bridging the gaps

Sustainable Urban Spaces

Sense and Sustainability

Networked Sensors for the Urban Environment

UCL

Energy Efficiency for the Computer Cloud

Kyle Jamieson, Venus Shum (Computer Science), **Muki Haklay** (Civil, Environmental and Geomatic Engineering) and **John Mitchell** (Electronic and Electrical Engineering)

Background and Objectives

Cloud computing provides flexibility to cloud providers in allocating computation to data centers. However, the cost of electricity required for running these data centers is very high. About one third of this electricity is spent only on the cooling mechanism. Therefore, cloud providers are under financial pressure to become “green”.

How can a cloud provider become green by saving energy and using environmentally-friendly energy resources, while simultaneously reducing the risk associated in shifting the responsibility to far-flung data centres?

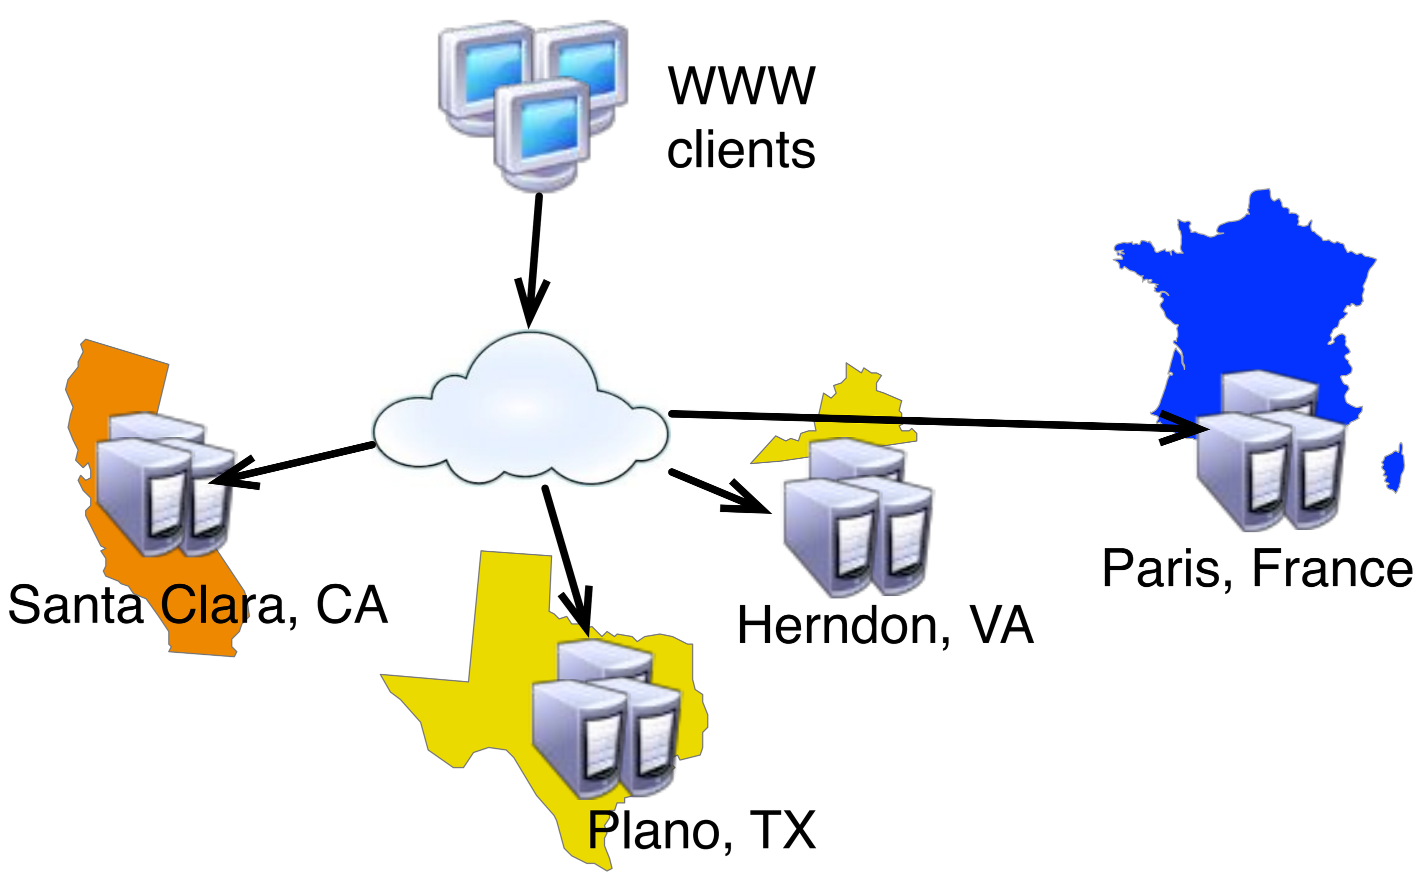

Figure 1: 1998 World Cup Web Site load balancing for traces taken from 30th April to 26th July.

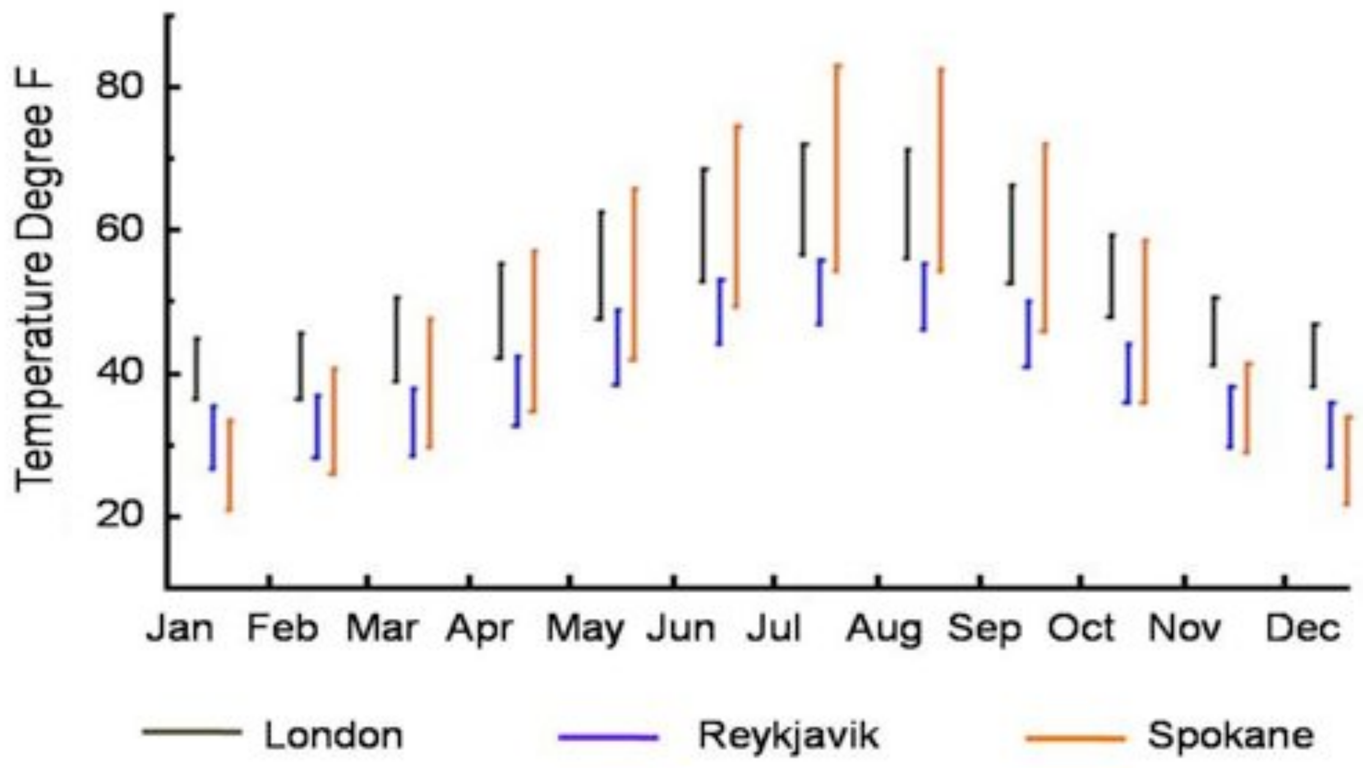

Figure 2: Average daily high and low temperatures in London, UK (GMT); Reykjavik, Iceland (GMT); Spokane, WA (GMT+8).

Approach

The opportunity we identify is that the cost of performing computation at a given location is not uniform across time and space, for the following reasons. (1) Seasonal variations in temperature favour certain locations on month-long timescales, as shown in Figure 2: Iceland lows prevail during the months from March to October, while Spokane lows prevail from November to February. (2) Daily variations in temperature, coupled with time zone differences (shown by the ranges of the bars in Figure 2) point to adaptation between data centers on different continents at a timescale of hours, during some months of the year.

(3) Spot electricity markets introduce hourly electricity price variations at different locations in the US, as Qureshi et al. have observed [ACM SIGCOMM ‘09]. (4) Factors unique to different countries, e.g. at many times in the year, Iceland offers lower-cost services in an environment with low ambient temperature, but possibly higher potential for impact from natural disasters.

Results

As countries roll out energy sources in differing proportions (such as wind, hydro, or geothermal power), the cloud should adapt as shown in Figure 3, which simulates our workload running on representative data centres in London with data centres in Reykjavik.

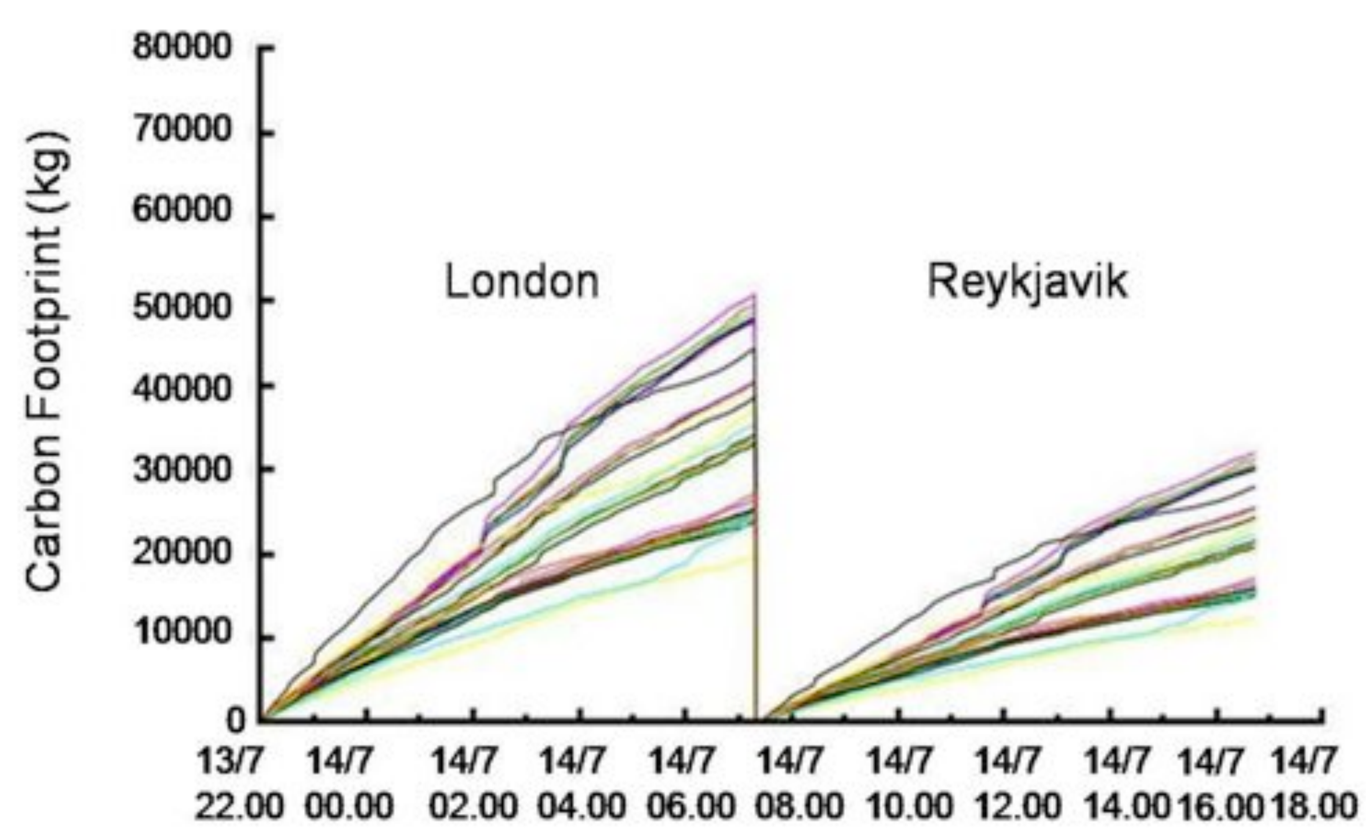

Figure 3: Cumulative kg CO2 produced by a web server workload at London (before 07:00) and Reykjavik (after 07:00) -- representative data center efficiencies.