Exploiting Temporal Diversity of Water Efficiency to Make Data Center Less "Thirsty"

M. A. Islam, K. Ahmed, **Shaolei Ren**, and G. Quan

Florida International University

A massive data center

Google's data center in Mayes County, Oklahoma

Something we know

- Data centers are energy hogs
  - Combined energy usage of all data centers would rank 5th in the world if data centers were a country
  - Tens of millions of $$$ in annual operational costs
  - Responsible for 2% electricity usage and 1% greenhouse gas emissions
  - ……

Something we may **not** know

- Data centers are very **"thirsty"**
  - 2013 Uptime Institute survey shows that a majority of large data centers use cooling towers to remove heat from server rooms to the environment
  - AT&T's data center facilities consume **1 billion** gallons of water in 2012
    - Taking up ***30%*** of the entire company's water consumption

Google’s cooling towers

Cooling towers in Google’s Belgium data center

Google is “drinking” water

Water evaporation in Google’s Oregon data center

Something we may **not** know

- Data centers also consume water **"remotely"**
  - Electricity production consumes water
    - The **#1** water withdrawal among all sectors
    - U.S. national average **1.8 L/kWh**, excluding hydropower
  - Data centers account for approx. **1%** water withdrawal in the U.S.
    - Data centers consume 1.7—2.3% electricity
    - Electricity accounts for 40+% water withdrawal

Electricity production consumes water

http://www.globalenvision.org/

Why should we care about water?

- Extended droughts are becoming a norm

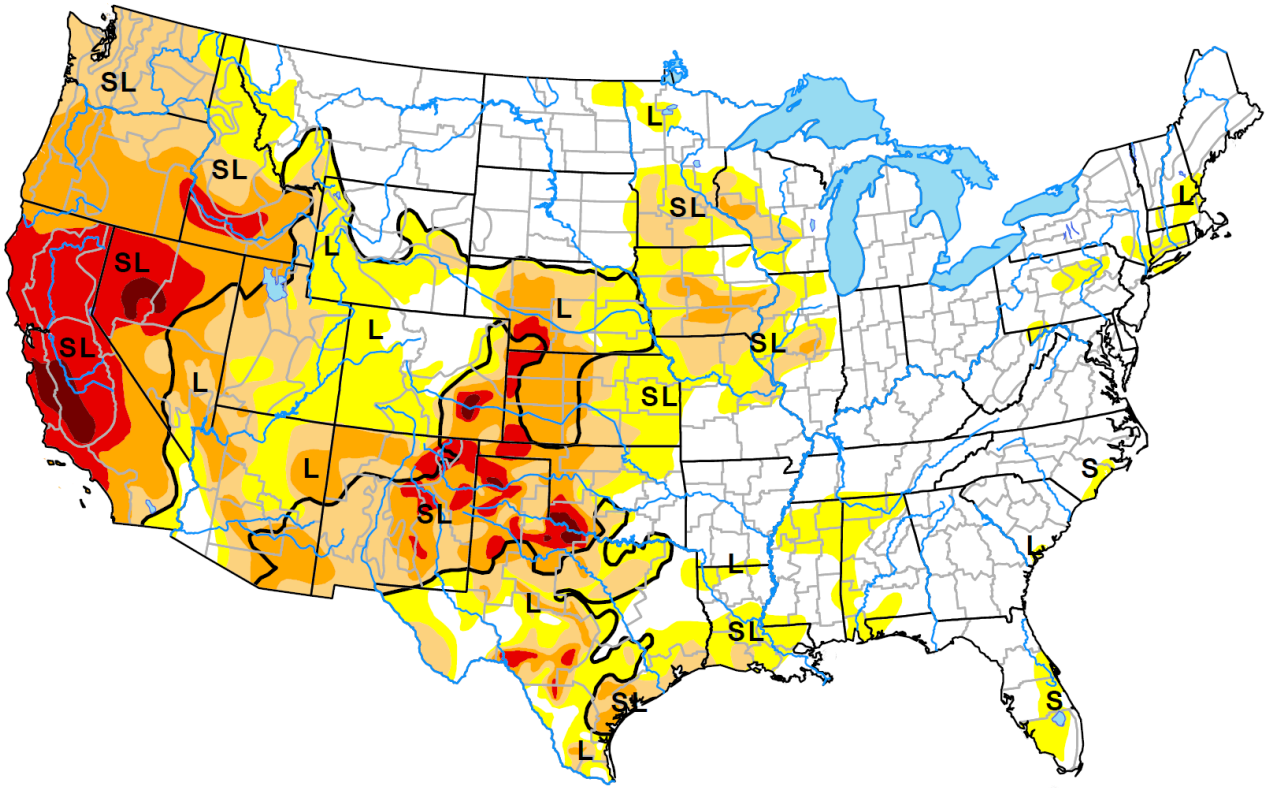

http://droughtmonitor.unl.edu/

Why should we care about water?

- Extended droughts are becoming a norm
  - California declared state-wide drought emergency on Jan. 17, 2014, urging all residents and businesses to cut back water usage by **20%**

What if drought is far away?

- Water conservation gets you…
  - Green certifications (e.g., LEED program)
    - Water reduction is a prerequisite
  - Tightening water compliance codes
    - U.S. government mandates 2% water efficiency improvement for federal facilities each year through 2020
  - Corporate social responsibility
    - AT&T, Google, Facebook, etc.

Industry has realized the water issue

- Google uses recycled water/seawater in some of its facilities
- Facebook uses air economizer (a.k.a. “free” cooling)
  - Recently released source code for monitoring water usage
- In 2013, LBNL recommends **water** and **energy** as two key considerations for U.S. federal data center consolidation

Industry has realized the water issue

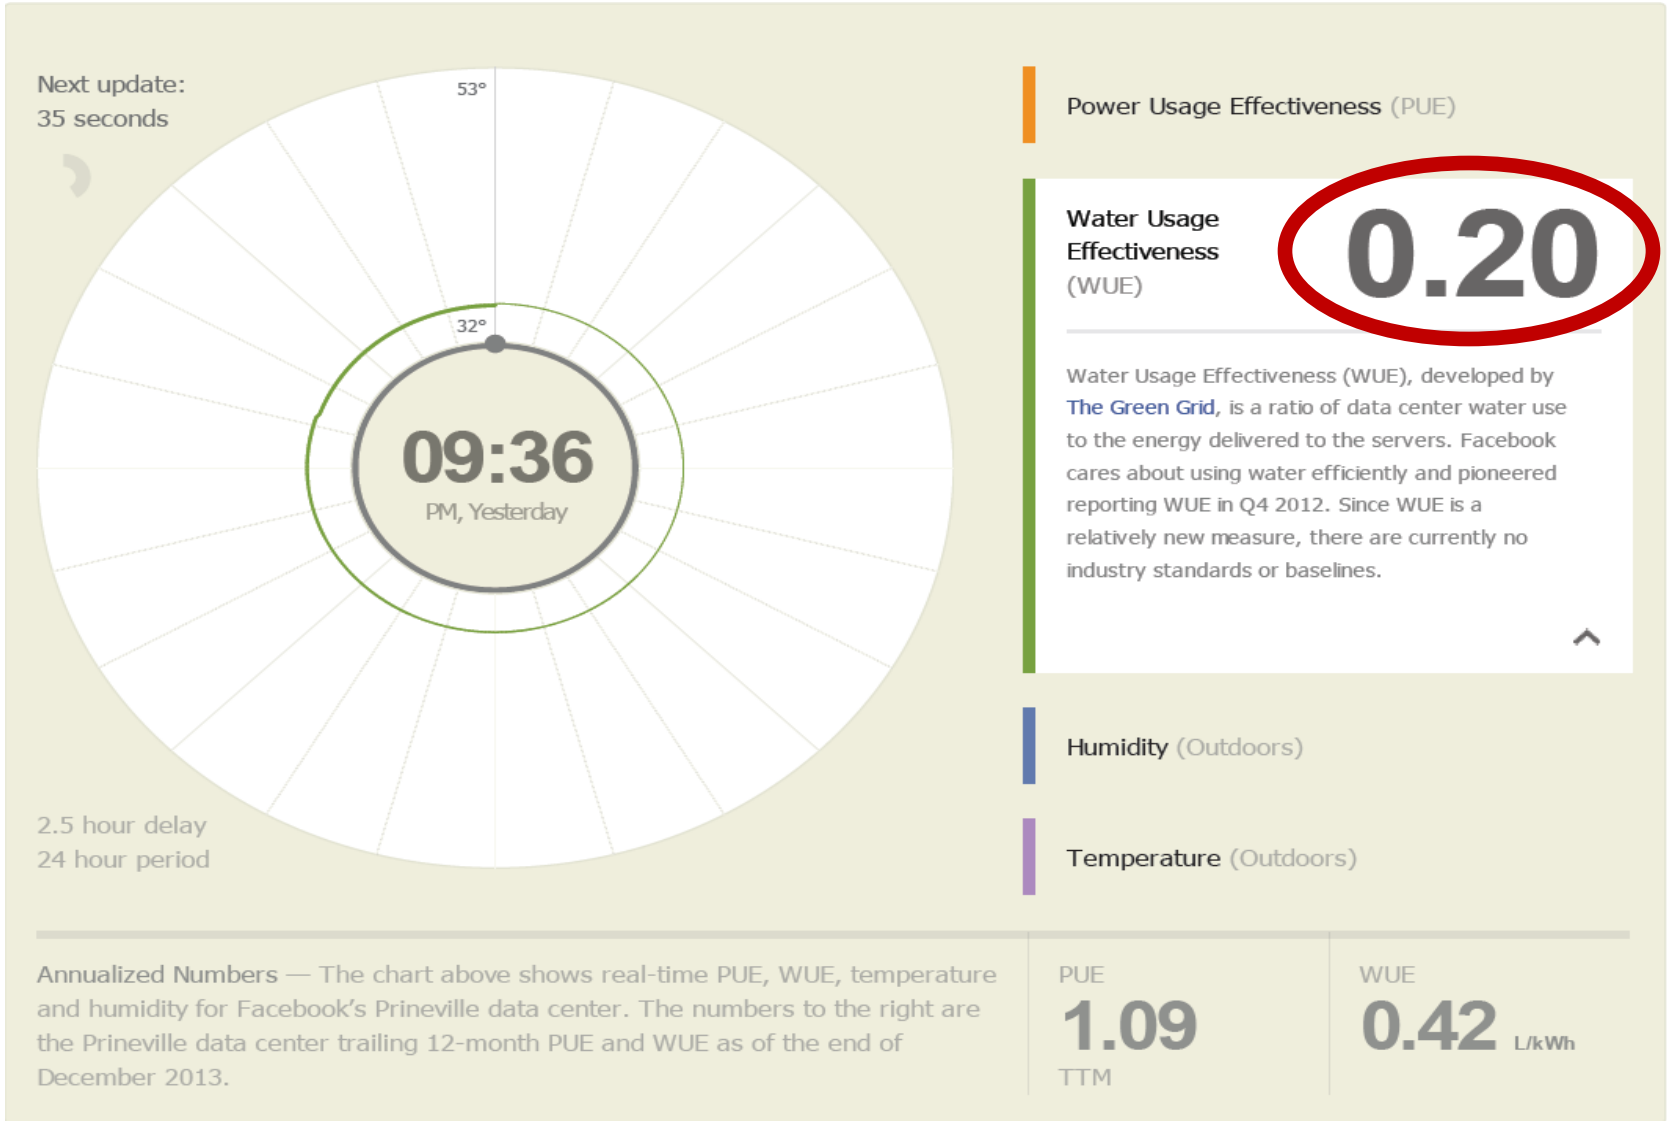

https://fbpuewue.com/prineville

Limitations of existing approaches

- The existing approaches primarily focus on improved "**engineering**"
  - High upfront capital costs
  - Suitable climate/locations
  - Offsite water is not explicitly addressed

Water is equal to energy?

- Not really…
  - Water efficiency changes over time
  - Just like minimizing electricity cost with time-varying electricity prices, we need to consider **"when"** to consume energy

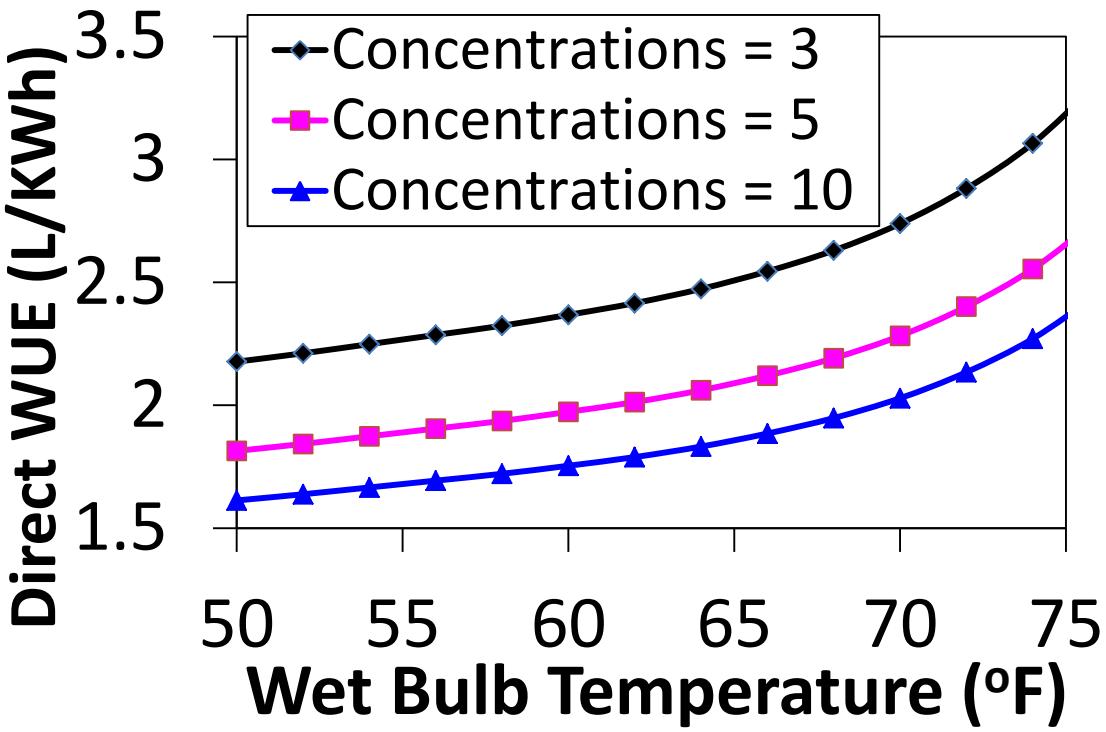

Time-varying (offsite) water efficiency

- Energy fuel mix changes over time for electricity
  - Different energy fuel types have different water efficiencies
  - In general, Hydro > Nuclear > Thermal > Renewables

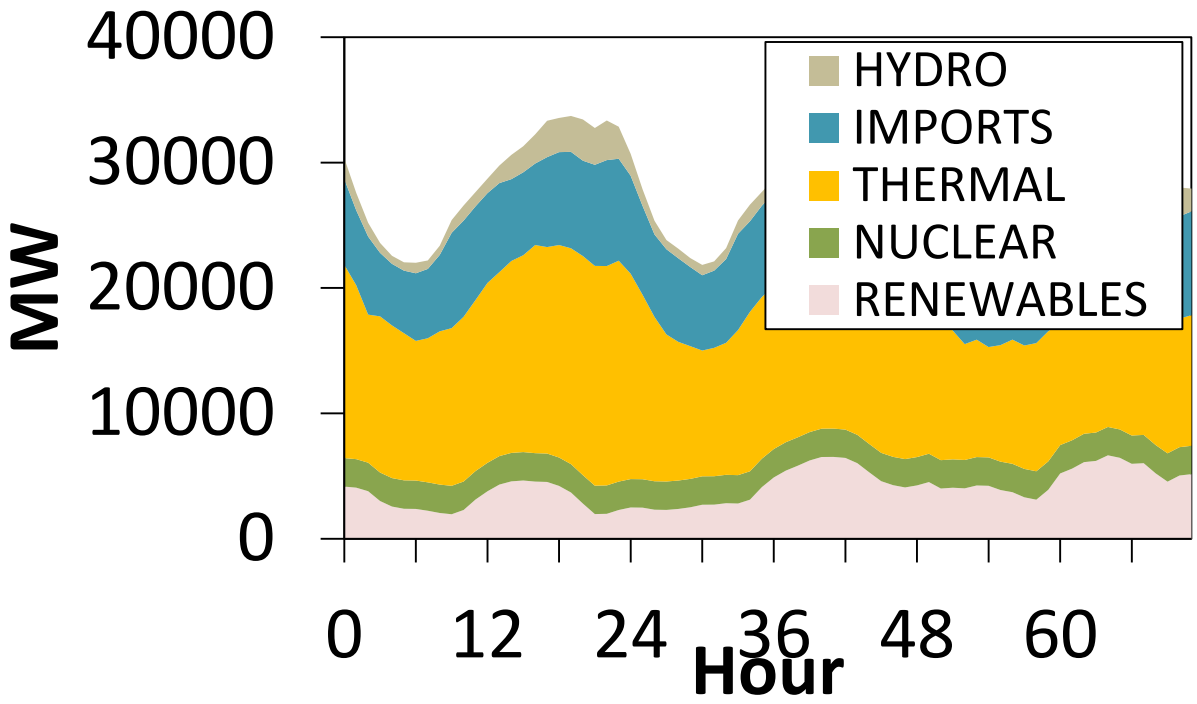

How can we exploit time diversity of water efficiency?

We propose…

- We can turn off some servers and defer some delay-tolerant workloads to time periods that are more water-efficient
  - But how to determine if a time period is **efficient**?
  - We also need to consider electricity cost (and carbon footprint if applicable)
- WACE: opportunistically schedule batch jobs
  - First, queue up the job arrivals
  - If the queue length is too large and/or water efficiency is already good, then process waiting jobs

How does WACE work?

- **Input:**
  - Online information: WUE, electricity price, etc.
- **Solve online optimization:**
  - Minimize $\mathrm{V} \cdot g(t) - q(t) \cdot b(t)$
    - $g(t)$: weighted sum of water footprint, electricity cost and carbon footprint
    - $q(t)$: batch job queue length
    - $b(t)$: processed batch jobs
- **Update batch job queue length**

Provably-efficient while bounding queue length

Case study

How to evaluate WACE?

- Let's consider a real-life trace simulation
  - MSR workload; data center in Mountain View, California
- Benchmarks
  - SAVING: Optimize for electricity cost
  - CARBON: Optimize for carbon footprint
  - ALWAYS: Optimize for delay performance

WACE reduces water footprint

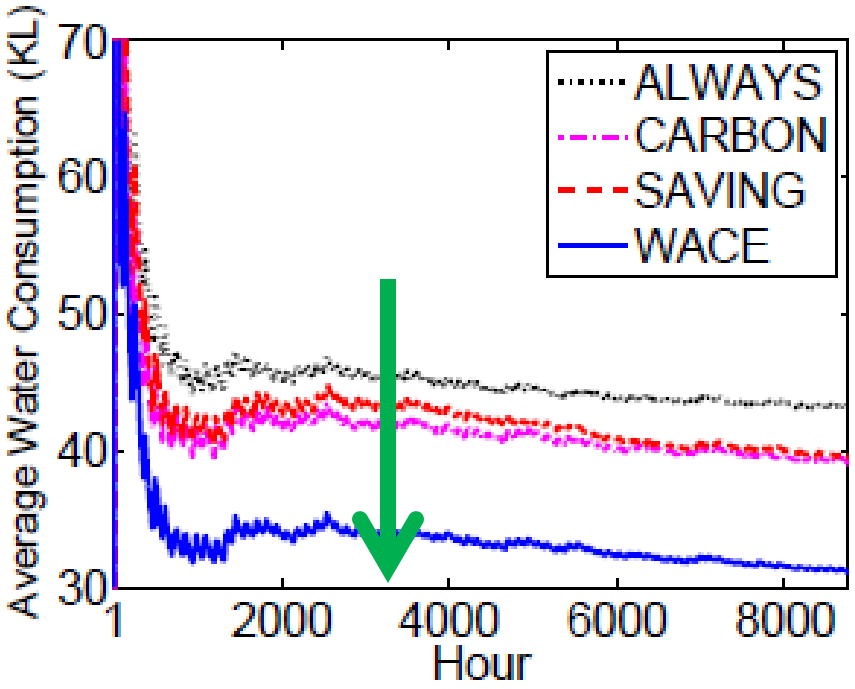

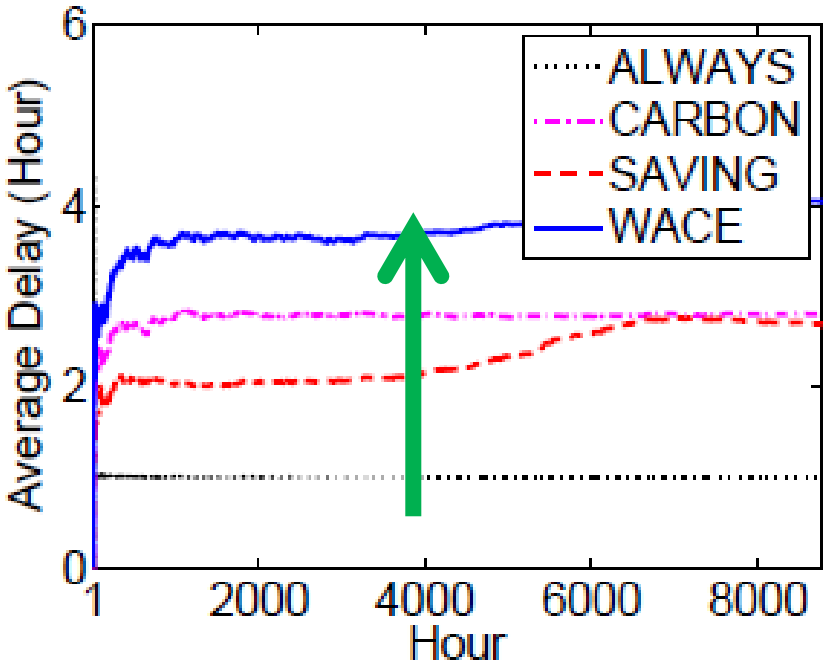

- Big water saving at (a small) delay increase

WACE reduces total cost

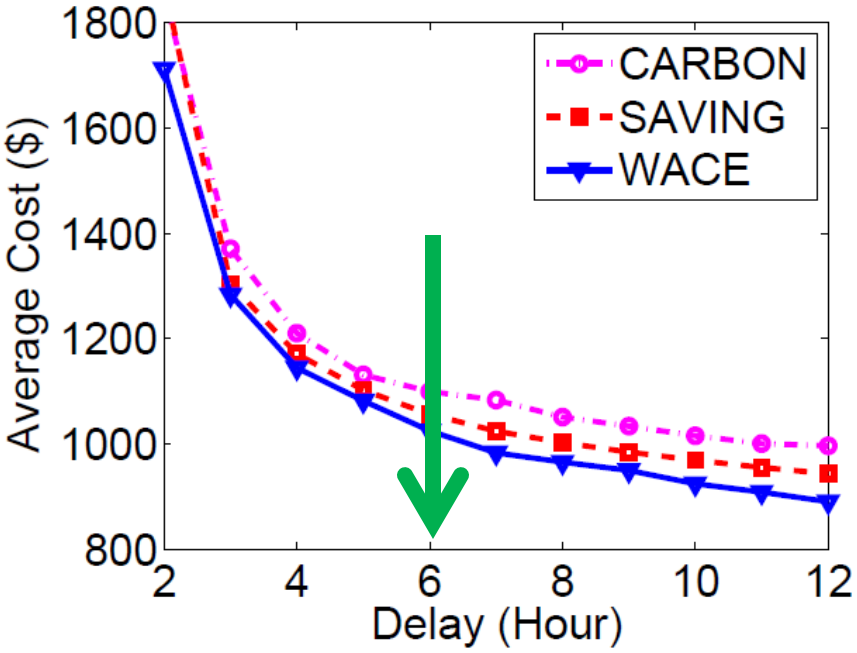

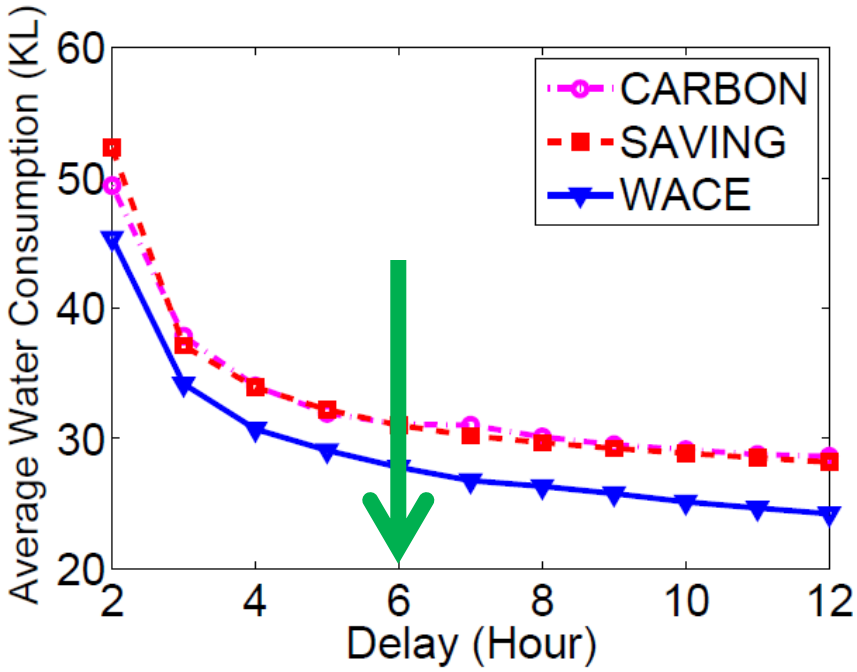

- WACE reduces the total cost without increasing delay
- Minimizing electricity cost or carbon emission is **not** sufficient for minimizing water footprint!

Three messages

- Data centers are “thirsty”
- Water footprint **!=** energy or electricity cost or carbon emission
- Time diversity of water efficiency can be exploited for reducing water footprint


Thanks!