Leveraging Renewable Energy in Data Centers: Present and Future

Ricardo Bianchini

Department of Computer Science

Collaborators: Josep L. Berral, Inigo Goiri, Jordi Guitart, Md. Haque, William Katsak, Kien Le, Thu D. Nguyen, Jordi Torres

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Motivation

- Data centers = machine rooms to giant warehouses
- Consume massive amounts of energy (electricity)

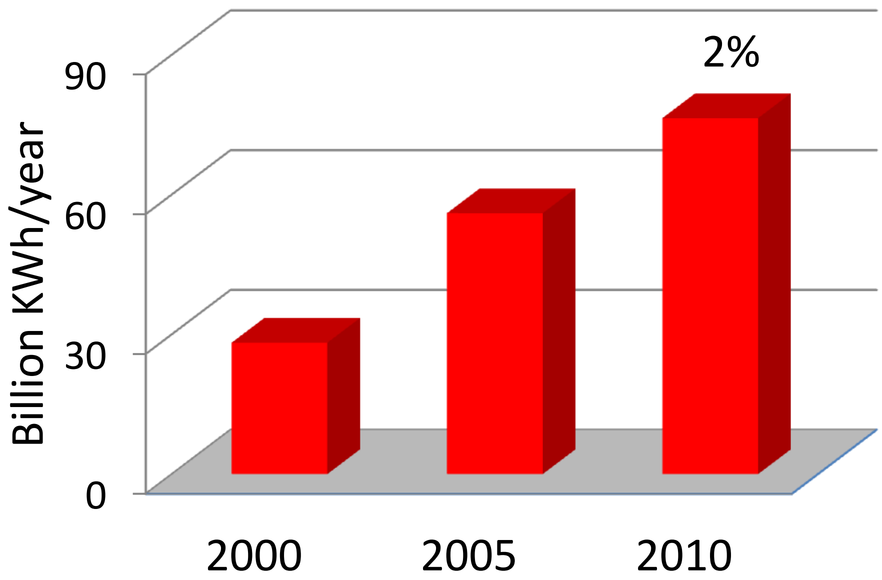

Electricity consumption of US DCs [JK'11]

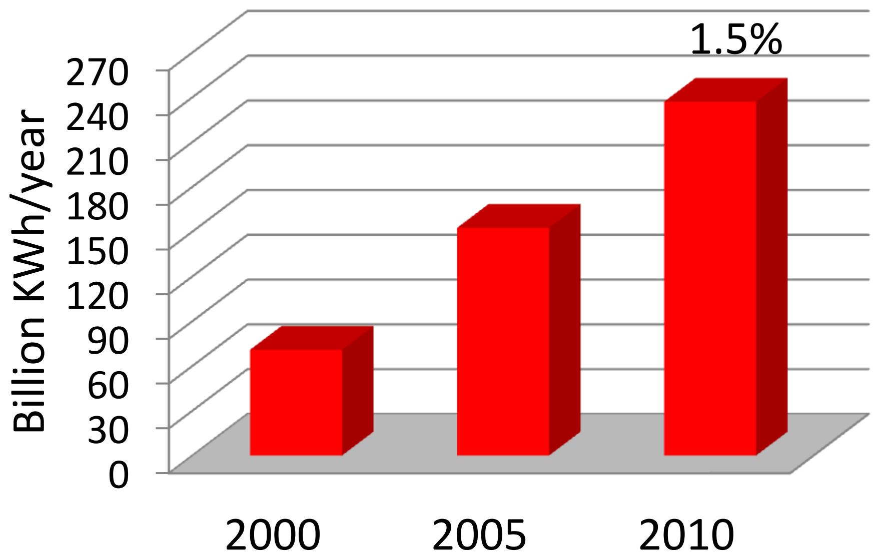

Electricity consumption of WW DCs [JK'11]

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Motivation

- Electricity comes mostly from burning fossil fuels

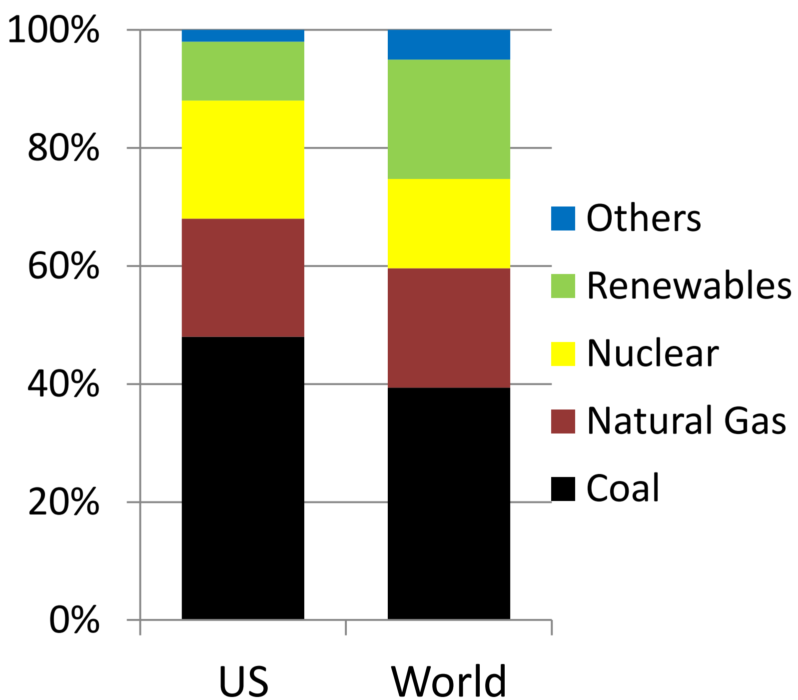

Electricity sources in US & WW [DOE'10]

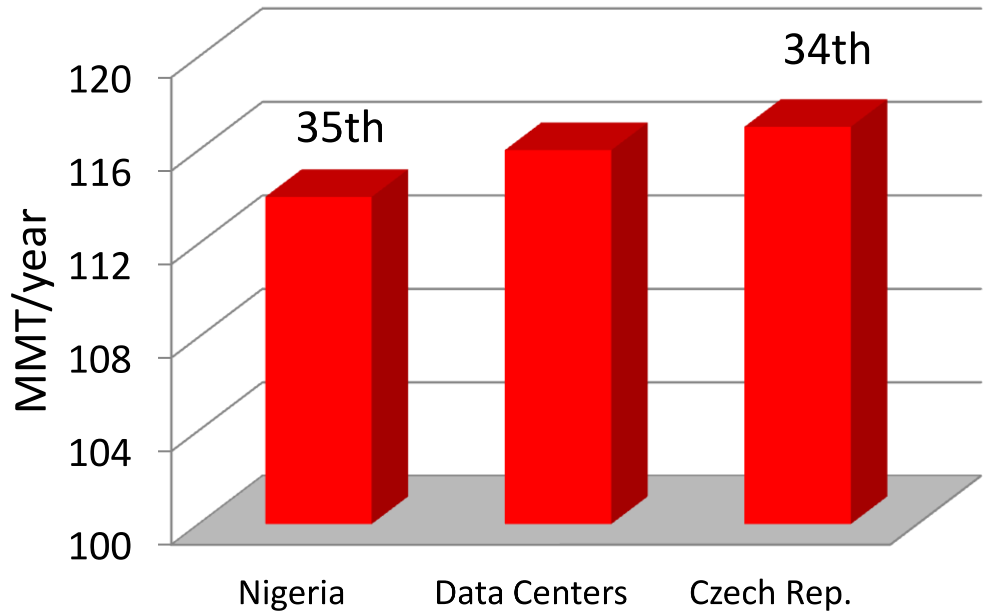

CO_2 of world-wide DCs [Mankoff'08]

Can we use renewables to reduce this footprint?

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Outline

- DC energy usage and carbon footprint
- Reducing carbon with renewables: 2 approaches
- Our target and research challenges
- Software for leveraging solar energy
- Parasol: our solar micro-data center
- Current and future works
- Conclusions

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Greening DCs: Grid-centric approach

- Pump renewables into the grid
- Pros:
  - If the grid is available, power is available
  - DC operator need not worry about renewable plants
  - Plants can be placed at the best possible locations
- Cons:
  - Energy losses of ~15% [IEC'07]
  - Dependence on the power grid or diesel generators
- Example: Google buys wind power from NextEra

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Greening DCs: Co-location approaches

- (1) Build DC near a renewable plant or (2) self-generate
- Pros:
  - Reduced energy losses: ~5%
  - No dependence on the grid
  - Lower peak-power/energy costs, after amortization period (2)
- Cons:
  - Location may not be good for DC (1) or renewable plant (2)
  - Energy may have already been committed (1)
  - Need to install and maintain renewable plant (2)
- Examples: Microsoft built DC near hydro plant in OR (1)
  Apple is building 20MW solar array in NC (2)

Outline

- DC energy usage and carbon footprint
- Reducing carbon with renewables
- Our target and research challenge
- Software and hardware for leveraging solar energy
- Current and future works
- Conclusions

RUTGERS
THE STATE UNIVERSITY
OF NEW JERSEY

Our target

- No approach is perfect
  - Different DC operators may take different approaches
- Co-location or self-generation with solar and/or wind
  - Pros: Clean and available
  - Cons: Space and cost

Solar and wind are clean

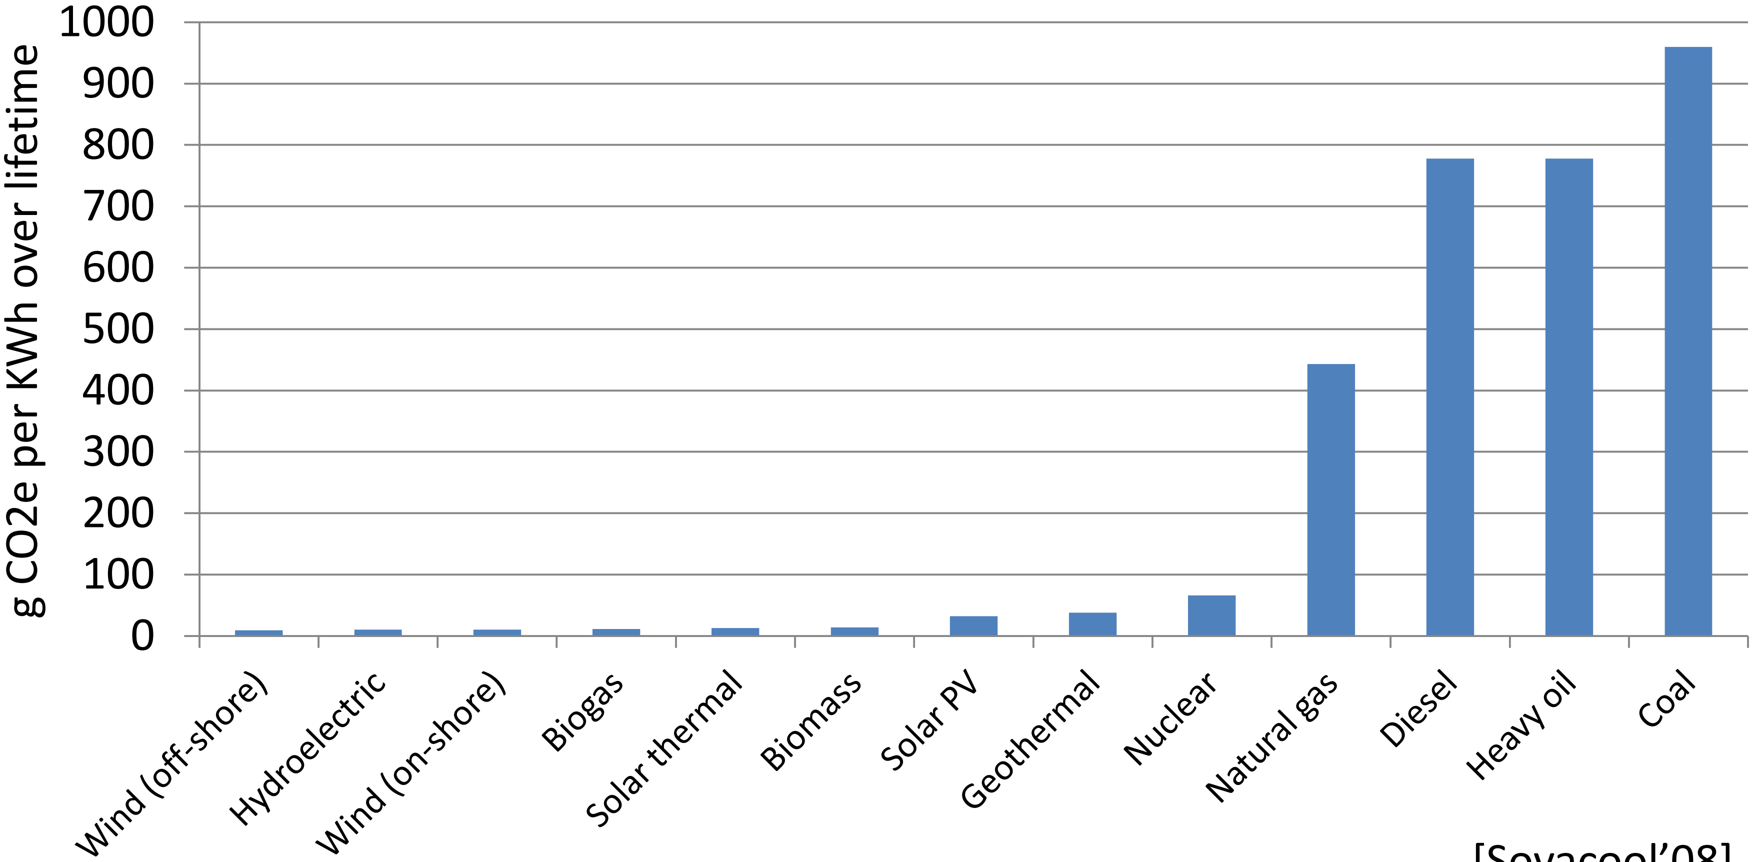

[Sovacool'08]

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Solar and wind are clean

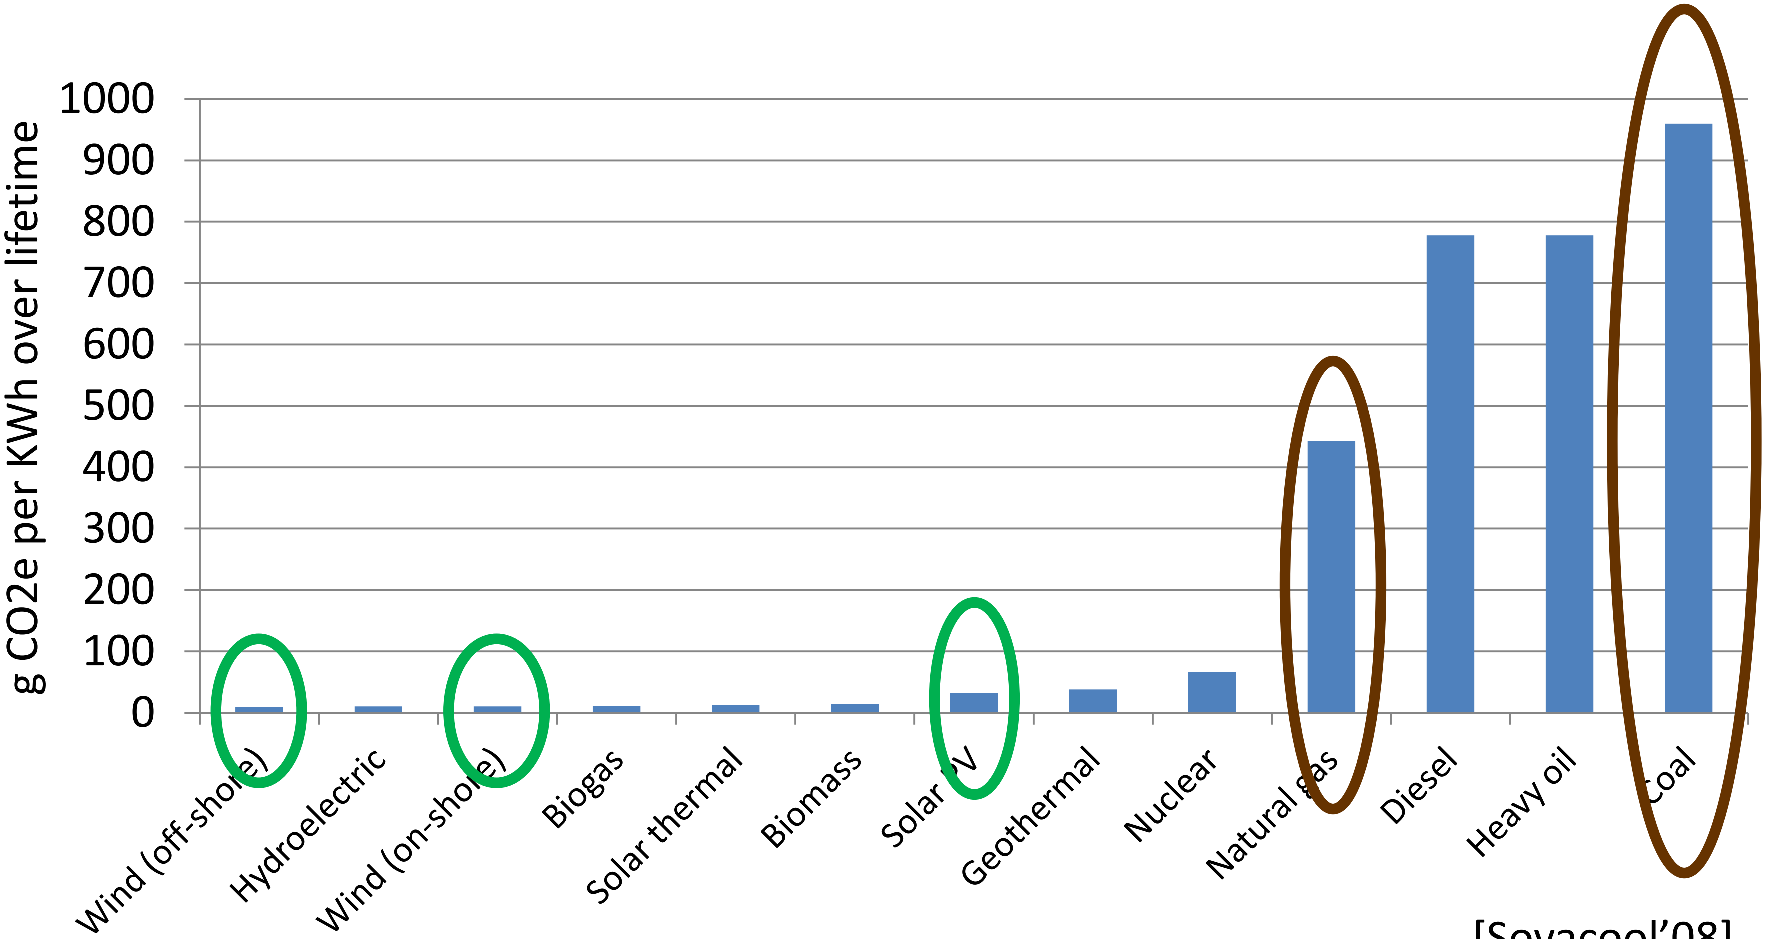

[Sovacool'08]

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Solar is more available in the US

Wind

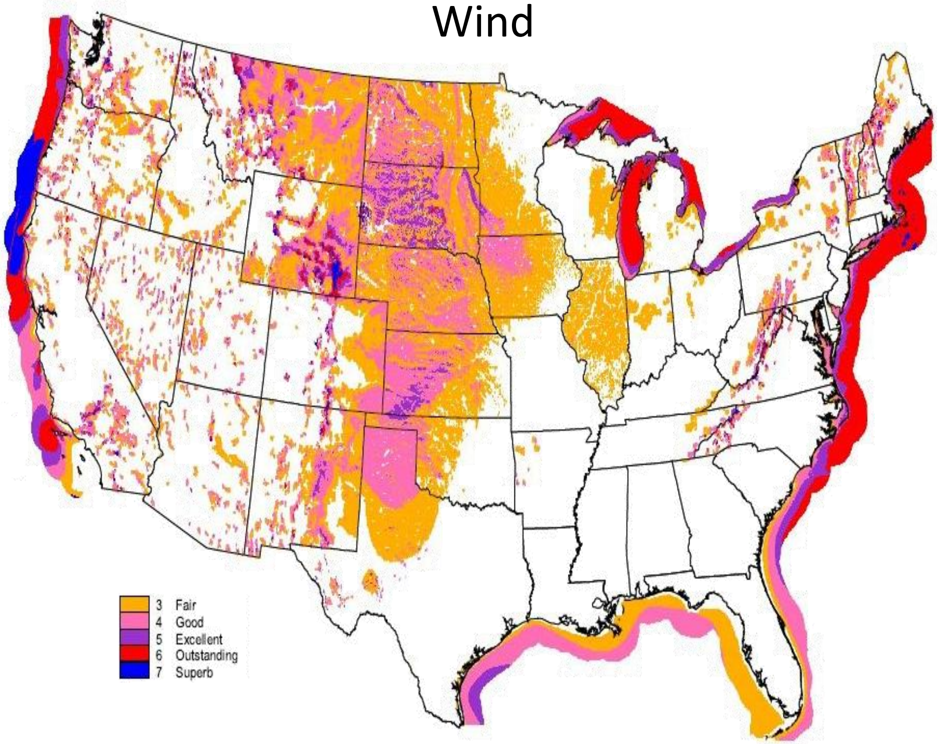

Solar

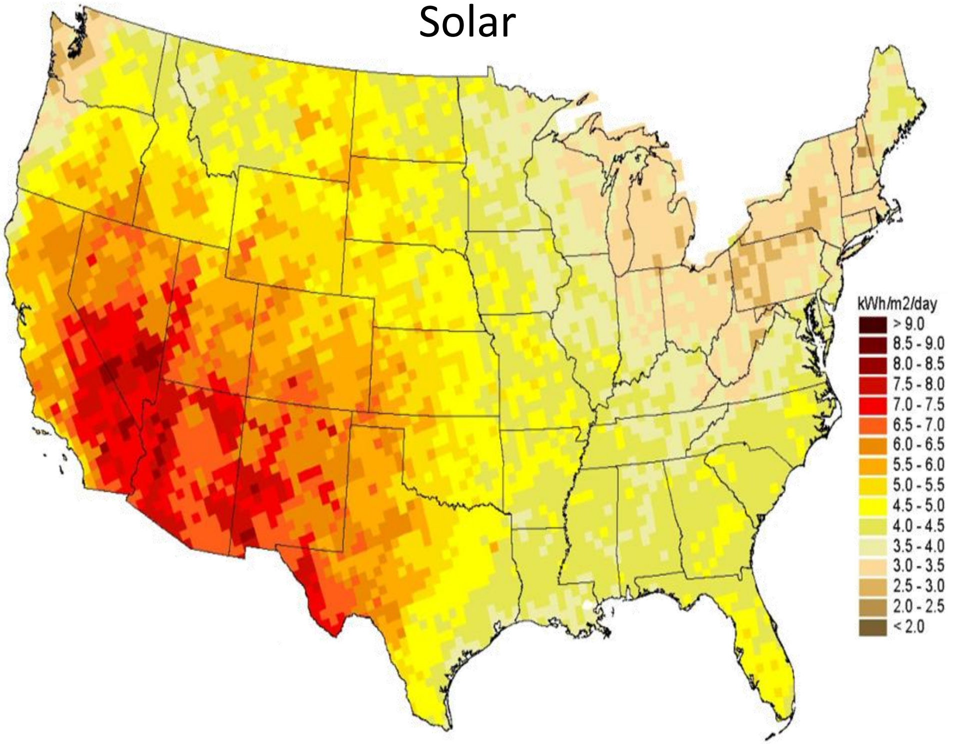

[NREL'12]

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Space: Solar PV efficiencies are increasing

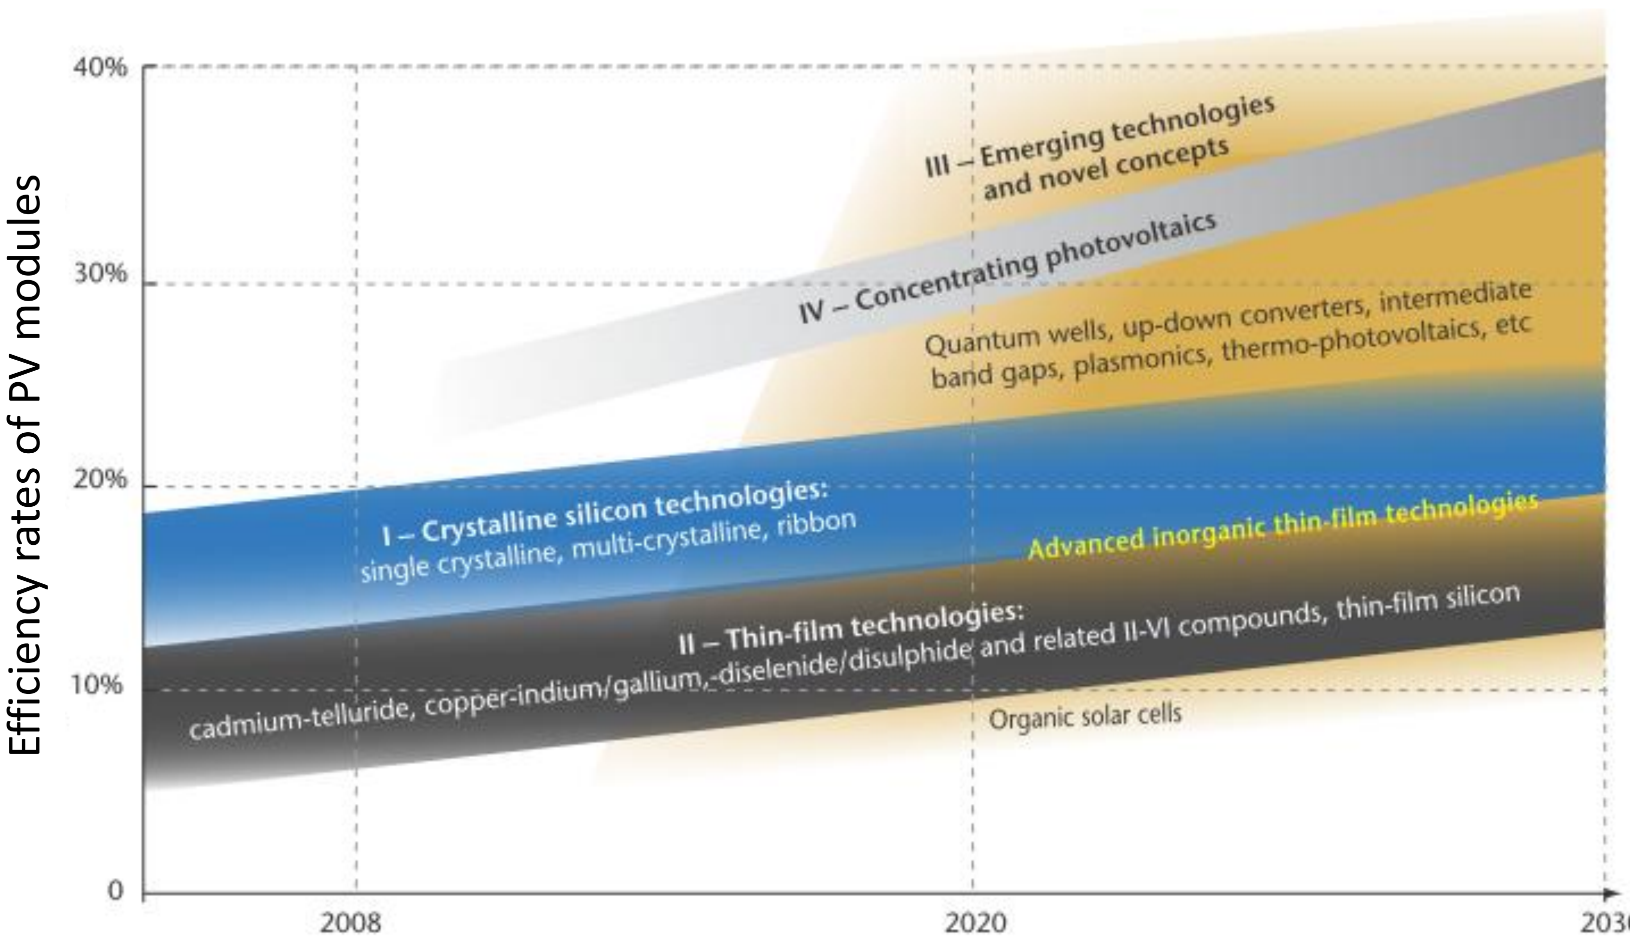

[IEA'10]

Space: Solar PV capacity factors today

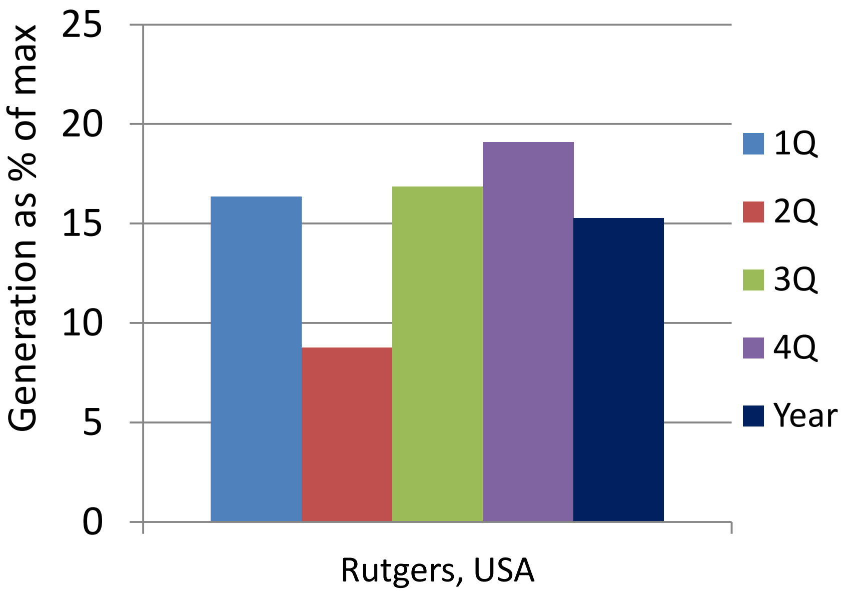

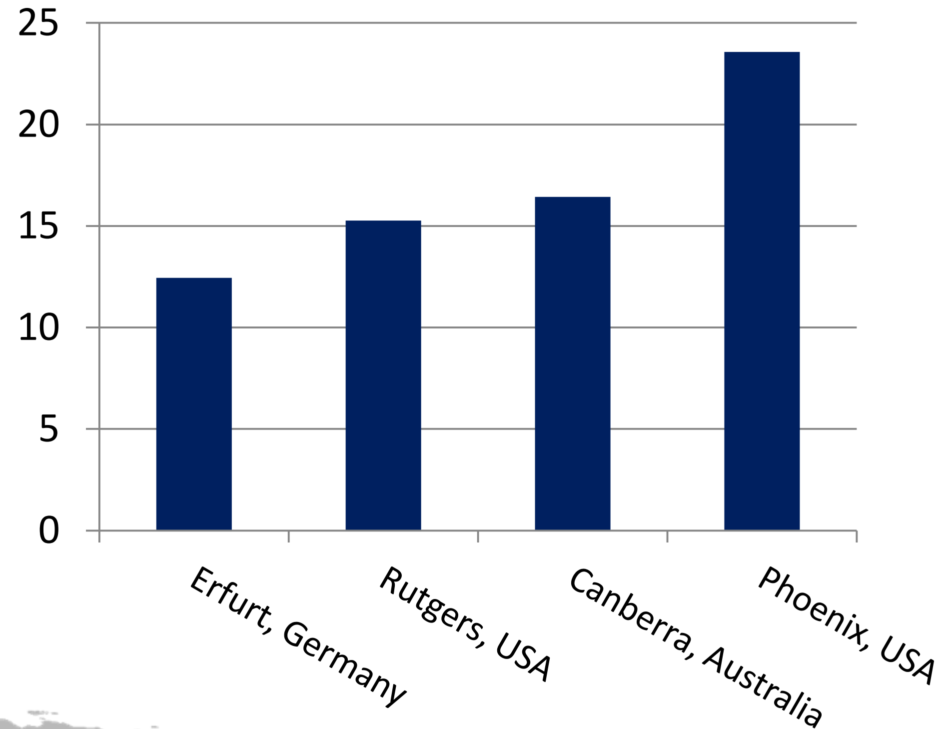

[PVOutput'12]

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Cost of solar PV energy is decreasing

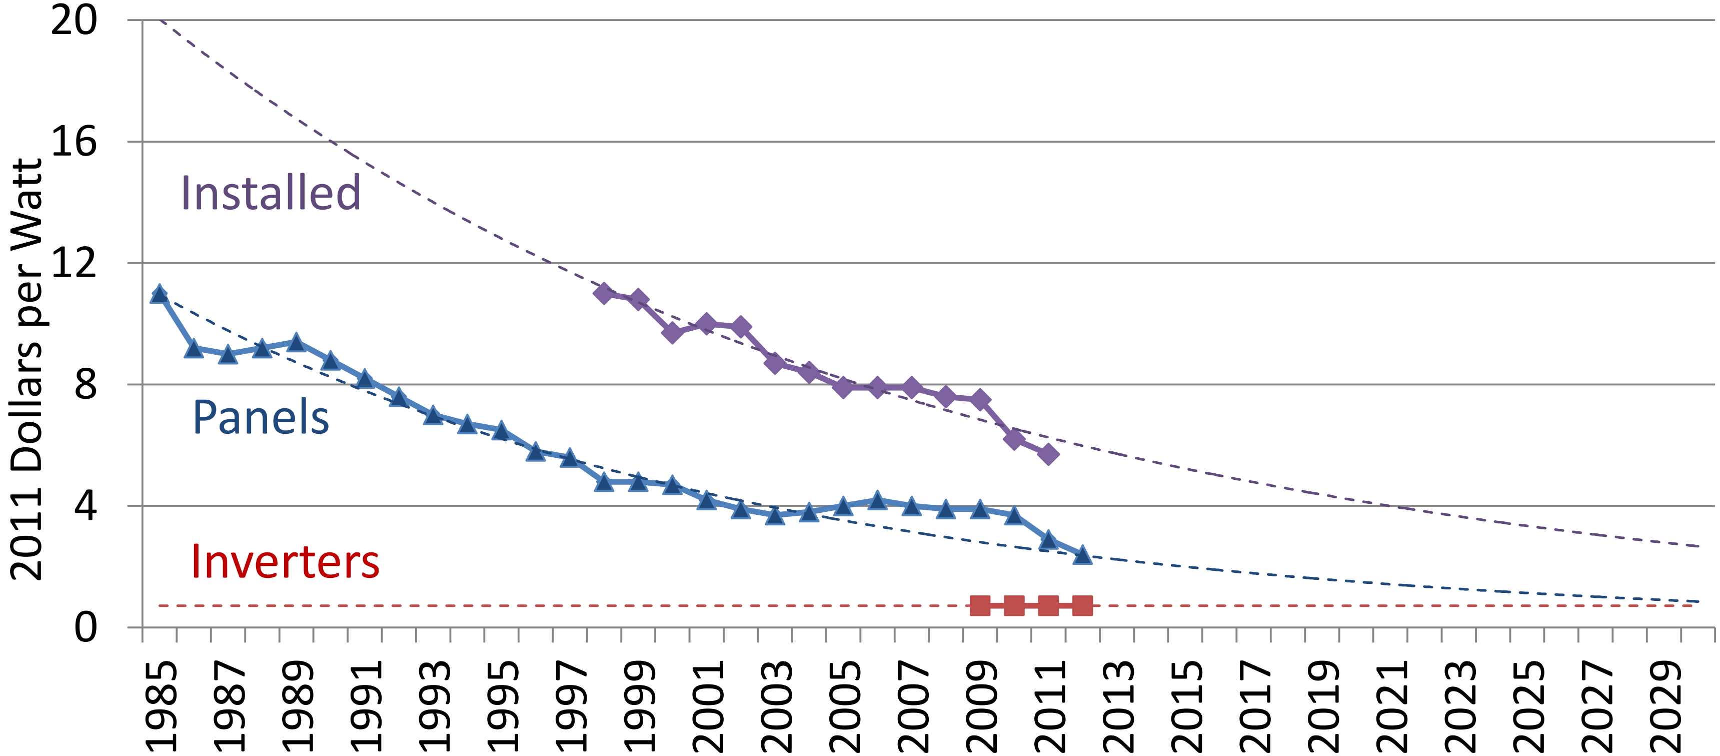

[DOE'11,Solarbuzz'12]

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

Cost of solar PV energy is decreasing

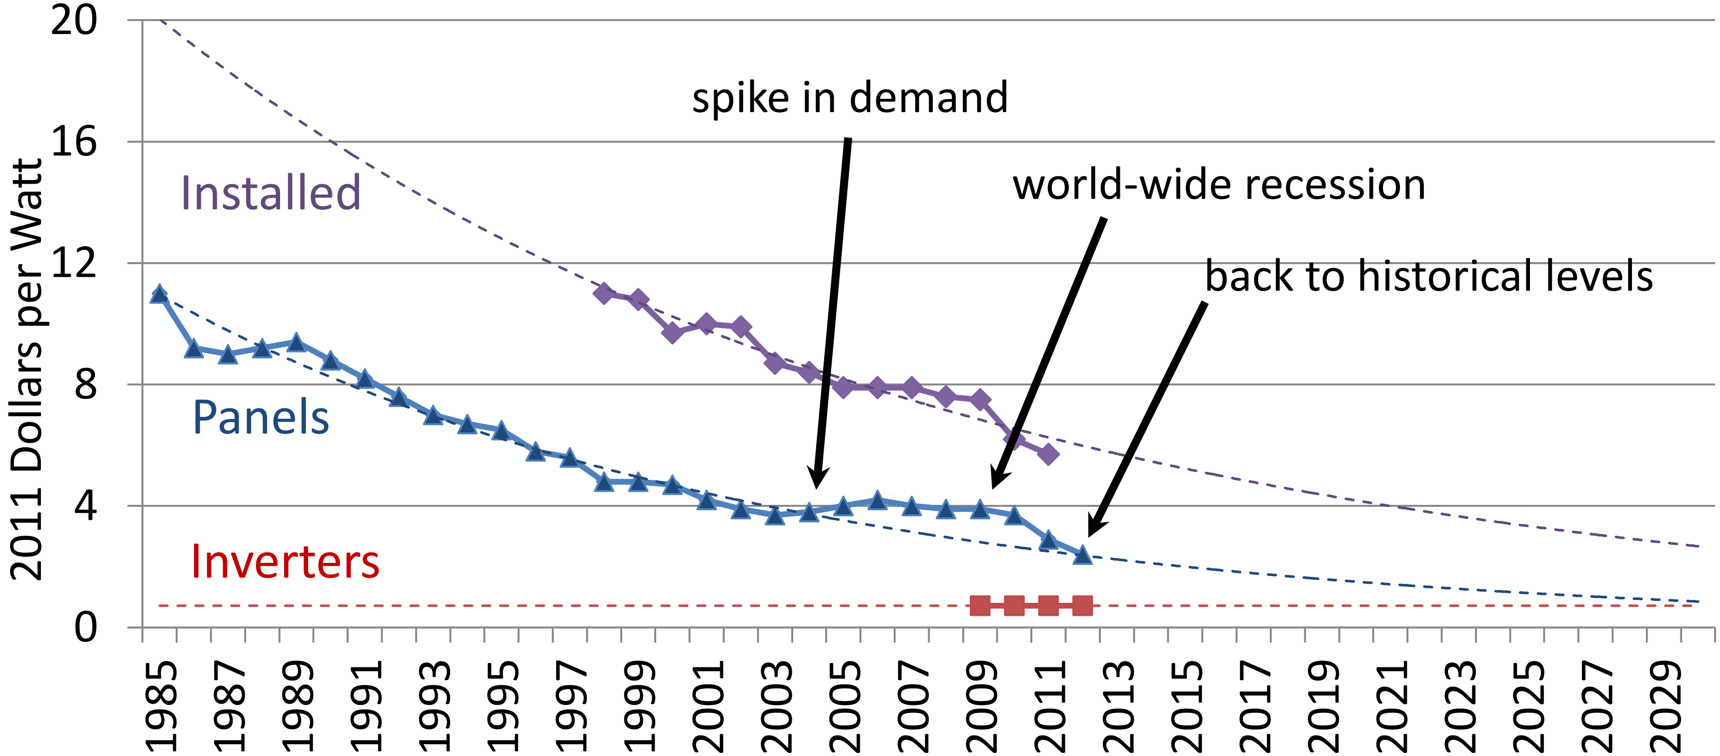

[DOE’11,Solarbuzz’12]

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Cost of solar PV energy is decreasing

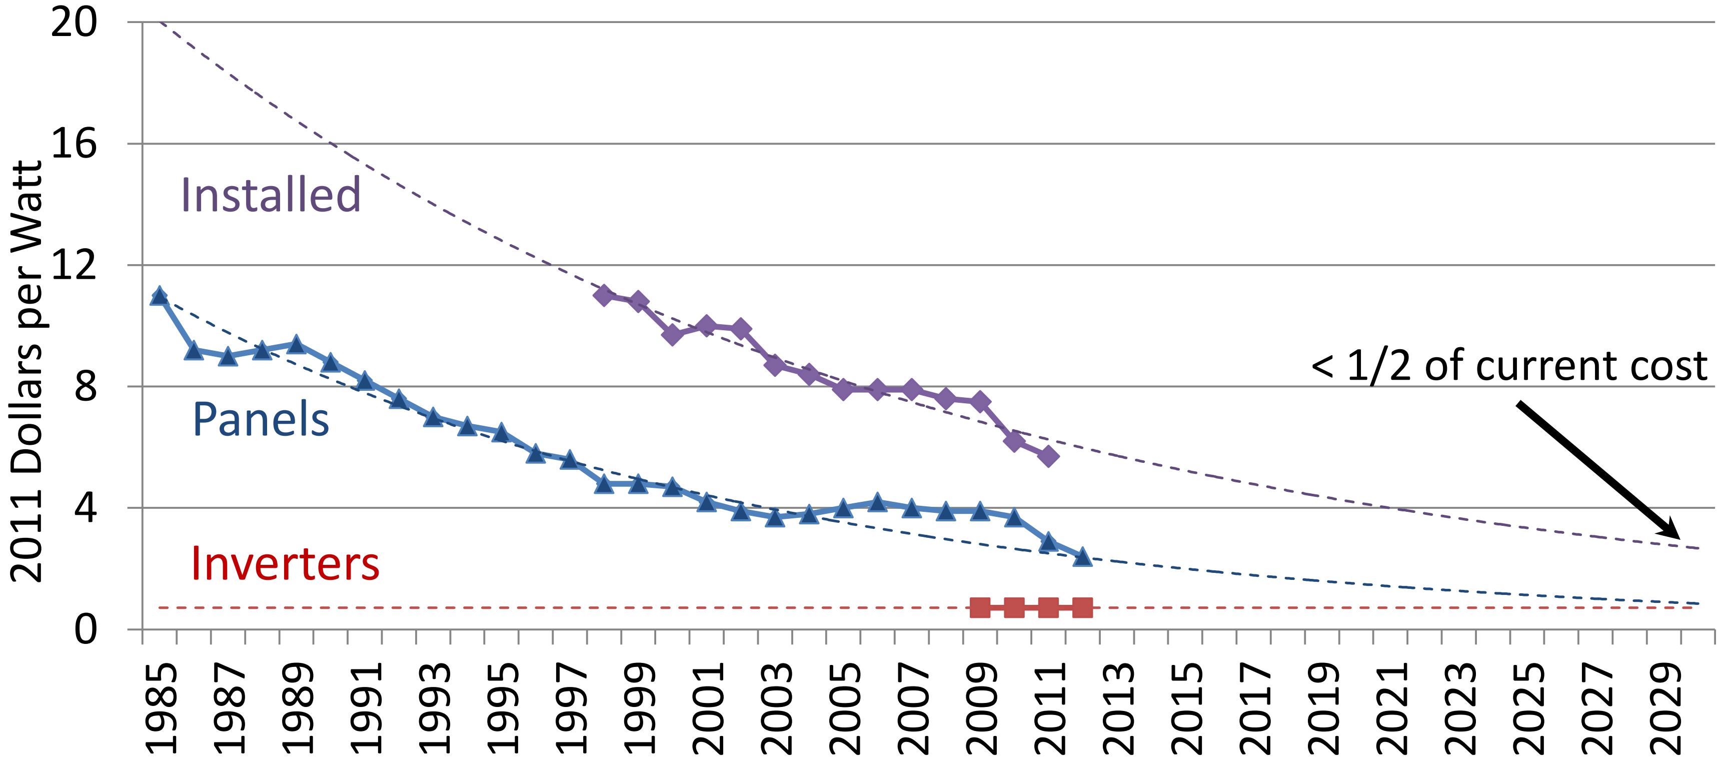

[DOE'11,Solarbuzz'12]

With incentives, the installed price can go down by another 50-60%

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

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

Assuming self-generation and federal + NJ incentives

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

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

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

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

Assuming self-generation and federal + NJ 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

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Main challenge: Supply of power is variable!

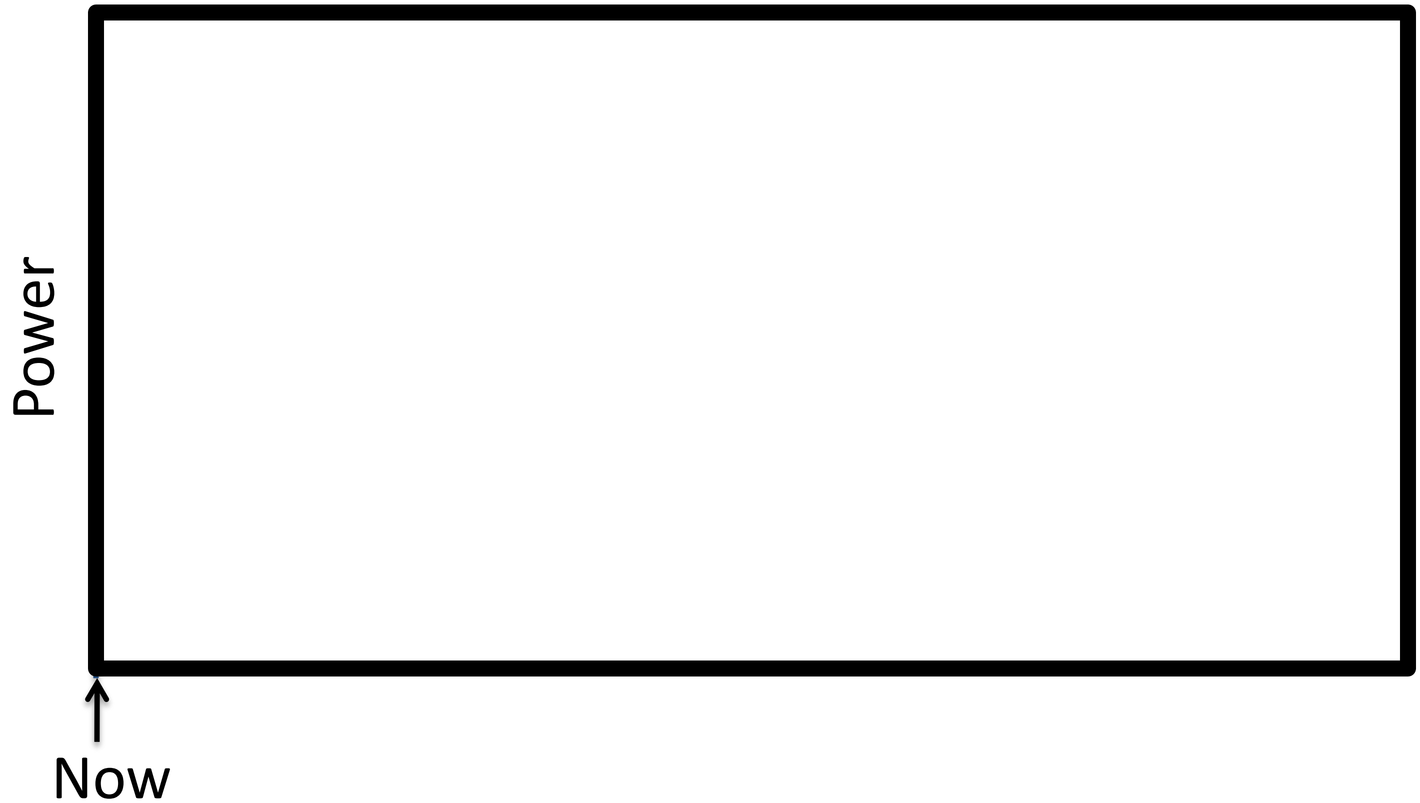

- Batteries and net metering are not ideal
- We need to match the energy demand to the supply

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Main challenge: Supply of power is variable!

- Many research questions:
  - What kinds of DC workloads are amenable?
  - What kinds of techniques can we apply?
  - Should we allow programmers to specify what can be done?
  - How well can we predict solar 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 with renewables
- Our target and research challenges
- Hardware and software for leveraging solar energy
- Current and future works
- Conclusions

Green DC software

- Follow the renewables [HotPower'09, SIGMETRICS'11]
- Duty cycle modulation with sleep states [ASPLOS'11]
- Quality degradation for interactive loads [UCB-TR'12]
- Adapt the amount of batch processing [HotPower'11]
- Delay jobs while respecting deadlines
  - GreenSlot [SC'11], GreenHadoop [Eurosys'12]

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Overall “delay-until-green” approach

- Predict green energy availability
  - Weather forecasts
- Schedule jobs
  - Maximize green energy use
  - If green not available, consume cheap brown electricity
- May delay jobs but must meet deadlines
- Send idle servers to sleep to save energy
- Manage data availability if necessary

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

GreenHadoop scheduling

Estimate the energy required by jobs (EWMA)

Job5

Job1

Job3

Job2

Job4

Job6

GreenHadoop scheduling

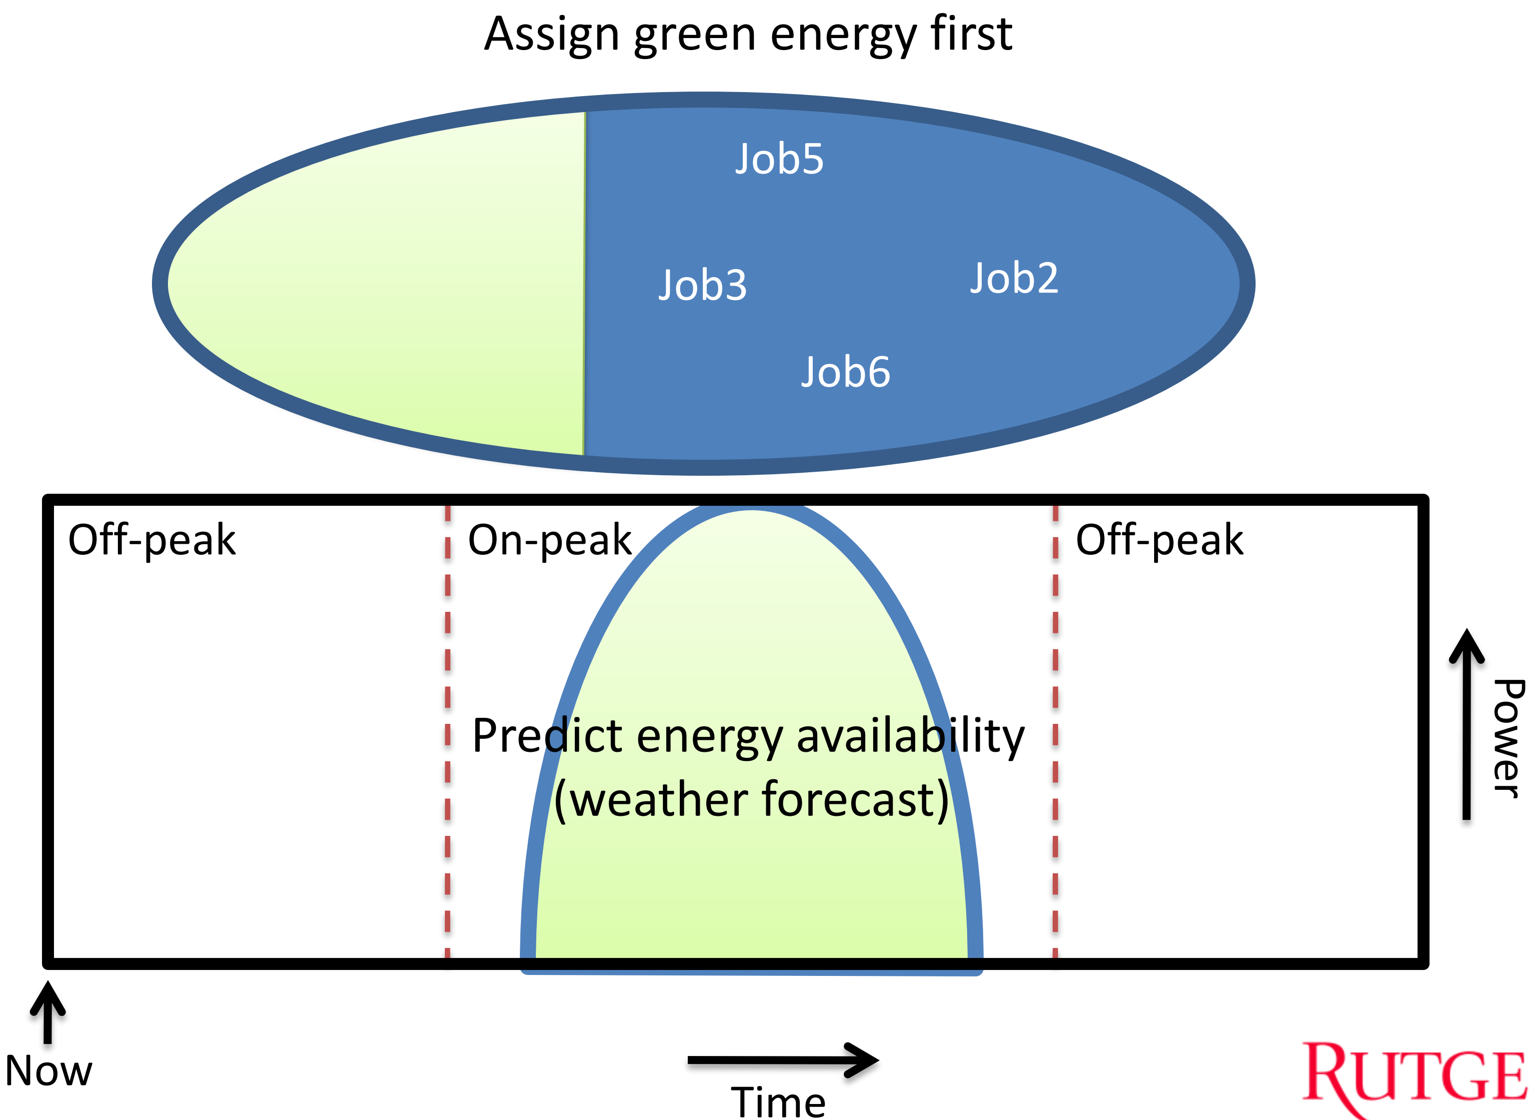

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

GreenHadoop scheduling

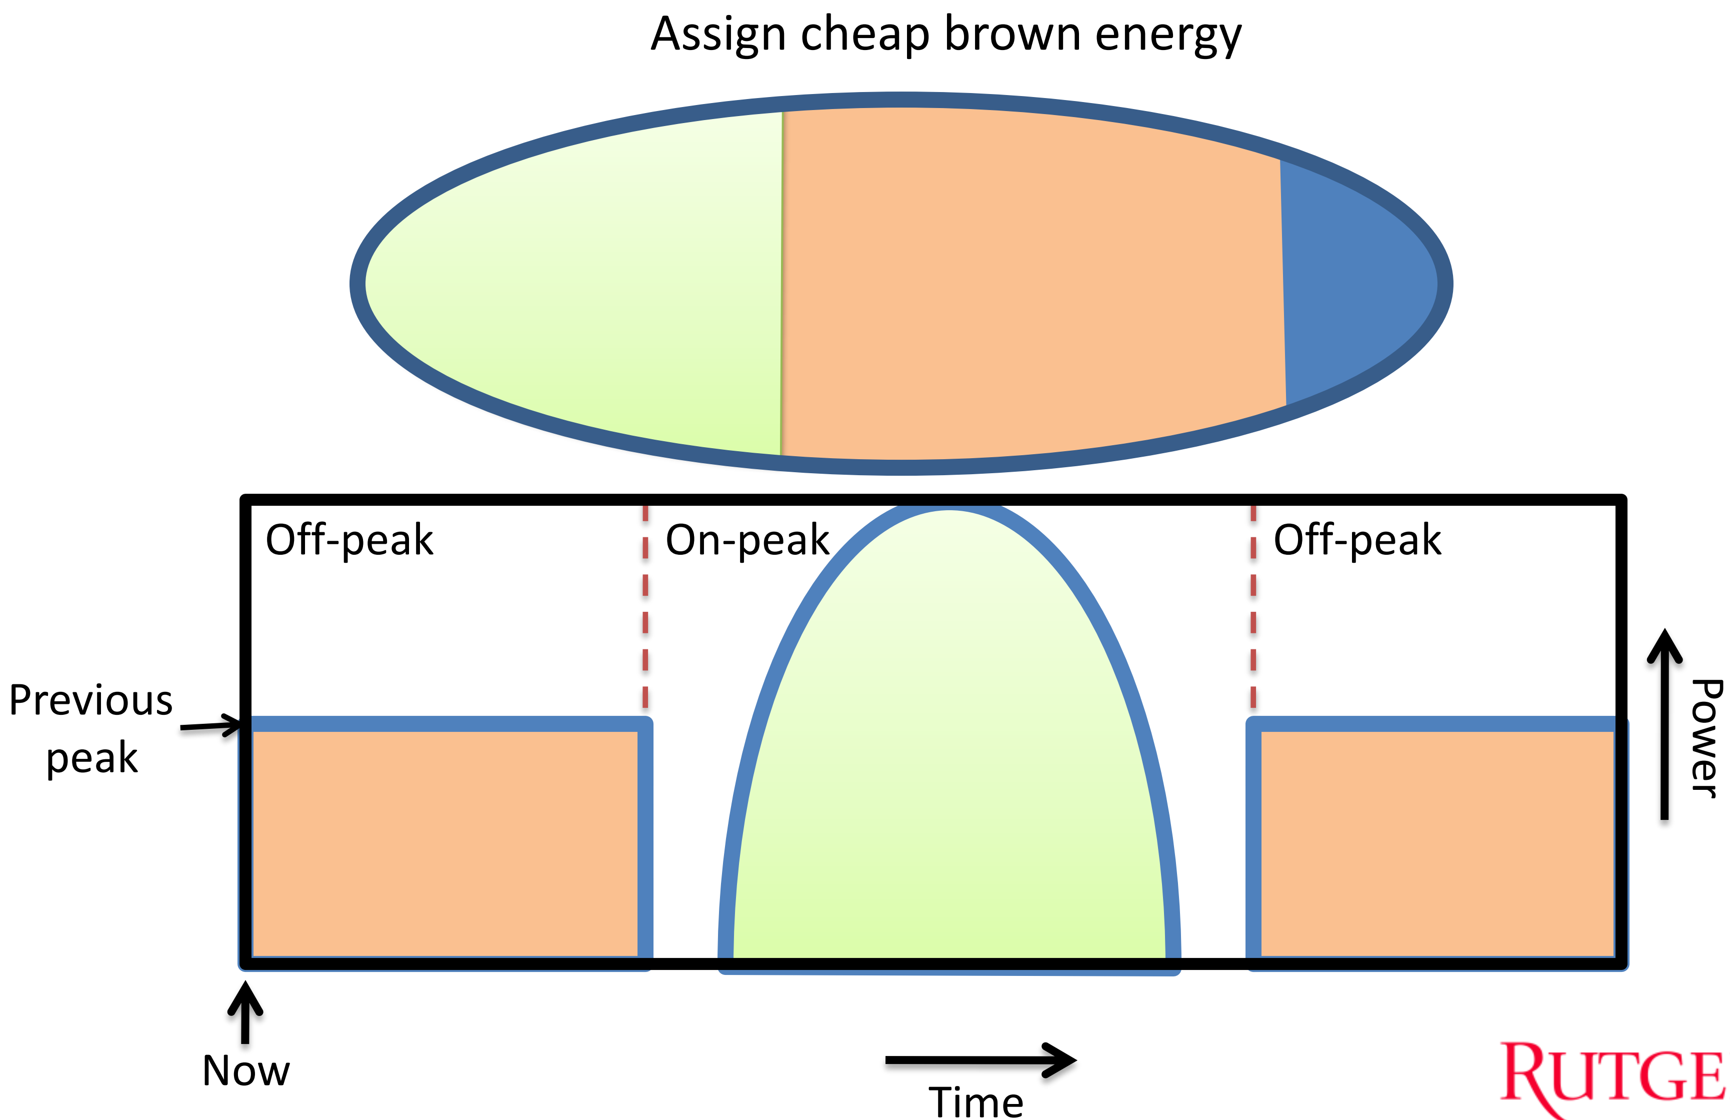

GreenHadoop scheduling

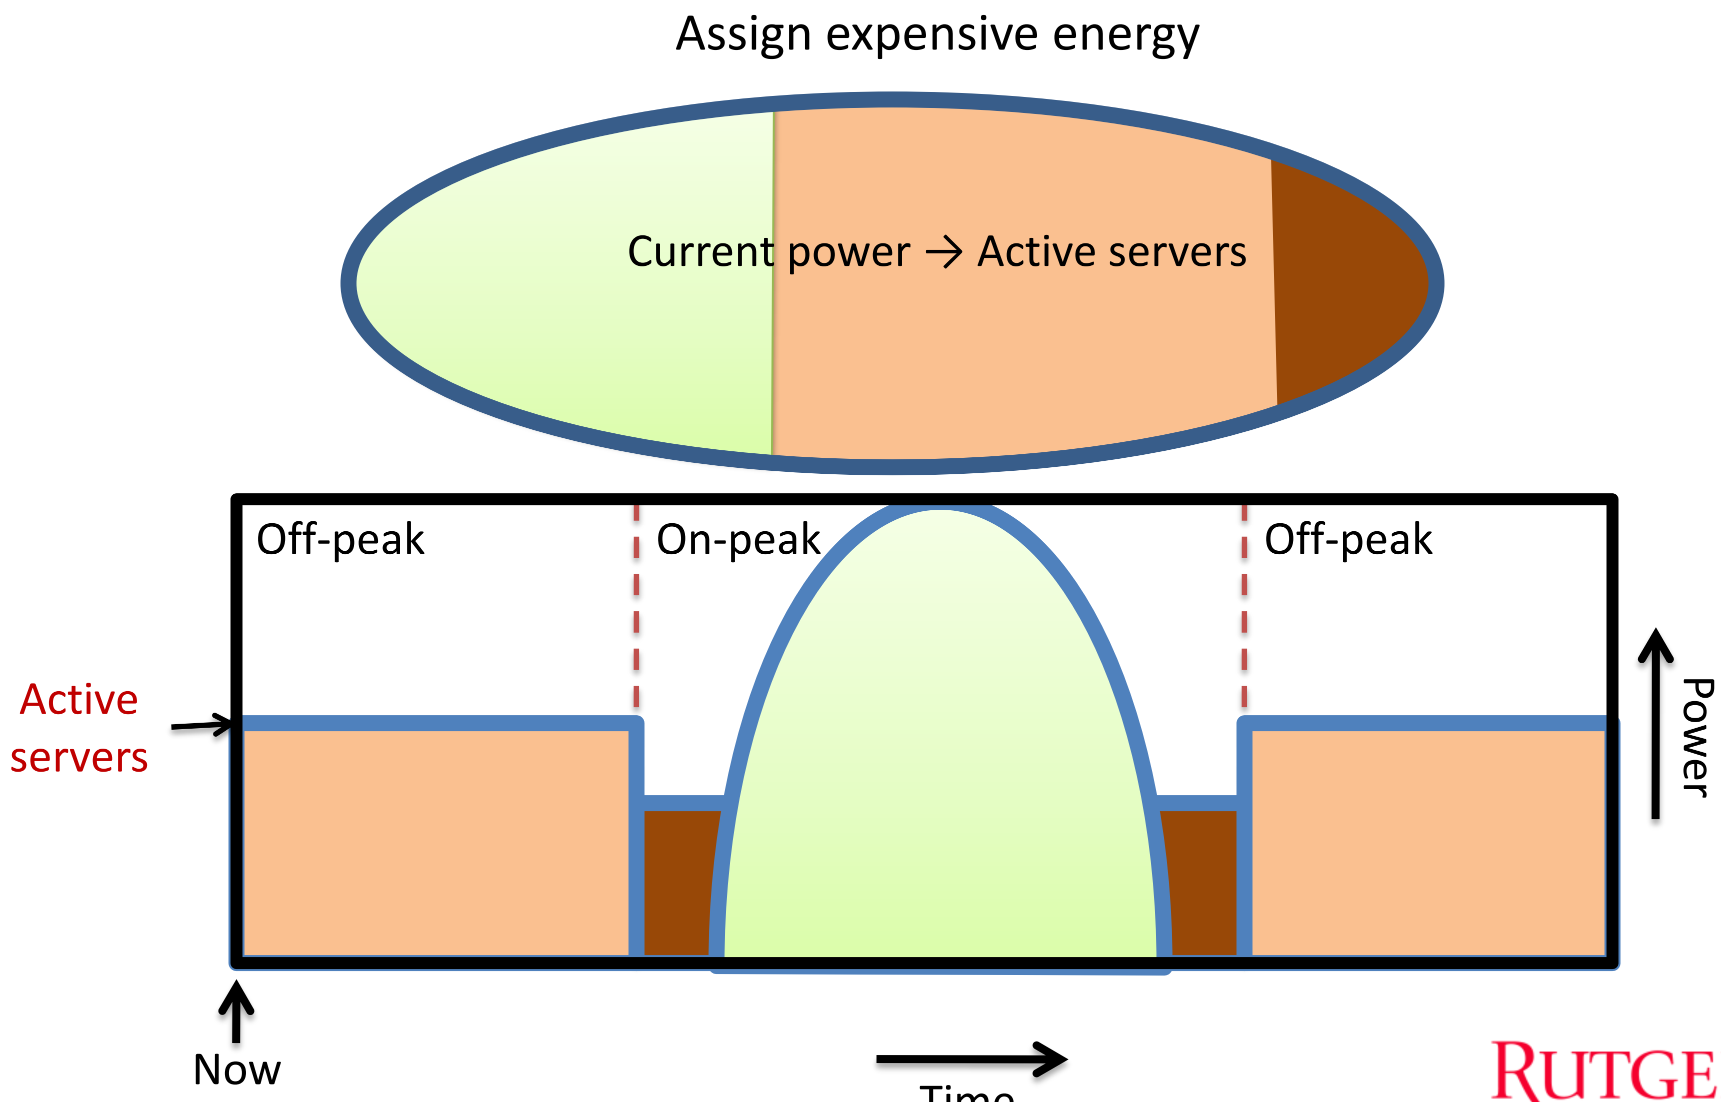

GreenHadoop scheduling

As time goes by…

the number of active servers changes

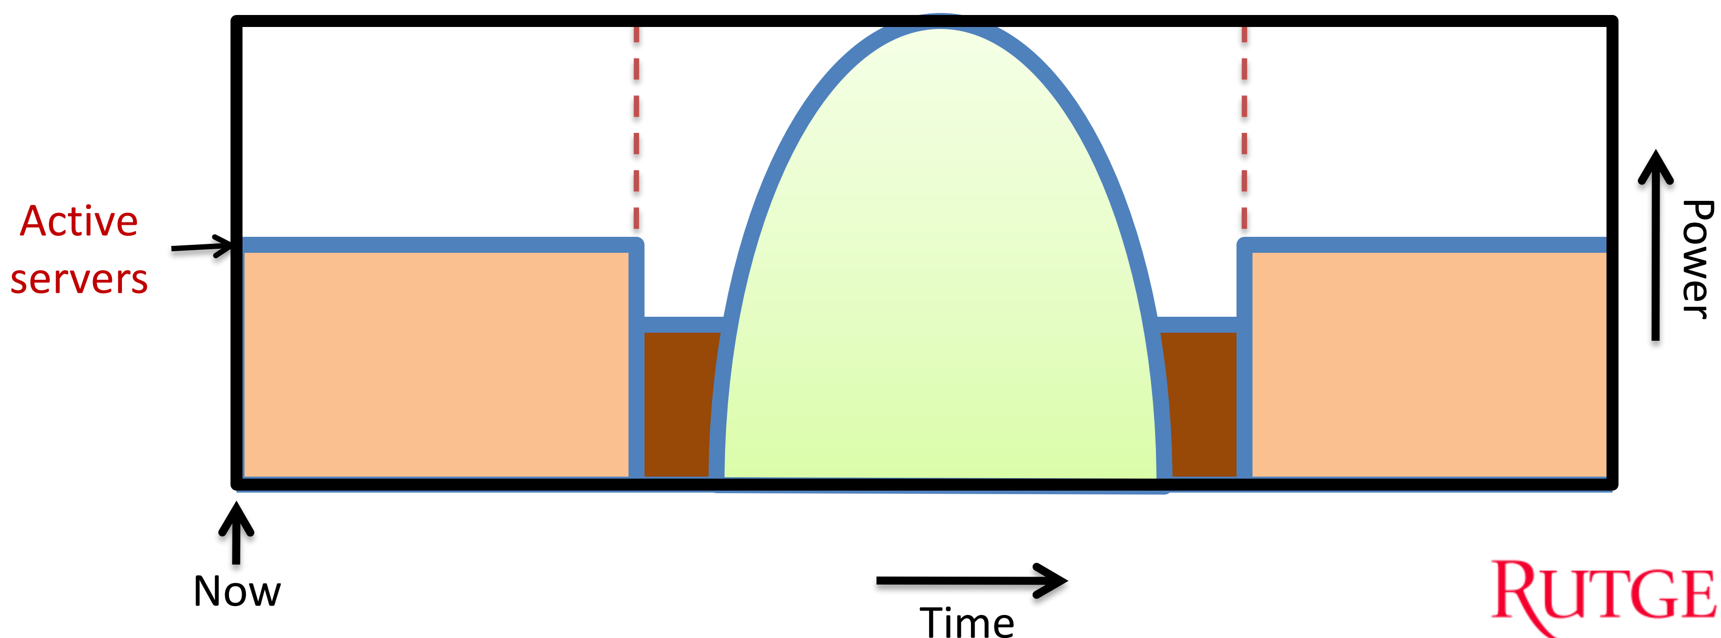

Energy prediction vs actual

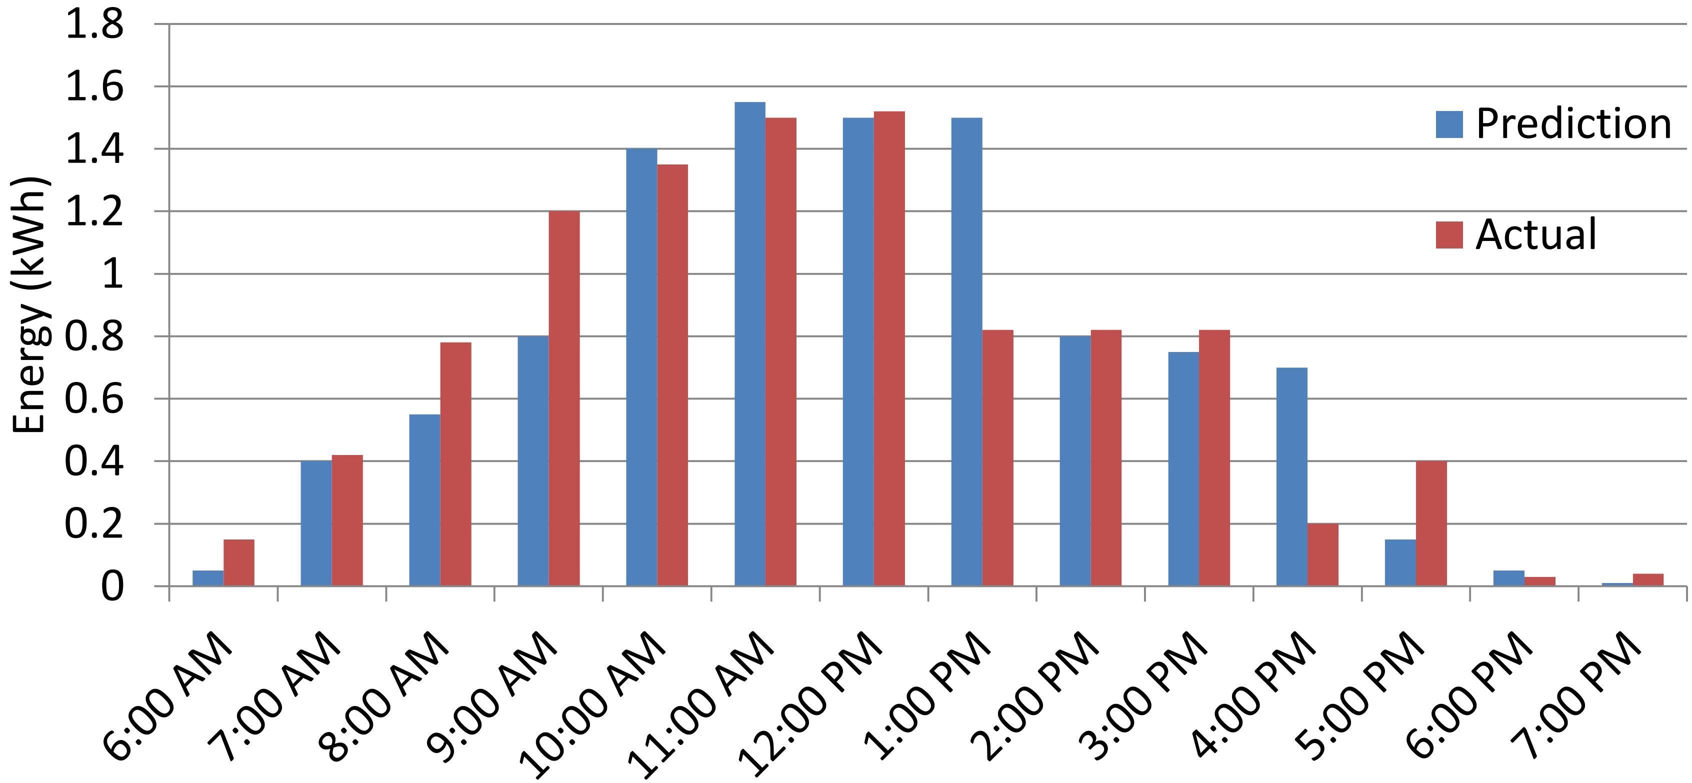

Actual data from the Rutgers solar farm (scaled down to our 16-node cluster)

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Energy prediction vs actual

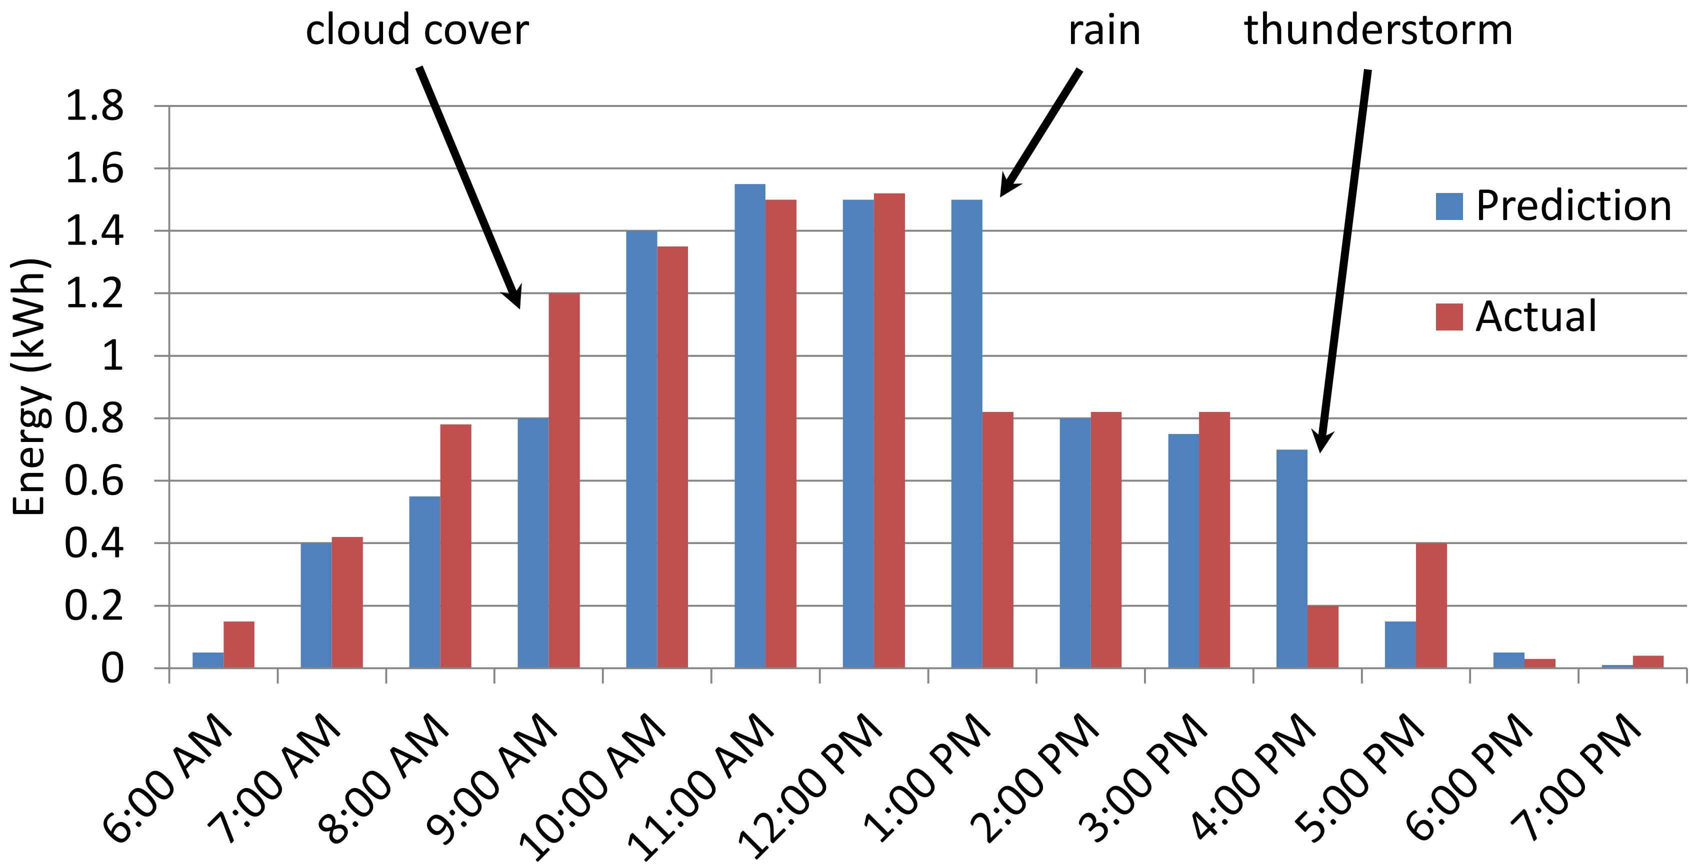

Actual data from the Rutgers solar farm (scaled down to our 16-node cluster)

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

GreenHadoop for Facebook workload

Hadoop

GHadoop

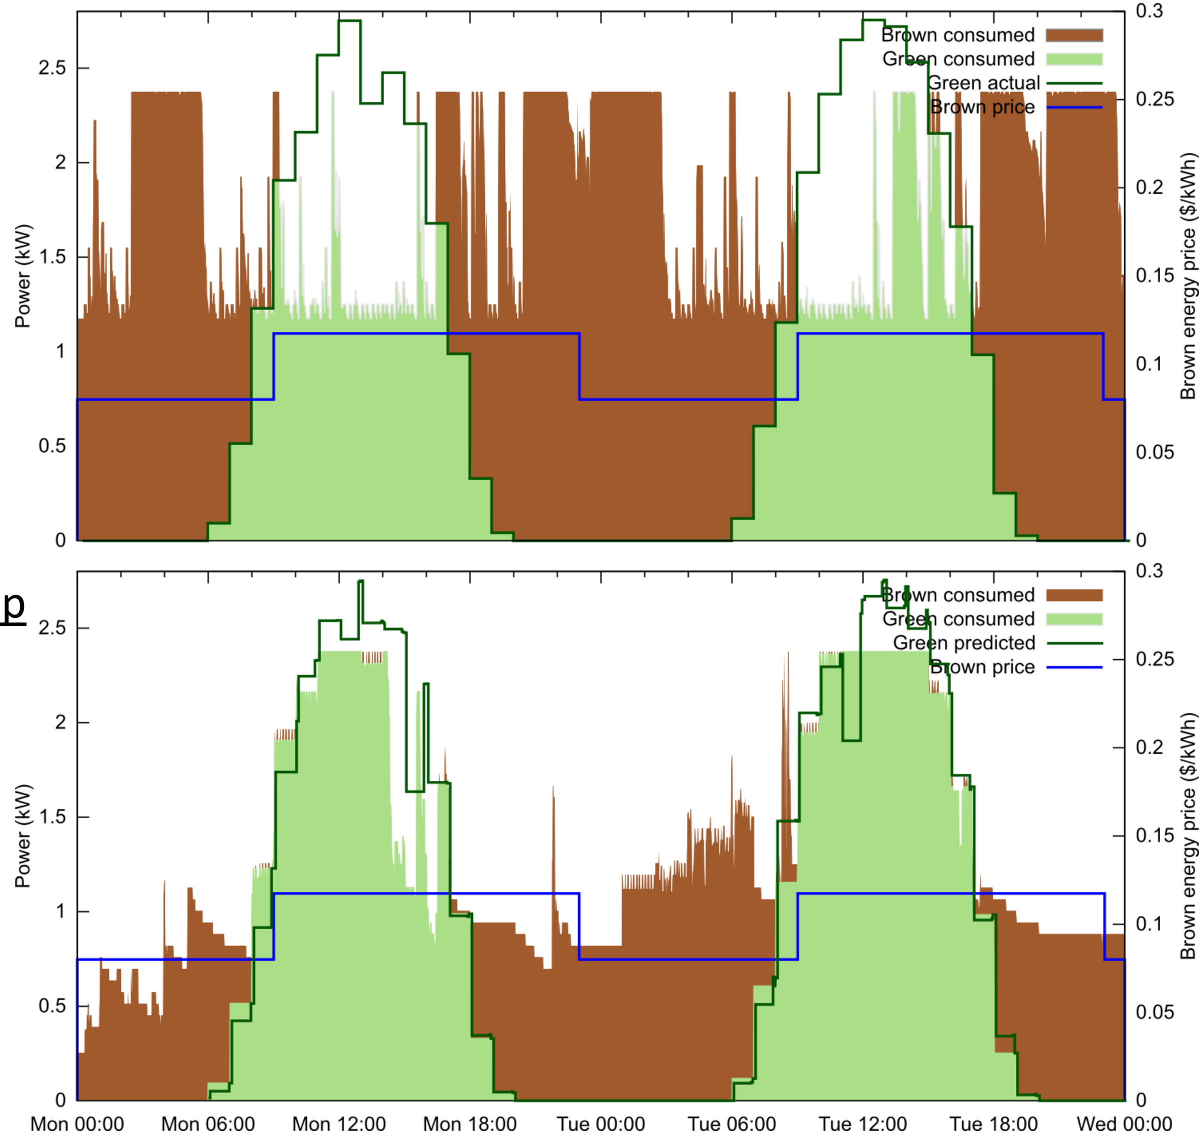

Brown consumed
Green consumed
Green produced
Brown price

31% more green
39% cost savings

Brown consumed
Green consumed
Green predicted
Brown price

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Outline

- DC energy usage and carbon footprint
- Reducing carbon with renewables
- Our target and research challenges
- Software and hardware for leveraging solar energy
- Current and future works
- Conclusions

The Rutgers Parasol Project

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

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
  - Full monitoring: resources, power, temperature, air

Parasol details

- Installed on the roof
- Steel structure
  - Container to host the IT
  - 16 solar panels: 3.2 kW peak
- Backup power
  - Batteries: 32 kWh
  - Power grid
- IT equipment
  - 2 racks
  - 64 Atom servers (so far): 1.7 kW
  - 2 switches and 3 PDUs
- Cooling
  - Free cooling: 110 W or 400 W
  - Air conditioning: 2 kW
  - Heating: 3 kW

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Outside and inside Parasol

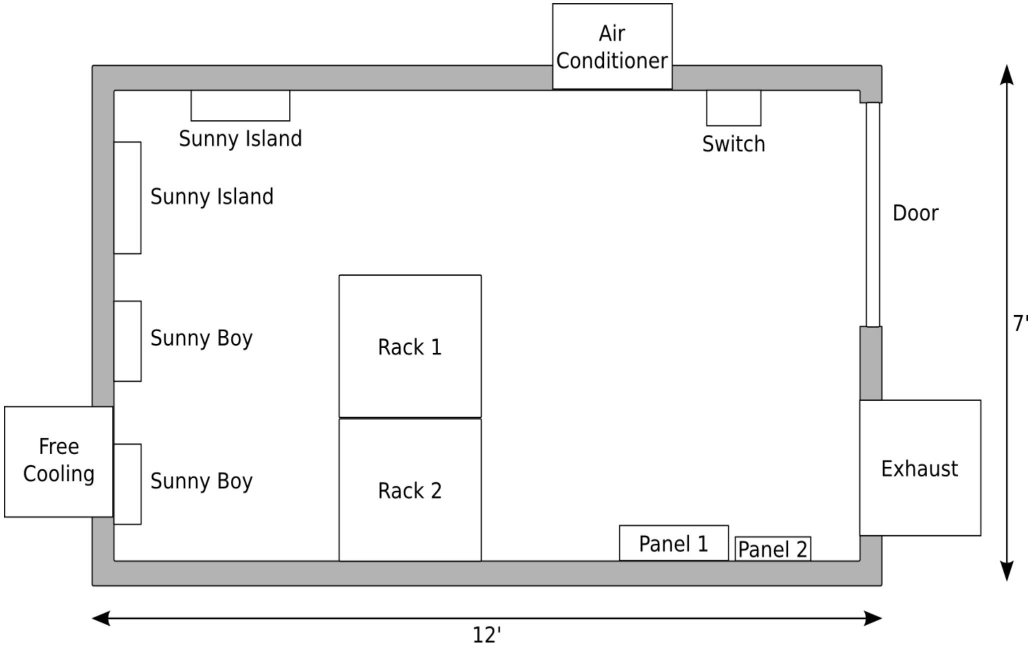

Outline

- DC energy usage and carbon footprint
- Reducing carbon with renewables
- Our target and research challenges
- Software and hardware for leveraging solar energy
- Current and future works
- Conclusions

Current and future works

- DC placement with probabilistic guarantees
- GreenNebula
- Smart management of energy sources
- Green SLAs
- Tradeoff between performance and green energy use
- Collect and make sense of the monitoring data

Conclusions

- Reduce the carbon footprint of ICT, data centers
- Topic is interesting and has societal impact
- Lots left to do…

More info -- http://parasol.cs.rutgers.edu