Distributed Computing & Stochastic Control for Demand Response In Mass Markets

Federal Energy Regulatory Commission

Washington DC, June 27, 2018

Alex Papalexopoulos, Ph.D.
CEO & Founder, ECCO International, Inc.
CEO & Chairman of the Board, ZOME Energy Networks, Inc.
San Francisco, CA
alexp@eccointl.com
alexp@zomepower.com

ECCO INTERNATIONAL
An Energy Consulting and Software Company

Outline

- Current DR Programs
- Key Challenges of a New Approach
- High Level Process of the New Approach
- Technology and the Load Controller
- ColorPower ™ Algorithm
- Formal Control Problem & Important Constraints
- Control Design Issues
- Cloud-Based Architecture
- ColorPower ™ Energy Token
- Conclusions

ECCO INTERNATIONAL
An Energy Consulting and Software Company

Multi-Market Integrated Electricity Framework

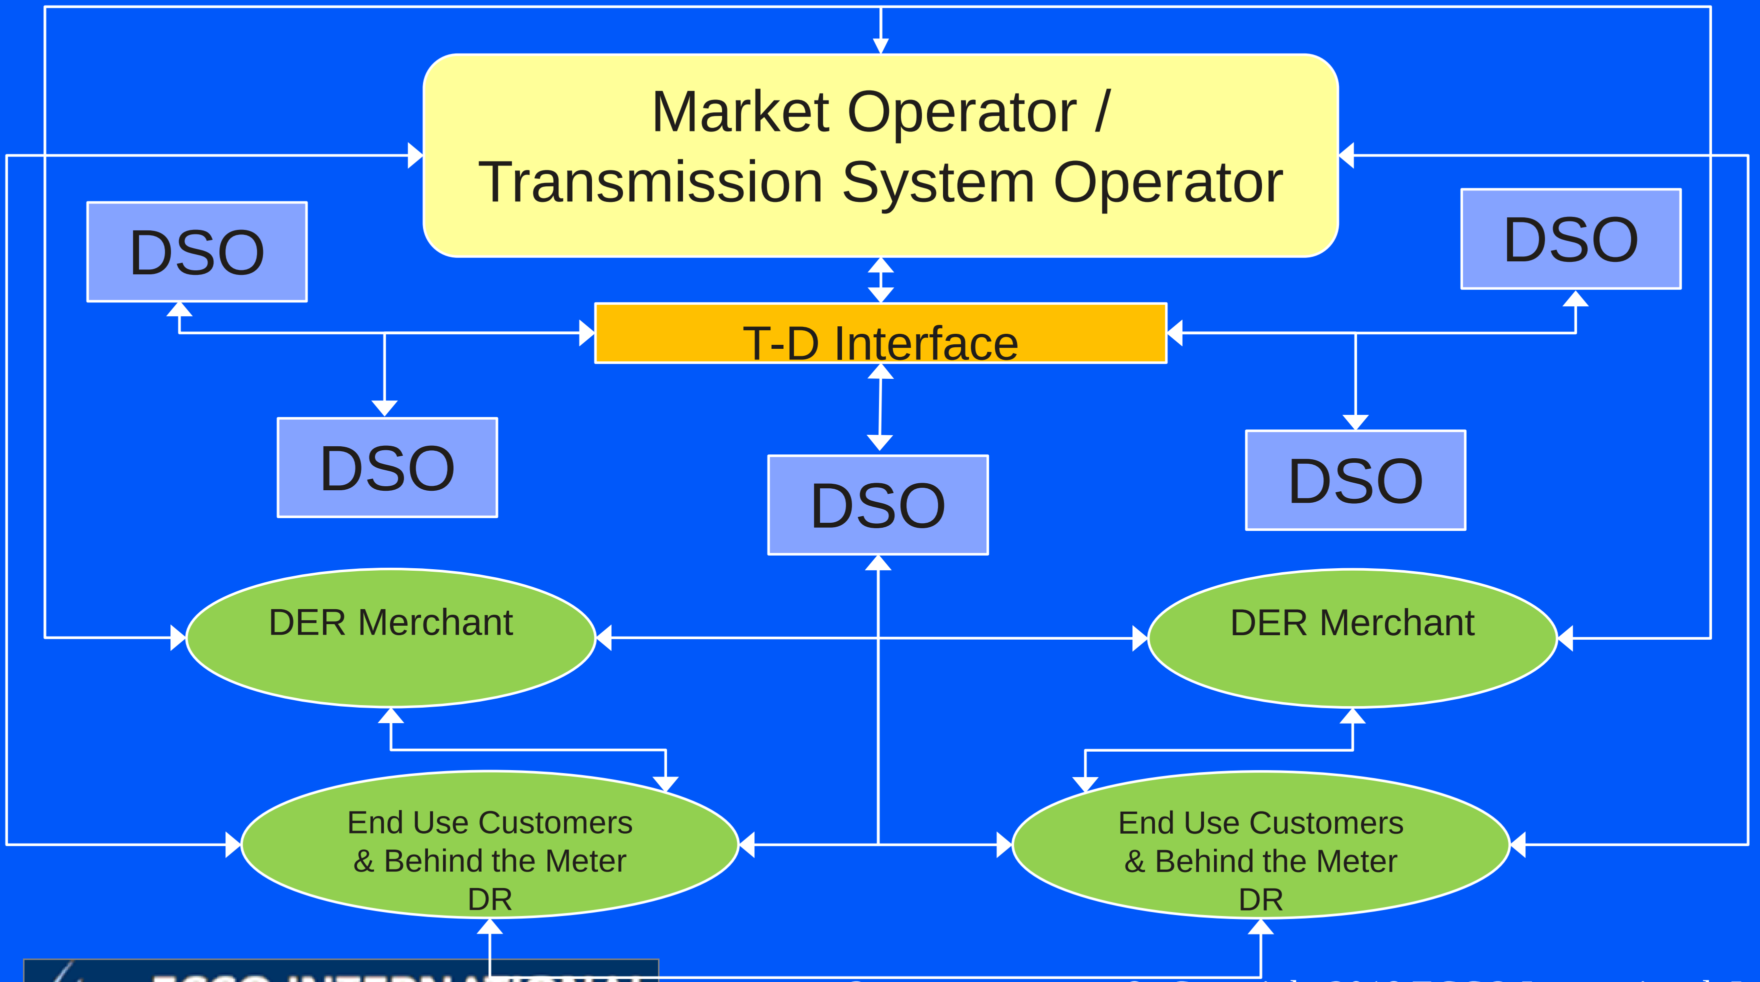

Where is the Problem
(The Era of Coercion Should Come to an End)

- Current DR programs are based on command and control approaches; programs are grouped in 4 groups:
- Customers submit their appliances to direct utility on/off load control
- Customers are exposed to price volatility—a concept called "prices-to-devices"; this is the "holy grail" today for activating DR in wholesale organized markets
- DR aggregators pay people for remote shutoff options; growth has stalled because customers see no other value than trading inconvenience for cash
- Finally, some programs rely on advanced analytics to predict customer behavior and drive messaging and pricing; they try to outguess what customers will do instead of asking them for their preferences

Key Challenges

- Scalability:
  - Safe, reliable coordinated response from millions of devices in < 2-4 minutes
- Consumer interface:
  - High benefit, low “annoyance factor”
  - Eliciting useful information (preferences)
  - Privacy concerns (detailed data and devices should remain private) – This means computations should be performed on consumer aggregates)
- Deployability:
  - Technology alignment with market & regulatory structure
  - Market fragmentation across grid & in home
- Fairness

High Level Process

- ColorPower agent (say customer’s meter) aggregates device flexibility information which is further aggregated across a network
- This forms a model of the overall system flexibility
- This system flexibility, along with a demand shaping target provided by the Utility or the Aggregator, is redistributed to every device in the system
- The devices then execute a distributed control algorithm (like flipping weighted coins) to determine if they respond or not

ColorPower ™ Algorithm

- Challenge: fast, private, robust, non-intrusive
- Approach: randomized distributed control
  - Aggregate flexibility information to shared model
  - Disseminate control signals
  - Local decision; coin-flip for fractional color
  - Weight for availability, over-damped control

Control problem: long timeouts on state changes

ColorPower ™ State Transitions

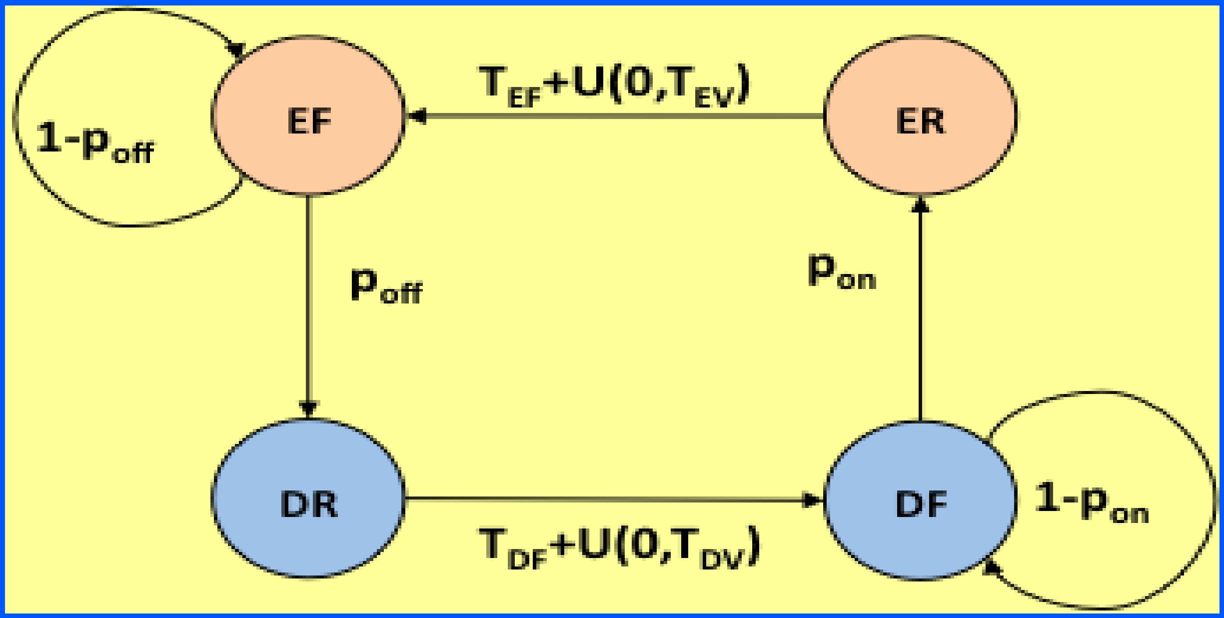

- Within each color, each device: (E)nabled vs. (D)isabled
- (R)efractory (it cannot switch states) vs. (F)lexible (eligible to switch)

ColorPower ™ State Transitions

- The evolution of each device is modeled like a modified Markov process
- In each round devices in state EF randomly switch off to state DR
- Once in DR device waits for certain rounds before transitions to state DF; the waiting time is a fixed number PLUS a uniform random addition to feather the distribution (so not many devices switch states at once)
- The other two distributions are complementary

Formal Control Problem

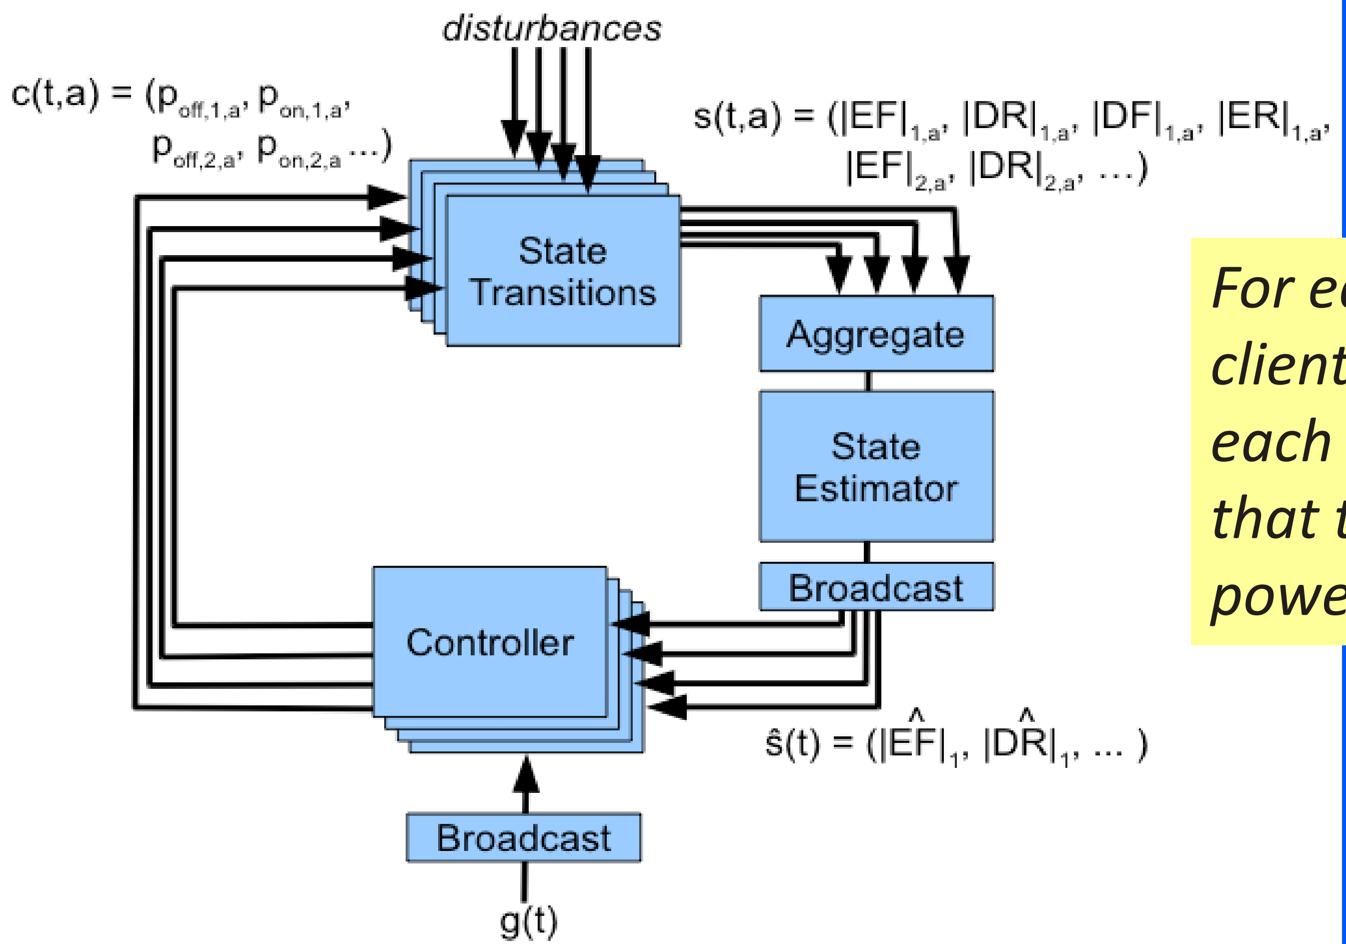

For each ColorPower client, set p_{on}, p_{off} for each device group, such that the total enabled power in s(t) tracks g(t)

Formal Control Problem

- The control problem is to set the transition probabilities such that the total Enabled Demand tracks the target as closely as possible, subject to the constraints
- Device with lower numbered colors are shut off first
- If a color has devices that are Enabled and Disabled, then every device is equally likely to be disabled
- No device is unfairly burdened by its initial bad luck in becoming Disabled

Control Problem Goals

- Continuous tracking (continuously track target power curtailment as loads change)
- User control (color priority)
- Fairness
- Cycling & Limited disruption

Control Problem Constraints

- **1. Goal tracking:** *shape power demand*

$$g(t) = \sum_i |EF_i| + |ER_i|$$

- (Sum of Enabled Demand over all colors i is equal to the goal)
- Typical transient response 2 to 4 minutes
- Simultaneous tracking of DR events and other rules applied to groups of facilities

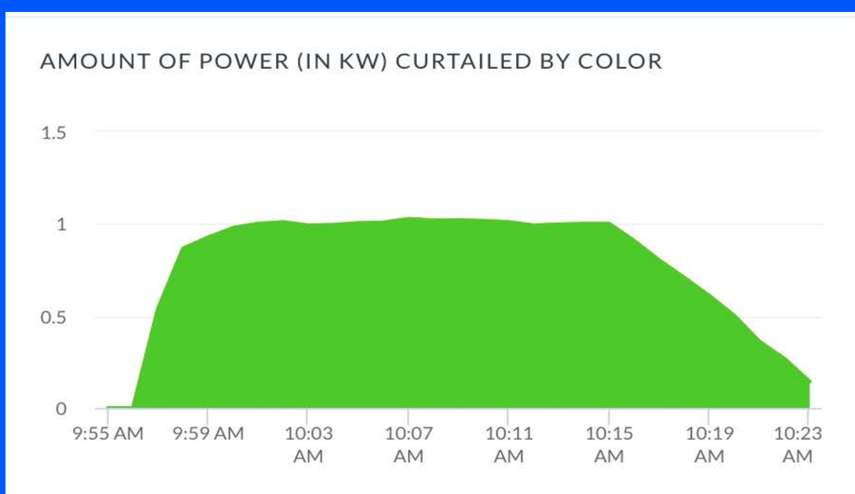

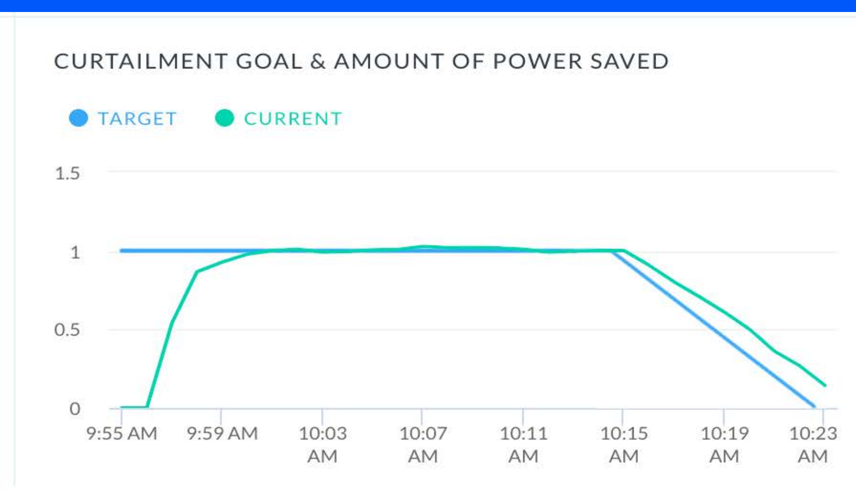

ECCO INTERNATIONAL
An Energy Consulting and Software Company

Control Problem Constraints

- **2. Color priority:** *respect user preferences*

$$|EF_i| + |ER_i| = \begin{cases} D_i - D_{i+1} & \text{if } D_i \leq g(t) \\ g(t) - D_{i+1} & \text{if } D_{i+1} \leq g(t) < D_i \\ 0 & \text{otherwise} \end{cases}$$

$$D_i = \sum_{j \geq i} |EF_j| + |ER_j| + |DF_j| + |DR_j|$$

- Demand D_i is the demand for the ith color and above
- Devices are Enabled from the highest color down until the goal is reached
- Users choose a ‘color’ for each device indicating willingness to be flexible in DR events
- User controllable from anywhere in the world via web or mobile app

Anytime Flexible
Peak Power
Emergencies Only
Uncontrolled

Control Problem Constraints

- **3. Fairness:** *no devices are favored*

$$\forall_{a,a'} c(t,a) = c(t,a')$$

- Meaning that the control state is identical for every agent
- Within the same user-selected color, treat all devices equally on average
  - Use randomized algorithms to ensure average equal treatment
- Use load information only in aggregate
  - Individual device load does not affect DR behavior
- Balance curtailment across different device types

Control Problem Constraints

- **4. Cycling:** *don't keep the same devices off*

$$\forall_{a,a'} c(t,a) = c(t,a')$$

- This means that as long as there are both Enabled and Disabled devices, some of them should be changing from Enabled to Disabled to vice versa

$$(|EF_i| > 0) \cap (|DF_i| > 0) \implies (p_{on,a,i} > 0) \cap (p_{off,a,i} > 0)$$

- **5. Limited Disruption:** Spread the curtailment as broadly as possible
  - Smart plugs—gradually cycle through which devices are off
  - HVAC—as small as possible temperature change across more homes
- Avoid frequent switching of loads

ECCO INTERNATIONAL
An Energy Consulting and Software Company

Controller Design Issues

- It is possible that not all constraints can be satisfied; some of them are more important than others
- Customer preferences are the most important ones
- Goal tracking is the second most important
- Least important is the Cycling constraint
- The Fairness constraint is the easiest to satisfy (simply the same stochastic algorithm on all clients is executed)
- We view the controller as having a “budget” of flexibility to spend with each color offering up to $|EF|_l$ of potential reduction in demand

ECCO INTERNATIONAL
An Energy Consulting and Software Company

Controller Design Issues

- Flexibility builds up as Refractory devices finish their time outs and move to the Flexible state
- The controller is formulated as a cascade of priorities of how to spend the “Flexibility budget” indicated by the state s(t)
- As the controller considers each constraint in turn, it allocates flexibility to satisfy that constraint (as much as possible)
- Then it attempts to satisfy the rest of the constraints with whatever flexibility remains unallocated
- Any unallocated flexibility is allowed to accumulate as a reserve improving future controllability

ECCO INTERNATIONAL
An Energy Consulting and Software Company

ColorPower™ Energy Demand Cloud

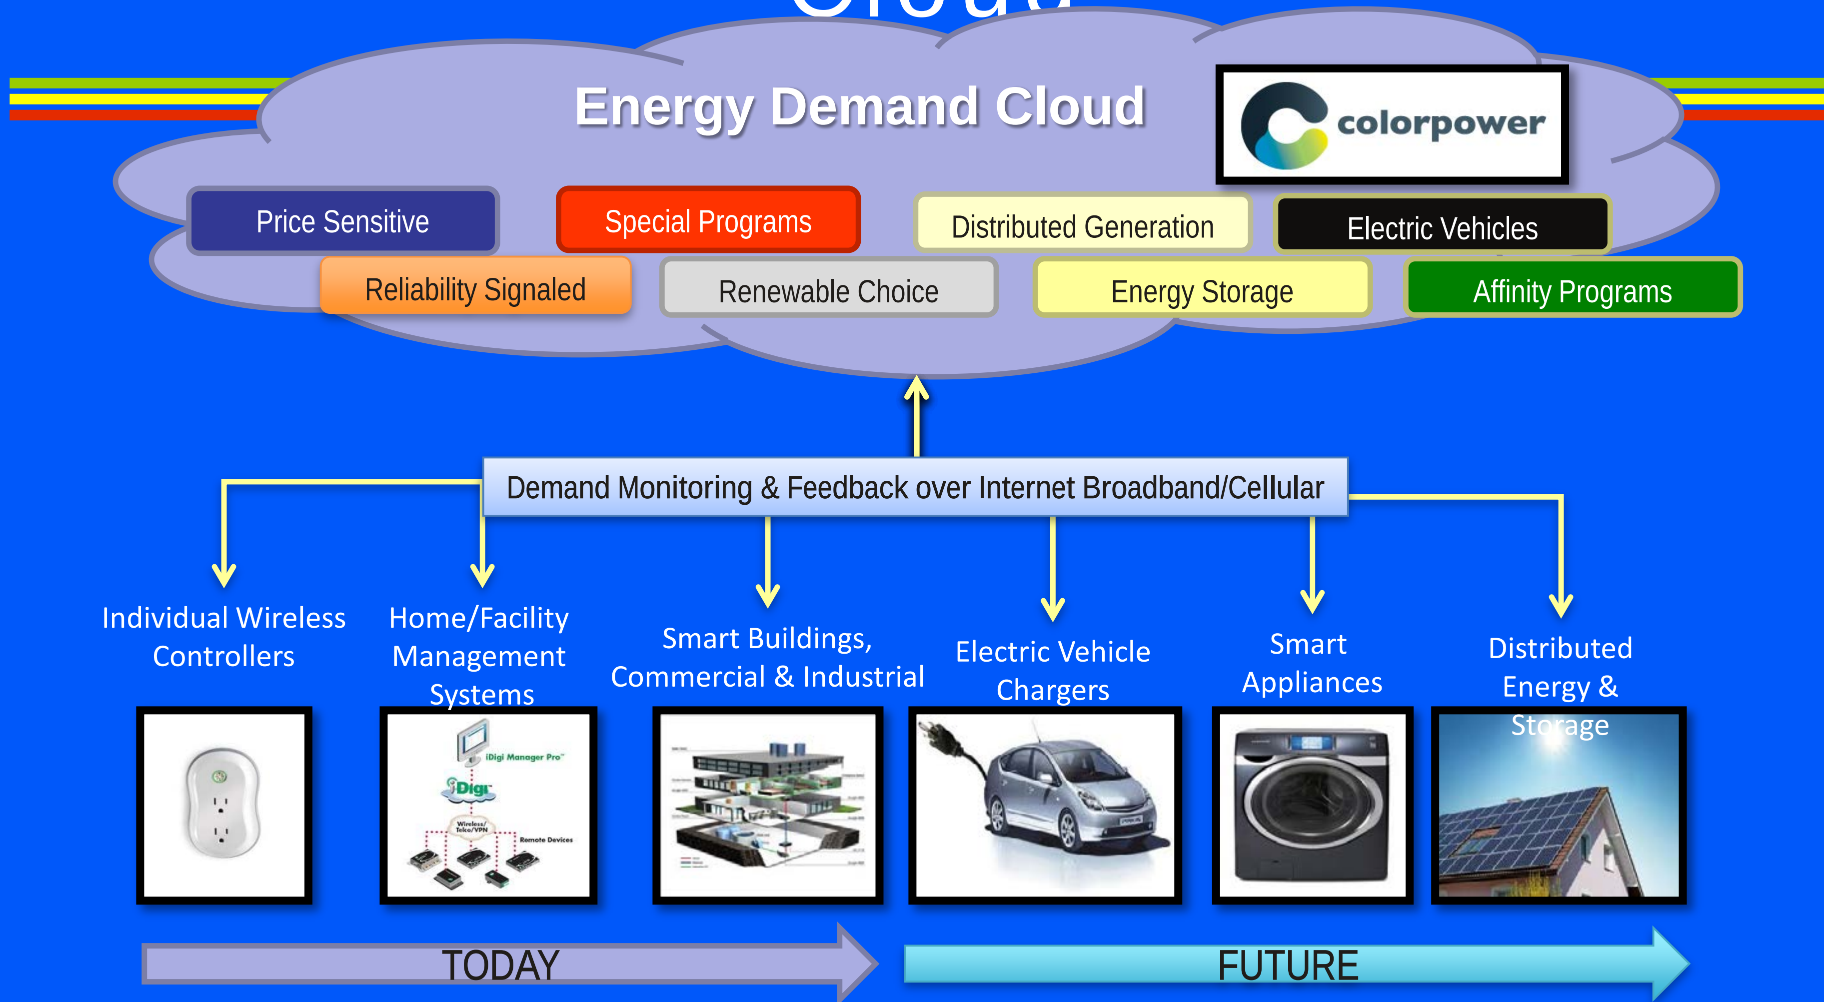

ColorPower Cloud-Based Implementation

- Cloud based—leverage modern cloud infrastructure
- Leverage smart devices people already have in their homes—smart thermostats, smart plugs, etc.
- Combine measurement and modeling
- Micro-Service Architecture
- Single API for everyone
- Independently Scalable
- Extendable with High Performance

Cloud Architecture

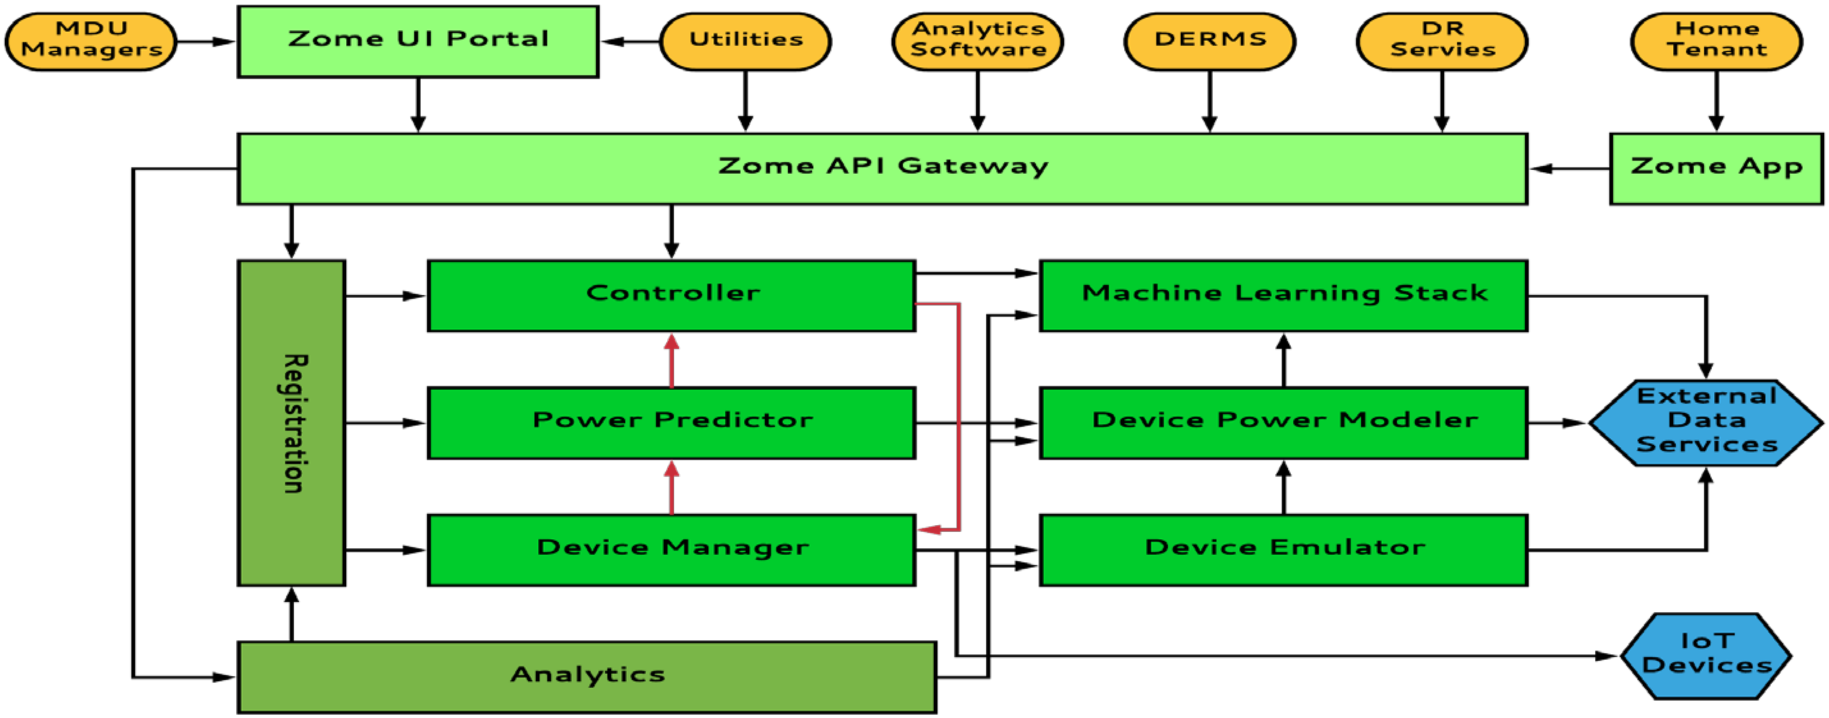

Arrows start at request origin and points to request destination.
Black Arrows – Request/Response type requests.
Red Arrows – Event type requests.

All greens represents components owned by Zome.
Mint – externally facing components.
Mantis – core algorithmic and CnC virtual stack.
Moss – Supplemental micro-services.

Yellow – customer agents / services.
Blue – third party and supplemental APIs.

Cloud Architecture

- Device Manager – Responsible for all device communications and state maintenance
- Power Predictor – Projects power use of the device when data is not available (thermostats) using behavioral models
- Power Modeler – Uses data gathered from the devices to generate device behavior models that will be used by the Predictor
- Controller – Heart of the system, location of the ColorPower algorithm; tasked with maintaining state of the grid and giving orders to devices to curtail power
- Emulator Micro Service is responsible for emulating devices; this allows ColorPower to run complex simulations on various “what if” scenarios

ColorPower Energy Token

- We are in the process of building a blockchain technology and we plan to introduce a utility crypto token: ZENT
- CP platform contributors (device owners) will be rewarded for participation in power saving events with ZENTs
- CP services with time will be purchasable only with ZENT tokens, to enable a robust and healthy energy token ecosystem
- Modularized cloud architecture allows CP to quickly integrate virtually any connected device into our system
- Adding support for new device takes days
- Today CP supports:
  - HIVE Thermostats
  - BOSS Smart Plug
  - Majority of generic power intensive z-wave devices (thermostats, power plugs)

ECCO INTERNATIONAL
An Energy Consulting and Software Company

Smart Devices

- Some example devices from our last deployment

Combine Measurement and Modeling

- For dynamic tracking of power use, we combine both measured power, when available, with model-based estimates
  - Some smart plugs provide power measurements when on
  - Thermostats do not indicate power of HVAC
- Physics-based models to estimate power when not available
- Combined information from many sources
  - Weather, location, local home construction statistics, etc.
  - Leverage AI machine learning techniques

ECCO INTERNATIONAL
An Energy Consulting and Software Company

Conclusions

- Current DR programs are not successful
- The new cloud-based proposed algorithm is based on a distributed computing based stochastic control algorithm allows fast, accurate and robust control of thousands to millions of devices
- Performance can be accurately predicted from stochastic model analysis; performance is robust against fluctuations, errors, and variation between devices
- We are in the process of rolling out three (3) products in San Francisco, Chicago and New York: CP for MDUs, CP for neighborhoods and HVAC Analytics
- We'll aggregate capacity and offer it in the ISO markets for various grid services and products by end of this year