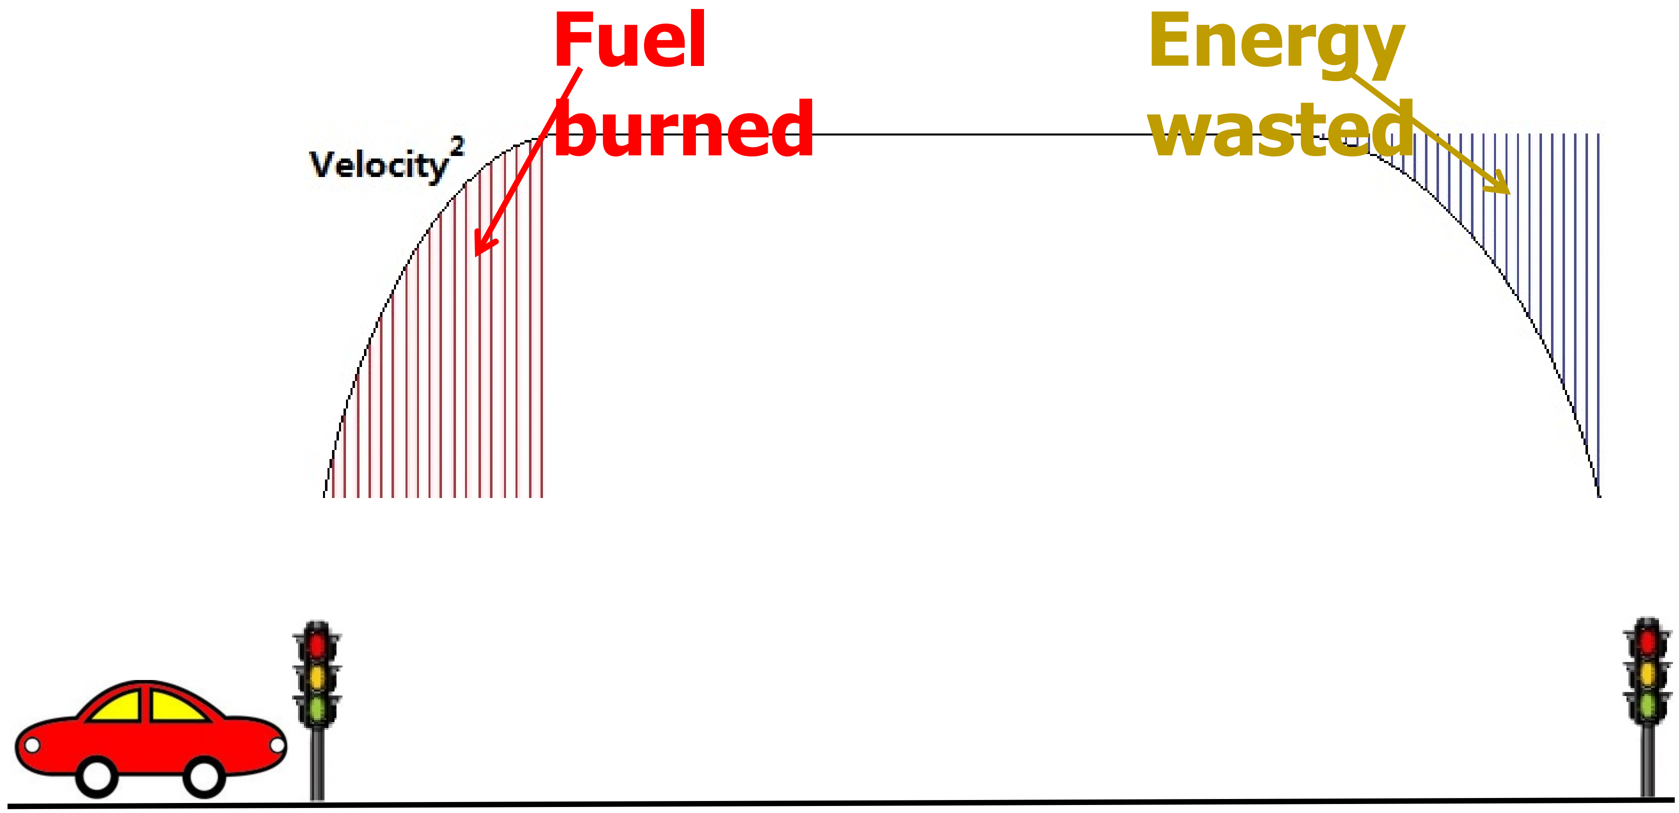

Slides by Yiran Zhao, Infocomm 2014 (with minor re-formatting)

Smartphone-based Prediction

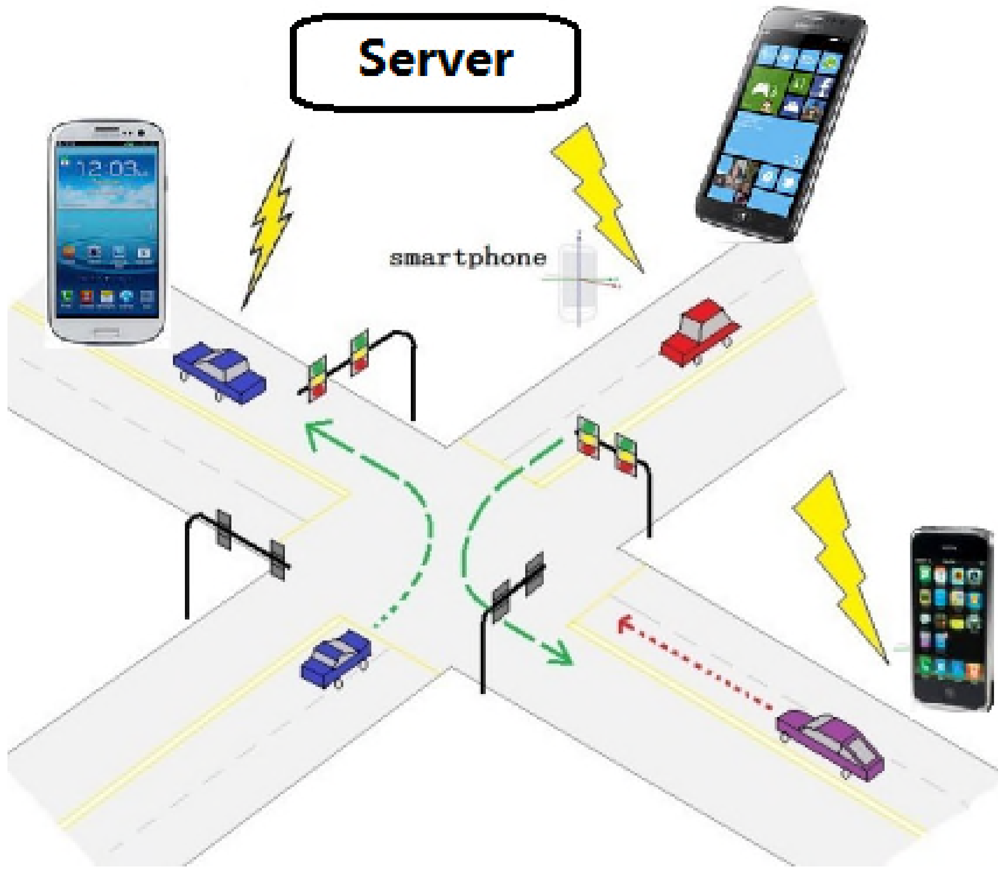

Slides by Yiran Zhao, Infocomm 2014 (with minor re-formatting)

Locating Intersections

- Car stops for more than 10 seconds before accelerating? → Intersection

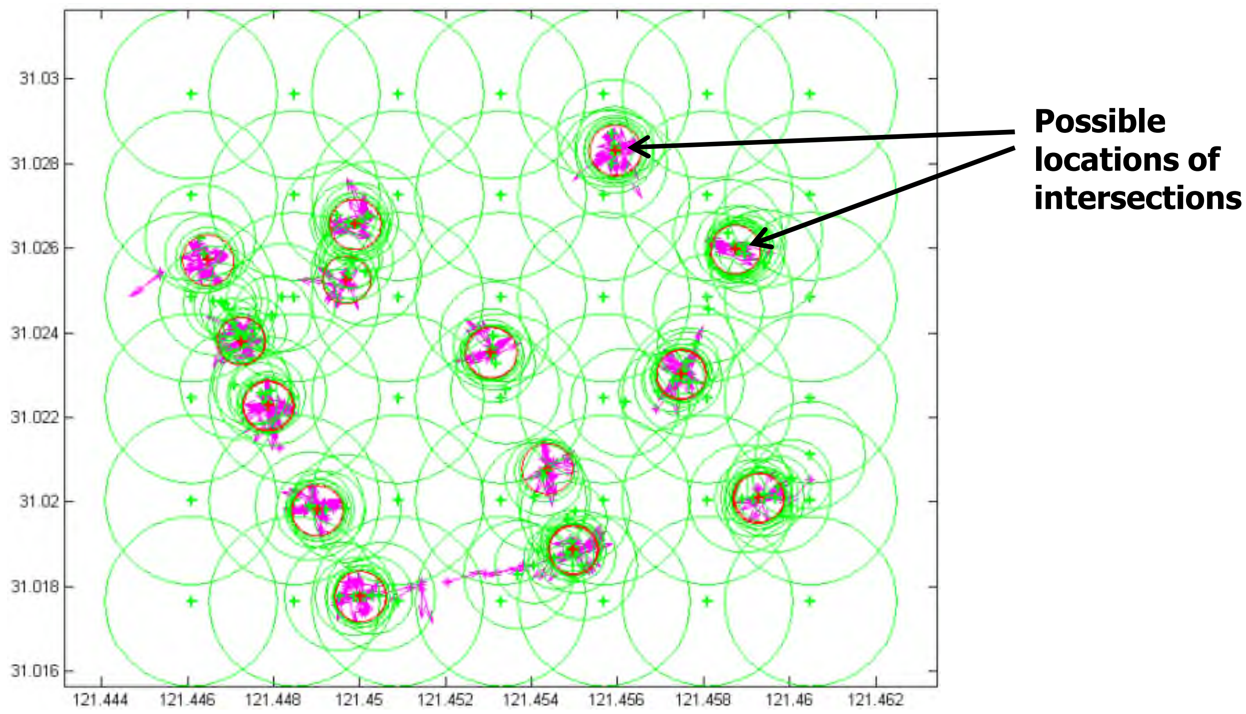

Slides by Yiran Zhao, Infocomm 2014 (with minor re-formatting)

Removing False Positives

- Group acceleration vectors with approximately the same direction into one cluster. Each cluster should represent one branch of the intersection.
- Intersections with less than 3 or greater than 5 clusters are removed.
- Intersections with outgoing vectors, are removed.

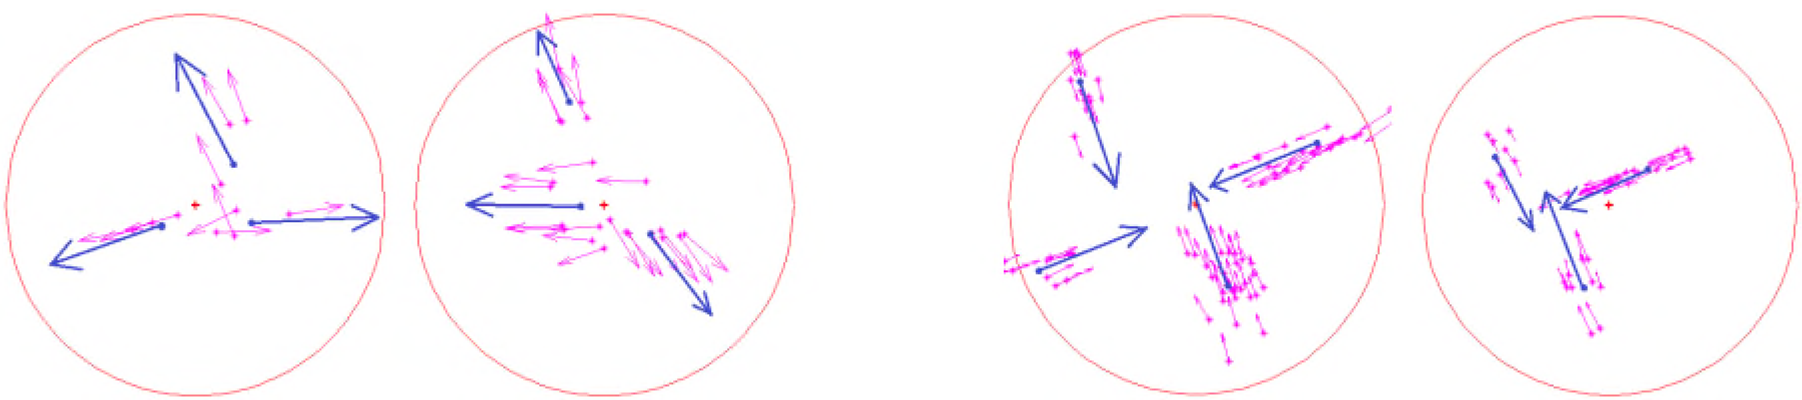

(a) Invalid intersection (b) Valid intersection

Slides by Yiran Zhao, Infocomm 2014 (with minor re-formatting)

Linking Intersections

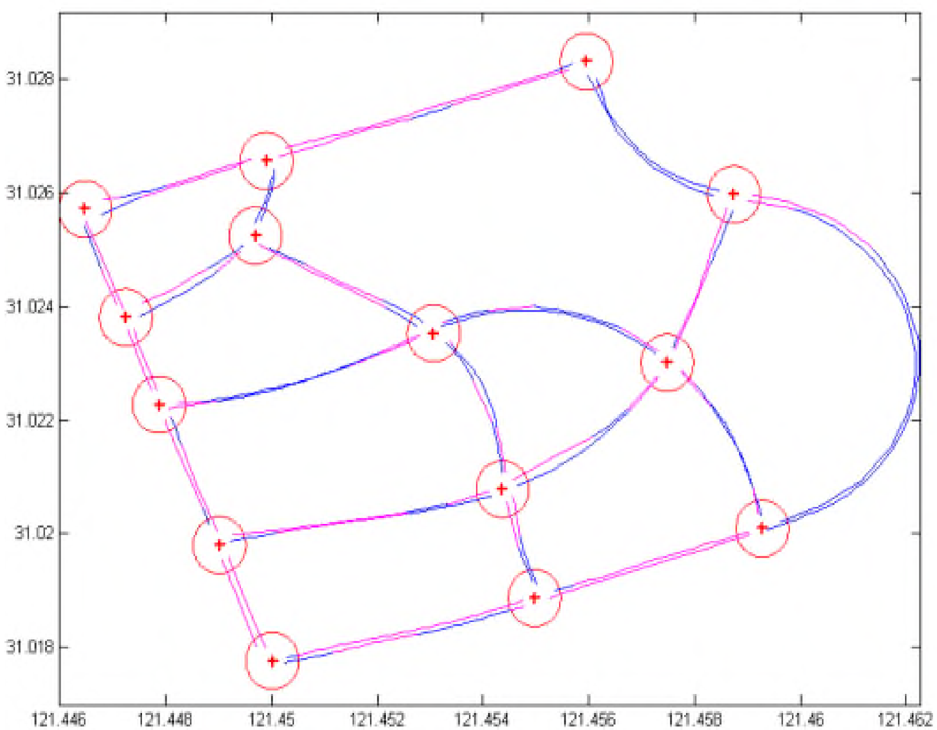

Slides by Yiran Zhao, Infocomm 2014 (with minor re-formatting)

Detect Cycle Length

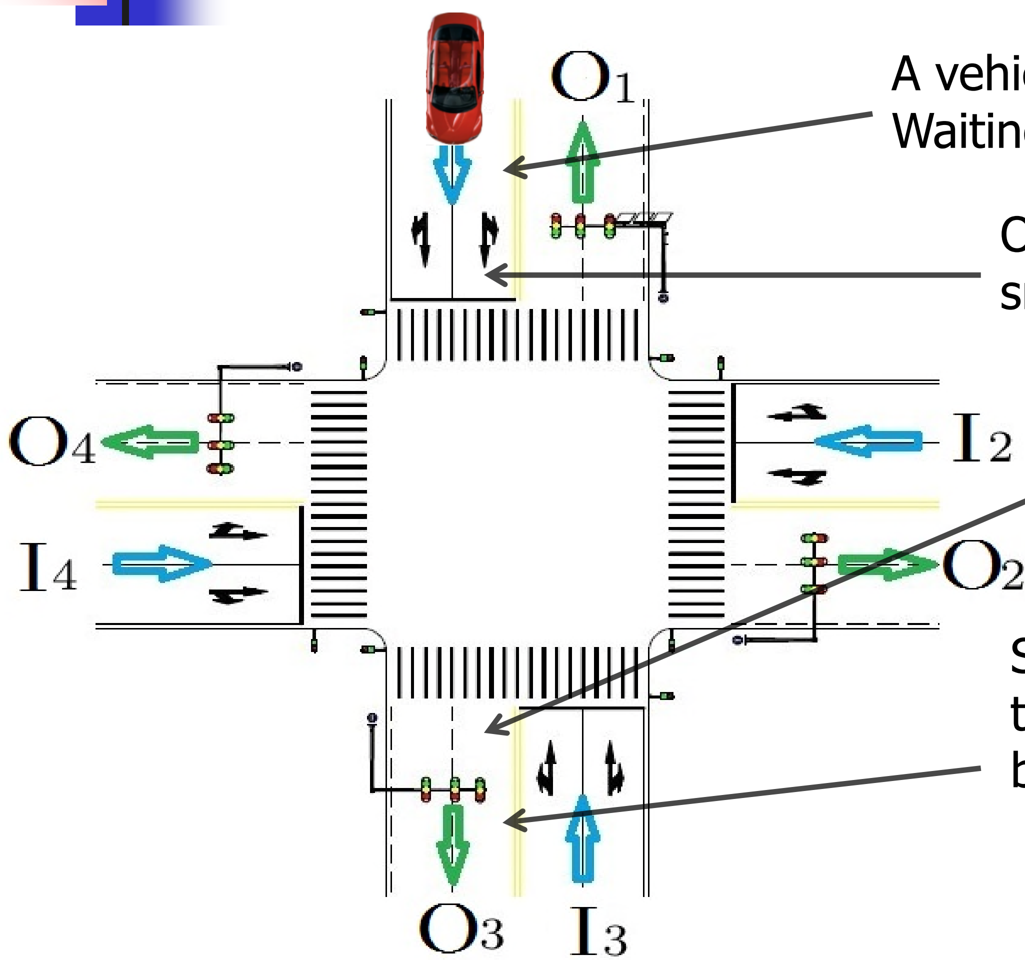

A vehicle stops and waits for green light. Waiting at time T0.

On detection of acceleration, smartphone records time T1.

On detection of entering another branch, smartphone records time T2.

Smartphone sends {wait time:(T1-T0); acc. time: (T2-T1); branch: I1,O3; ID} to the server.

Phase Sequence Inference

- Once the traffic signal cycle length Tc is obtained, the sequence of states $\{S_i, i = 1,2, \ldots N\}$ are to be determined.
- Ideally, and typically, there are four states after complete merging:
- And the above four states should happen in sequence:

$$S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow S_4 \rightarrow S_1 \ldots$$

- We give a new name to states that are in sequence and in closed loop: phase.

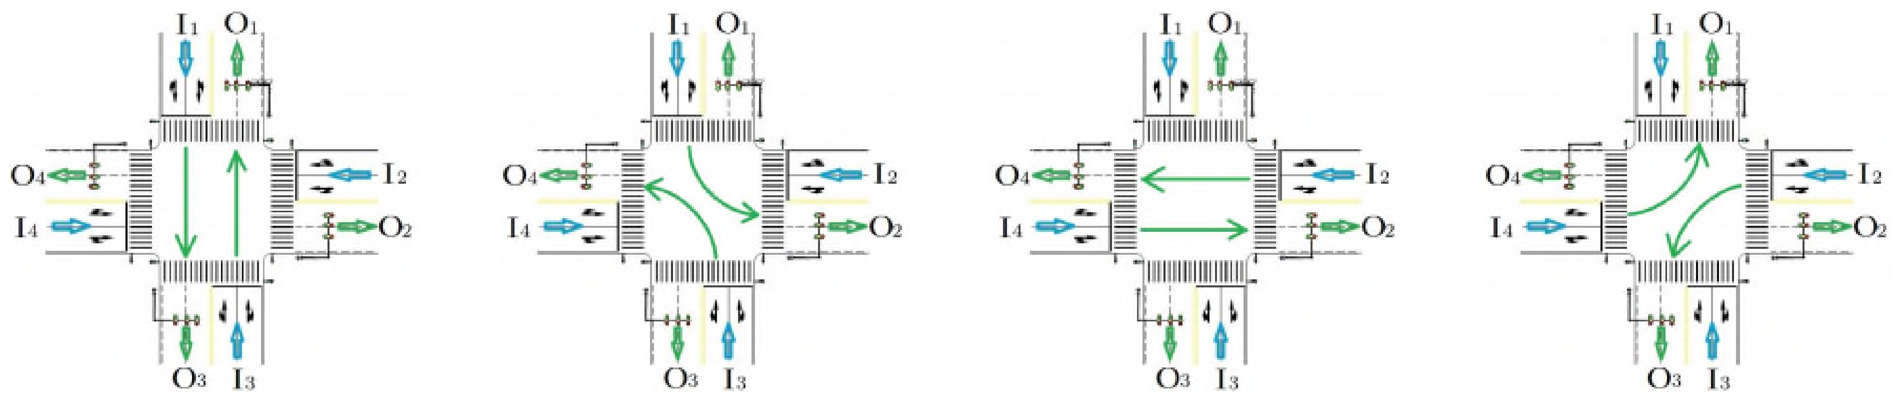

Slides by Yiran Zhao, Infocomm 2014 (with minor re-formatting)

Results

- Less getting stuck in red lights

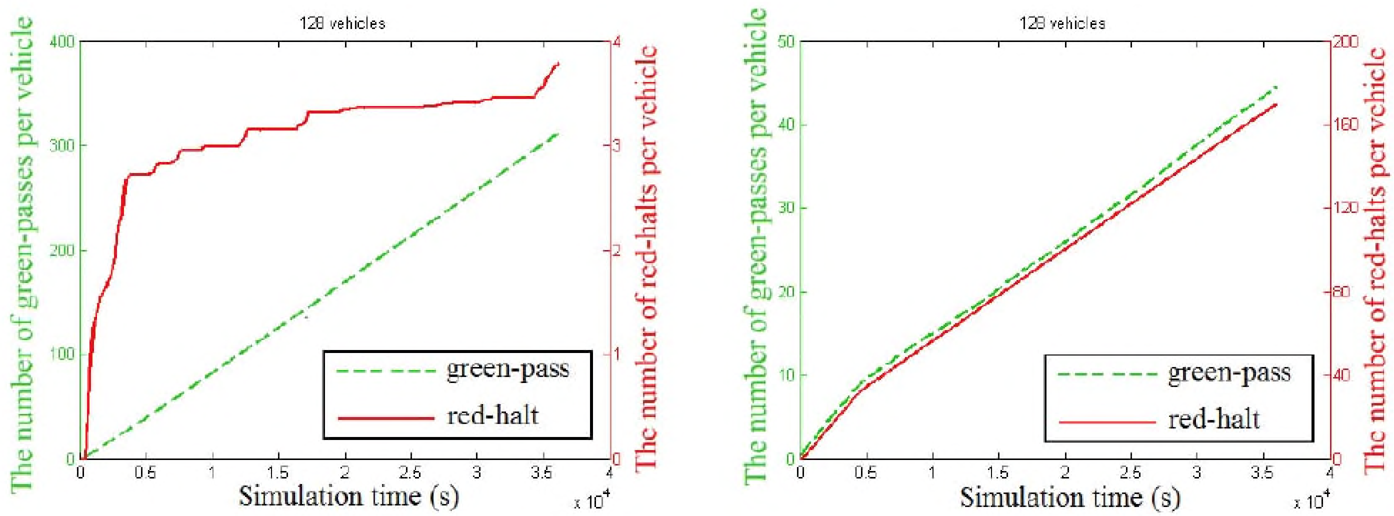

(a) With CityDrive (b) Without CityDrive

	With CityDrive	Without CityDrive
Energy consumption	13.9	33.7
Number of acceleration	0.09	1.67

Slides by Yiran Zhao, Infocomm 2014 (with minor re-formatting)

Traffic Light Phase Prediction

- Can we predict when the light is going to turn green?

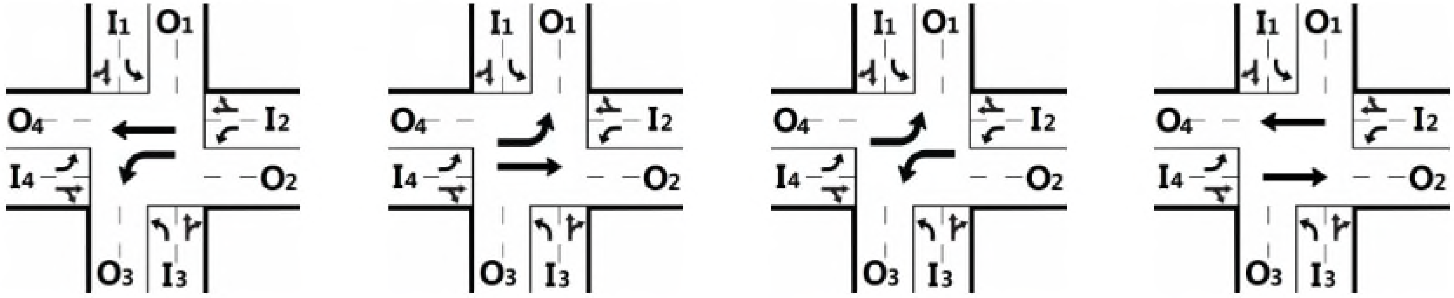

(1) Phase 1 (S_1). (2) Phase 2 (S_2). (3) Phase 3 (S_3). (4) Phase 4 (S_4).

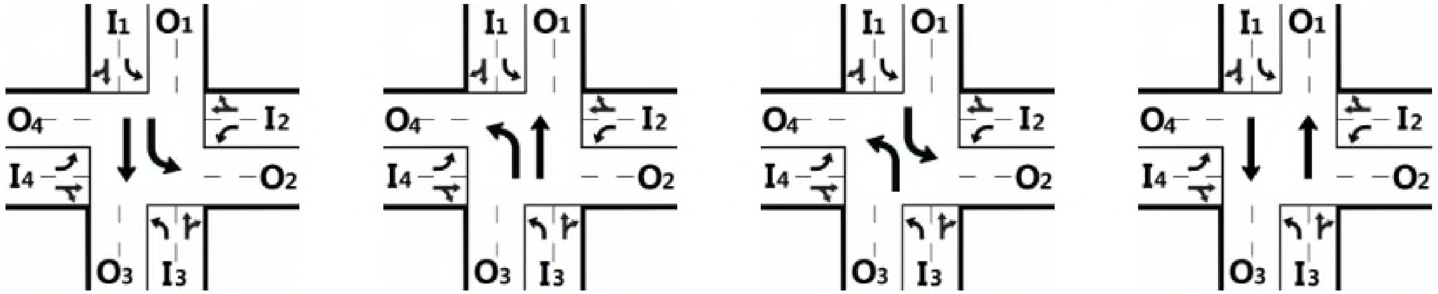

(5) Phase 5 (S_5). (6) Phase 6 (S_6). (7) Phase 7 (S_7). (8) Phase 8 (S_8).

Traffic Light Timing Prediction

- Maximum likelihood algorithm: reconstructs maximum likelihood estimate of phase durations and sequence given acceleration patterns of vehicles.

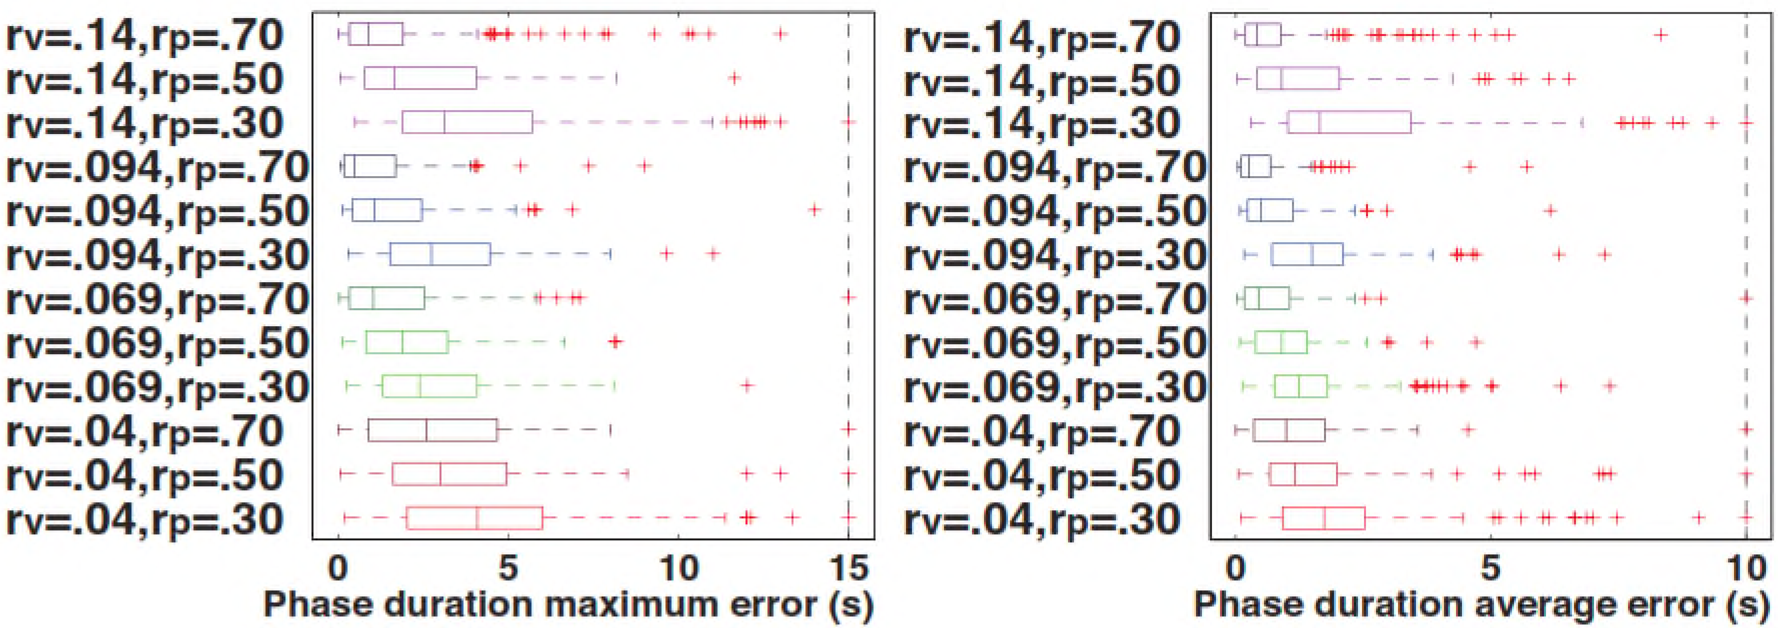

Energy Delay Trade-off

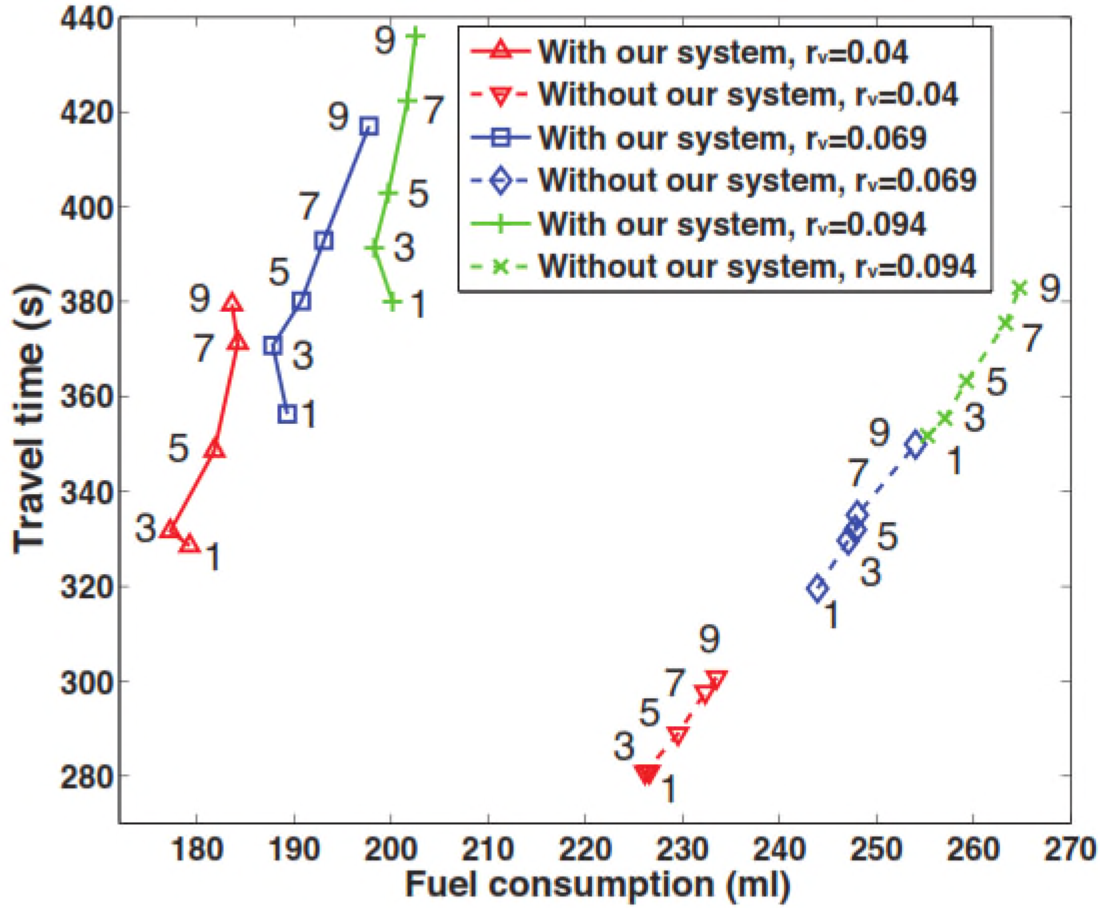

Another Example of Modeling

- Understanding urban pollution profiles

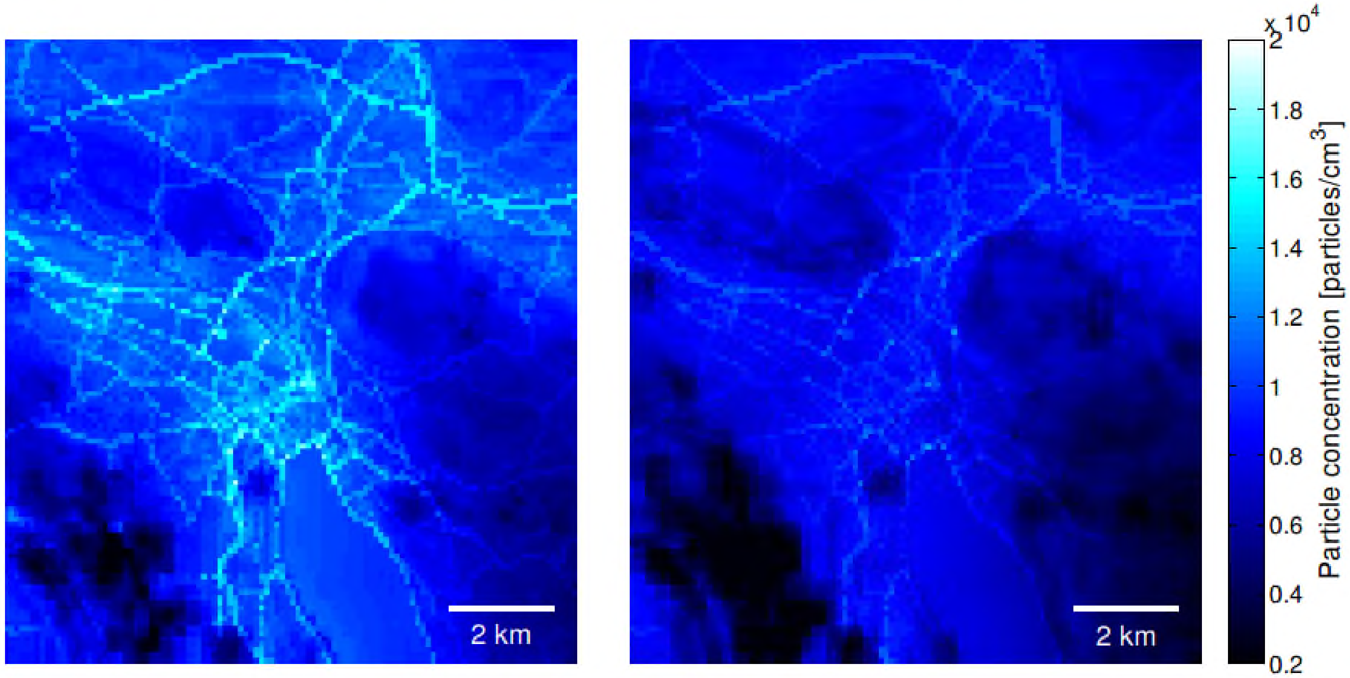

(a) Monday–Saturday. (b) Sunday.

Measuring Urban Pollution

- 10 sensor nodes on public buses

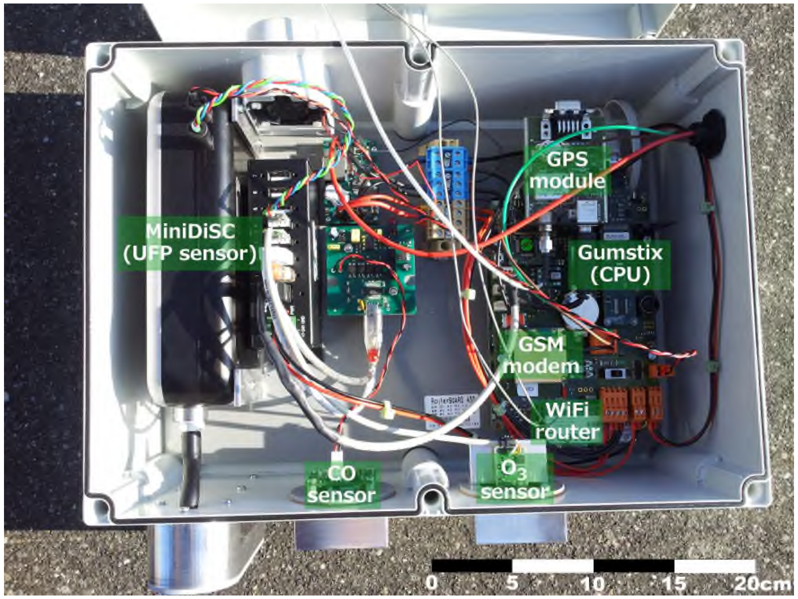

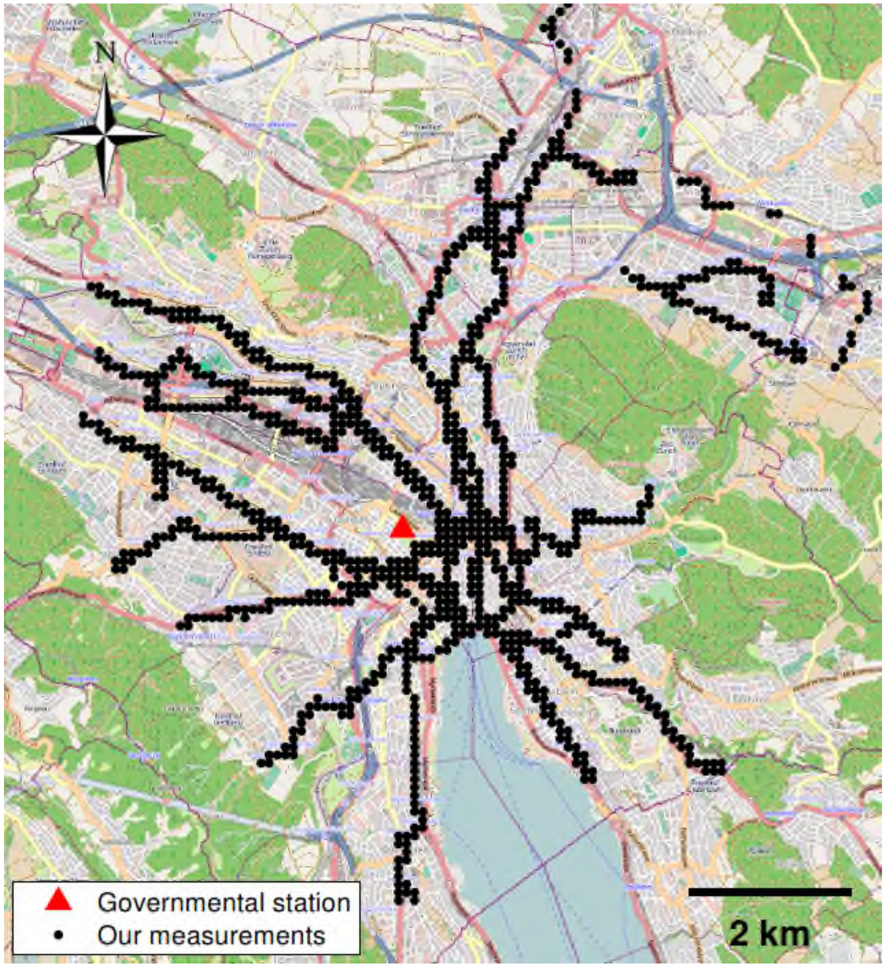

Selecting Regression Parameters

- What factors might impact pollution?

Variable [unit]	Variable [unit]
Population [inhabitants/ha]	Industry [industry buildings/ha]
Building height [floor levels/ha]	Heating [oil and gas heatings/ha]
Terrain elevation [average m/ha]	Road type [busiest road type/ha]*
Distance to next road [m]	Distance to next large road [m]†
Terrain slope [average degree/ha]	Terrain aspect [average degree/ha]
Traffic volume [vehicles per day/ha]	Distance to next traffic signal [m]

*Five road types: residential, tertiary, secondary, primary, and freeway.
†Road types classified as large: secondary, primary, and freeway.

Selecting Regression Parameters

- What factors might impact pollution?

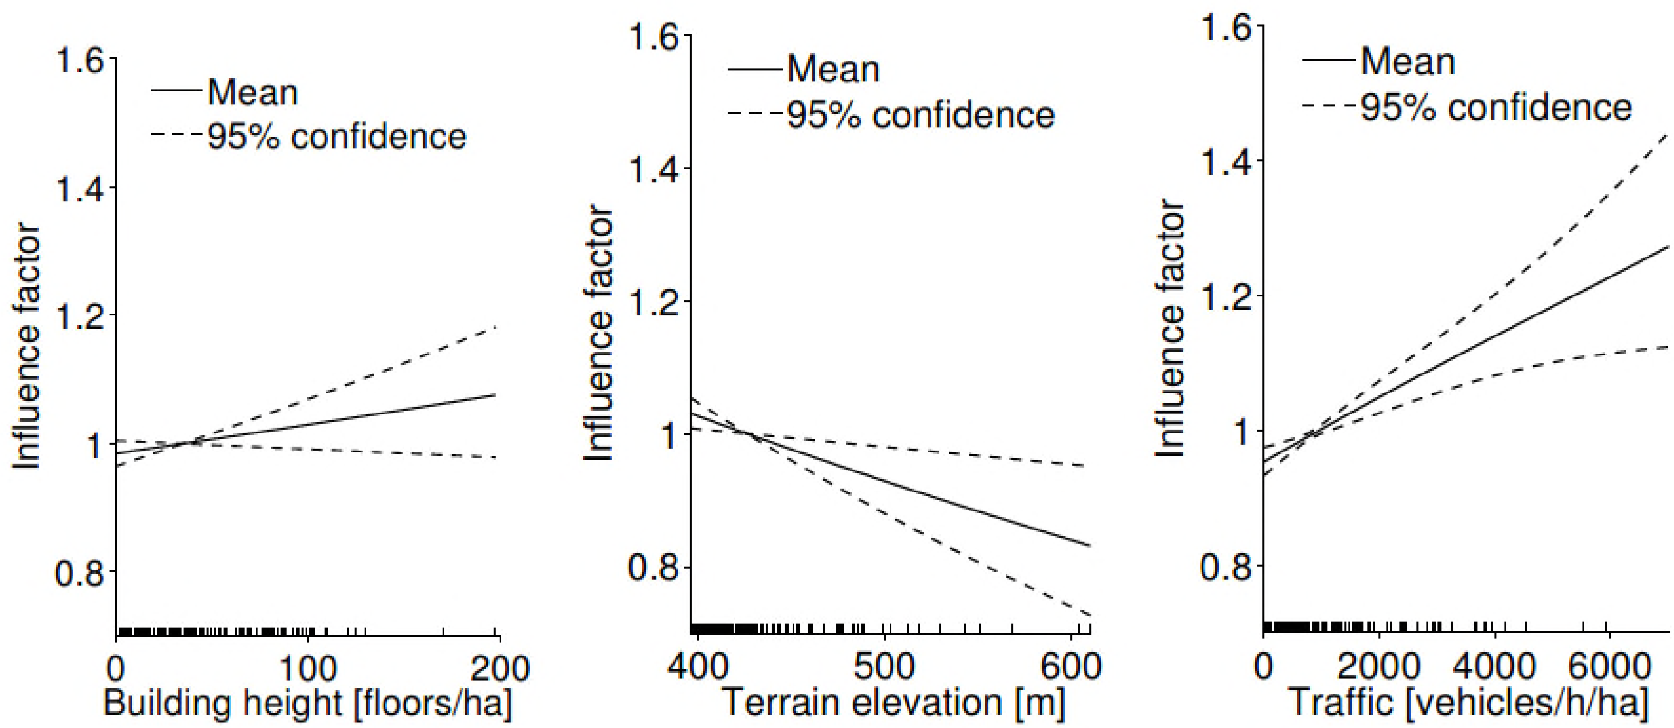

Prediction Results

- Measured versus predicted pollution level

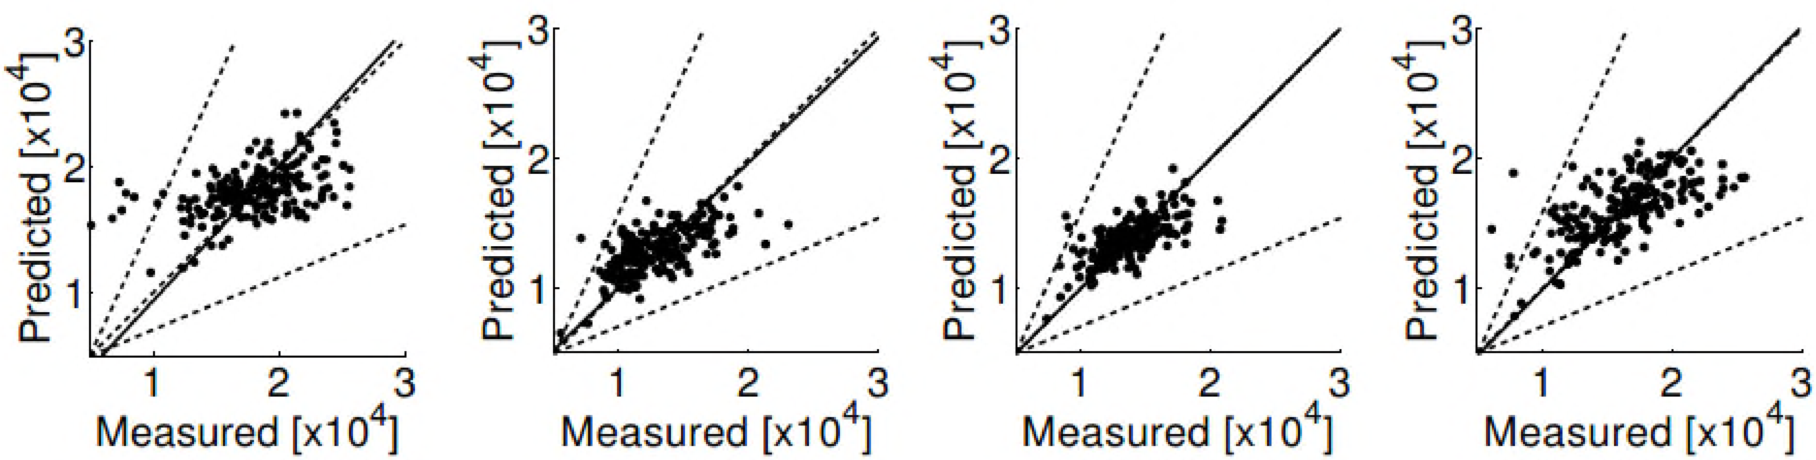

(a) Winter. (b) Spring. (c) Summer. (d) Fall.

Prediction Results

- Quantifying the Error

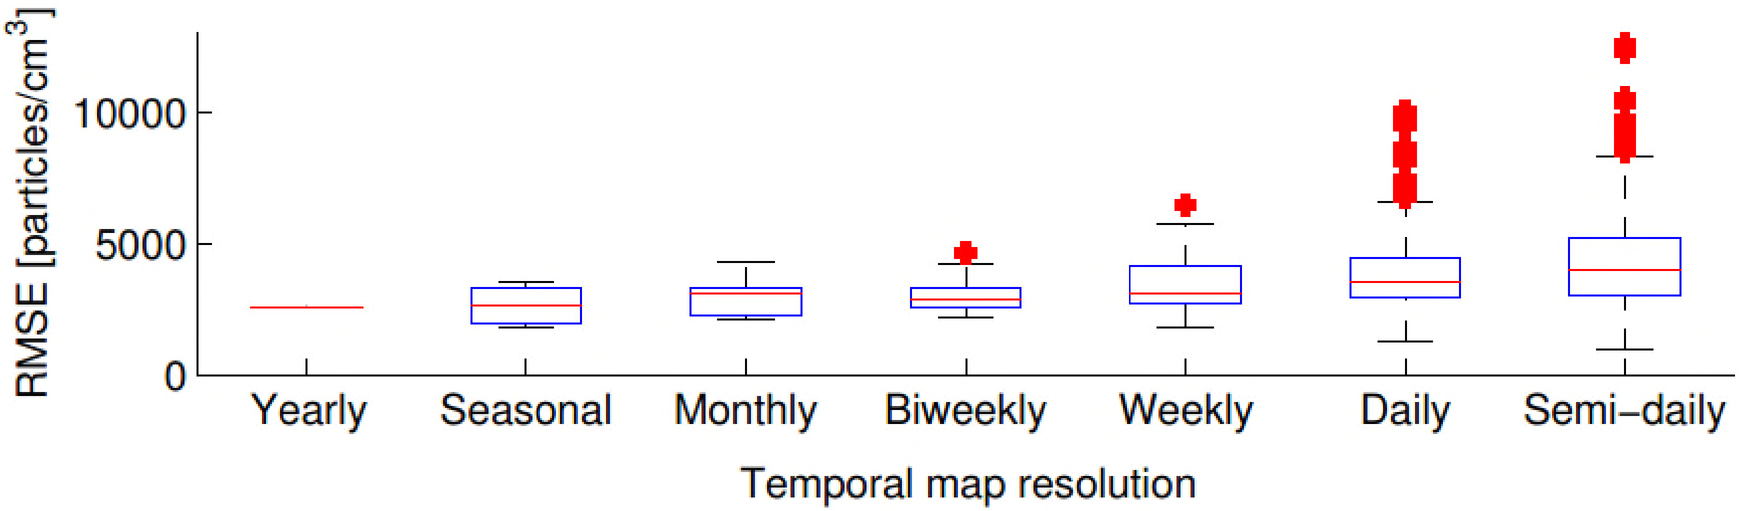

Combining Models and Statistics

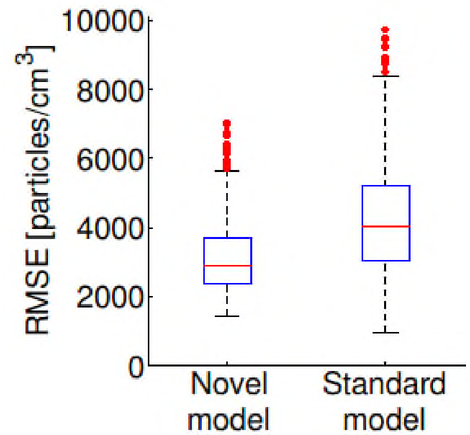

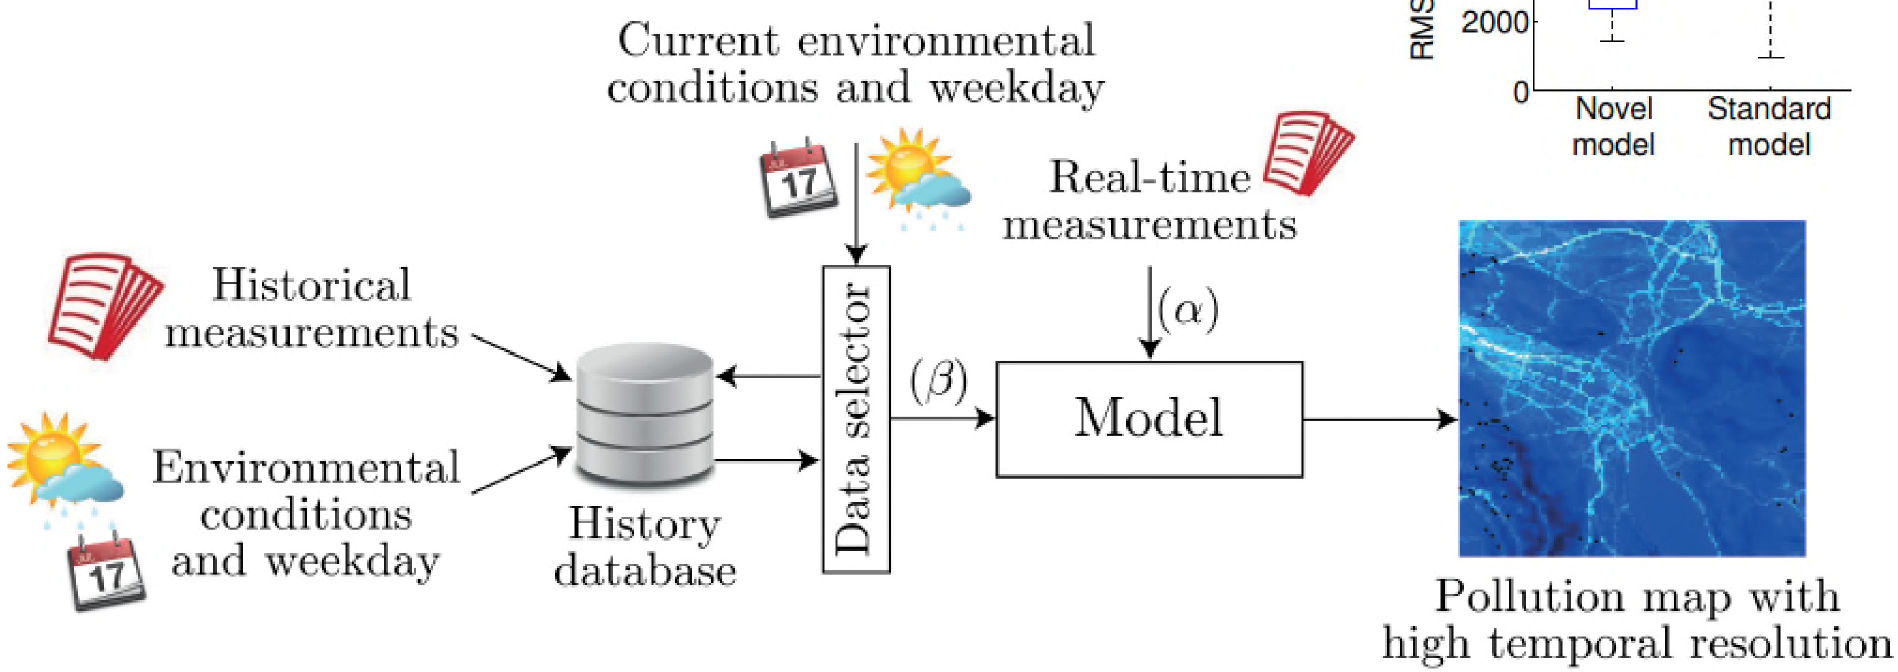

The Electric Grid:
Yesterday and Today

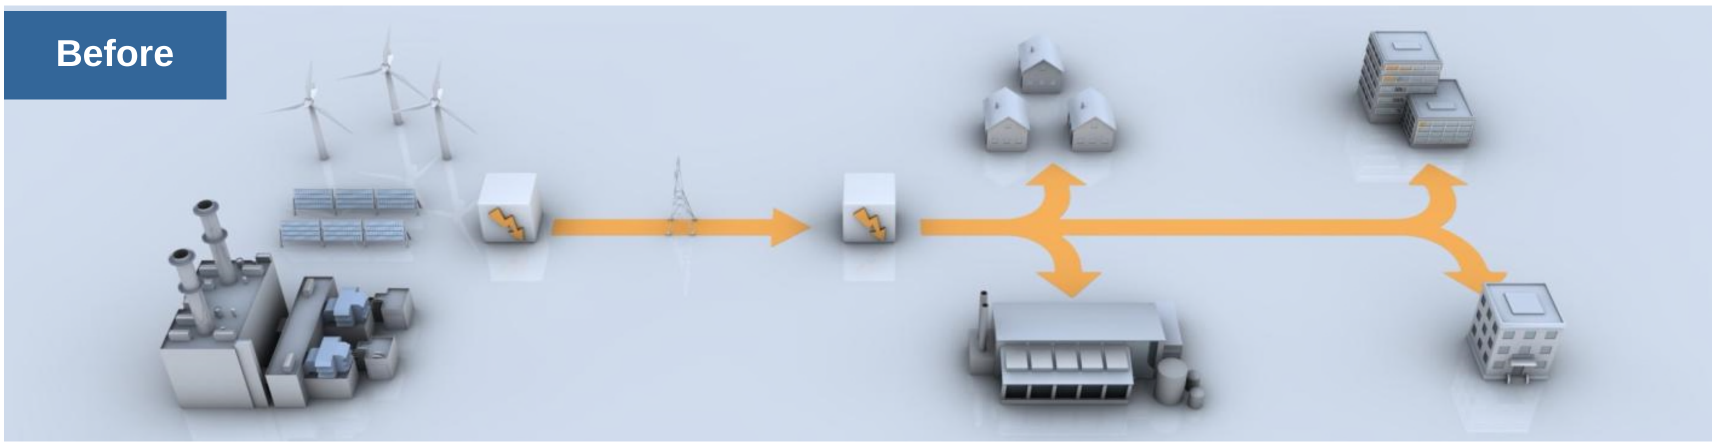

- One-way limited communication
- One-way power flow
- Centralized generation
- No electric vehicles
- Few sensors and analog control
- Little to no consumer choice
- Reactive maintenance
- Limited usage transparency

The Electric Grid:
Yesterday and Today

- **Bi-directional and instantaneous communication and metering**
- **Bi-directional power flow**

- Pervasive monitoring and digital control
- Self-monitoring & high visibility
- Many consumer choices

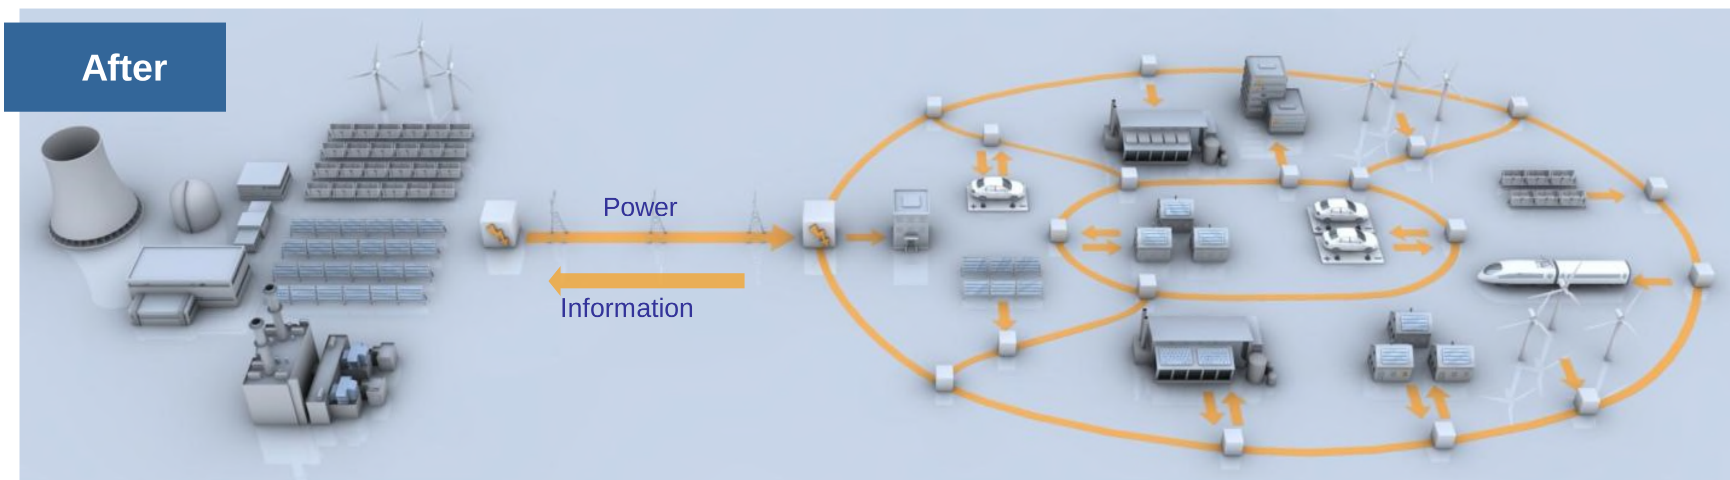

TODAY The transition has begun, with peak-demand management (demand response, ILM), and dynamic pricing (e.g. critical peak pricing programs)

Supply-following Loads

Optimizing

- Energy storage
- Pricing
- CO_2 reduction
- Energy efficiency
- E-car integration

Intelligent Load Management (ILM)

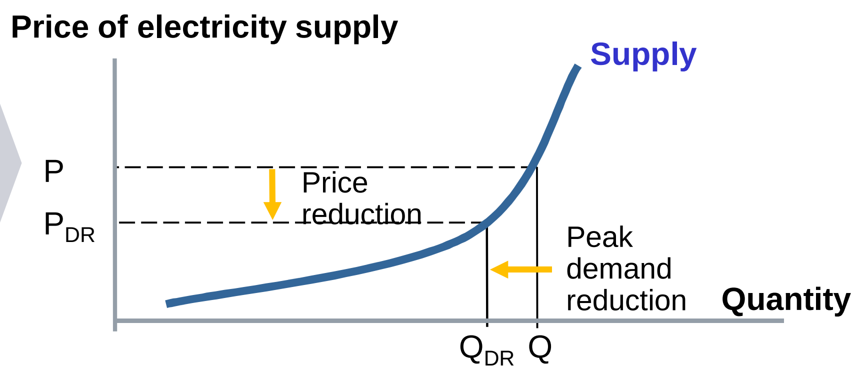

Balancing the grid

- Avoid investments in new power plants
- Increase power quality
- Integrate volatile renewable energy
- E-Car charging

Smart Consumption

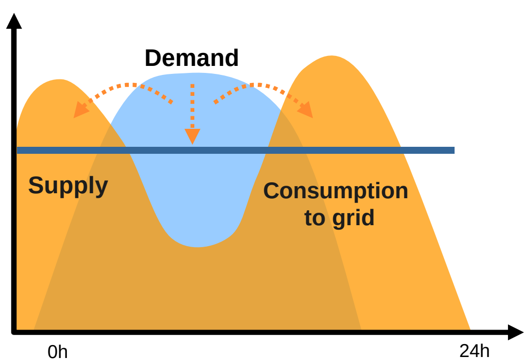

Smart Buildings

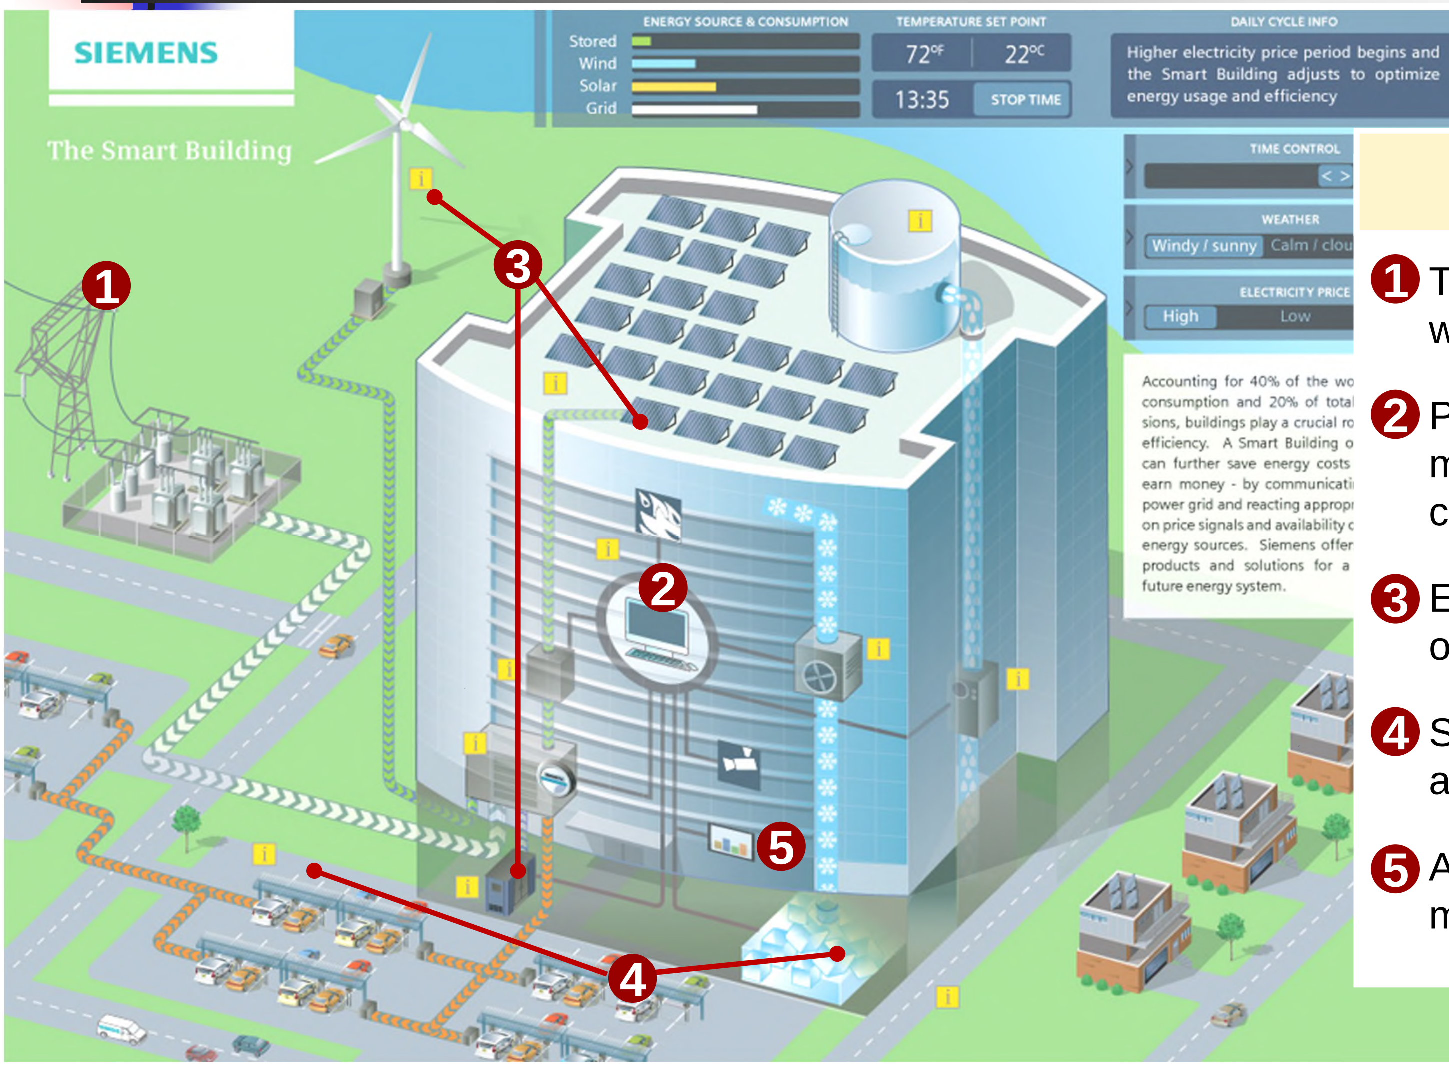

Applications

1. Two-way communication with utilities
2. Proactive energy management / smart consumption
3. Energy sources with onsite generation assets
4. Storage capacity for added flexibility
5. Active carbon management

Why do we need Demand Response?

An oversold or undersold flight is similar to the electrical grid at capacity….

Every Person who gives up their seat is a "Negawatt" and will receive compensation for giving up a seat.

In other words the utility company will pay you to reduce your load during peak demand

Intelligent Load Management

- Leverages existing BAS equipment to generate cash payments through automated load management

- Allows building operators to participate in Demand Response, Critical Peak Pricing and Smart Grid programs through local utilities

- Balances multiple factors:
  - Corporate standards
  - Efficiency
  - Financial
  - Site conditions

- The technology leader in multi-site load aggregation with ***proven financial results***

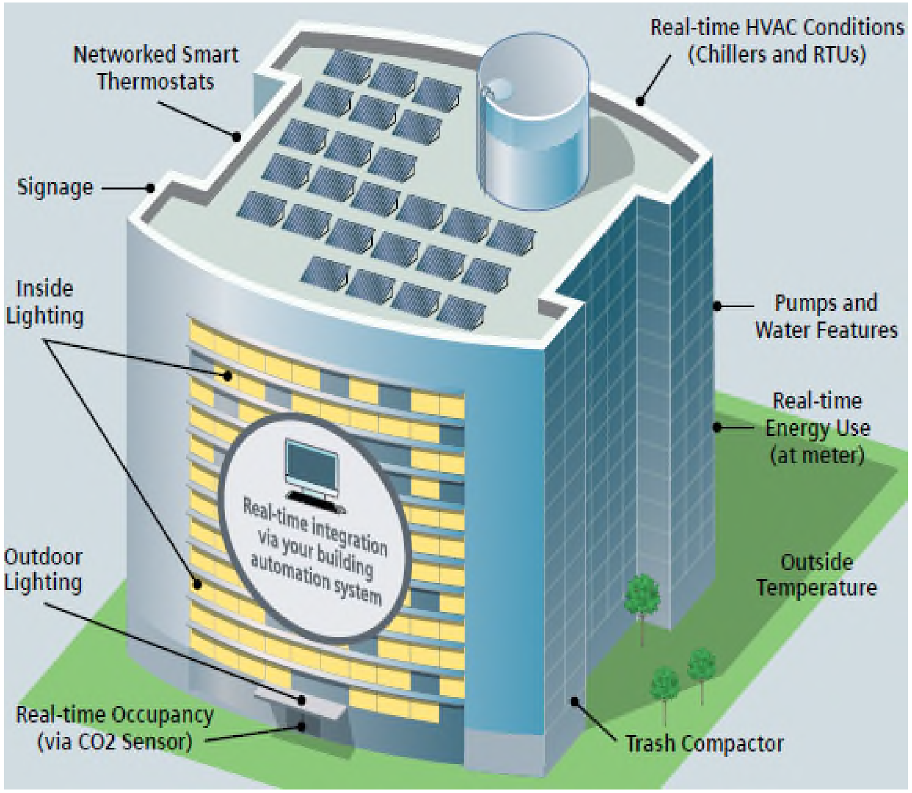

Example: Siemens Demand-Response System

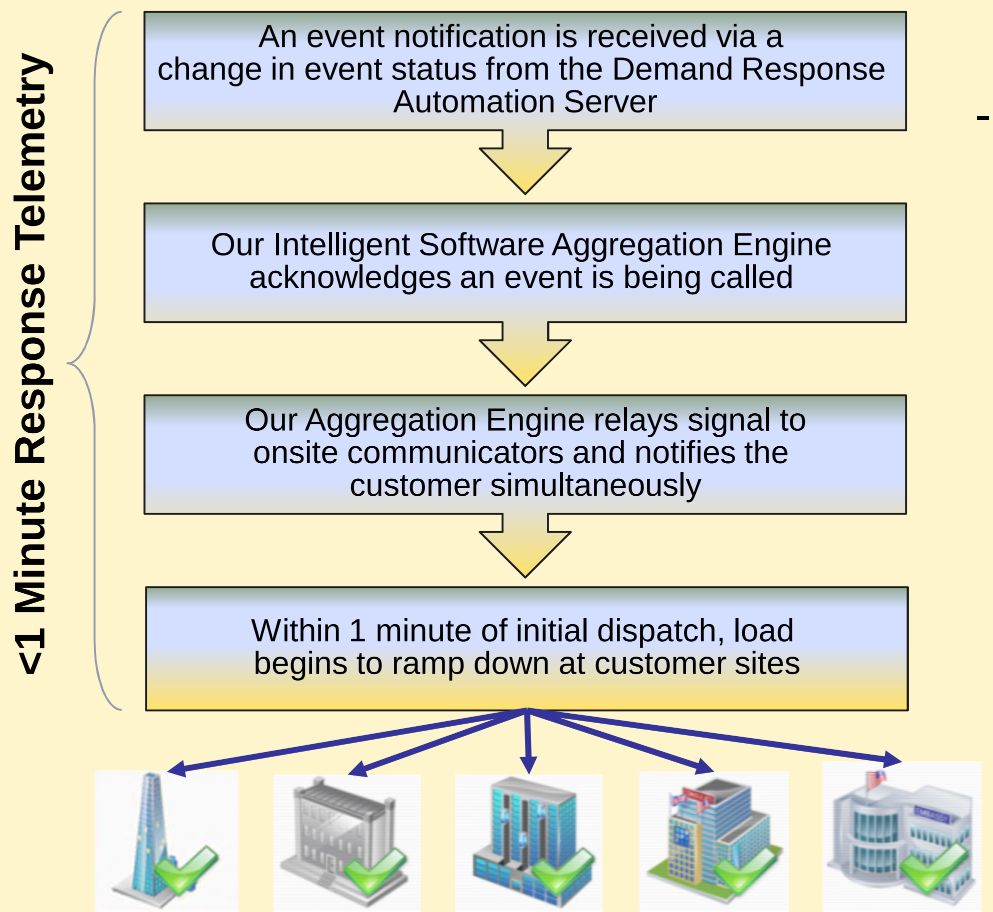

- **Approach facilitates reliable participation in short notice programs.**

Demand Response

- How much capacity/flexibility is available to respond to supply variations?
  - A study of HVAC systems

The Dataset

- Examples of dataset measurements
- Data used to build model relating external temp, internal temp, set point, AC power, etc.

Table 1. Data field characteristics

Data Field	Resolution	Sampling Rate
Temperature setpoint, $T_{set}(t)$	1 $[^oF]$	5 mins
Internal Temperature, $T_{int}(t)$	1 $[^oF]$	
Duty Ratio, $d(t)$	0.001	
Average Power, P_{avg}	0.01 $[W]$	
External Temperature, $T_{ext}(t)$	0.1 $[^oC]$	1 hour
Solar Insolation, $\phi_{sol}(t)$	0.01 W/m^2	

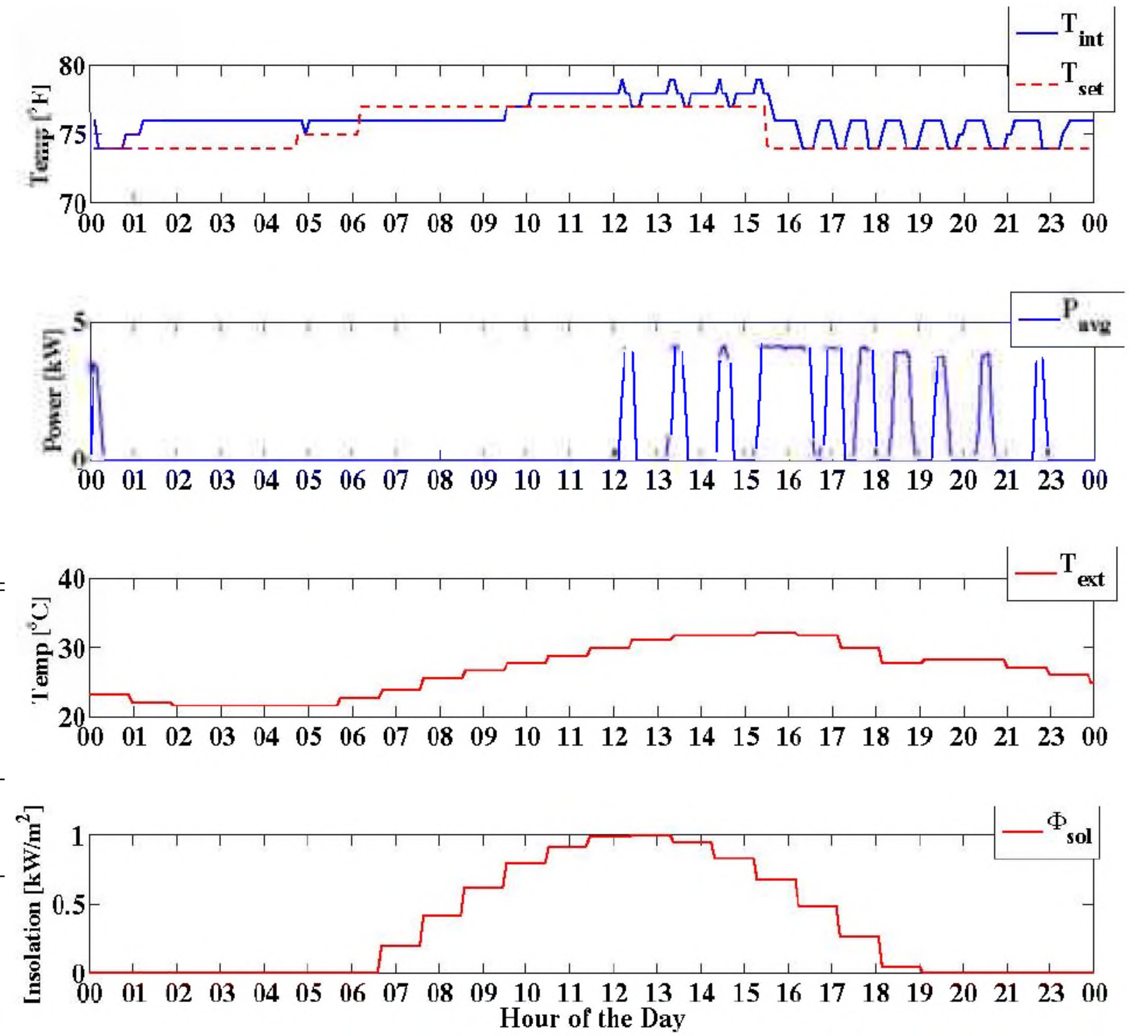

Model Validation

- Predicting how internal temperature will change in response to large set point changes

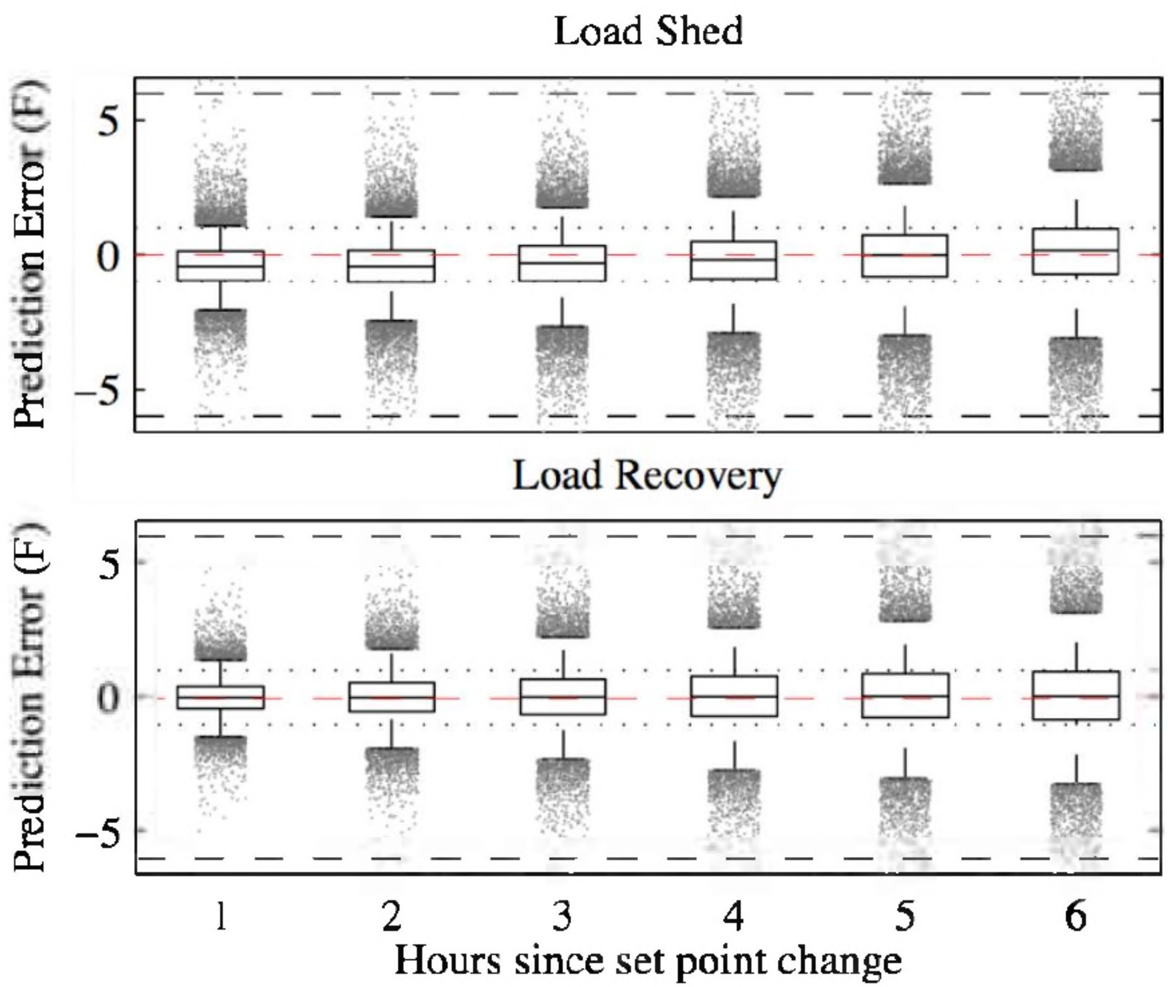

Demand Response Potential

- Adapting temperature set point to external supply variations

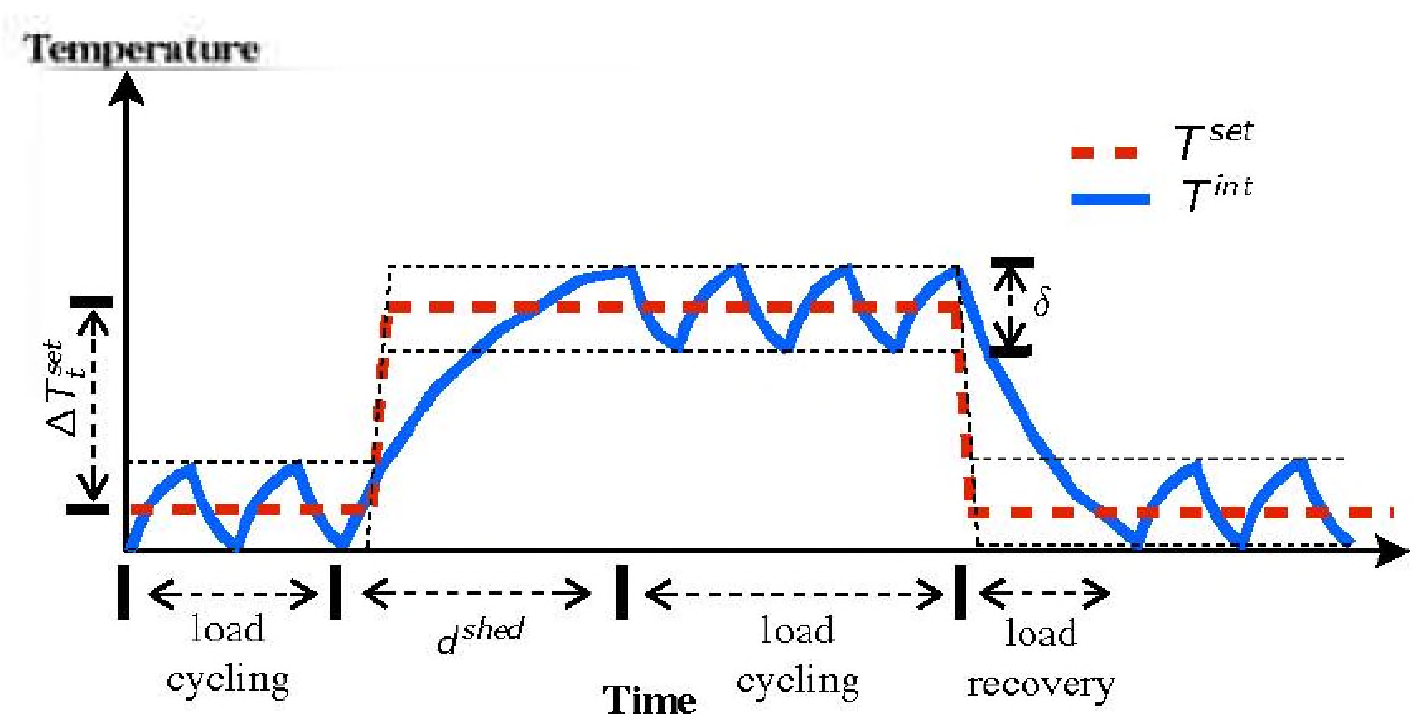

Demand Response Potential

- Several "What If" Scenarios

Table 2. Characteristics of different scenarios and case studies

Case Study	External Temp.	Setpoint Profile	Setpoint Change $\Delta T^{set}_{t_{event}}$	Scenarios	Event Start Hour, t_{event}, [hour of the day]	DR event Duration, D_{DR}, [hours]
CS-1	82^oF, Constant	76^oF, Constant	2^oF	A	12	[1,2,3,4] hours
			2^oF	B	15	
			4^oF	C	12	
			4^oF	D	15	
CS-2	Measured, Thu, 21/06/2012	Avg. Weekday Profile	2^oF	A	12	[1,2,3,4] hours
			2^oF	B	15	
			4^oF	C	12	
			4^oF	D	15	
CS-3	Measured, Sun, 17/06/2012	Avg. Weekend Profile	2^oF	A	12	[1,2,3,4] hours
			2^oF	B	15	
			4^oF	C	12	
			4^oF	D	15	

Impact on Power

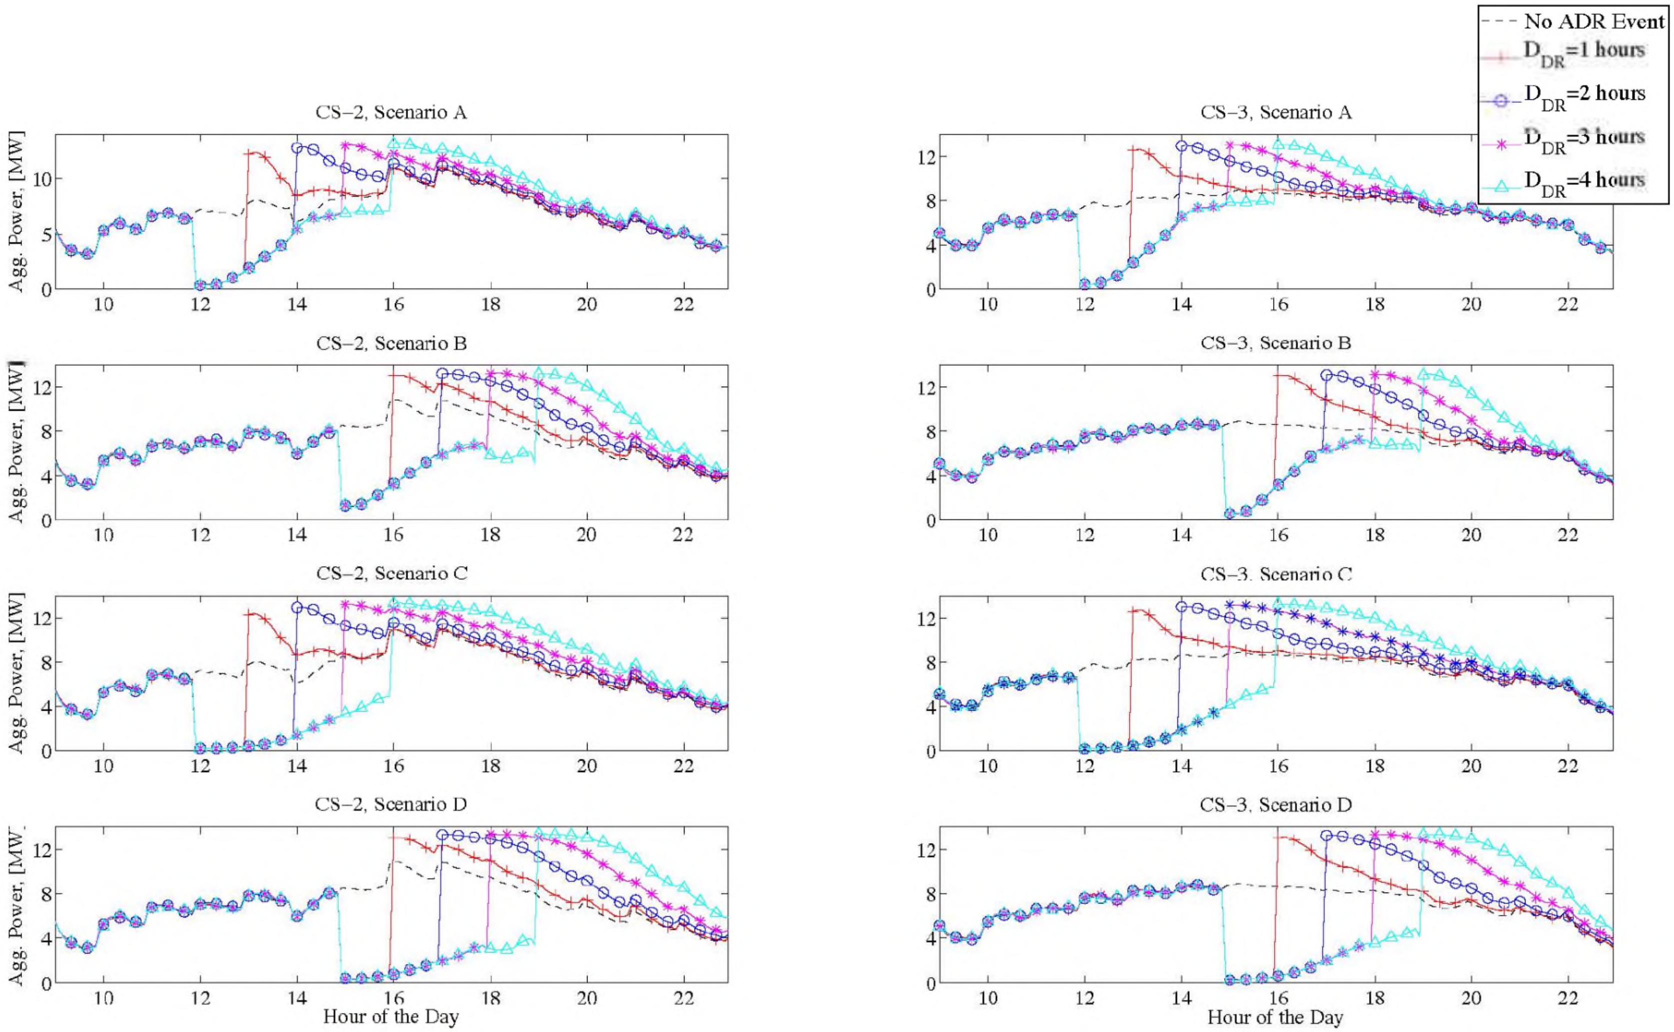

Impact on Energy Savings