Energy-Delay Tradeoffs

Brent Horine

March 28, 2011

Citation

Moo-Ryong Ra, Jeongyeup Paek, Abhishek B. Sharma, Ramesh Govindan, Martin H. Krieger, Michael J. Neely, "**Energy-Delay Tradeoffs in Smartphone Applications**", In Proc. Of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys), 2010, pp. 255-270.

Motivation

* Reduce energy consumption of uploading videos in Tomography project, while providing a user controllable tradespace w.r.t. delay
* 5000 videos (as of writing)
* http://tomography.usc.edu/

Applications of Tomography

* Surveillance of transportation hub in LA
* Behavior analysis of developmental disabilities in children
* Document construction in post-Katrina Mississippi for improving zoning regs and ordinances

Example Scenario

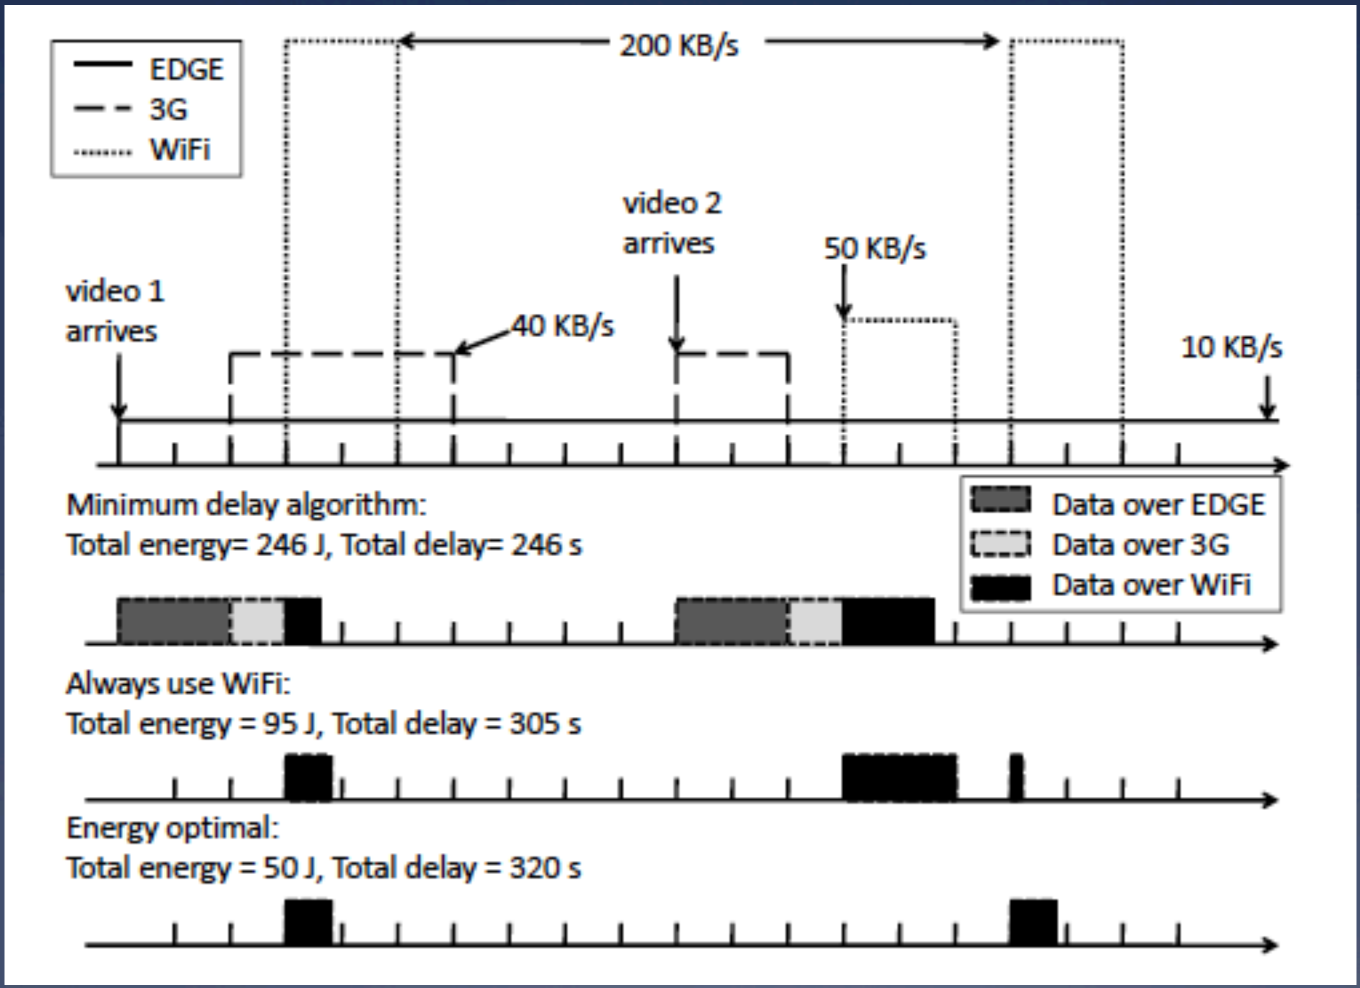

Strategies

* Link Selection Problem
* Minimum Delay (5X energy consumption)
* Always use WiFi (2X)
* Energy Optimal (1X)

Problem Formulation

* Link Selection Problem solved with a Lyapunov optimization framework
* Minimizes total energy expenditure subject to keeping the average queue length finite
* Control Algorithm: SALSA (Stable and Adaptive Link Selection Algorithm)

Implementation Details

* Adjust parameter V in Lyapunov control to tune the energy-delay tradeoff
* Derive link rate estimates empirically from RSSI measurements, then learn with use
* Simulated and experimental traces indicate an energy savings (in terms of battery charge life) of 10-40%

Generic Results and Benefits

* Using Lyapunov optimization techniques, they get arbitrarily close to a target Power consumption, while maintaining queue stability
* Cost of reduced power consumption is a larger delay
* Average queue backlog grows linearly in V
* Does not require prior knowledge of distributions of A[t] and SI[t], except that variances are finite

Time Average Power Consumption and Queue Backlog

$$\overline{P} = \limsup_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}\{P[\tau]\} \leq P^* + \frac{B}{V}$$

$$\overline{U} = \limsup_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}\{U[\tau]\} \leq \frac{B + VP^*}{\varepsilon}$$

$\varepsilon > 0$ constant describing distance between arrival pattern and the capacity region boundary
P^* = theoretical lower bound on time average power consumption
B = upper bound on the sum of variance of A[t] and μ [t]

V Parameter Control

- Although V controls the tradeoff, it is not a simple relationship
- Let α be slope of time averaged power consumption

$$\frac{d(P^* + B/V)}{dV} = \frac{-B}{V^2} = -\alpha$$

$$\Rightarrow V = \sqrt{\frac{B}{\alpha}}$$

Expressing V in User Units

$$V[t] = \sqrt{\frac{B[t]}{\alpha(D[t]+1)^{\alpha}}}$$

D[t] = instantaneous delay in data transfer (time that the bit at the head of the queue has been resident in queue)

B is now time varying

$\alpha \rightarrow 0$ for energy efficiency (V[t] decreases slowly)

α *large* for quick transfers (V[t] decreases rapidly)

Assumed Input Data Patterns

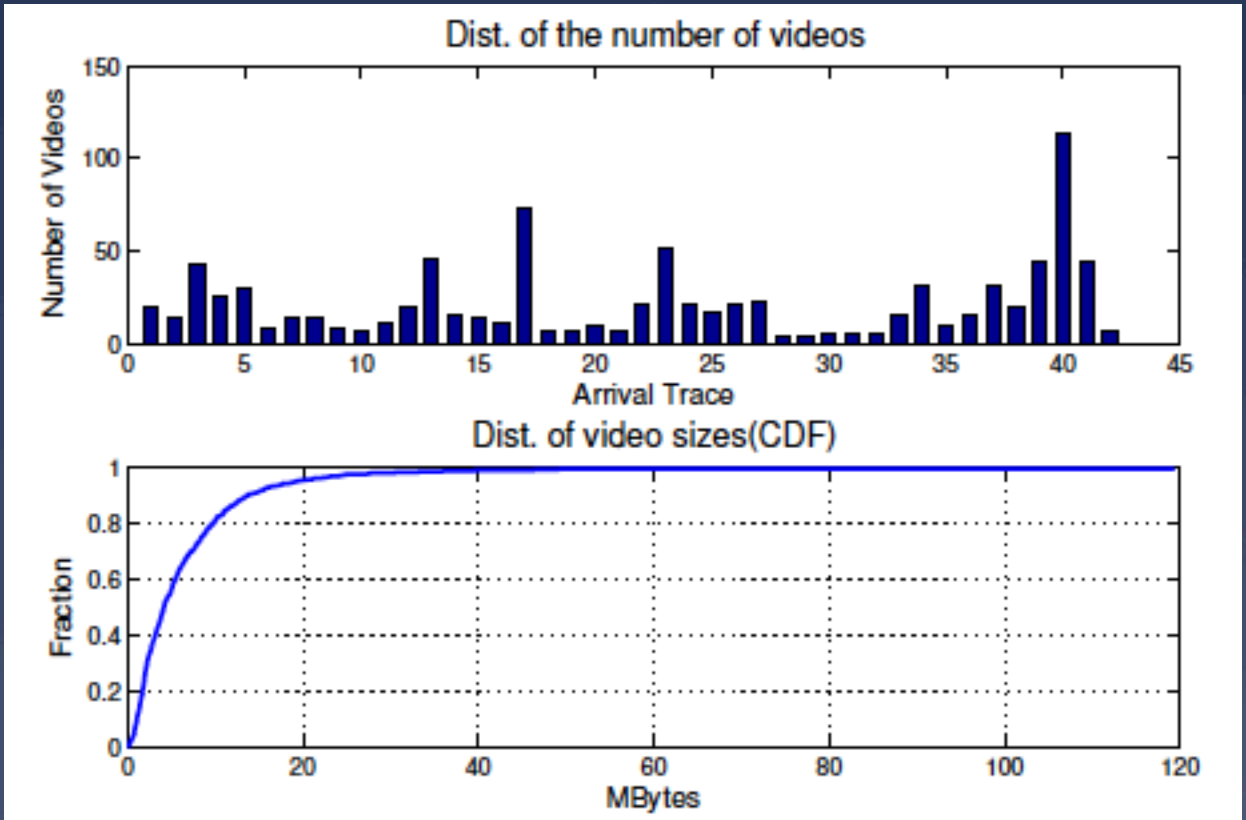

Link Availability w/Failure Probability (CDF)

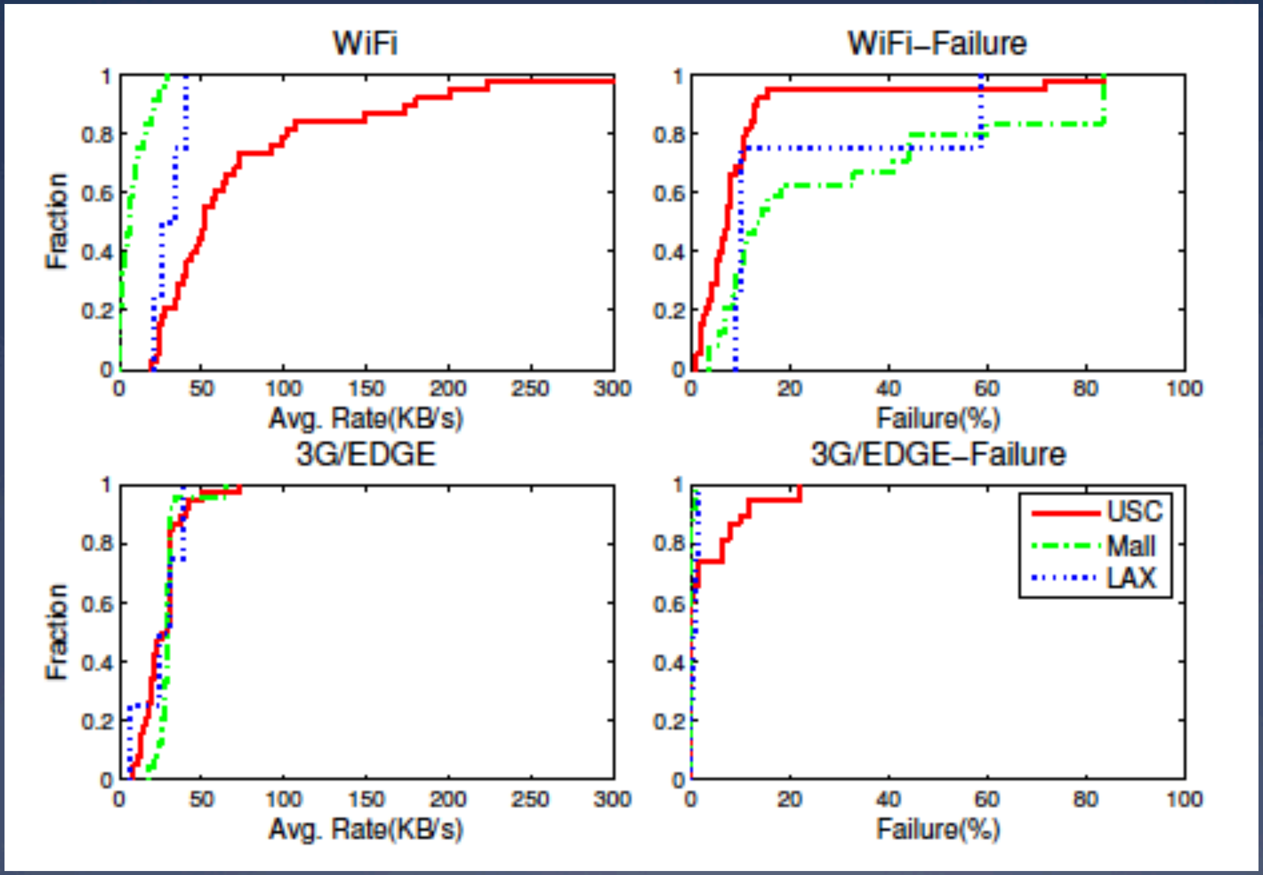

Simulation Results

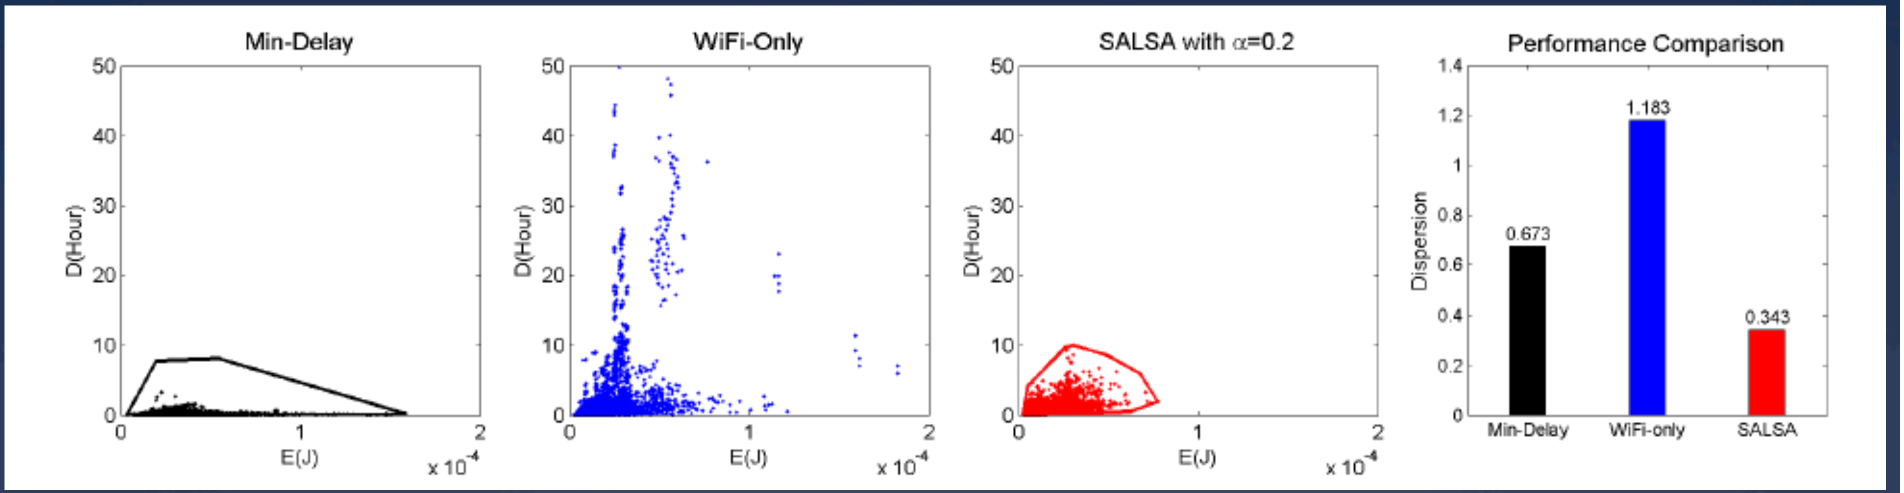

$$\overline{E} = \frac{\sum_{i=1}^{N} E_i}{\sum_{i=1}^{N} S_i}, \quad \overline{D} = \frac{\sum_{i=1}^{N} (D_i S_i)}{\sum_{i=1}^{N} S_i}$$

SALSA Energy Savings and Delay Penalty

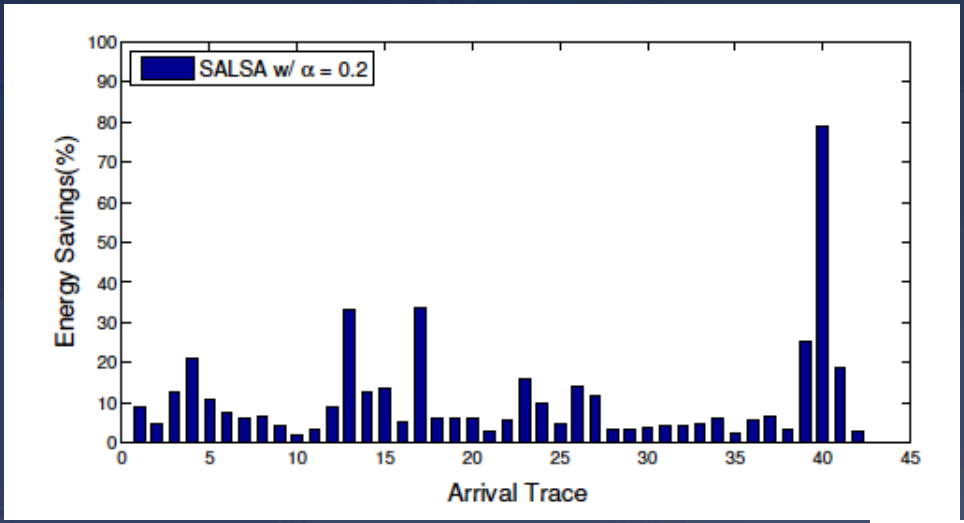

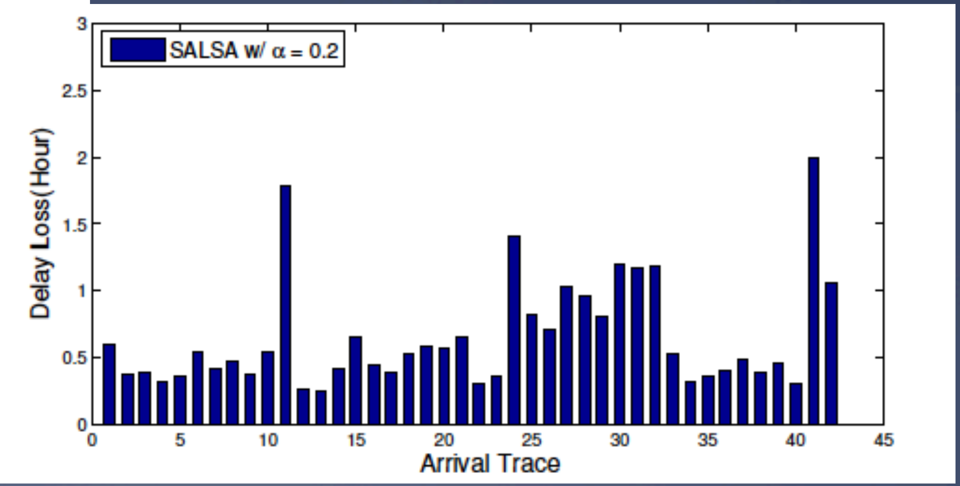

Dispersion versus Alpha

* Dispersion describes distance from "ideal"

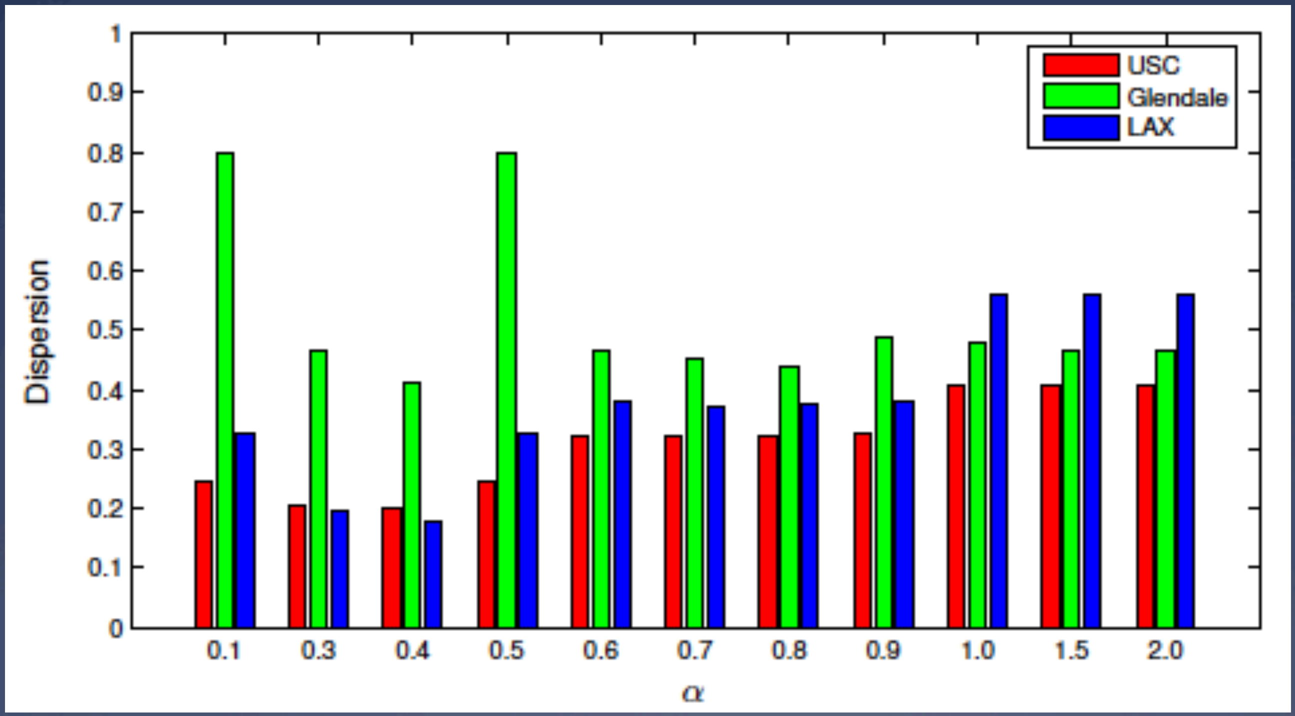

Simulation Results

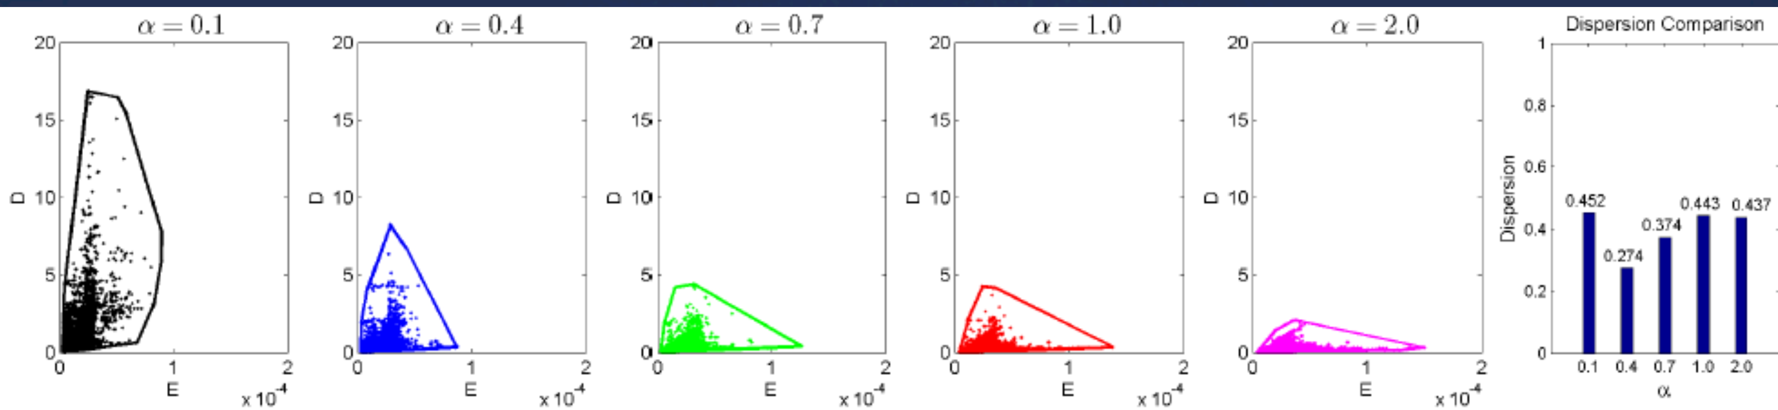

Figure 8: SALSA envelopes for different α

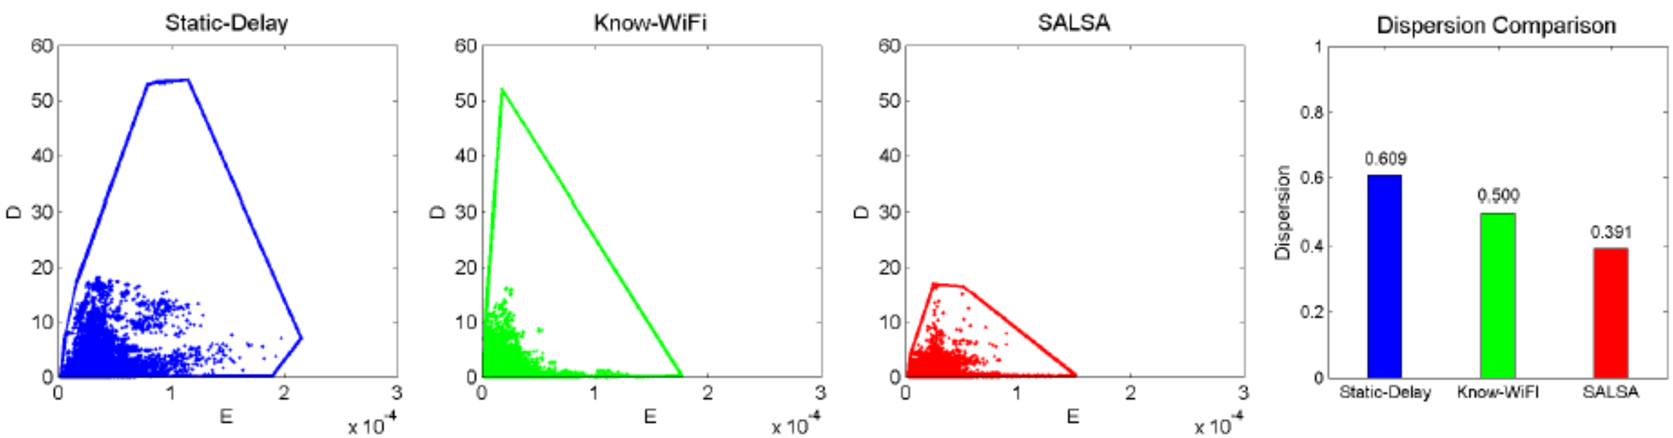

Figure 10: STATIC-DELAY vs KNOW-WIFI vs SALSA

Scan Interval Comparisons

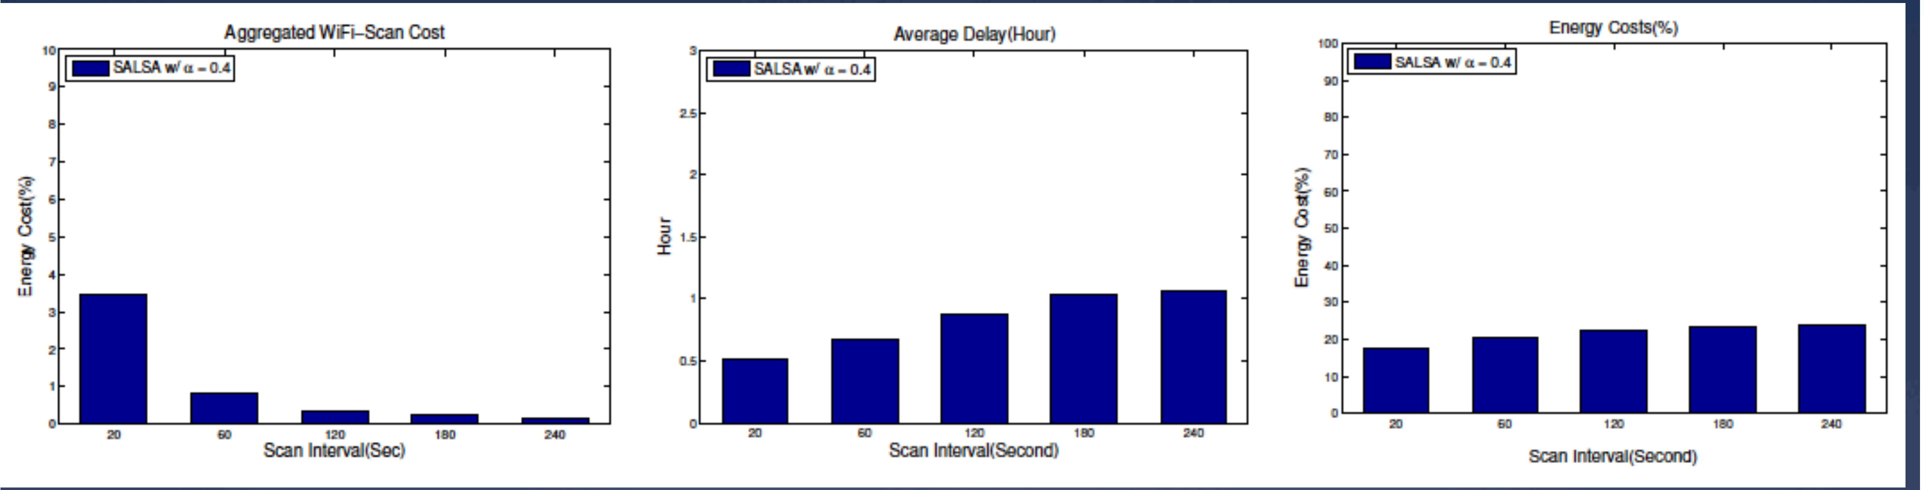

Experimental Results

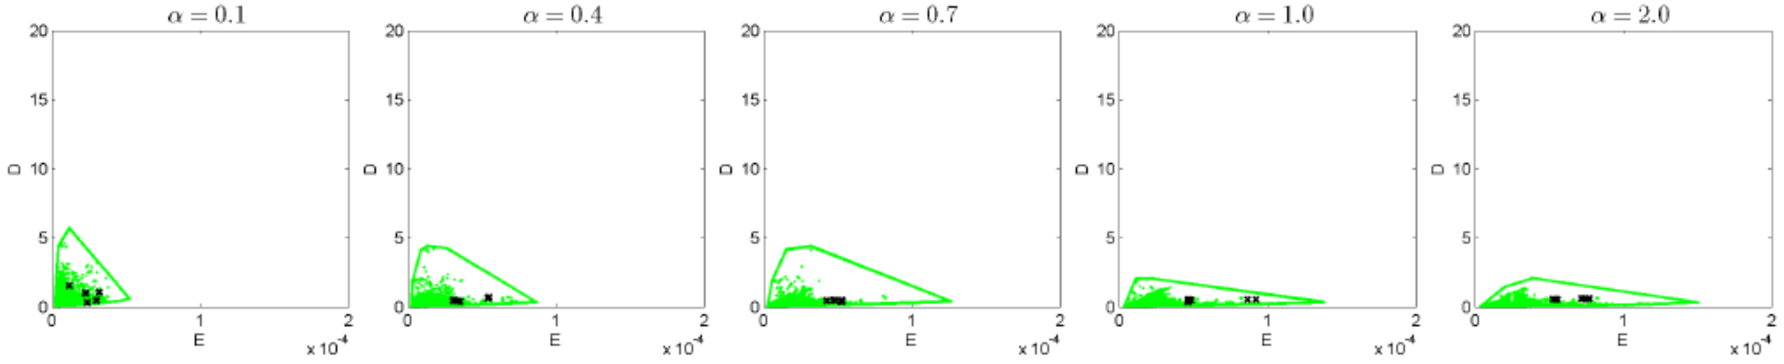

Figure 15: Experimental result at the USC Campus compared to simulation results

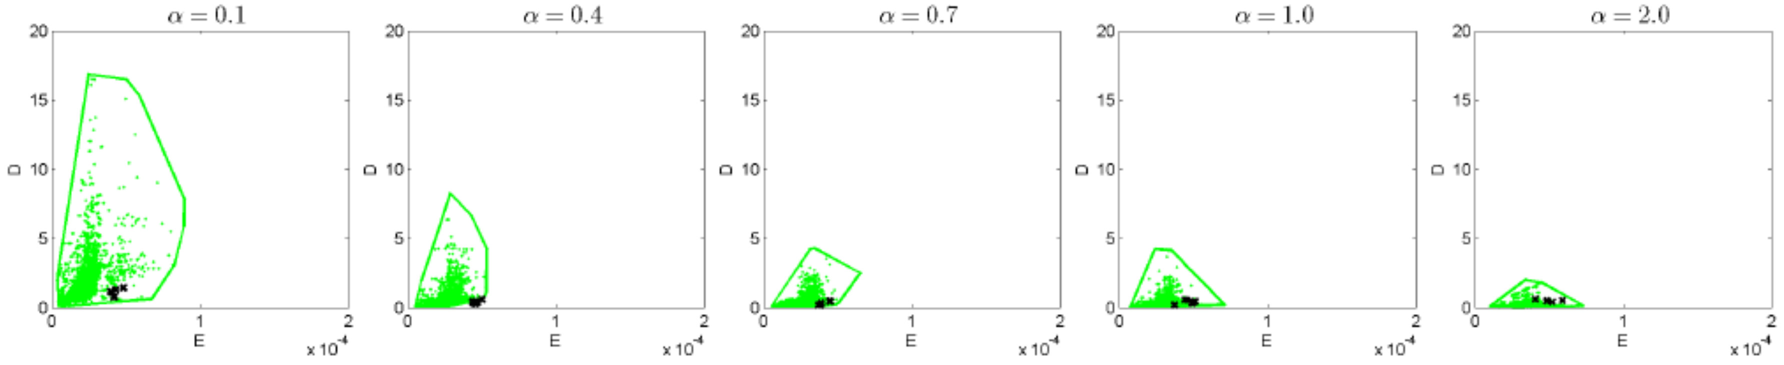

Figure 16: Experimental result at Shopping Mall compared to simulation results

Significance of Work

* A user controllable framework with a large tradespace between latency of video upload and battery life

Potential Improvements

- Explore a larger tradespace using the simulator rather than relying upon empirical traces
- Incorporate learning to predict when a repetitive user might arrive near a quality WiFi AP
- As the authors suggest, continue to work on relating α to user units
- Could use location rate info to guess if user is likely to move into a superior region soon