Submitted 31 July 2015
Accepted 3 November 2015
Published 9 December 2015

Corresponding author
Navin Sharma,

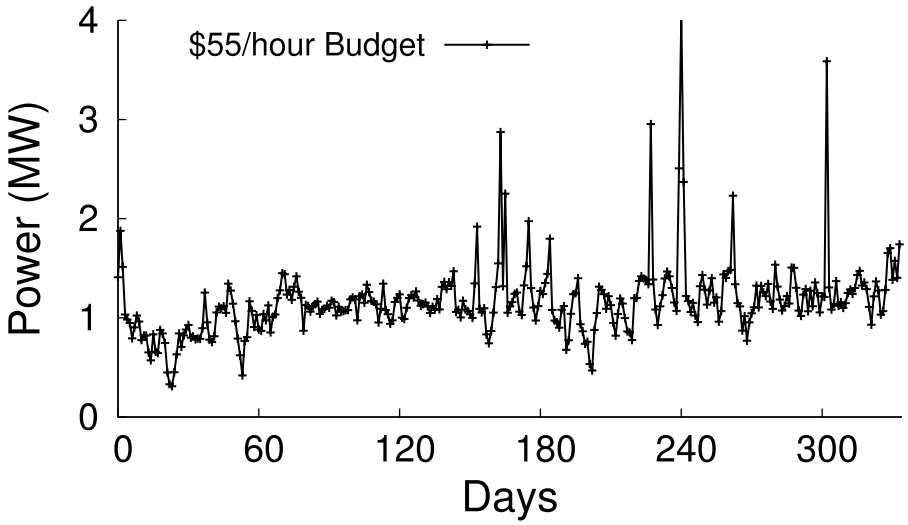

Figure 1 **Electricity prices vary every five minutes to an hour in wholesale markets, resulting in the power available for a fixed monetary budget varying considerably over time**

data center capital costs are enormous, maximizing the power delivery infrastructure's utilization by operating as many servers as possible is important. However, data centers

Energy sources

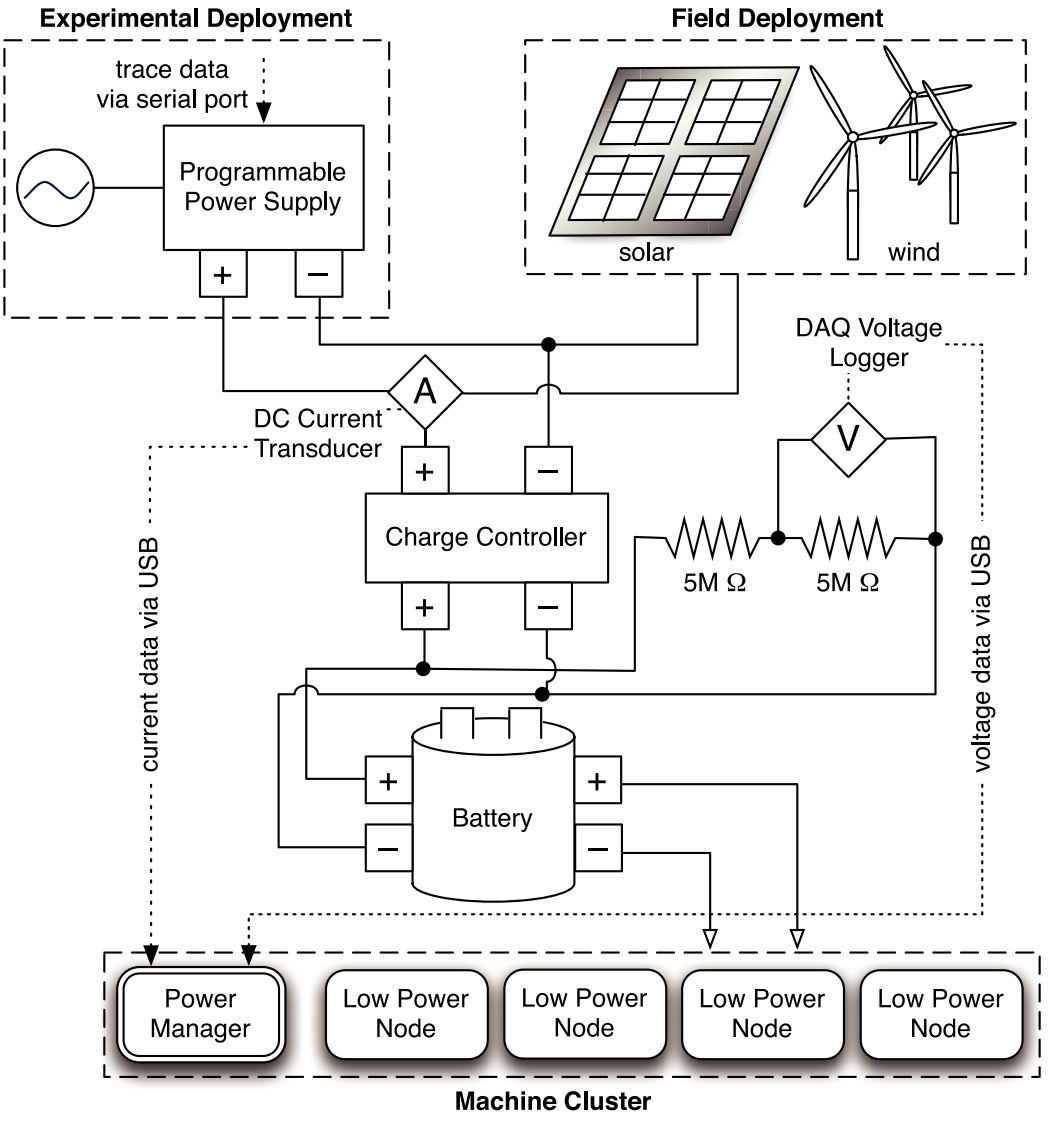
Experimental Deployment
Field Deployment
trace data
via serial port
Programmable
Power Supply
+
−
solar
wind
DAQ Voltage
Logger
A
DC Current
Transducer
V
+
−
Charge Controller
5M Ω
5M Ω
current data via USB
voltage data via USB
Battery
Power
Manager
Low Power
Node
Low Power
Node
Low Power
Node
Low Power
Node
Machine Cluster

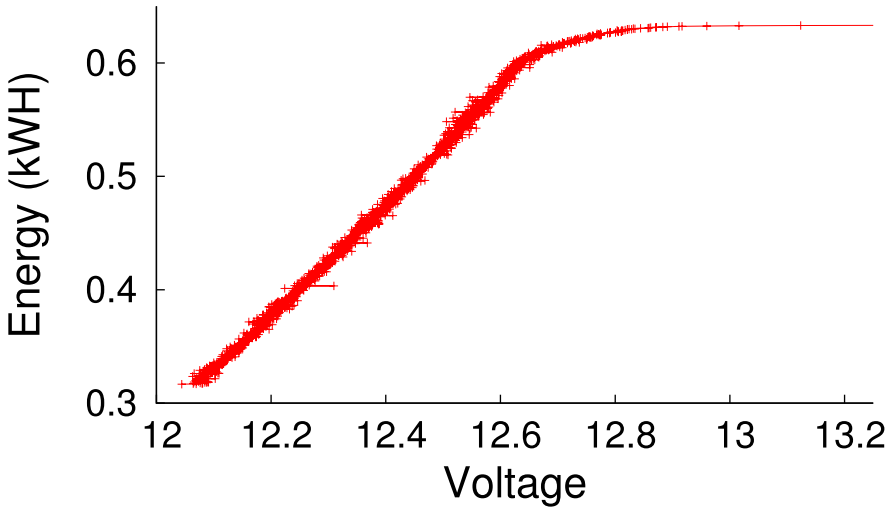

Figure 4

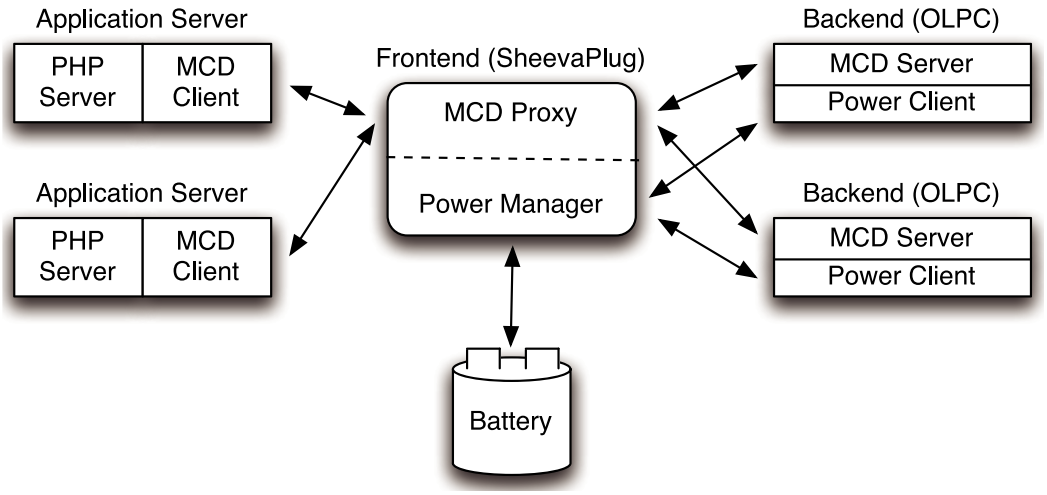
Application Server
PHP Server
MCD Client
Application Server
PHP Server
MCD Client
Frontend (SheevaPlug)
MCD Proxy
Power Manager
Battery
Backend (OLPC)
MCD Server
Power Client
Backend (OLPC)
MCD Server
Power Client

Figure 6

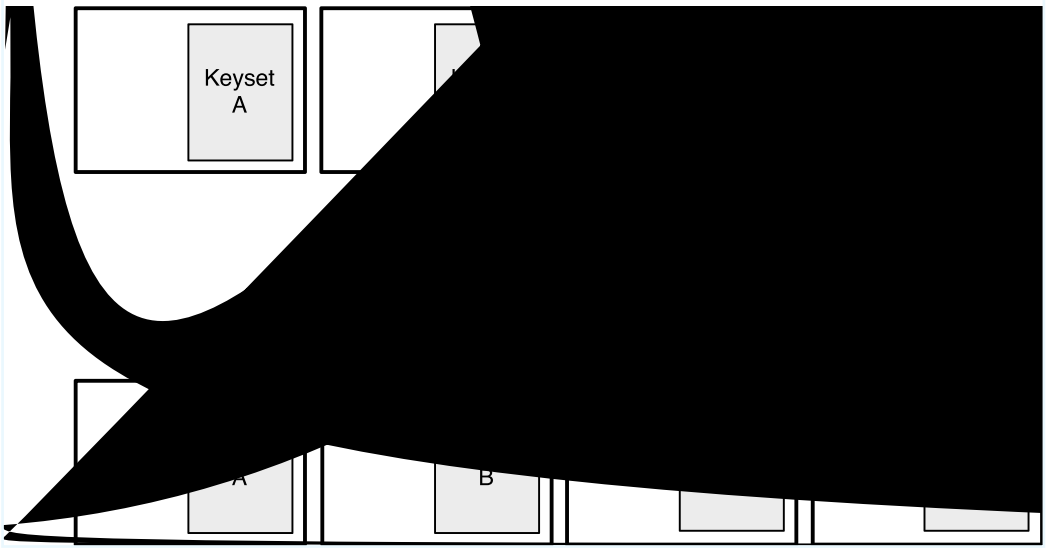
Keyset
A
B

Table 4 Summary of the best policy for a given performance metric and workload combination.

the characteristics of a distributed multimedia cache, which serves videos requested by

generally limited by its network card and disk I/O. So, the cache server can potentially reduce its energy consumption by sending data at its full capacity for a fraction of a time interval (usually few seconds) and going to a low-power state for the remaining period of the time interval, as shown in Fig. 8. In essence, the server could employ the blinking

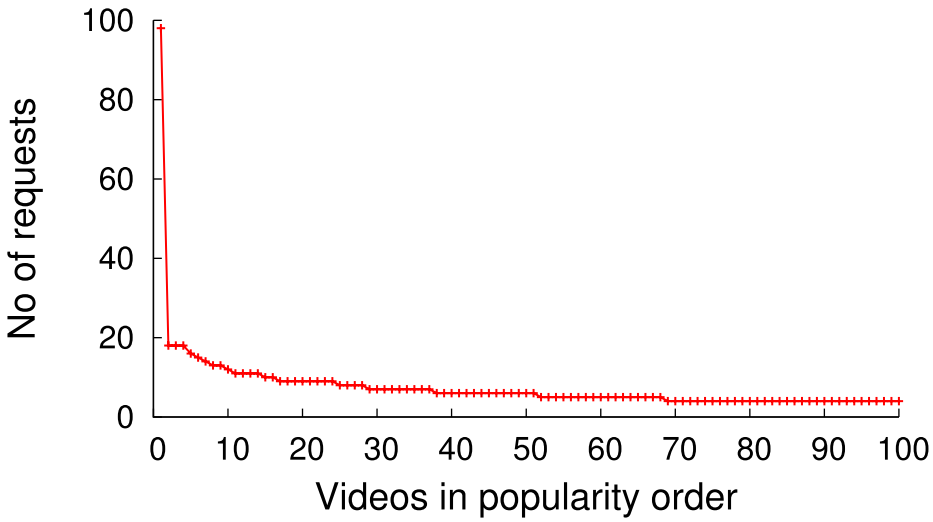
100
80
60
40
20
0
No of requests
0
10
20
30
40
50
60
70
80
90
100
Videos in popularity order

Figure 9

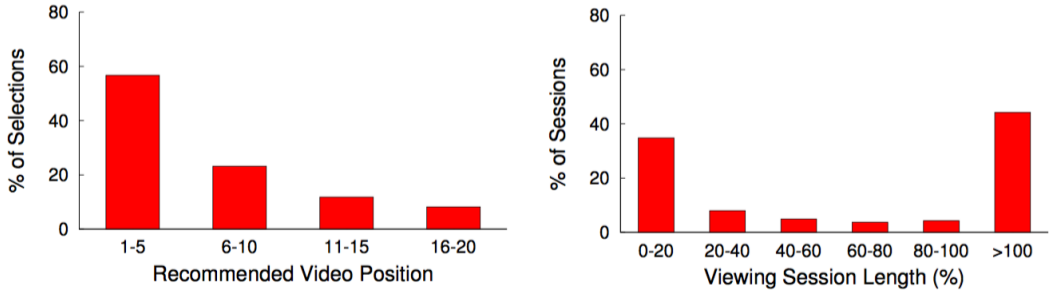
80
60
40
20
0
% of Selections
1-5
6-10
11-15
16-20
Recommended Video Position
A
80
60
40
20
0
% of Sessions
0-20
20-40
40-60
60-80
80-100
>100
Viewing Session Length (%)
B

while also controlling chunk server placement and eviction. Clients, e.g., web browsers,

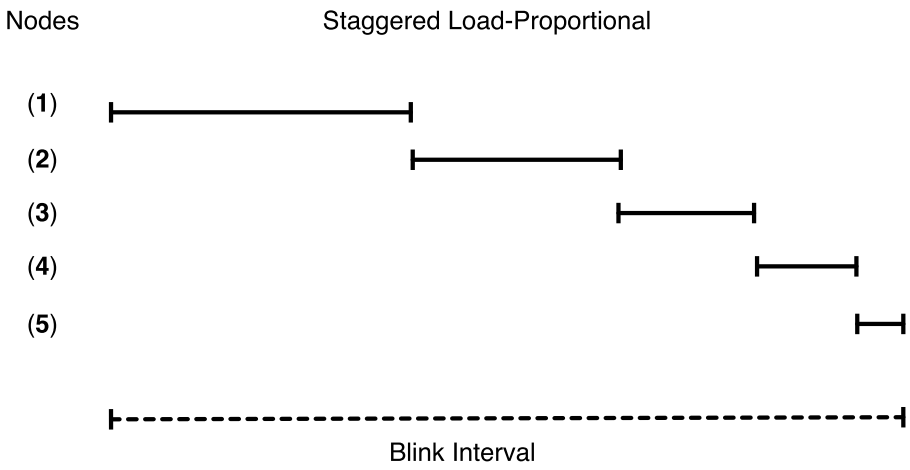
Nodes
Staggered Load-Proportional
(1)
(2)
(3)
(4)
(5)
Blink Interval

key requests at a configurable but steady rate according to either a Zipf popularity distribution, parameterized by , or a uniform popularity distribution. As in a typical

As discussed in 'Blink Prototype,' we use a cluster of ten Mac minis for GreenCache. We use one Mac mini to run the proxy and power manager, whereas we run a cache server and power client on other Mac minis. The proxy connects to a WiMAX base station (NEC Rel.1 802.16eBS) through the switch. We use a Linux laptop with a Teletonika USB WiMAX modem to run as a client. We also use a separate server to emulate multiple WiMAX clients.

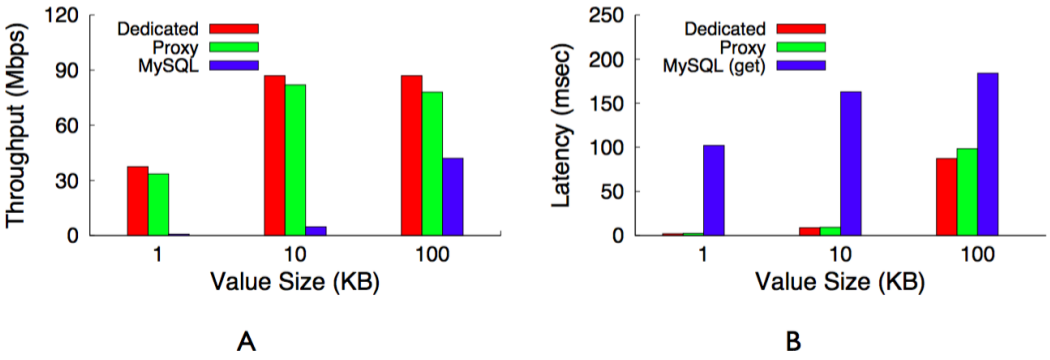

Figure 14 Maximum throughput (A) and latency (B) for a dedicated memcached server, our memcached proxy, and a MySQL server. Our proxy imposes only a modest overhead compared with a dedicated memcached server. (A) Throughput. (B) Latency.

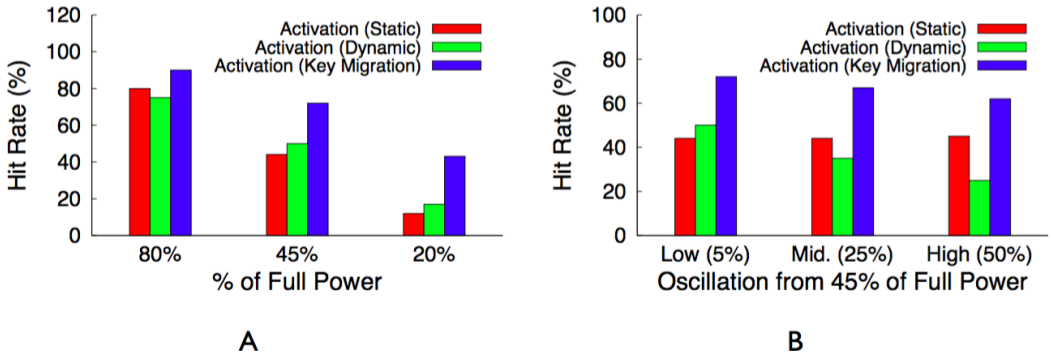
Hit Rate (%)
120
100
80
60
40
20
0
Activation (Static)
Activation (Dynamic)
Activation (Key Migration)
80%
45%
20%
% of Full Power
A
Hit Rate (%)
100
80
60
40
20
0
Activation (Static)
Activation (Dynamic)
Activation (Key Migration)
Low (5%)
Mid. (25%)
High (50%)
Oscillation from 45% of Full Power
B

Figure 15

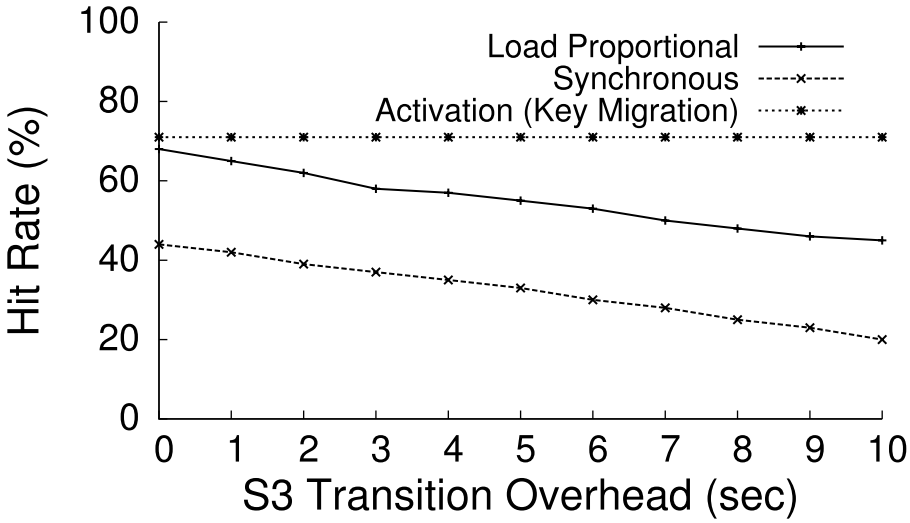
100
80
60
40
20
0
Hit Rate (%)
0
1
2
3
4
5
6
7
8
9
10
S3 Transition Overhead (sec)
Load Proportional
Synchronous
Activation (Key Migration)

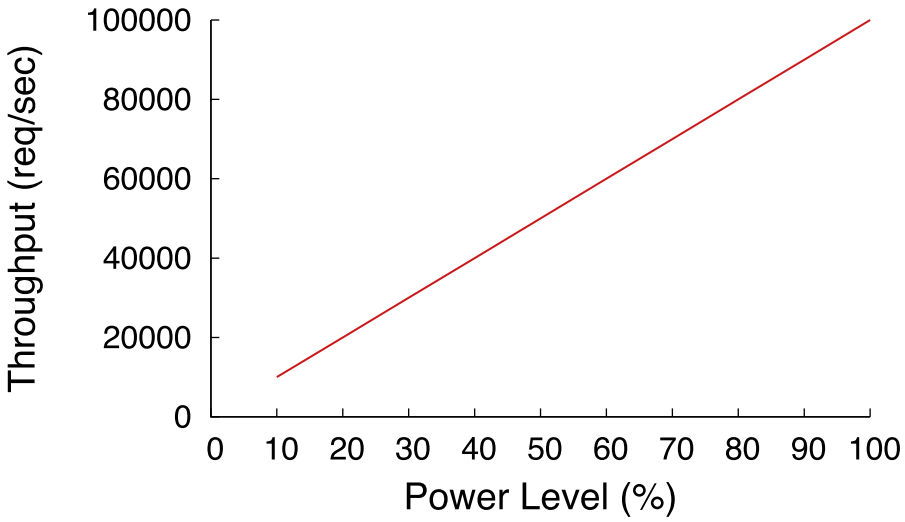
Throughput (req/sec)
100000
80000
60000
40000
20000
0
0
10
20
30
40
50
60
70
80
90
100
Power Level (%)

Table 5 Standard deviation, 90th percentile, and average buVering time at diVerent power levels and blink intervals.

BuVering time (s)	Power (%)				
+	20	40	60	80	100

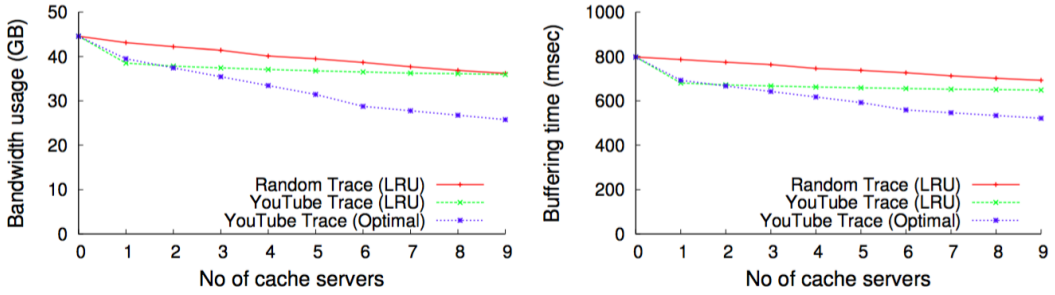
Bandwidth usage (GB)
50
40
30
20
10
0
0 1 2 3 4 5 6 7 8 9
No of cache servers
Random Trace (LRU)
YouTube Trace (LRU)
YouTube Trace (Optimal)
Buffering time (msec)
1000
800
600
400
200
0
0 1 2 3 4 5 6 7 8 9
No of cache servers
Random Trace (LRU)
YouTube Trace (LRU)
YouTube Trace (Optimal)

A B

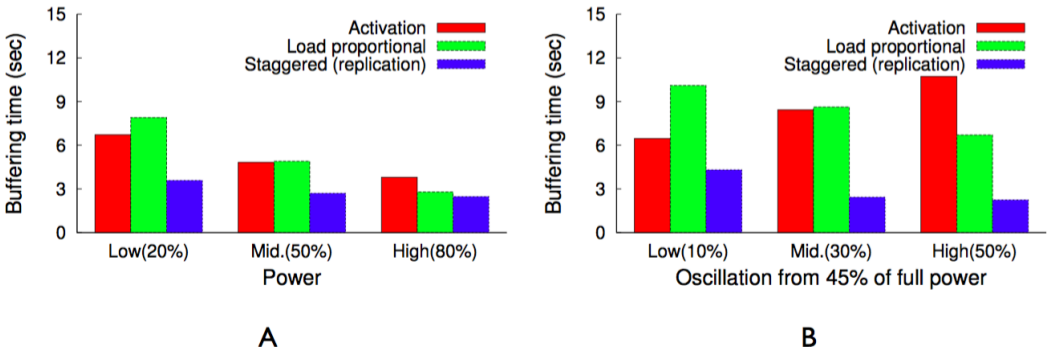

Figure 24 BuVering time at various steady and oscillating power levels. (A) Steady power. (B) Oscillating power.

Figure 24A shows the average buVering time at diV

Figure 25 **BuV**

center's power delivery infrastructure, including the possibility of using dense clusters of