Managing Carbon Footprints and Electricity Costs in Data Centers

Ricardo Bianchini

Department of Computer Science
Rutgers University

Collaborators: Ozlem Bilgir, Kien Le, Yogesh Jaluria, Margaret Martonosi, and Thu Nguyen

Motivation

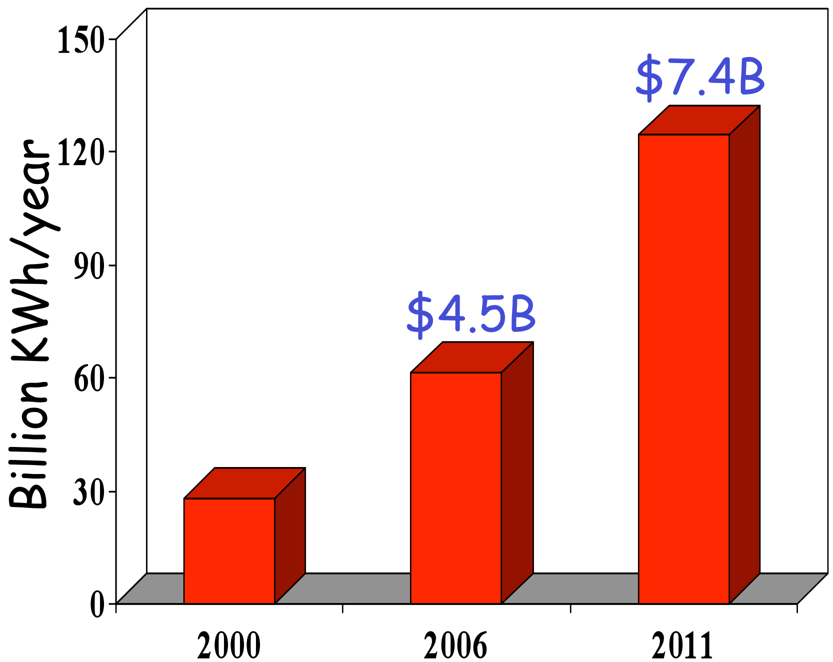

Electricity consumption of US data centers [EPA'07]

Electricity produced by carbon-intensive means

Carbon emissions of world-wide data centers [Mankoff'08]

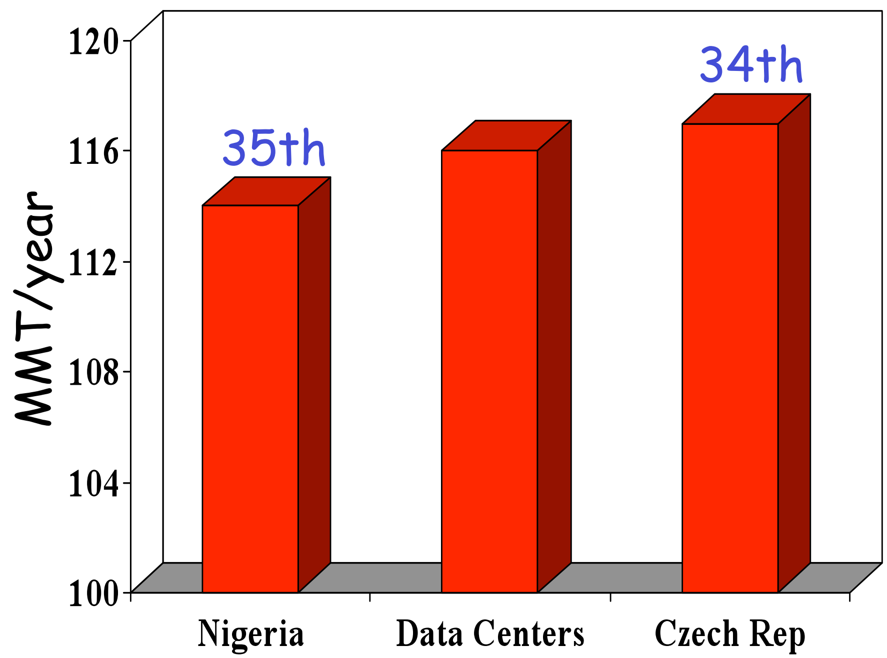

What To Do?

- Improve efficiency: HW, cooling, supply of power
- Limit carbon emissions: green energy, carbon offsets
- Educate people about the importance of greening IT
- Capping brown energy can help achieve these goals
  - Cap-and-trade: if cap exceeded, purchase carbon offsets
  - Cap-and-pay: if cap exceeded, pay higher brown energy price
  - Cap-as-target: if cap exceeded, must compensate somehow

Research Questions

- Can we cap brown energy without degrading performance or excessively increasing costs and overheads?
- How can we take advantage of distributed data centers, green energy, and variable electricity prices?
- As a first environment to study: Internet services
  - Examples: Google, iTunes, Yahoo!
  - Up to millions of servers and multiple data centers
  - Challenging in many ways, e.g. Service-Level Agreements (SLAs)

Multi-DC Internet Services

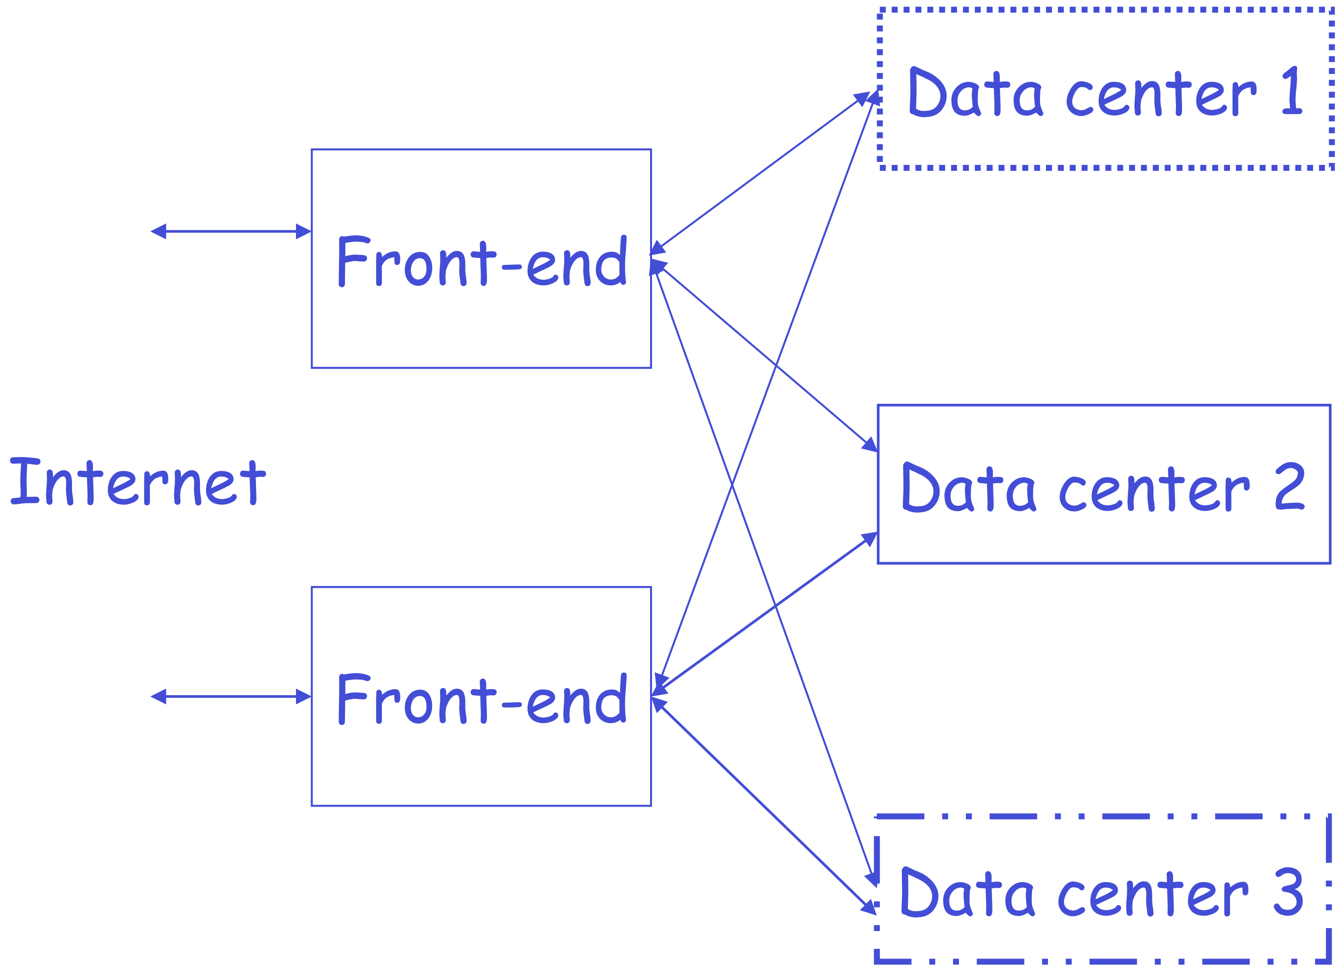

Multi-DC Internet Services

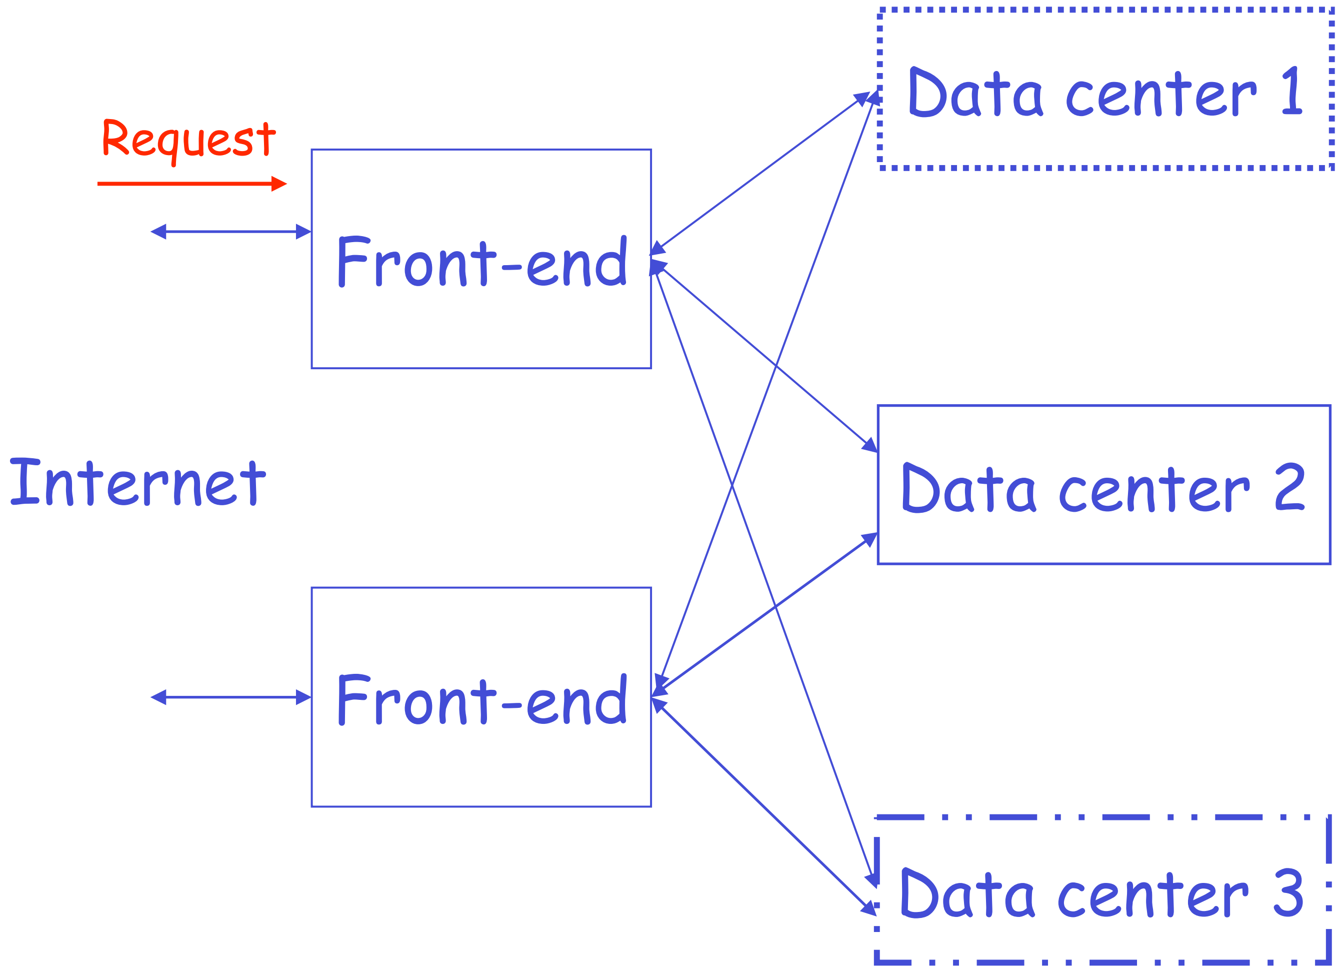

Multi-DC Internet Services

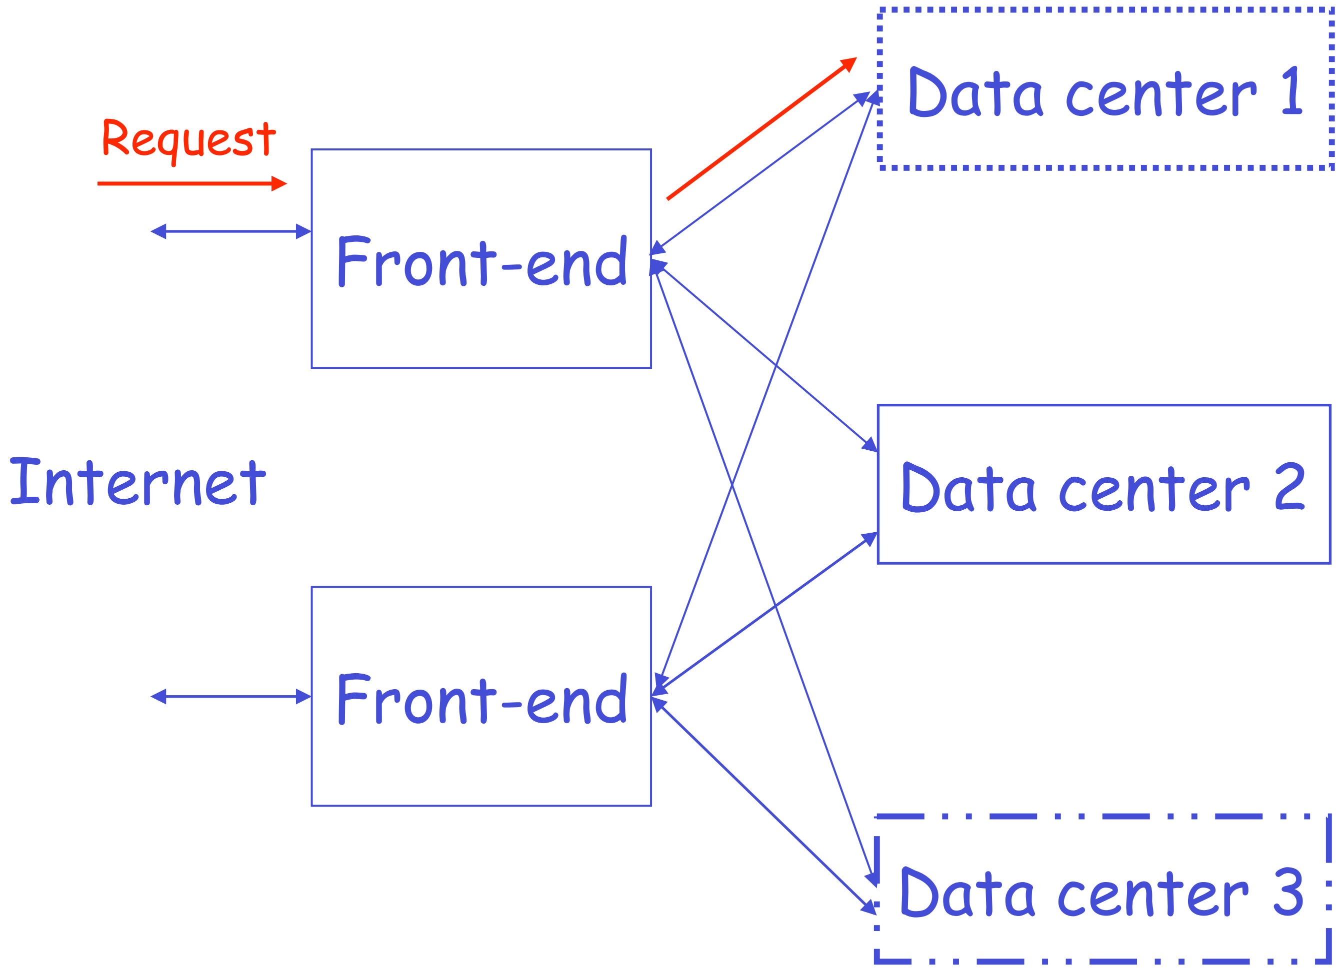

Multi-DC Internet Services

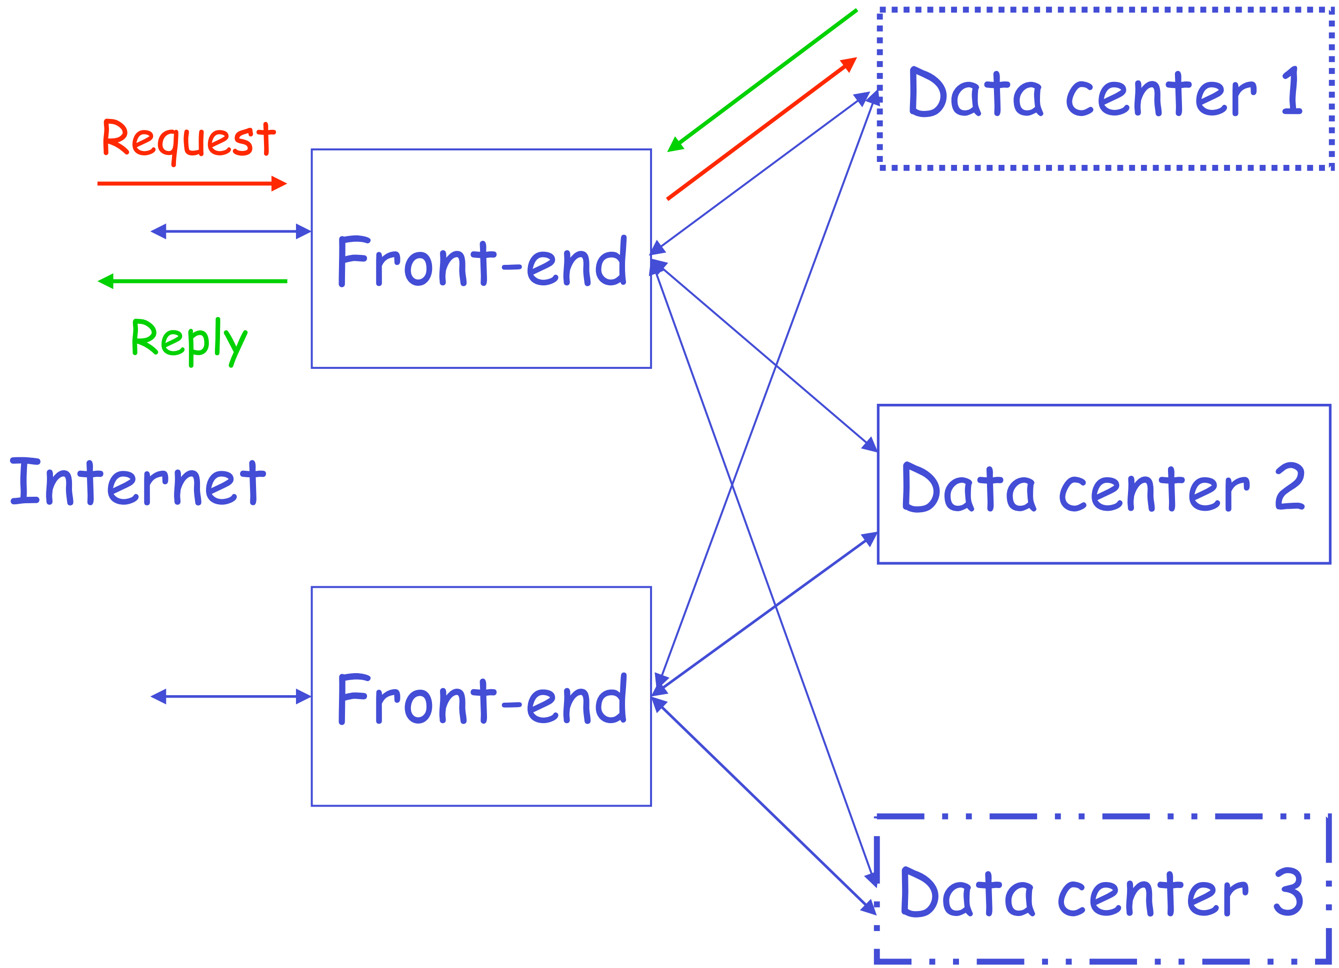

Multi-DC Internet Services

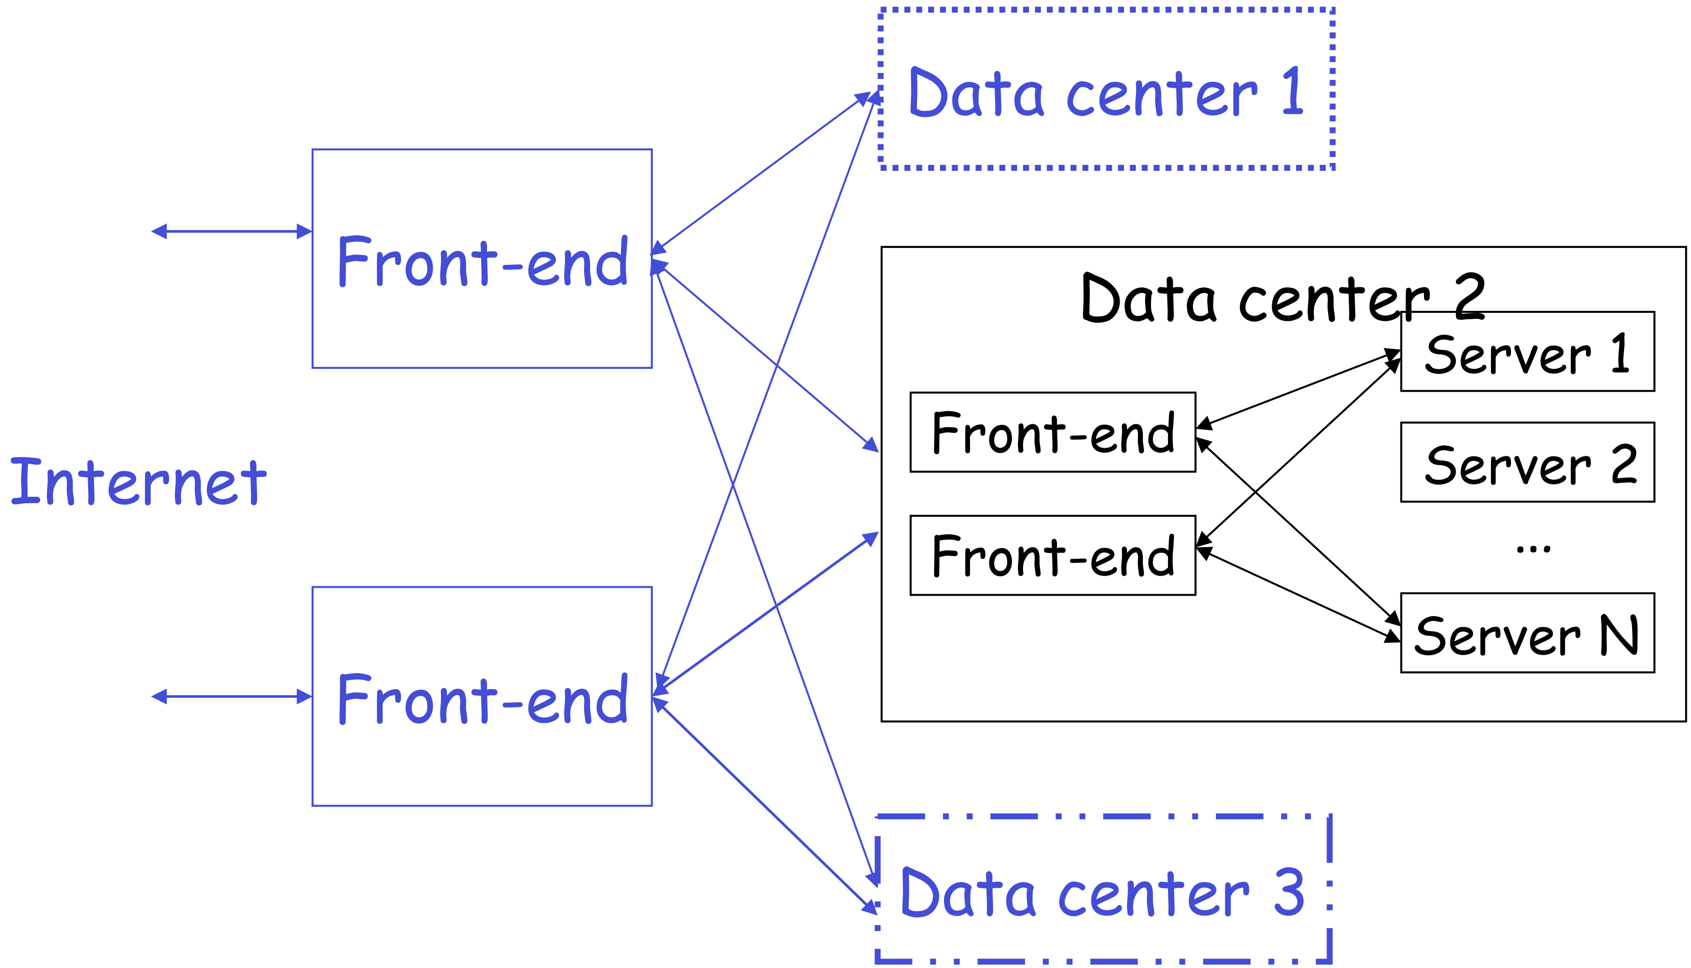

Research Activities and Approach

- Goal: Build SW to manage energy based on modeling & optimization
- Across-DC request distribution
  - Request distribution policies: time series analysis + mathematical optimization + statistical performance data
  - Time zones, electricity prices, cooling types, green energy availability
- Within-DC distribution, energy management, and enforcement
  - Hierarchical policies for cap and SLA enforcement: time series analysis + queuing theory + power modeling
  - Electricity price, cooling costs, effect on cooling behavior
- Putting it all together: integrated and coordinated decisions

First Study: Across-DC Request Distribution

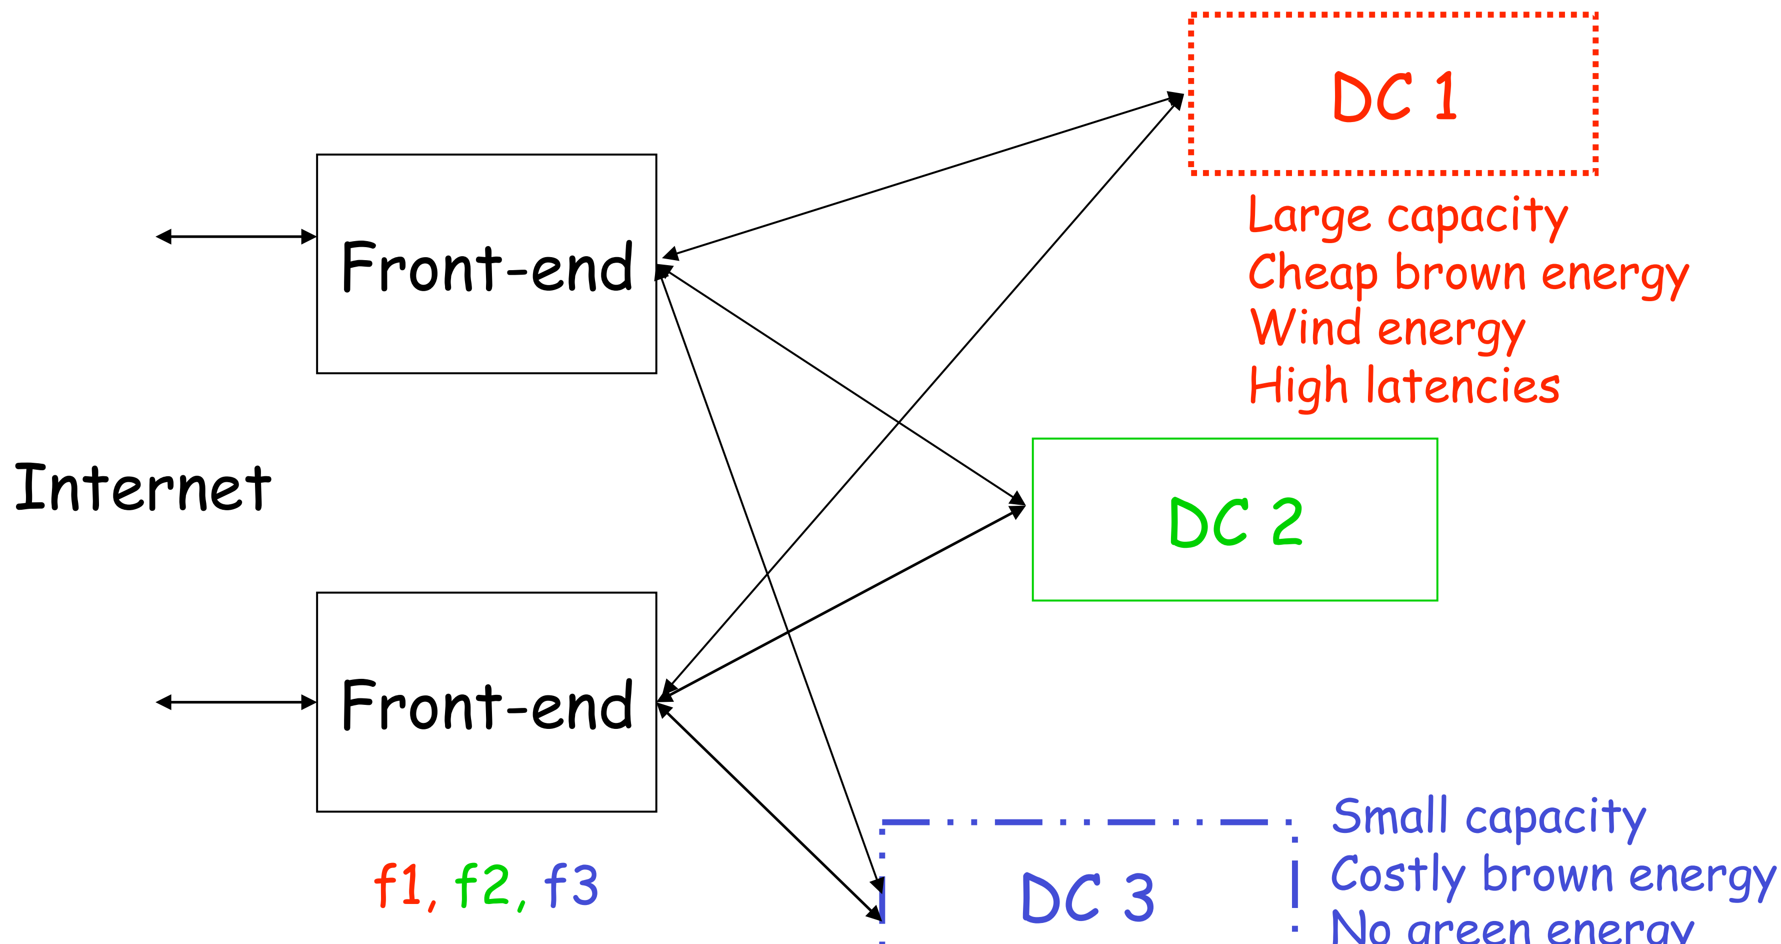

FEs compute fi's periodically

Example Distribution Policy for Cap-and-Pay

$$\text{Overall Cost} = \left(\sum_t \sum_i \sum_y f_i^y(t) \times p^y(t) \times LT(t) \times Cost_i^y(t)\right) + \left(\sum_t \sum_i Base_i(offered_i)\right)$$

$$Cost_i^y(t) = \begin{cases} m_i^{brown} \times b_i^y(t) + m_i^{green} \times g_i^y(t), \text{ if brown energy consumed so far } \leq BEC \\ fee_i^y(t) + m_i^{brown} \times b_i^y(t) + m_i^{green} \times g_i^y(t), \text{ otherwise} \end{cases}$$

1. $\forall i \; m_i^{brown} \text{ and } m_i^{green} \geq 0$
2. $\forall i \; m_i^{brown} + m_i^{green} = 1$
3. $\forall t \forall i \forall y \; f_i^y(t) \geq 0$
4. $\forall t \forall y \sum_i f_i^y(t) = 1$
5. $\forall t \forall i \sum_y (f_i^y(t) \times p^y(t) \times LR(t)) \leq LC_i$
6. $\sum_t \sum_i \sum_y (f_i^y(t) \times p^y(t) \times LT(t) \times CDF_i(L, offered_i)) / \sum_t LT(t) \geq P$

Statistical response time prediction

Load intensity prediction

Example Distribution Policy for Cap-and-Pay

Goal: Compute fi's that minimize the overall energy cost

$$Cost_i^y(t) = \begin{cases} m_i^{brown} \times b_i^y(t) + m_i^{green} \times g_i^y(t), \textit{ if brown energy consumed so far } \leq BEC \\ fee_i^y(t) + m_i^{brown} \times b_i^y(t) + m_i^{green} \times g_i^y(t), \textit{ otherwise} \end{cases}$$

1. $\forall i \; m_i^{brown} \textit{ and } m_i^{green} \geq 0$
2. $\forall i \; m_i^{brown} + m_i^{green} = 1$
3. $\forall t \forall i \forall y \; f_i^y(t) \geq 0$
4. $\forall t \forall y \sum_i \; f_i^y(t) = 1$
5. $\forall t \forall i \sum_y \; (f_i^y(t) \times p^y(t) \times LR(t)) \leq LC_i$
6. $\sum_t \sum_i \sum_y (f_i^y(t) \times p^y(t) \times LT(t) \times CDF_i(L, offered_i)) / \sum_t LT(t) \geq P$

Statistical response time prediction

Load intensity prediction

Example Distribution Policy for Cap-and-Pay

Goal: Compute fi's that minimize the overall energy cost

Cost/request: Pay fee if brown energy cap is exhausted

1. $\forall i \; m_i^{brown}$ and $m_i^{green} \geq 0$
2. $\forall i \; m_i^{brown} + m_i^{green} = 1$
3. $\forall t \forall i \forall y \; f_i^y(t) \geq 0$
4. $\forall t \forall y \sum_i \; f_i^y(t) = 1$
5. $\forall t \forall i \sum_y \; (f_i^y(t) \times p^y(t) \times LR(t)) \leq LC_i$
6. $\sum_t \sum_i \sum_y (f_i^y(t) \times p^y(t) \times LT(t) \times CDF_i(L, offered_i)) / \sum_t LT(t) \geq P$

Statistical response time prediction

Load intensity prediction

Example Distribution Policy for Cap-and-Pay

Goal: Compute fi's that minimize the overall energy cost

Cost/request: Pay fee if brown energy cap is exhausted

Constraints:
- Make sure that no DC is overloaded
- Make sure that SLA is not violated

Generality: same formulation applies even when there are no caps, no green energy, and no price variation

Complexity: formulation is non-linear and recursive

Using Caps to Encourage Conservation

3 DCs
Power mix
On/off
Ask.com

Simulated
Annealing

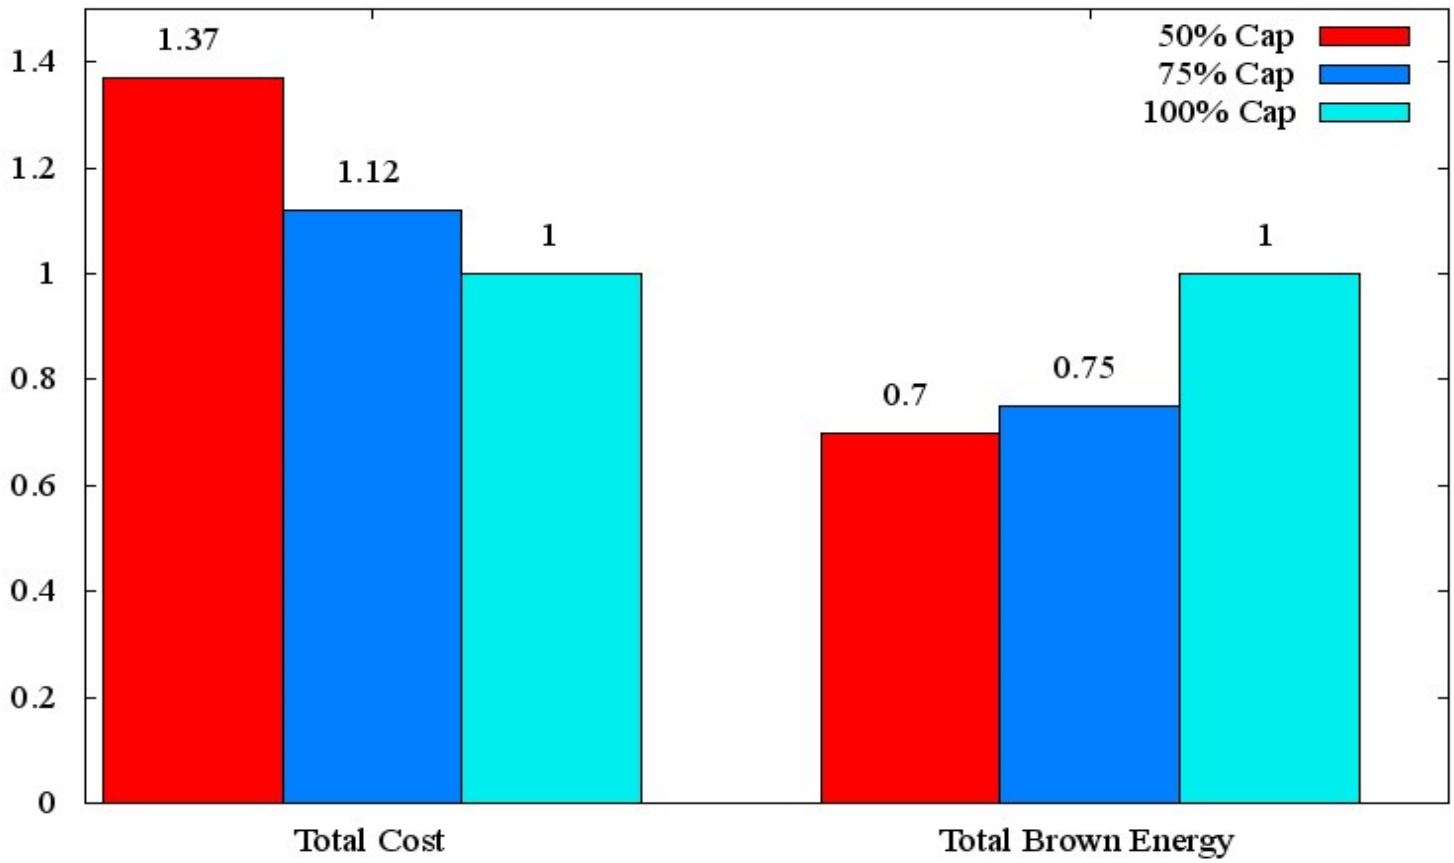

75% cap lowers brown energy by 25% for a 12% cost increase
Lowering the cap further produces small gains and high costs

Conclusions

- Many new ideas and directions
- Formal, scientific approach
- Very promising preliminary results
- Easy to broaden scope beyond services: HPC and enterprise DCs
- No related work on energy capping, across-DC distribution, time zones, green energy, variable electricity prices. Topics are new to us too! ☺
- Much more work to do…