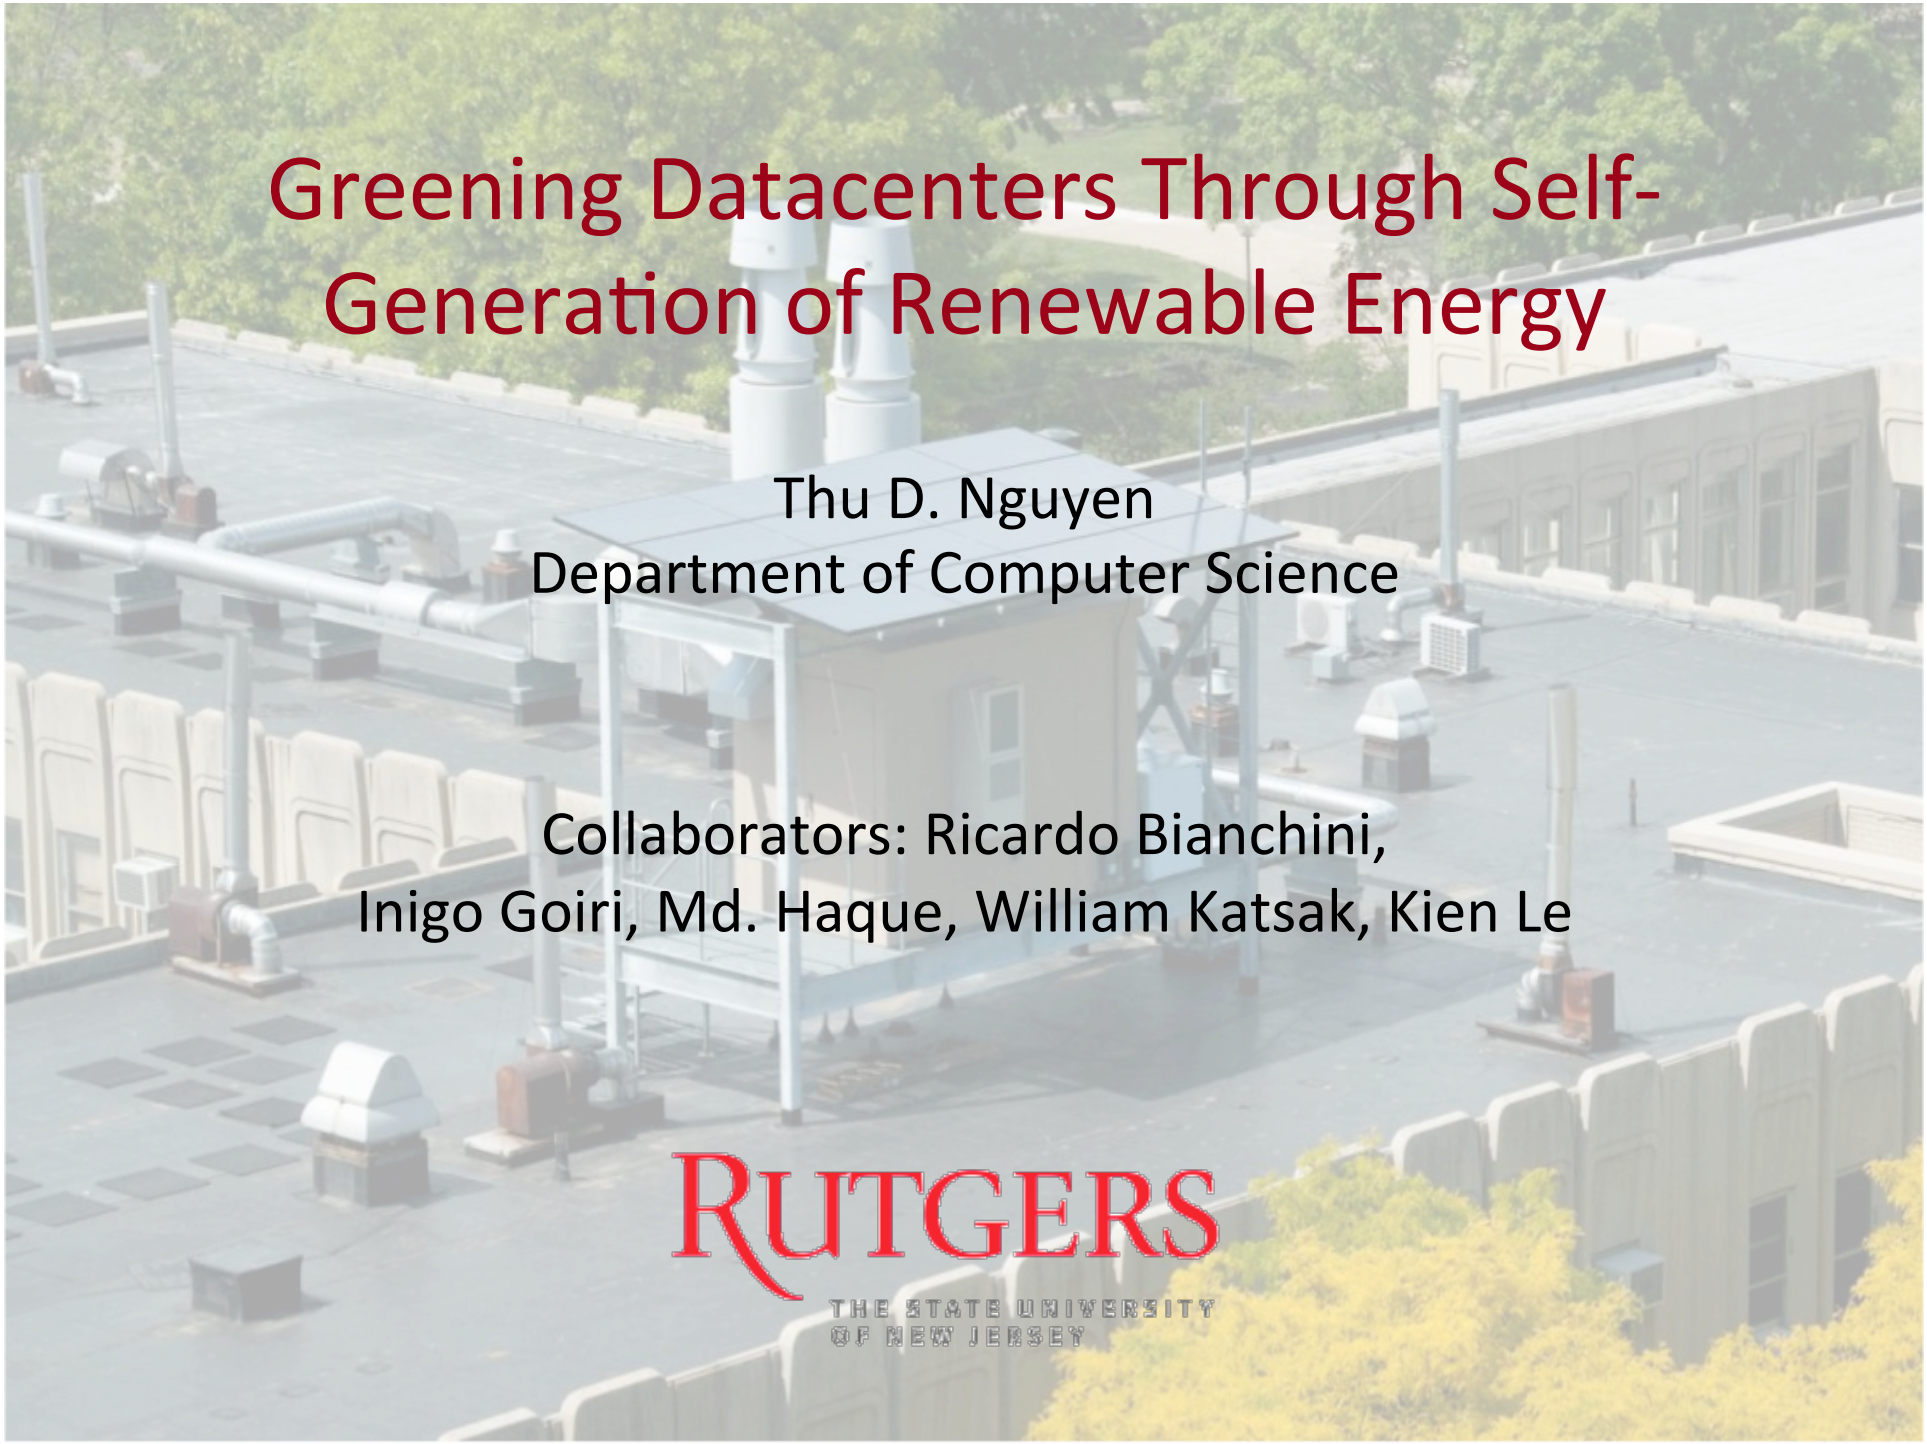

Greening Datacenters Through Self-Generation of Renewable Energy

Thu D. Nguyen
Department of Computer Science

Collaborators: Ricardo Bianchini,
Inigo Goiri, Md. Haque, William Katsak, Kien Le

RUTGERS
THE STATE UNIVERSITY
OF NEW JERSEY

Motivation

- Datacenters consume massive amounts of energy (electricity)
- Vast majority currently due to small and medium scale datacenters

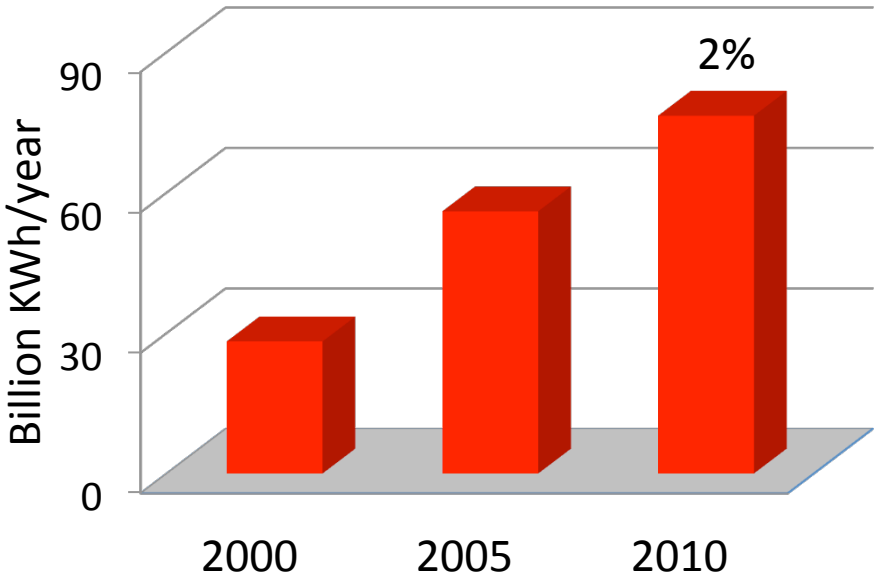

Electricity consumption of US DCs [JK'11]

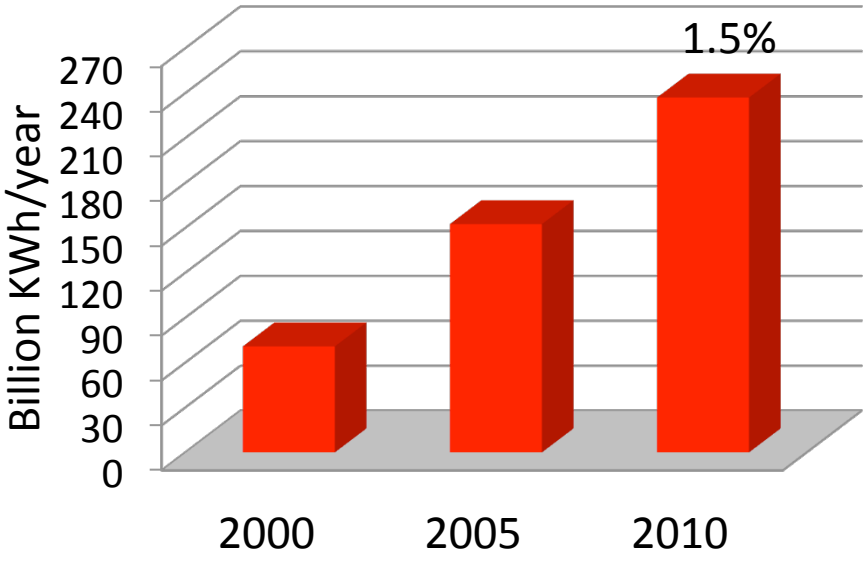

Electricity consumption of WW DCs [JK'11]

Motivation

- Electricity comes mostly from burning fossil fuels

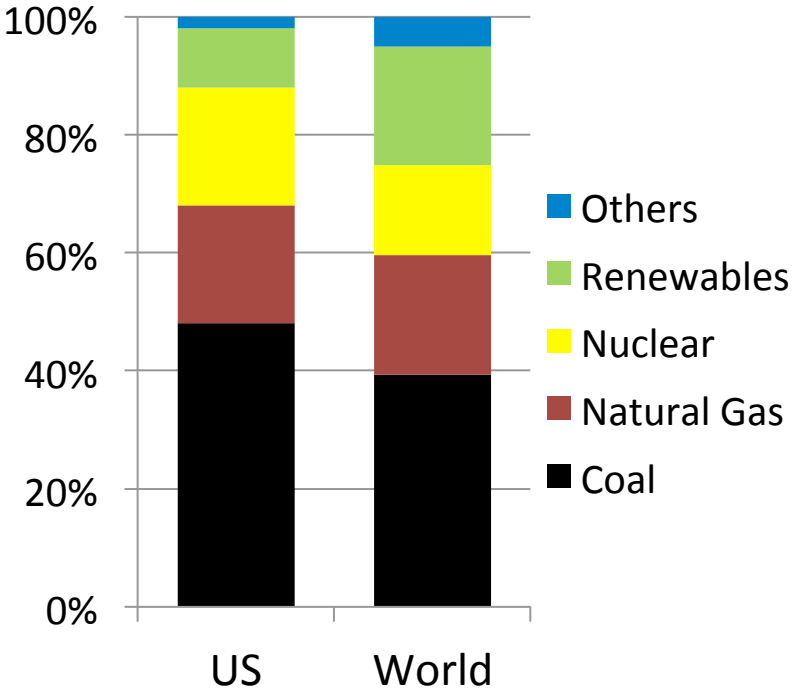

Electricity sources in US & WW [DOE'10]

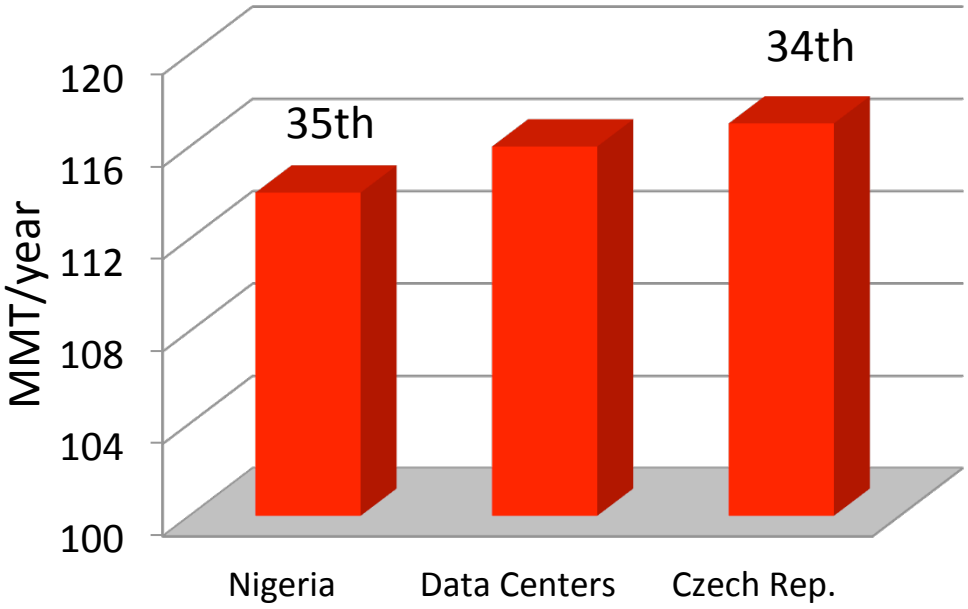

CO_2 of world-wide DCs [Mankoff'08]

Can we use renewables to reduce this footprint?

- Reducing electricity costs would be nice too

Outline

- Motivation: DC energy usage and carbon footprint
- Reducing carbon footprint & cost with renewables
- Our target and research challenges
- Parasol: our solar-powered micro-datacenter
- GreenSwitch: managing power sources
- Previous, current, and related work
- Conclusions

Greening Datacenters

- Power purchase agreement, off-site generation
  - Renewable energy produced at the best location
  - Requires transmission
    - Construction of transmission lines
    - Energy losses: ~15% [IEC'07]
    - Variability challenging for grid operators
  - Example: Google buys wind power from NextEra
- Self-generation, co-location
  - Lower peak power, energy costs with self-generation
  - Location may not be ideal for DC or renewable plant
  - Examples:
    - Microsoft placed DC near a hydro plant in OR
    - Apple built a 40MW solar array in NC
- No approach is perfect

Self-Generation Example: Apple NC

Outline

- DC energy usage and carbon footprint
- Reducing carbon footprint & cost with renewables
- Our target and research challenges
- Parasol: our solar-powered micro-datacenter
- GreenSwitch: managing power sources
- Previous, current, and related work
- Conclusions

Our Research Target

- Self-generation or co-location with solar and/or wind
  - Pros: Clean and available
  - Cons: Space and cost

Solar and Wind Are Clean

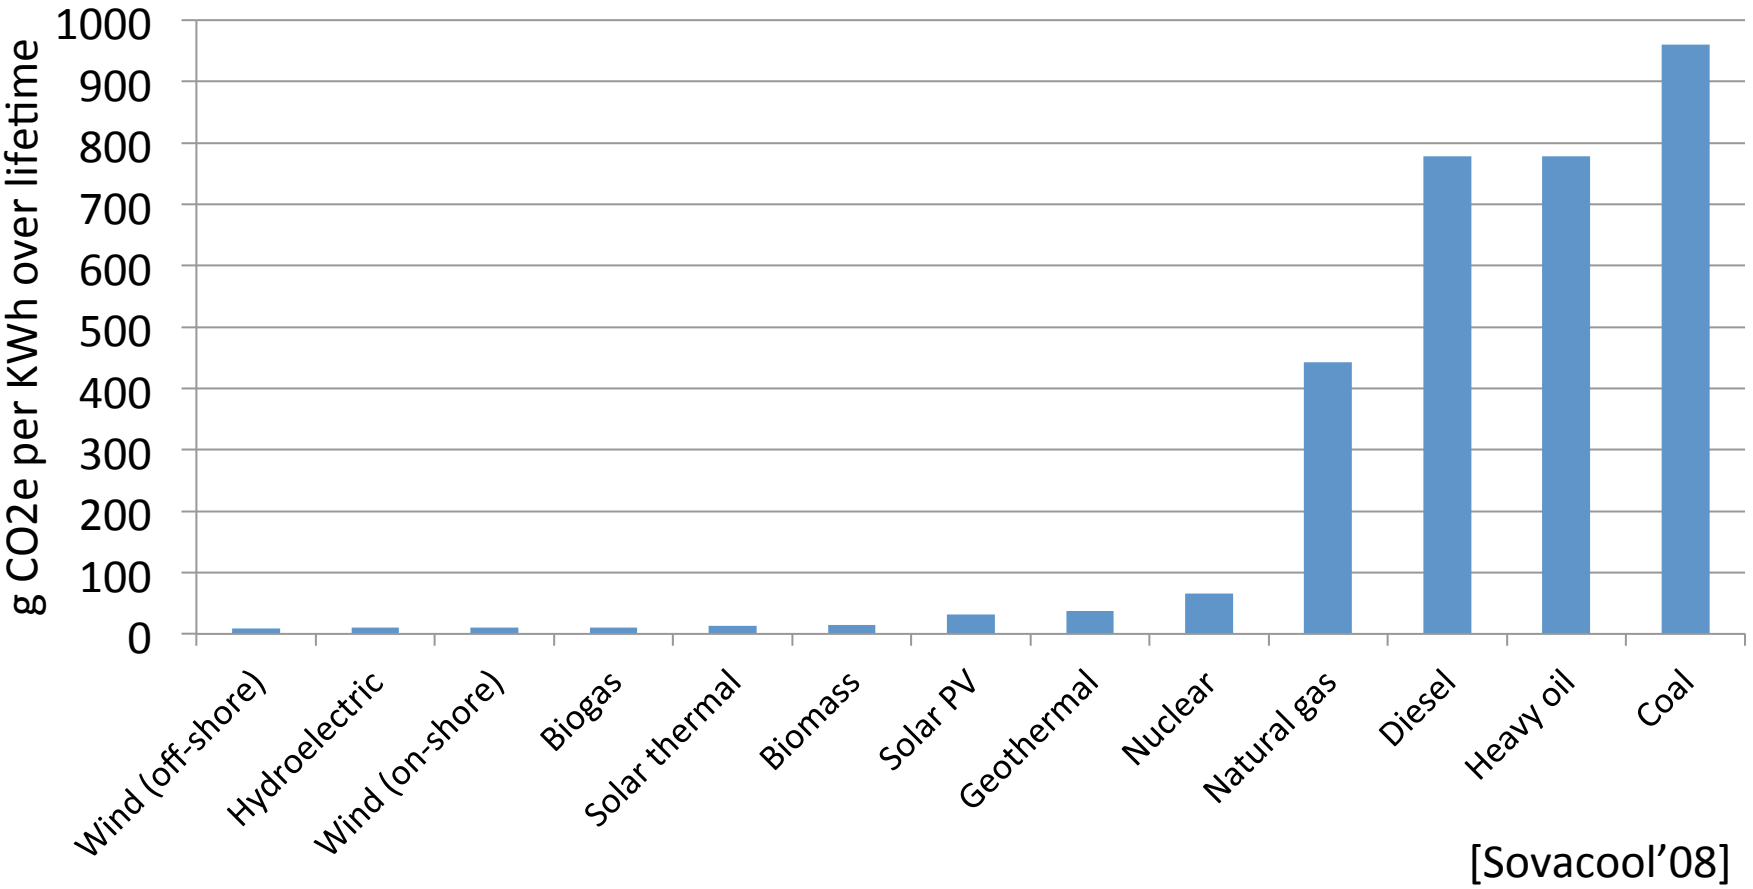

[Sovacool'08]

Solar and Wind Are Clean

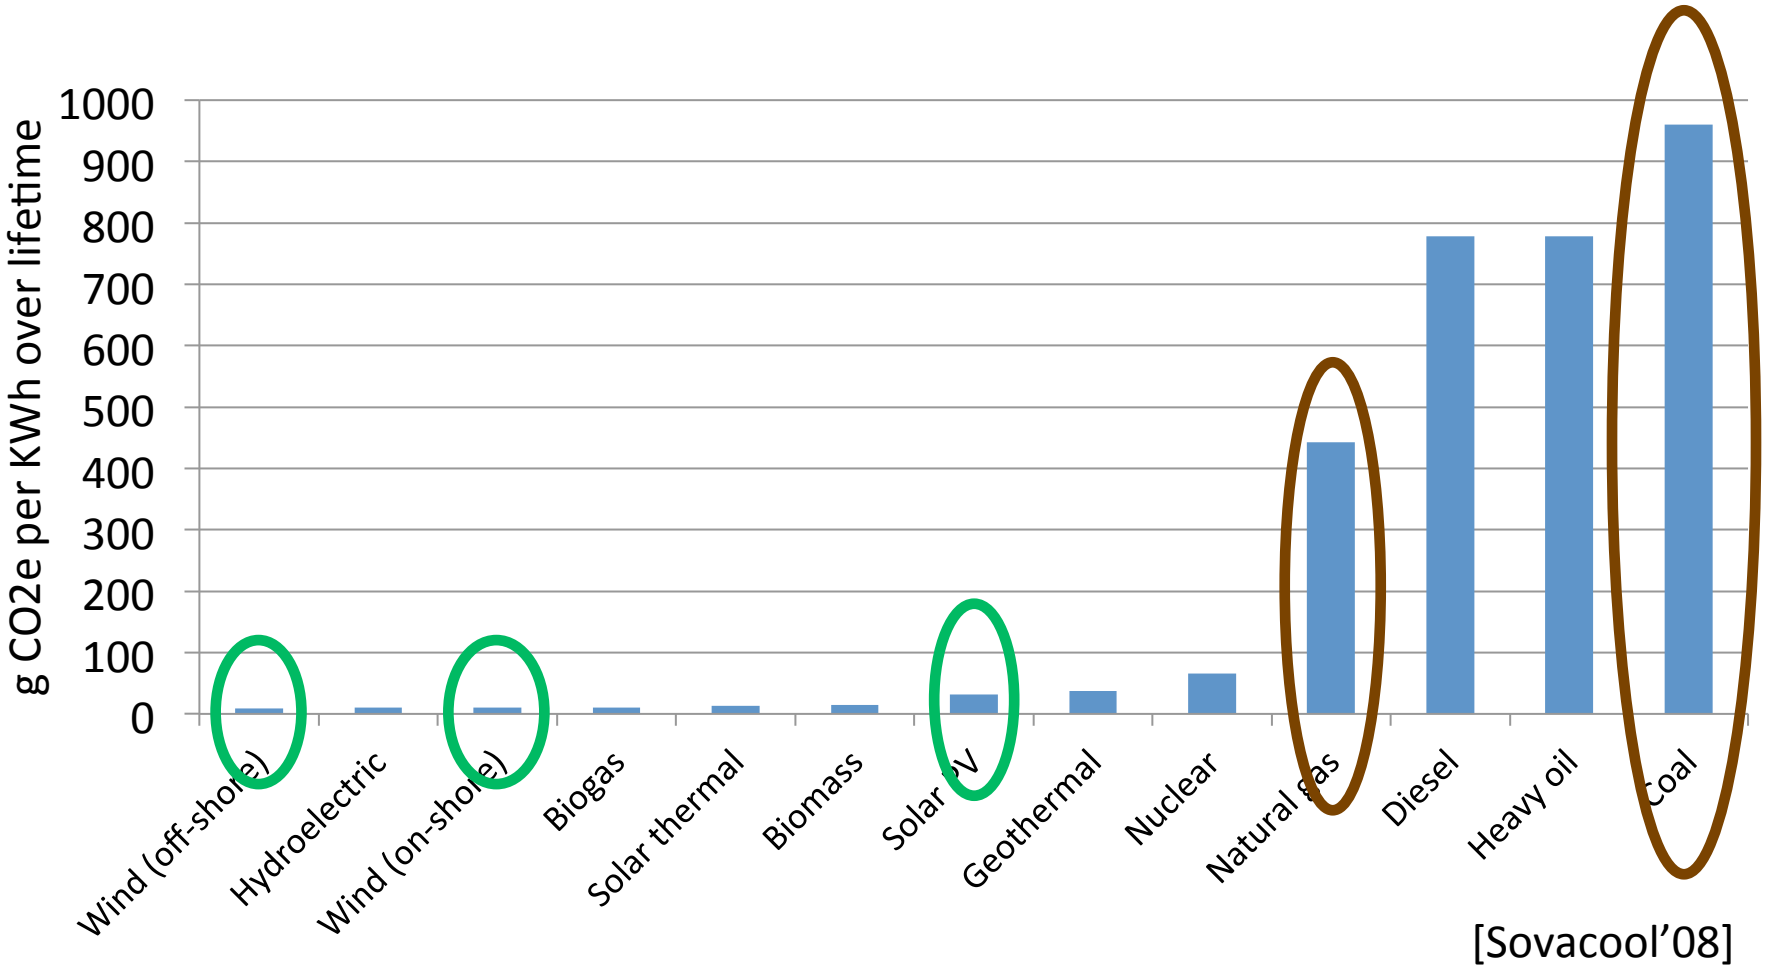

[Sovacool'08]

Solar More Available In US

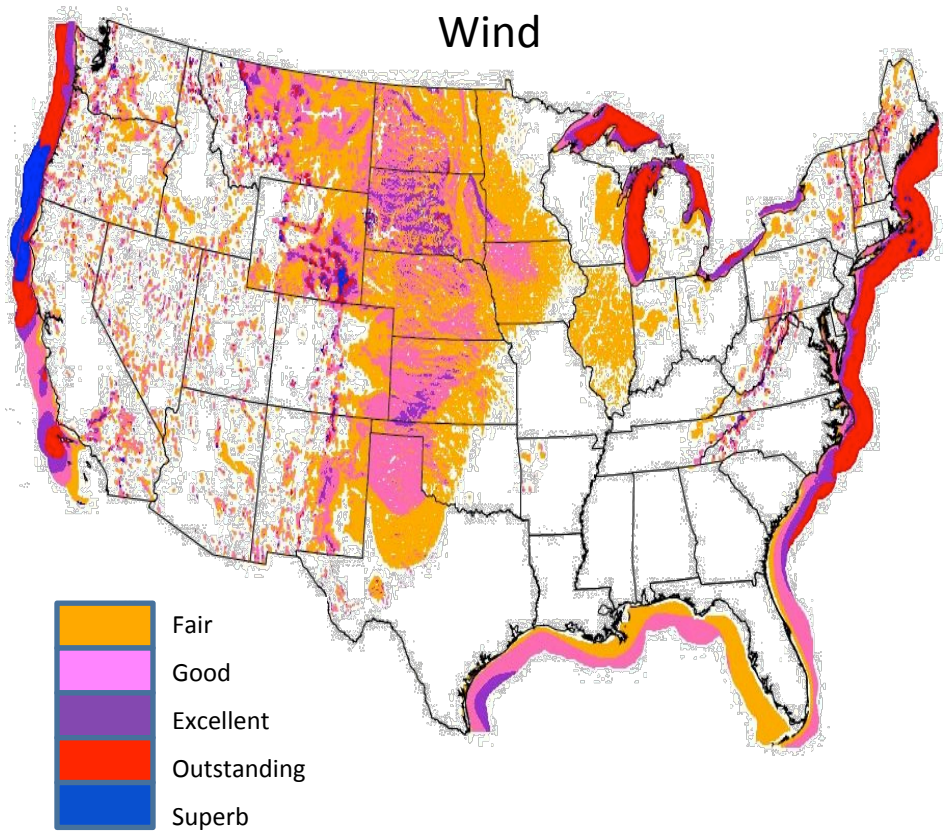

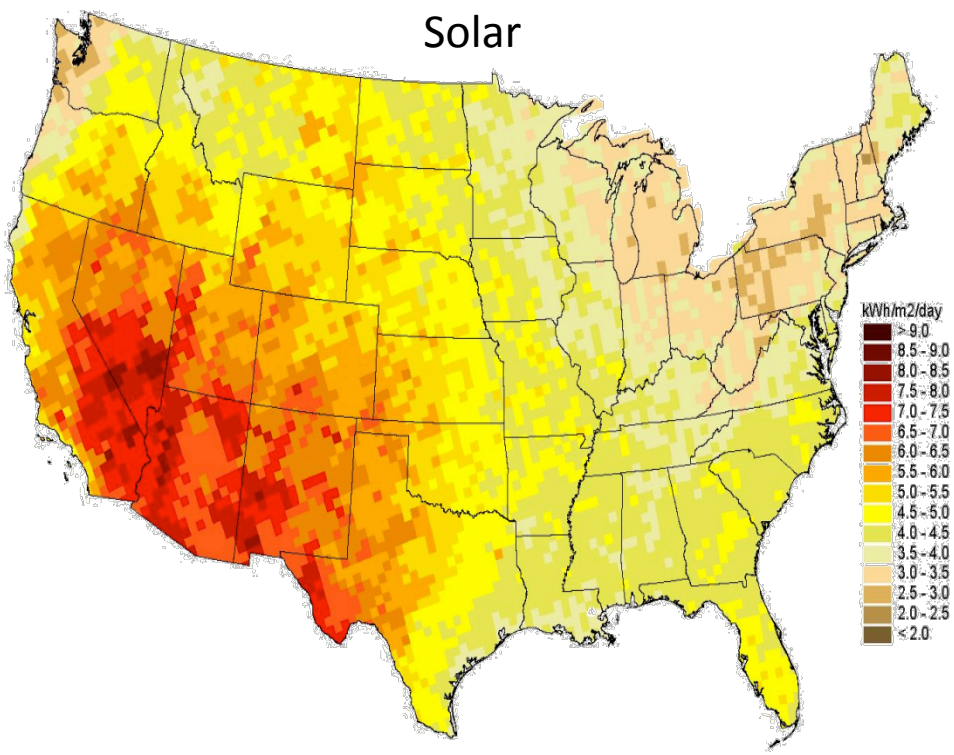

[NREL'12]

Solar PV Efficiencies Are Increasing

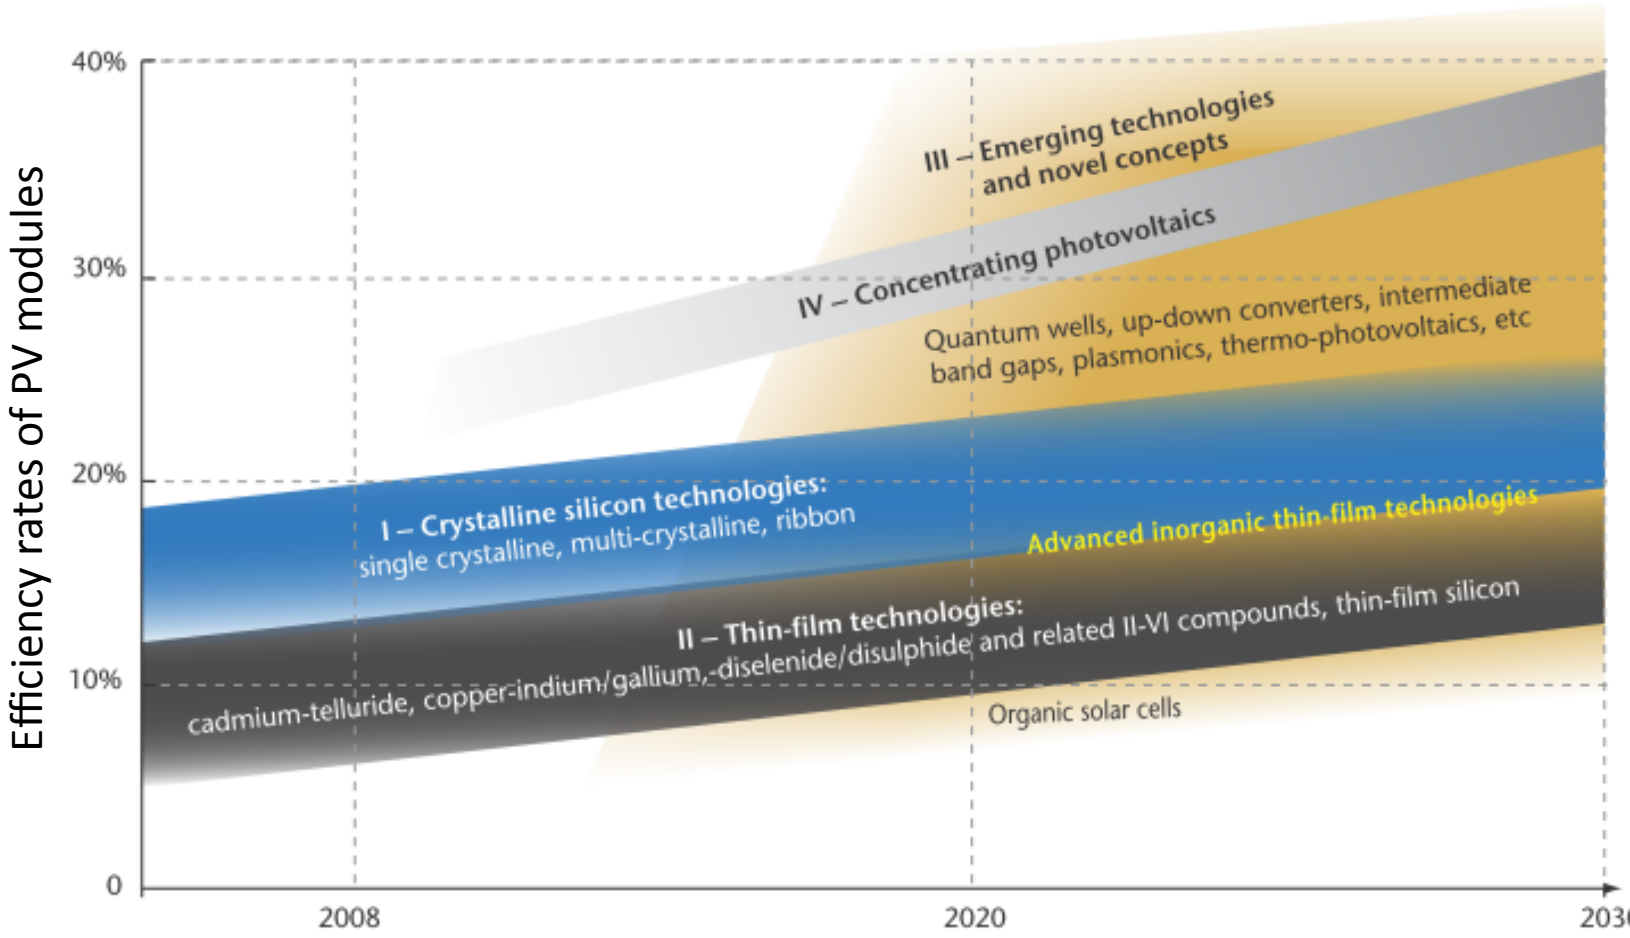

[IEA'10]

Solar PV Capacity Factors Today

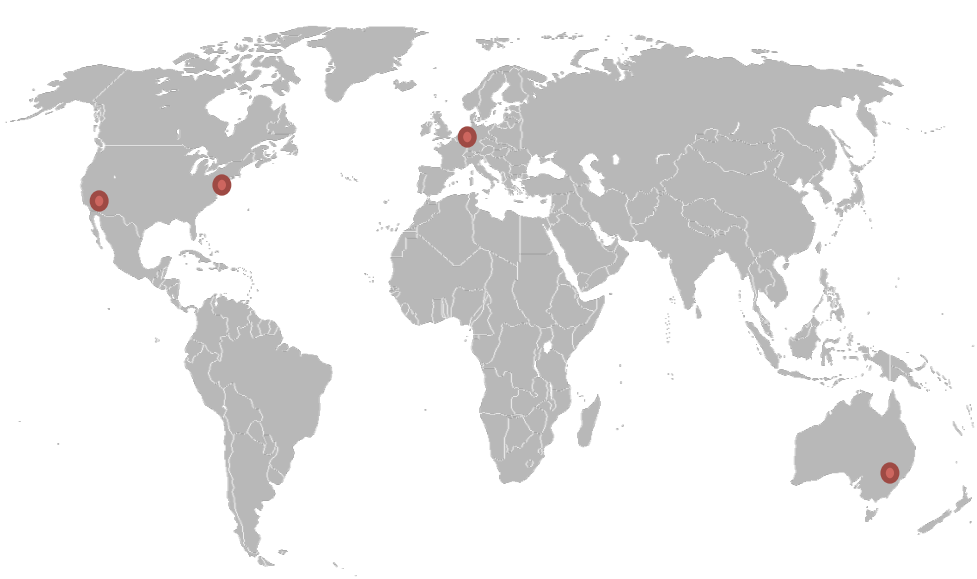

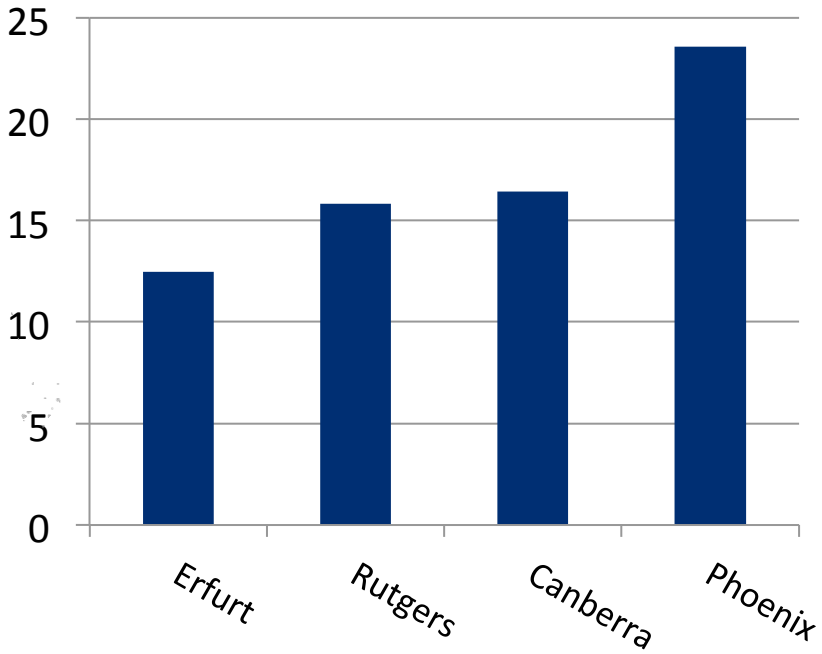

[PVOutput'12]

Cost of Solar PV Energy Decreasing

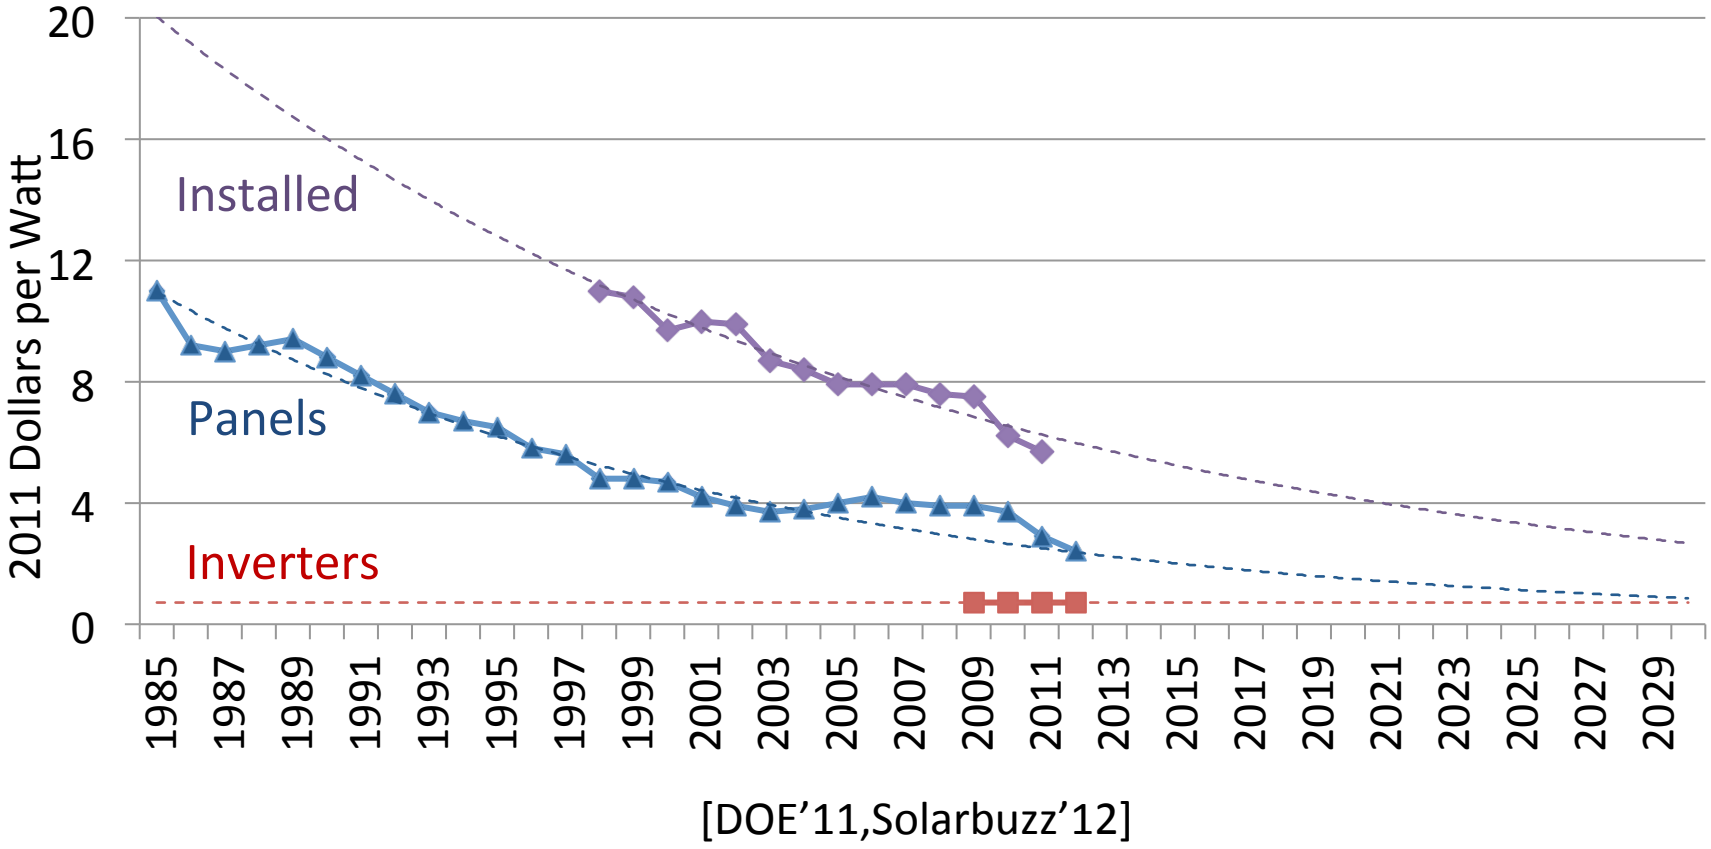

[DOE'11,Solarbuzz'12]

Grid electricity prices have been increasing: 30%+ since 1998 [EIA'12]

Cost of Solar PV Energy Decreasing

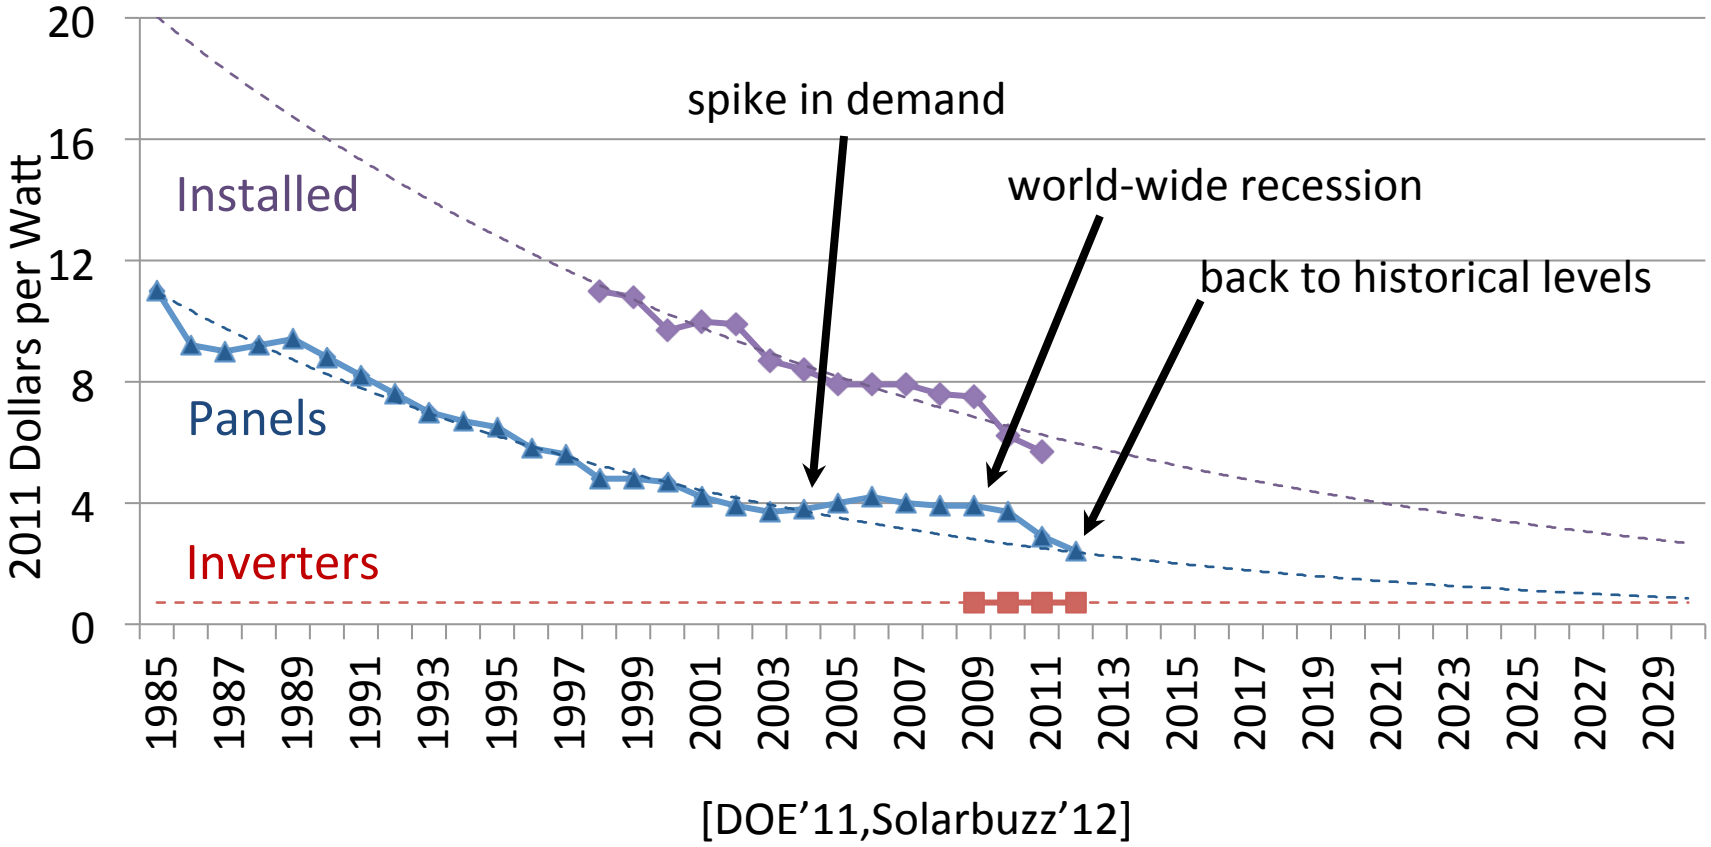

[DOE'11,Solarbuzz'12]

Cost of Solar PV Energy Decreasing

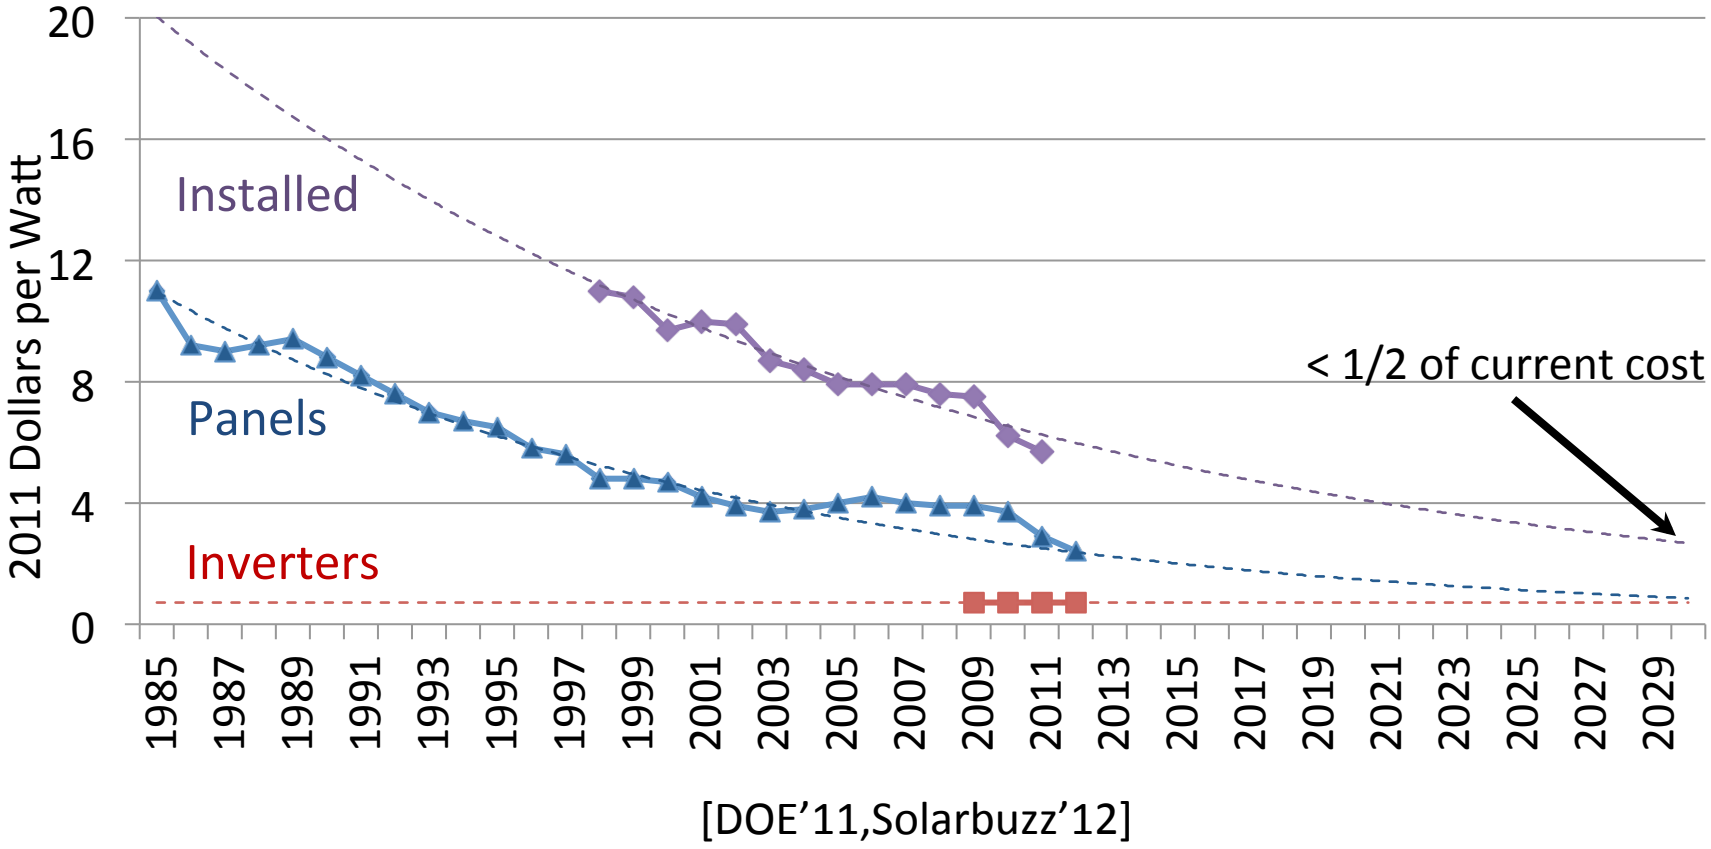

[DOE'11,Solarbuzz'12]

With incentives, the installed price can go down by another 40-60%

Solar Space and Cost: Present and Future

Space as a factor of rack area	Present	Future (2020-2030)
Density per rack		
8kW (200W 1U servers)	~47x	~24x
2kW (25W 0.5U servers)	~12x	~6x

Assuming 30% server utilization, 50% solar energy, NJ capacity factor, and 1 row of panels

Solar space and cost: Present and future

Space as a factor of rack area	Present	Future (2020-2030)
Density per rack		
8kW (200W 1U servers)	~47x	~24x
2kW (25W 0.5U servers)	~12x	~6x

Assuming 30% server utilization, 50% solar energy, NJ capacity factor, and 1 row of panels

Cost per AC Watt	Present	Future (2020-2030)
	~$2.30	< $1.20

Assuming self-generation and federal + state incentives

Time to amortize cost	Present	Future (2020-2030)
	~12 years	< 6 years

Assuming above costs, NJ capacity factor, and NJ grid energy prices

Solar space and cost: Present and future

Space as a factor of rack area	Present	Future (2020-2030)
Density per rack		
8kW (200W 1U servers)	~47x	~24x
2kW (25W 0.5U servers)	~12x	~6x

Assuming 30% server utilization, 50% solar energy, NJ capacity factor, and 1 row of panels

Cost per AC Watt	Present	Future (2020-2030)
	~$2.30	< $1.20

Assuming self-generation and federal + state incentives

Time to amortize cost	Present	Future (2020-2030)
	~12 years	< 6 years

Assuming above costs, NJ capacity factor, and NJ grid energy prices

Wind takes ~12x less space and is ~3x cheaper

Main Challenge: Power Supply is Variable

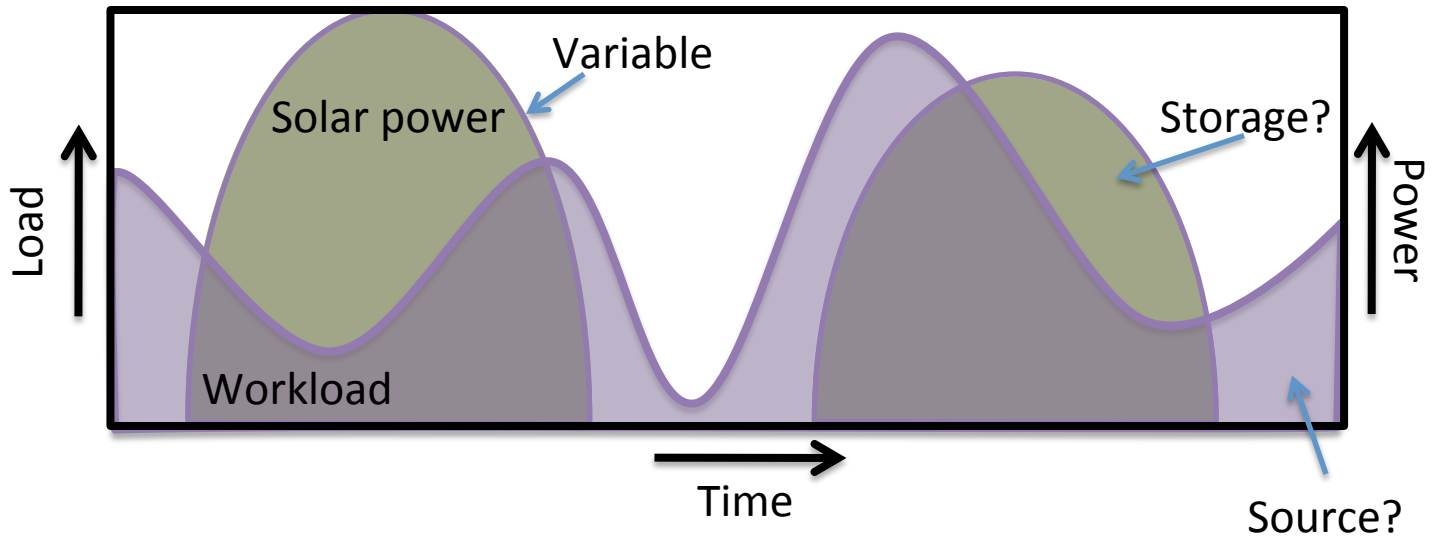

- Power generation is variable
  - Unlikely to match workload
- Match power demand and supply

Main Challenge: Power Supply is Variable

- Many research questions:
  - What kinds of DC workloads are amenable?
  - What kinds of techniques can we apply?
  - How well can we predict solar energy availability?
  - If batteries are available, how should we manage them?
  - Can we leverage geographical distribution?

- Building hardware & software to answer questions

Outline

- DC energy usage and carbon footprint
- Reducing carbon footprint & cost with renewables
- Our target and research challenges
- Parasol: our solar-powered micro-datacenter
- GreenSwitch: managing power sources
- Previous, current, and related work
- Conclusions

The Rutgers Parasol Project

Parasol: Our Hardware Prototype

- Unique research platform
  - Solar-powered computing
  - Remote DC deployments
  - Software to exploit renewables within and across DCs
  - Tradeoff between renewables, batteries, and grid energy
  - Free cooling, wimpy servers, solid-state drives

Parasol: Our Hardware Prototype

Parasol Details

- Steel structure on the roof
  - 16 solar panels: 3.2 kW peak
  - Container hosts 2 racks of IT
- Backup power
  - Batteries: 32kWh
  - Grid
- IT equipment
  - 64 Atom servers (so far): 1.9 kW
  - 2 switches
- Cooling
  - Free cooling: 10 – 400 W
  - Air conditioning: 2 kW

Electrical infrastructure

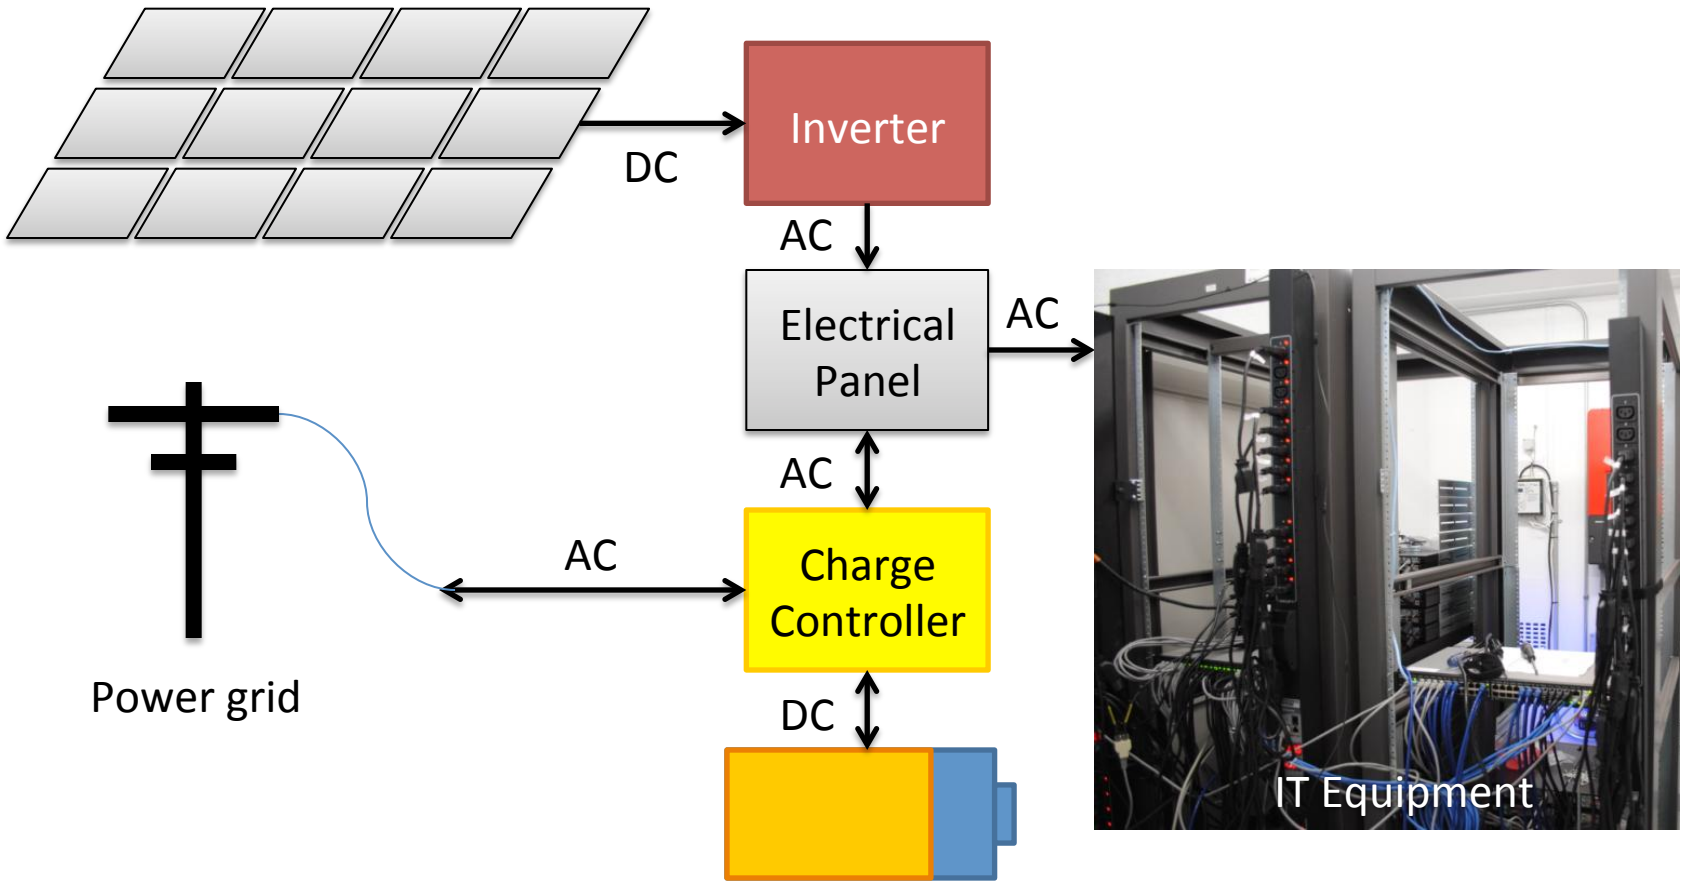

Outline

- DC energy usage and carbon footprint
- Reducing carbon footprint & cost with renewables
- Our target and research challenges
- Parasol: our solar-powered micro-datacenter
- GreenSwitch: managing power sources
- Previous, current, and related work
- Conclusions

Example Energy Source Management

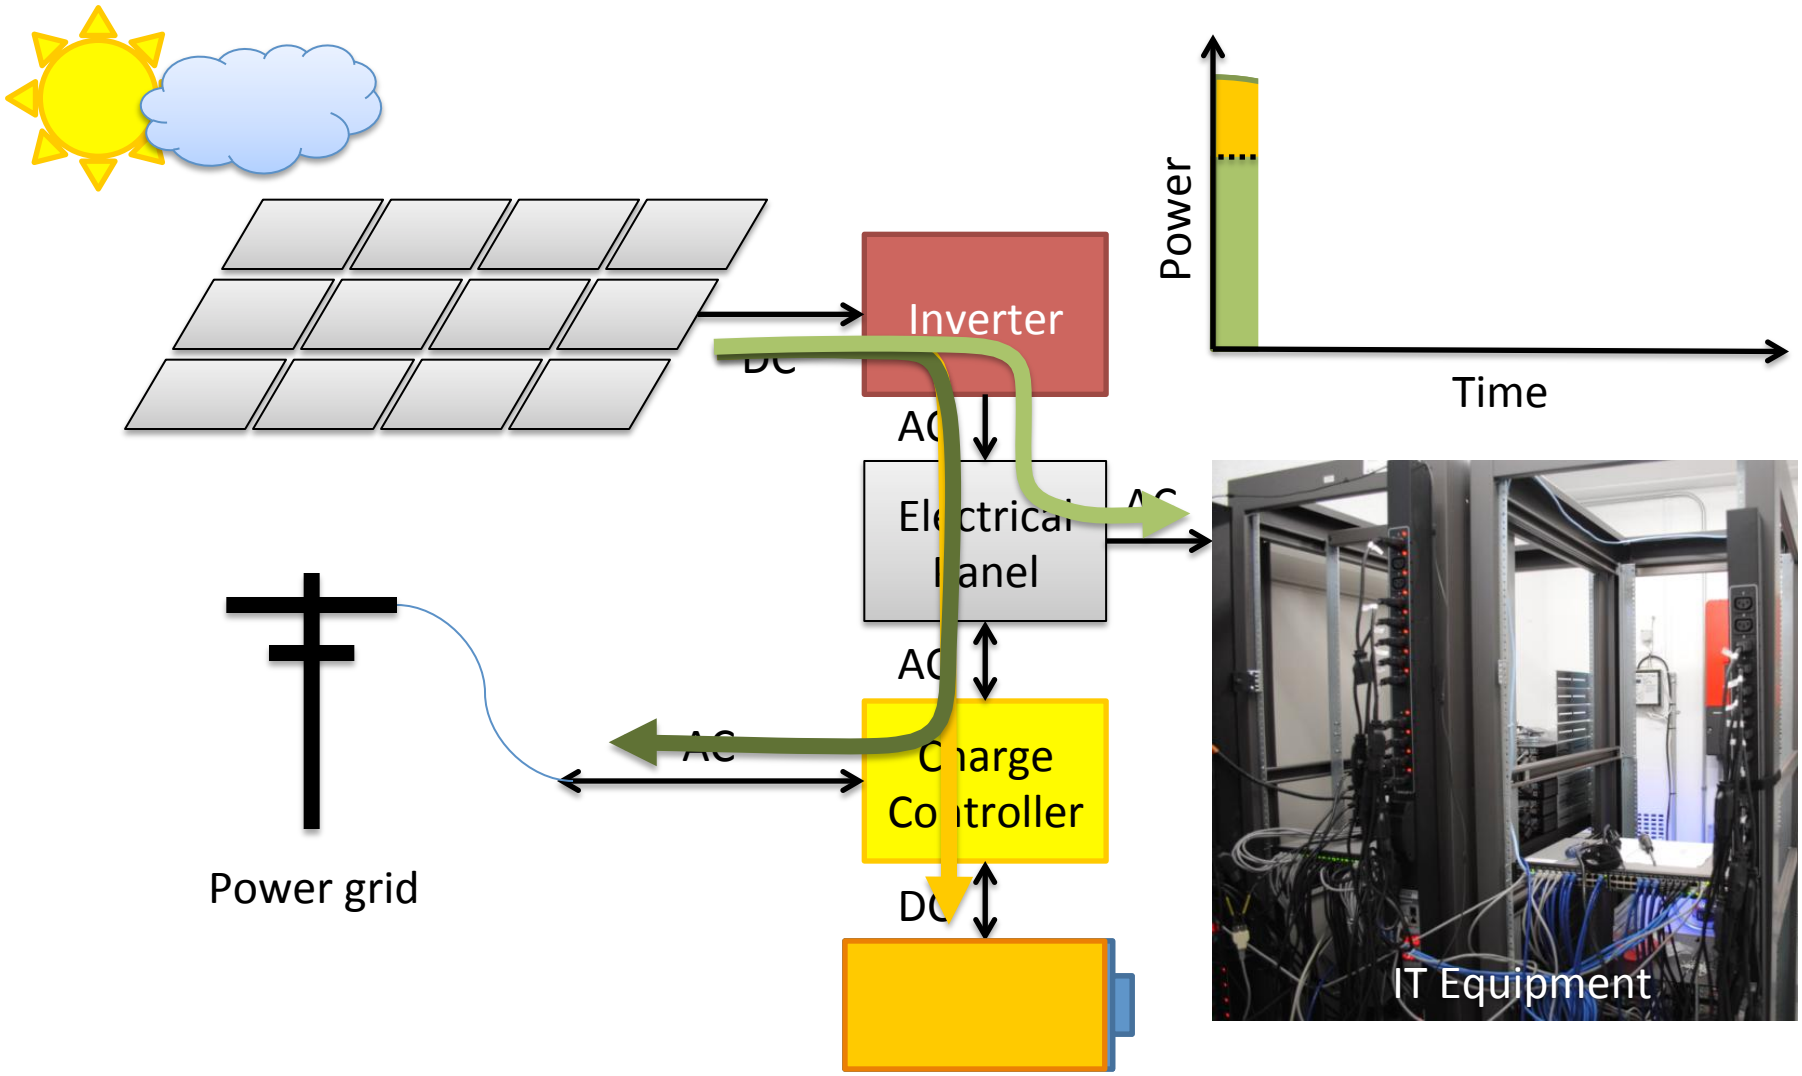

Example Energy Source Management

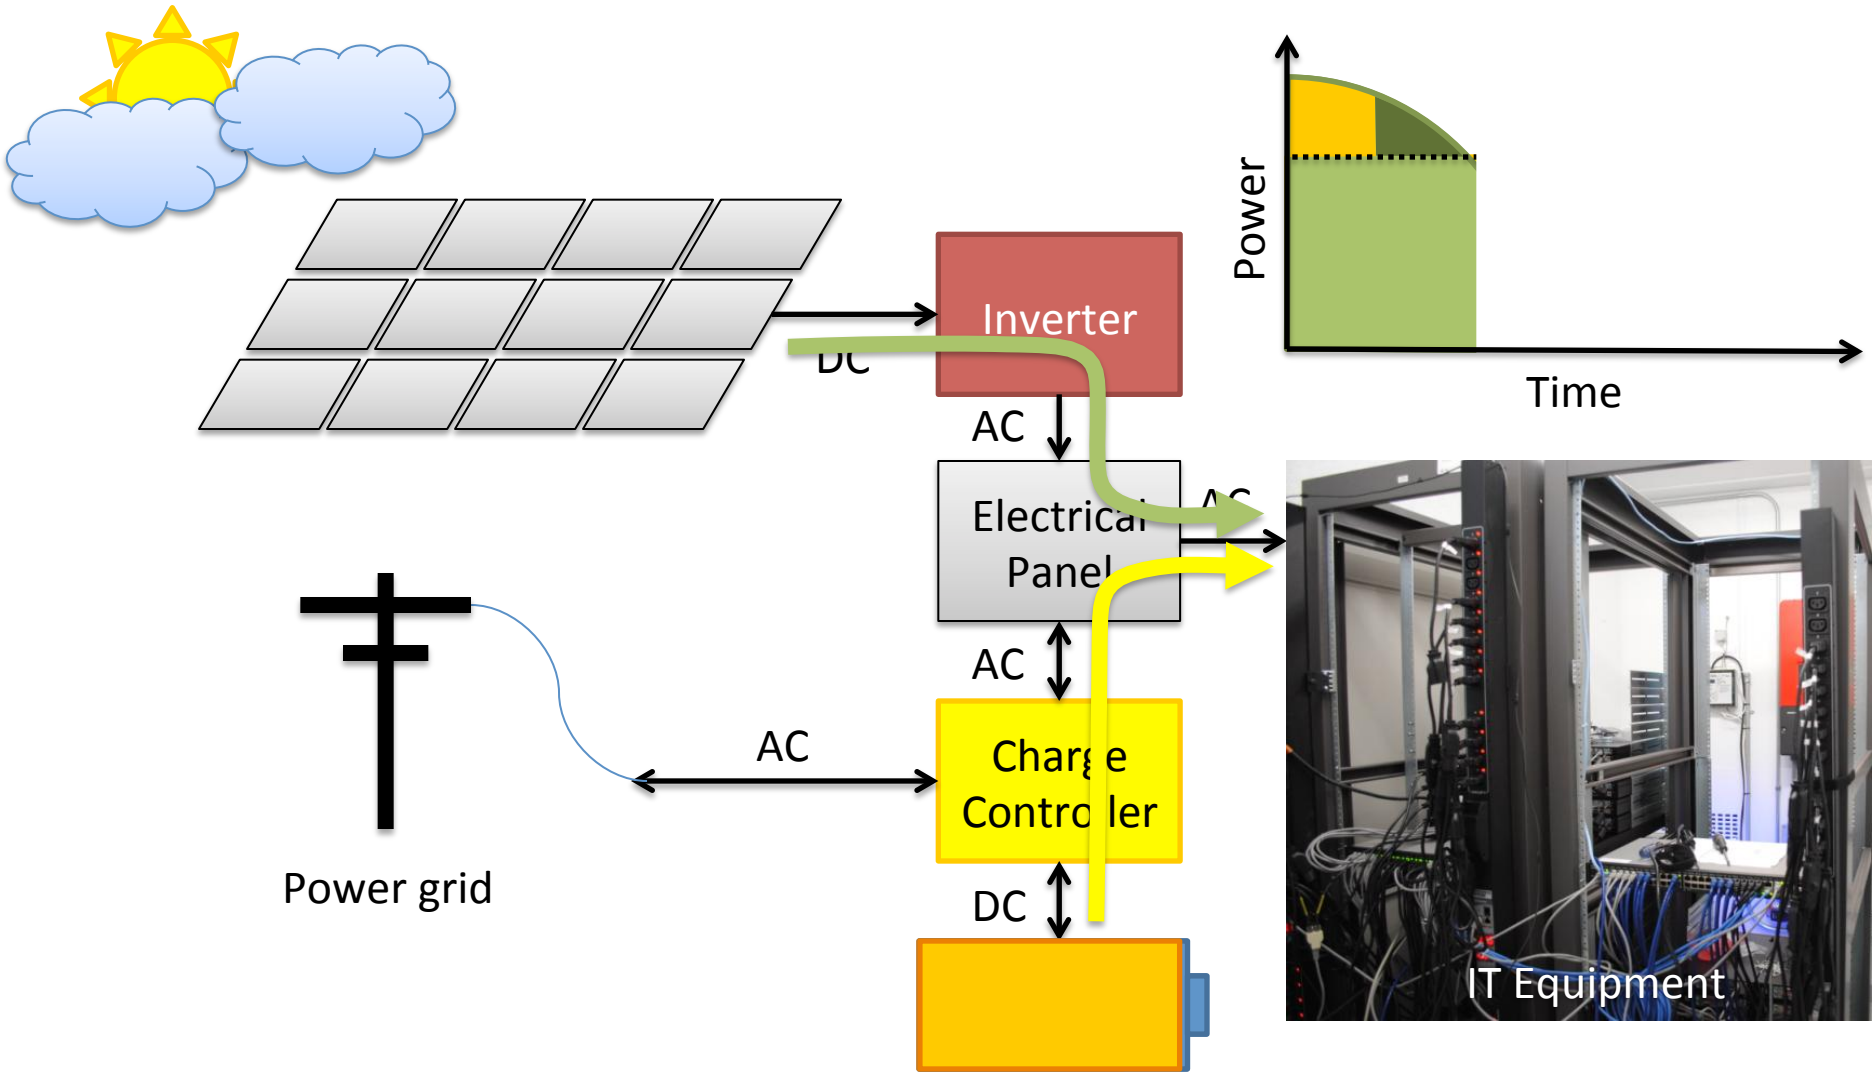

Example Energy Source Management

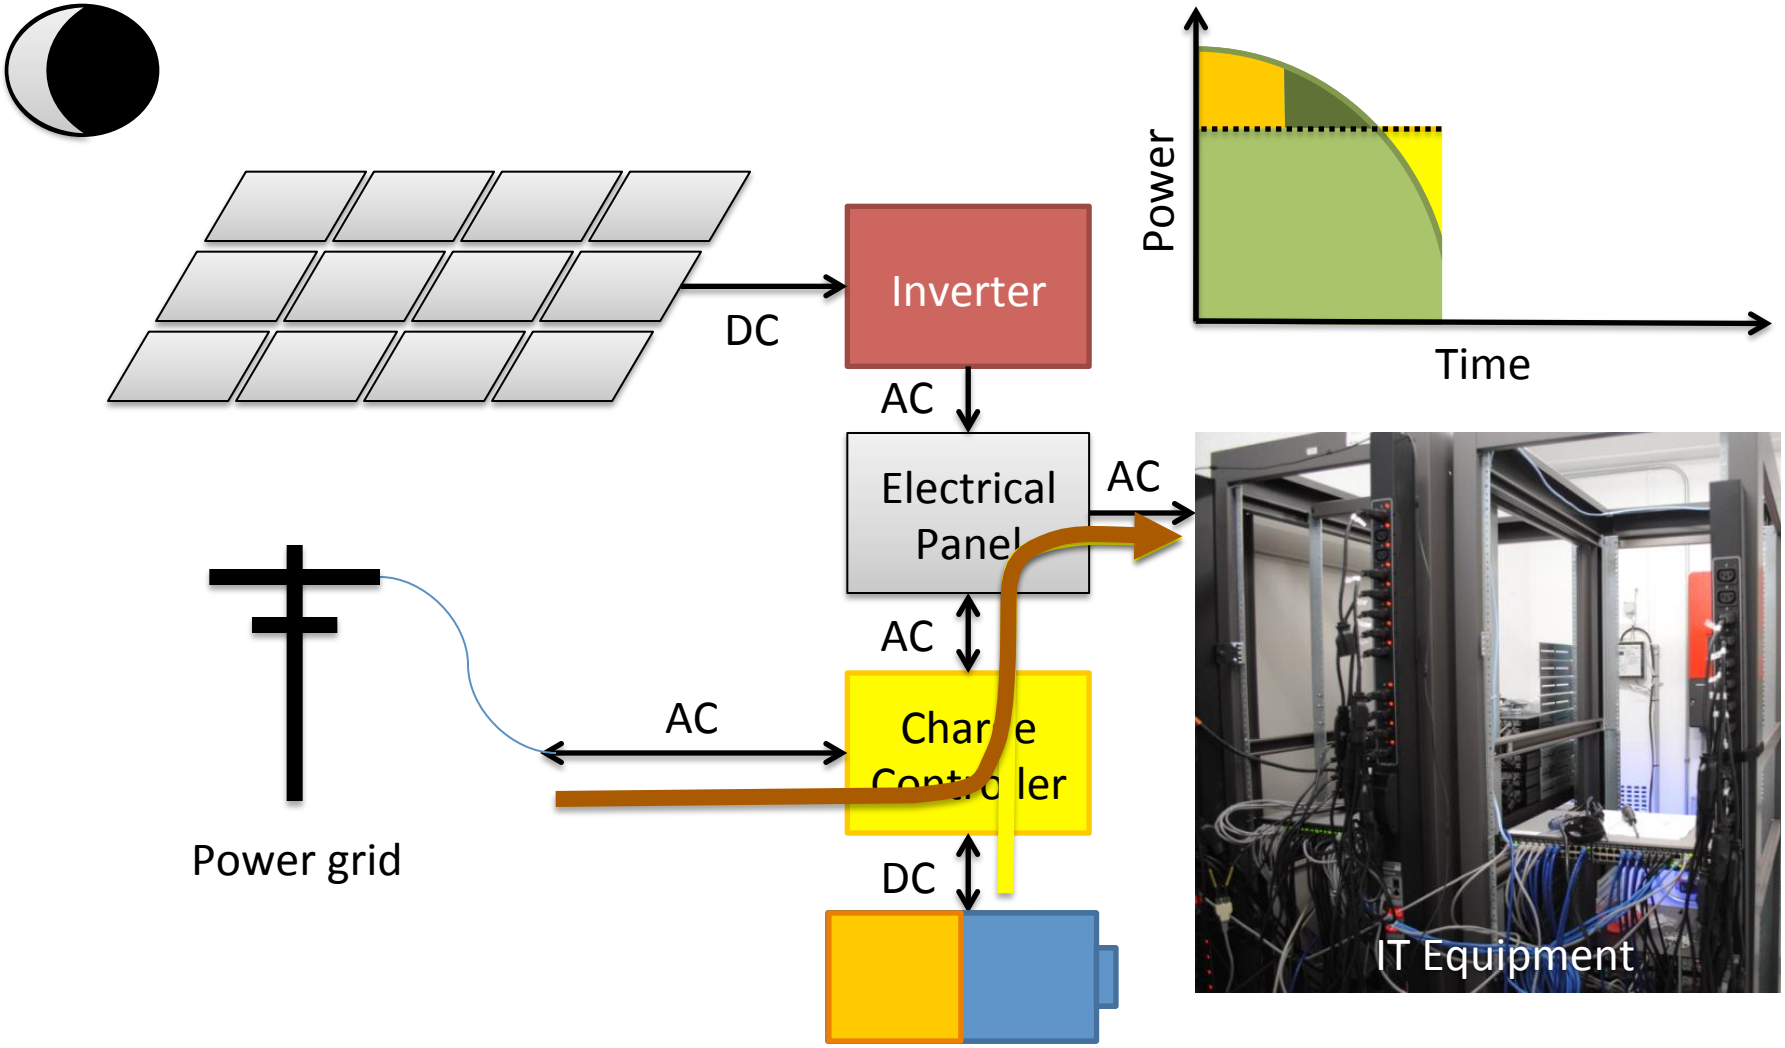

Possible Energy Source Management

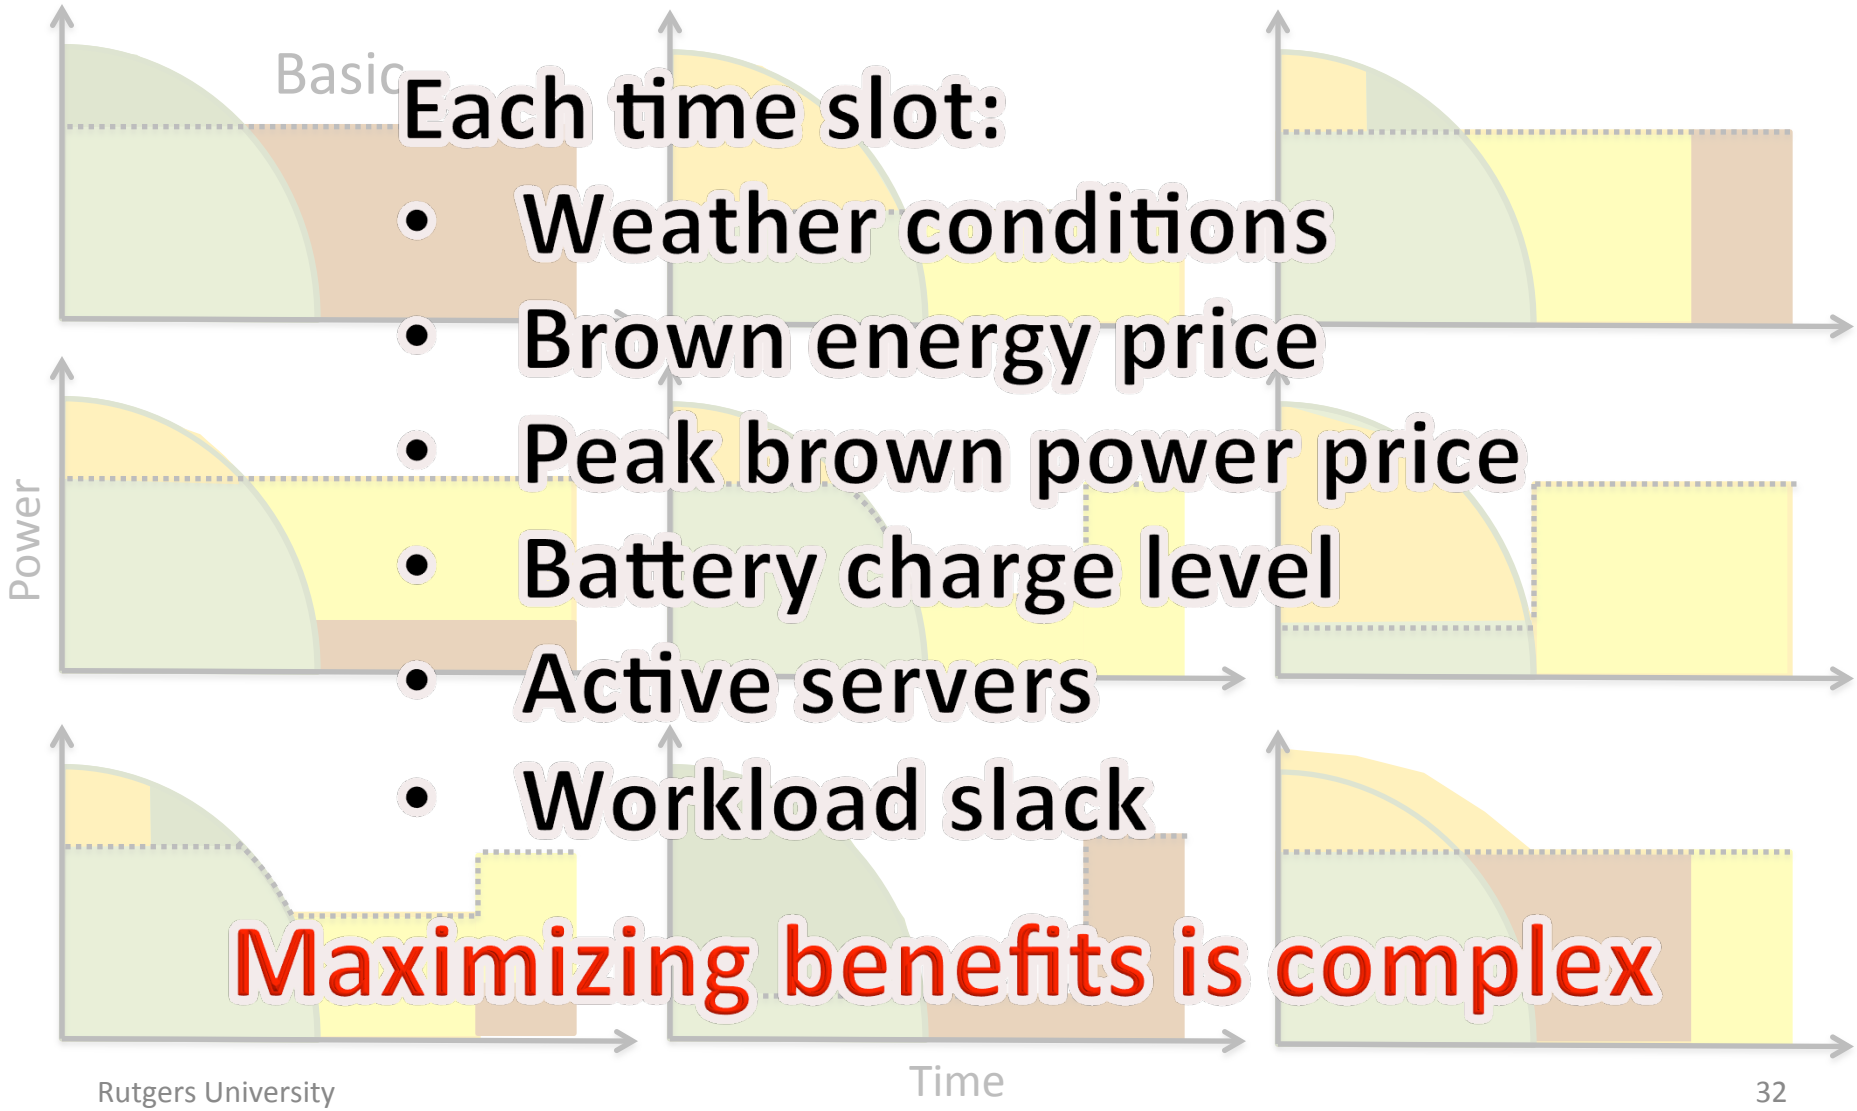

Each time slot:

- Weather conditions
- Brown energy price
- Peak brown power price
- Battery charge level
- Active servers
- Workload slack

Maximizing benefits is complex

GreenSwitch

- Minimize brown electricity cost
  - Brown energy
  - Peak brown power
  - Battery lifetime constraint
- Manage energy sources
  - Use solar/net metering
  - Charge/discharge battery
  - Limit brown peak power
- Manage workload
  - Turn servers on/off
  - Delay deferrable jobs

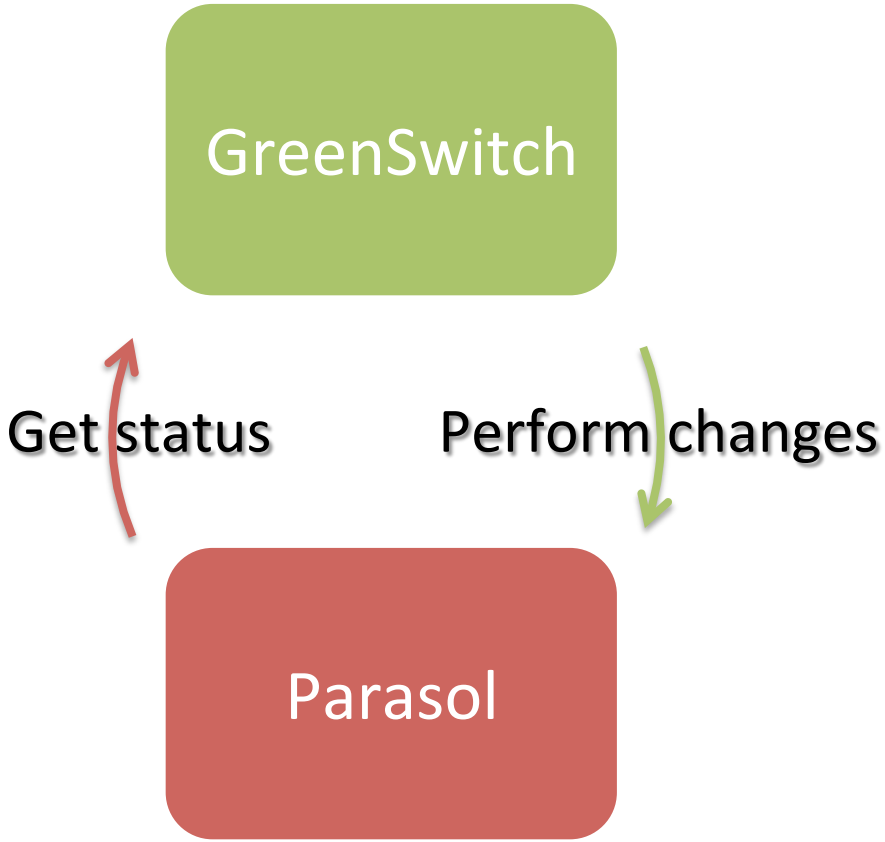

GreenSwitch Architecture

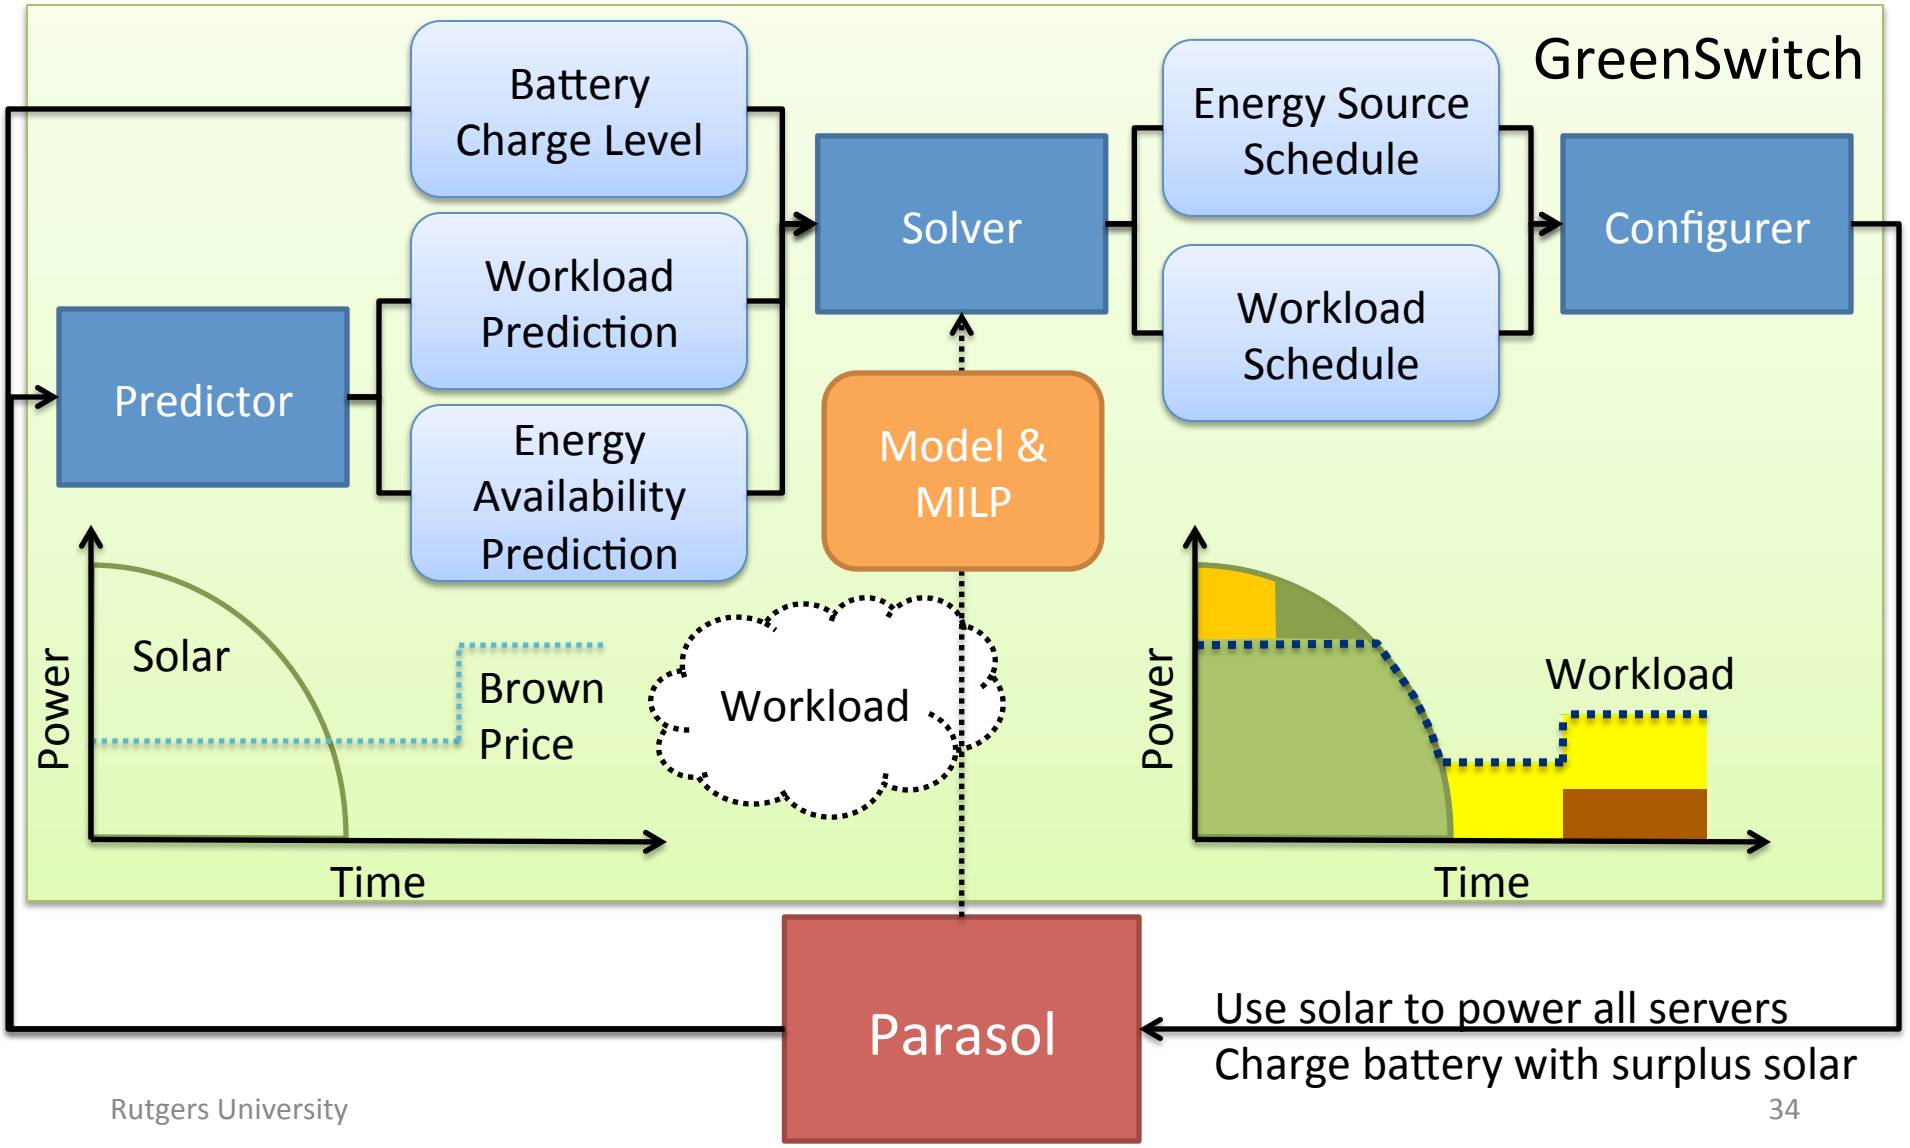

Experimental Environment

- Evaluation on 64 Parasol nodes
  - 12 one-day experiments
  - Deferrable vs. non-deferrable workloads
  - Baseline datacenter (no solar, no batteries, no delays)
- New Jersey brown electricity pricing
  - On/off-peak energy, peak power, net metering
- GreenSwitch for Hadoop (configurer)

Experimental Environment

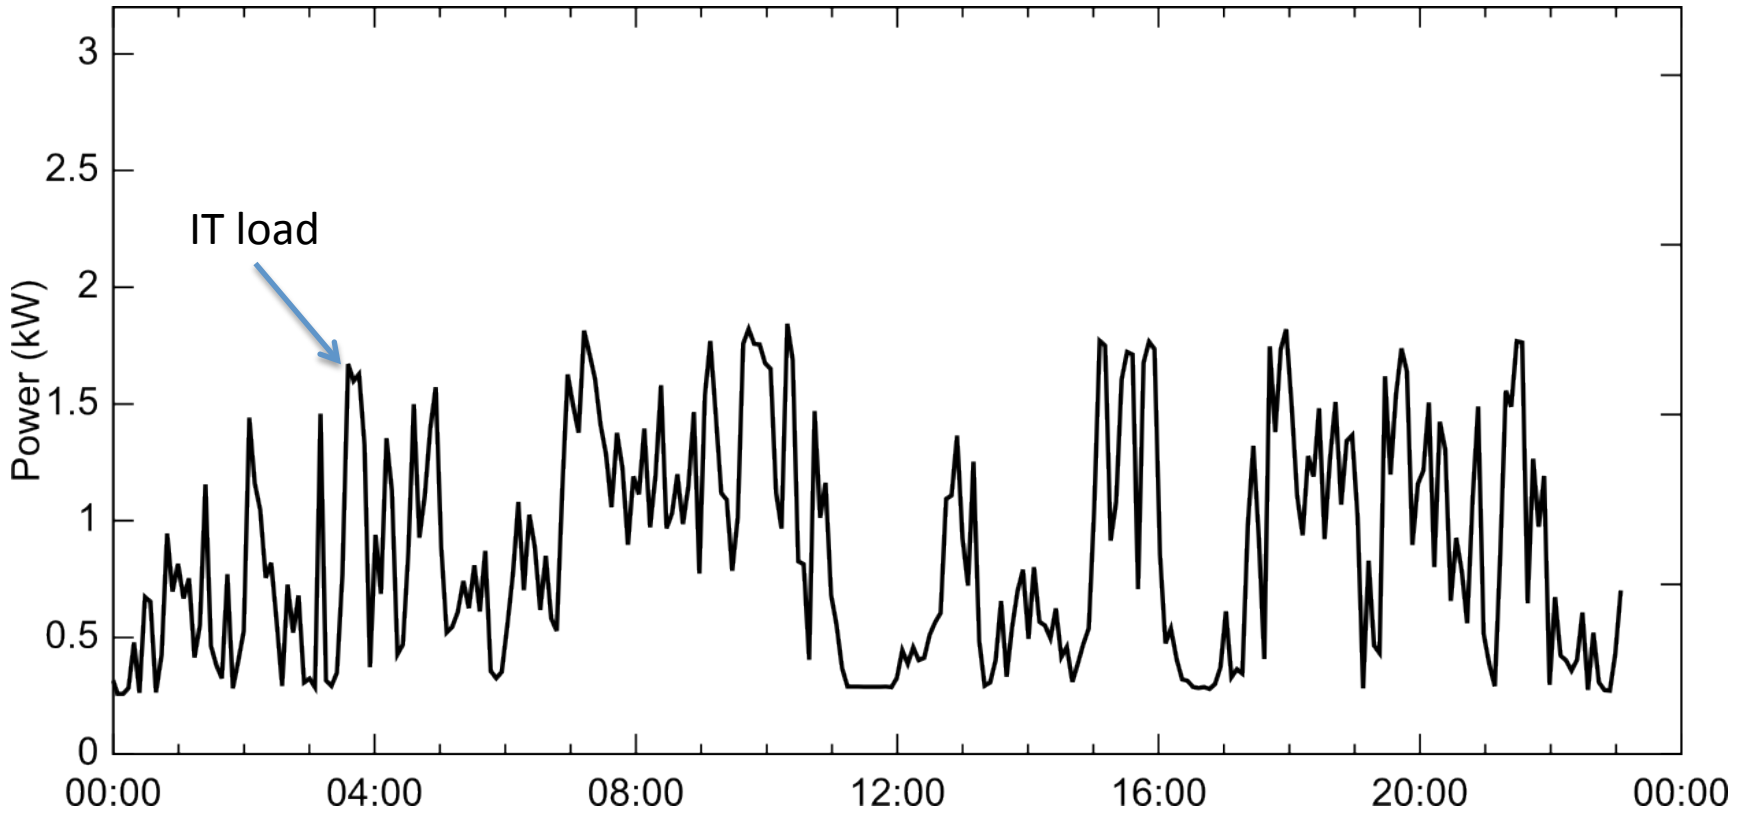

SWIM: Facebook based workload [MASCOTS'11]

Parasol Without GreenSwitch

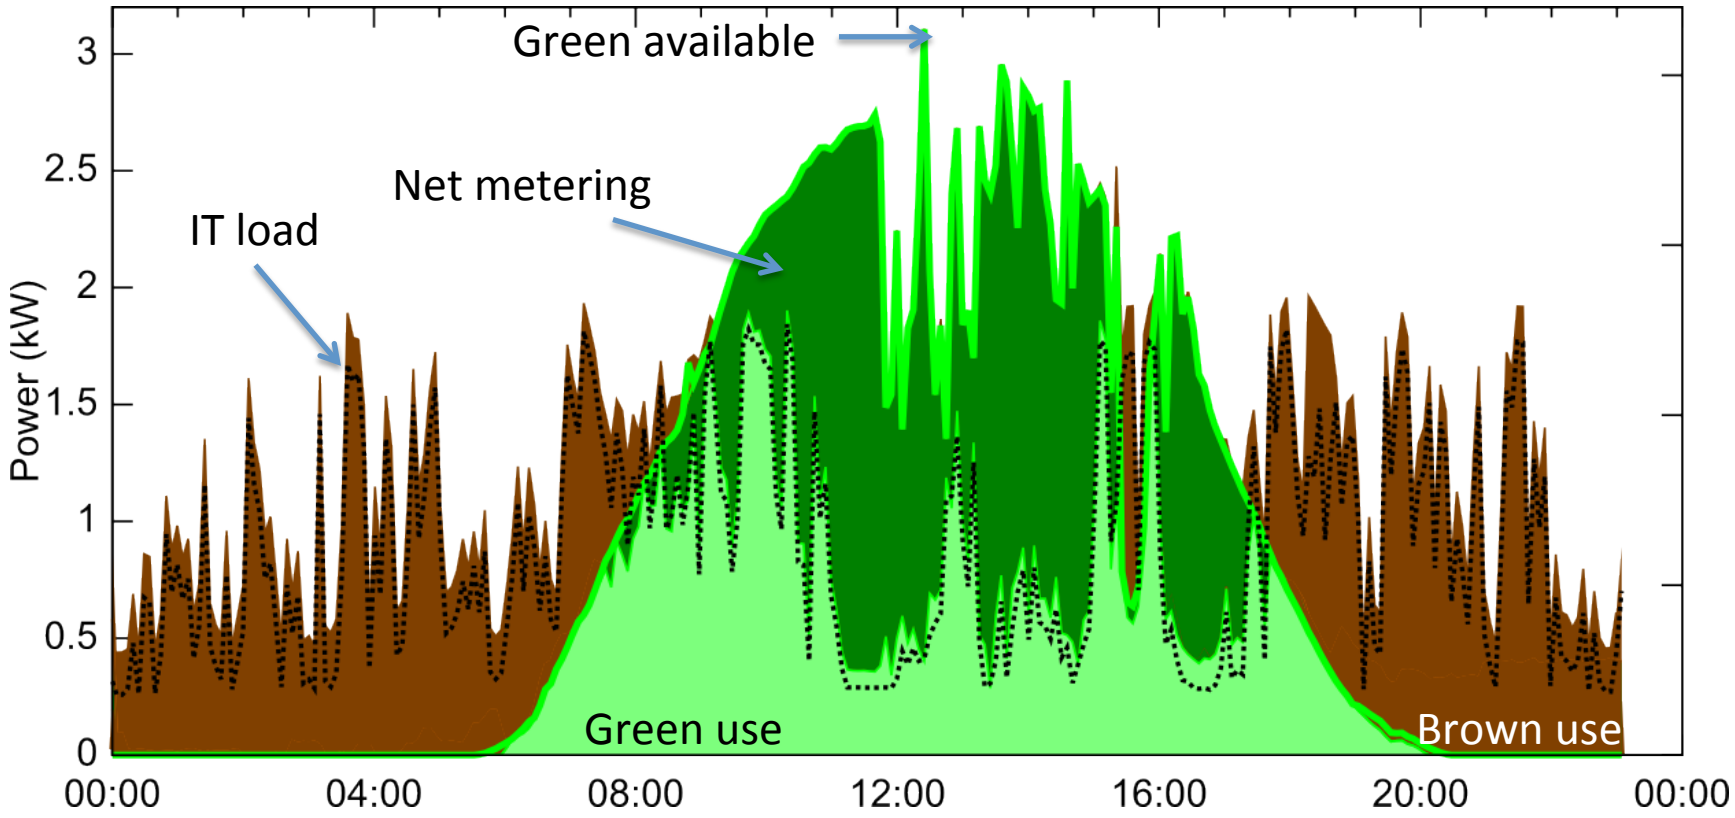

66% cost savings → Solar amortized in 7 years

GreenSwitch: Non-Deferrable Workload

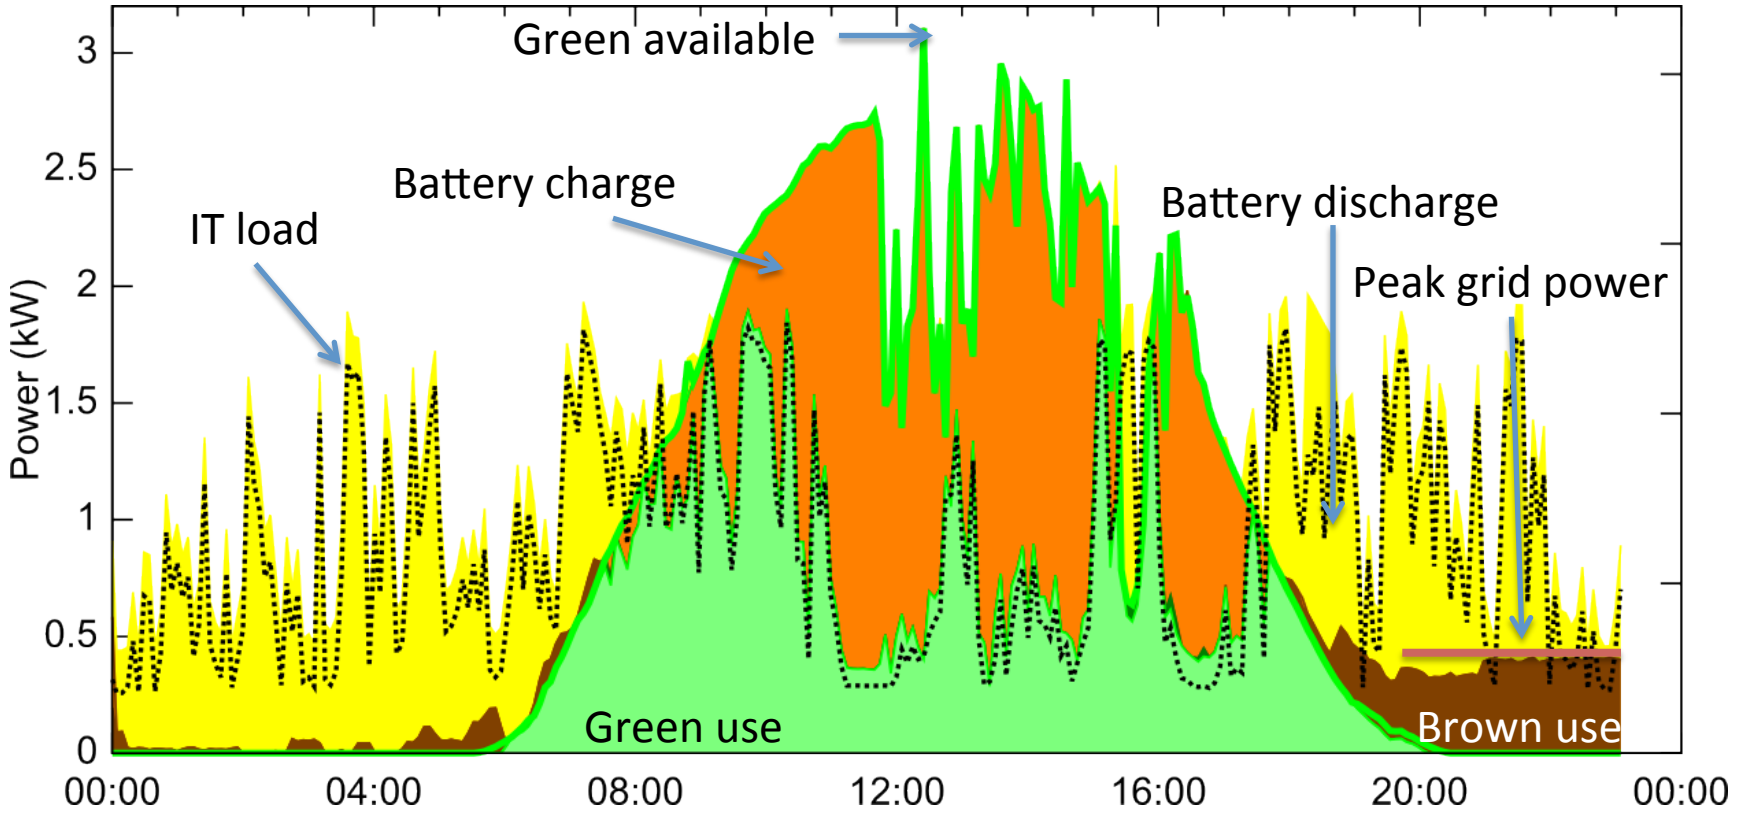

75% cost savings → Batteries cannot be amortized

GreenSwitch: Deferrable Workload

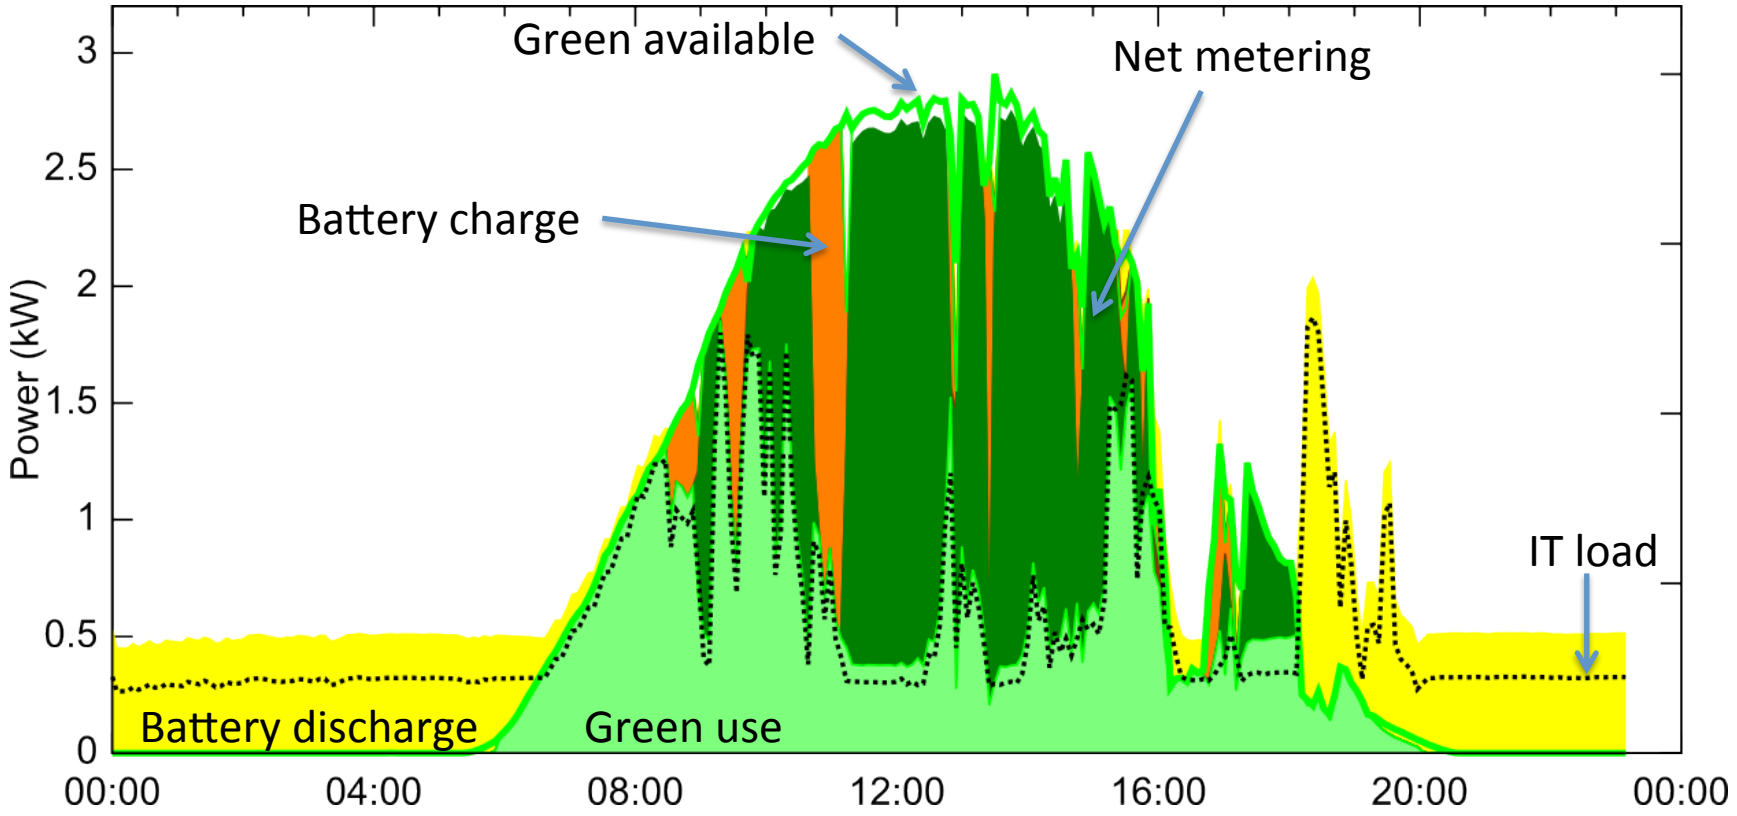

96% cost savings → Solar + batteries amortized in 7.6 years

Parasol: A Real System

- Real software running on real hardware
- Power losses
- Overhead of energy source switching
- System limitations
  - Net metering vs. battery charging
  - Use grid vs. net metering
  - Green battery charging vs. use grid

Outline

- DC energy usage and carbon footprint
- Reducing carbon footprint & cost with renewables
- Our target and research challenges
- Parasol: our solar-powered micro-datacenter
- GreenSwitch: managing power sources
- Previous, current, and related work
- Conclusions

Other Works on Green DC Software

- Follow the renewables [HotPower'09, IGCC'10]
- Delay batch jobs while respecting deadlines [SC'11, Eurosys'12]
- Power source management [SustainIT'12, ASPLOS'13]
- Green SLAs in HPC clouds [IGCC'13]

Current Works

- Temperature- and variation-aware management of free-cooled datacenters
- DC placement world-wide for cost-effective follow-the-renewables
- Matching power demand to power supply for non-deferrable workloads
  - Trading-off response time and durability for interactive workloads (GreenCassandra)

Related Work

- Blink, UMass, Amherst
- Algorithms for Sustainable IT, Caltech
- Sustainable datacenters, HP Labs
- Ren et al., MASCOTS 2012
- Li et al., ISCA 2012

Conclusions

- Greening datacenters
  - Challenges & opportunities
  - Hardware/software solution
- GreenSwitch benefits
  - Delaying load and solar gives the best results
  - Reduces amortization time by 1.8-2x
  - Flexibility: no batteries, workloads, wind …

http://parasol.cs.rutgers.edu