

PuRPLe: Predictive Reduction of Power and Latency

Matthew Craven

Ahmed Amer

Department of Computer Science

University of Pittsburgh

MSST2005



Conflicting Goals

- ▶ *Traditionally, system designers have had to compromise between power consumption and performance*
- ▶ *Reducing power leads to increased latency*
- ▶ *For disks, we apply a spindown policy to reduce power, but pay a price to bring the disk back to active state*

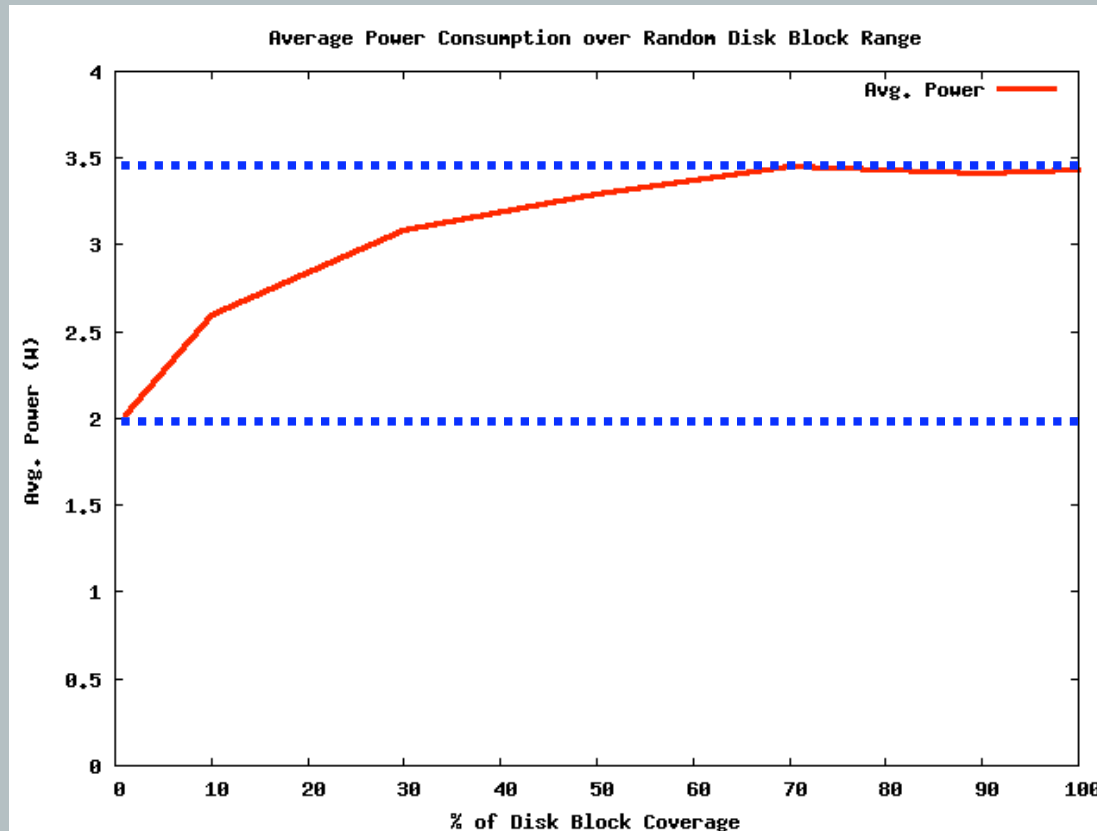


An End to Conflict

- ▶ *By predicting disk accesses over a short interval, we can prefetch those accesses and shut down the disk for the remainder of this interval*
- ▶ *This leads to substantial power savings*
- ▶ *We also reduce the number of times we spin down the disk, which improves performance*



Disk Power Consumption

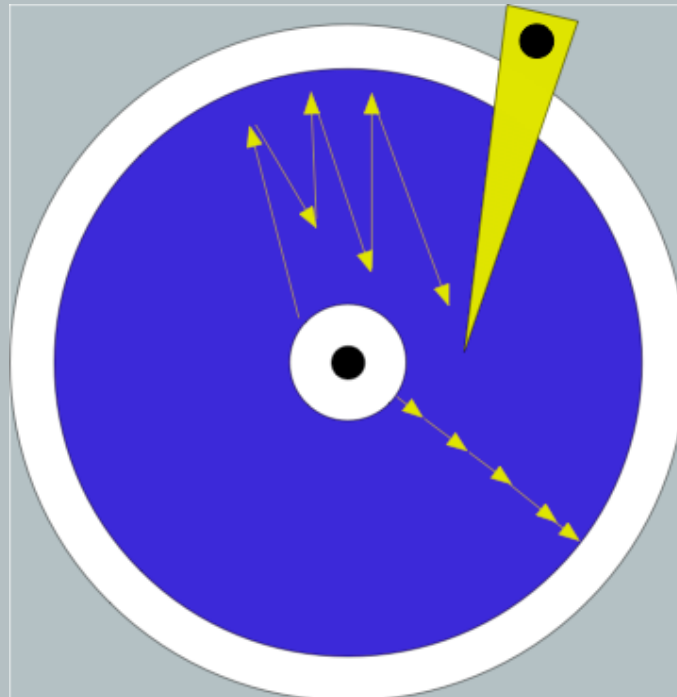


- Accurate disk power measurements demonstrate that random arm movements can increase power consumption up to 40%

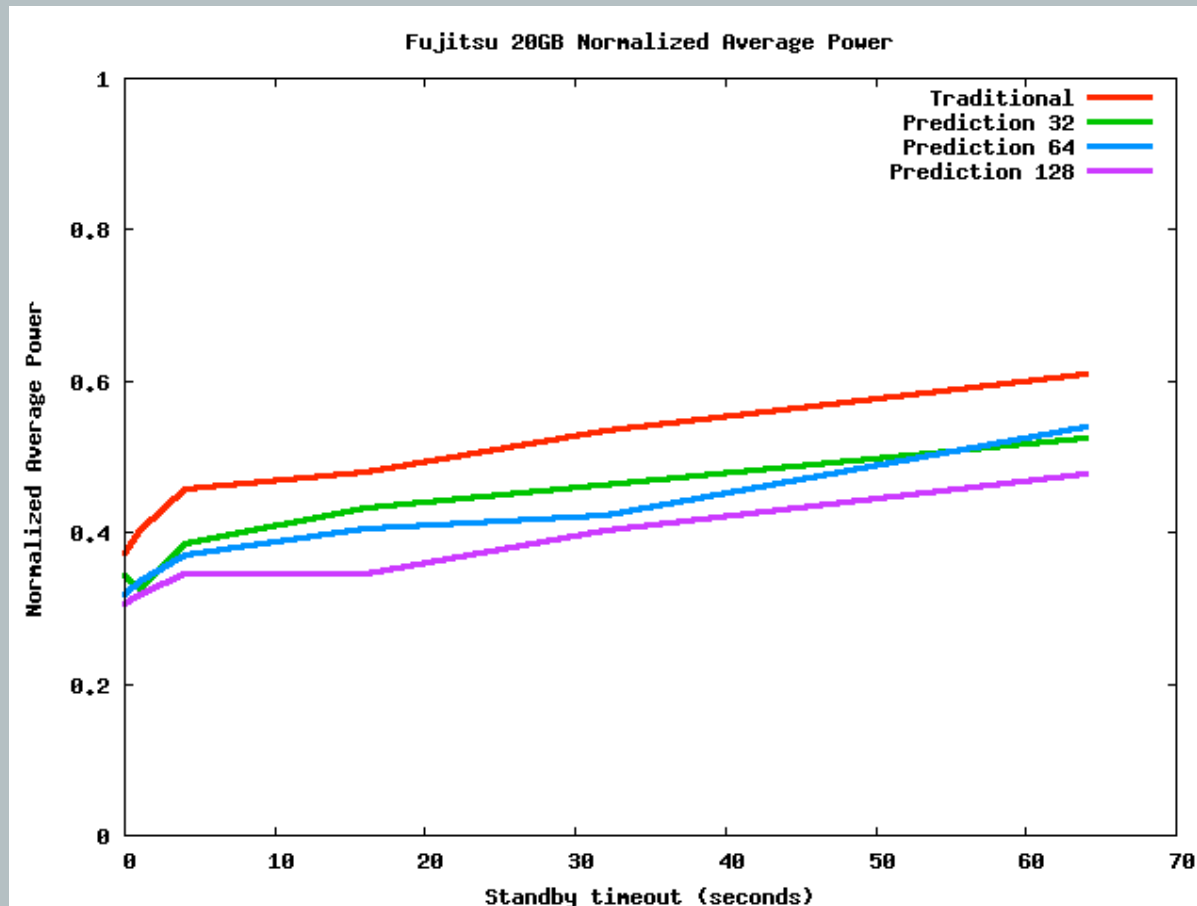


Reordering

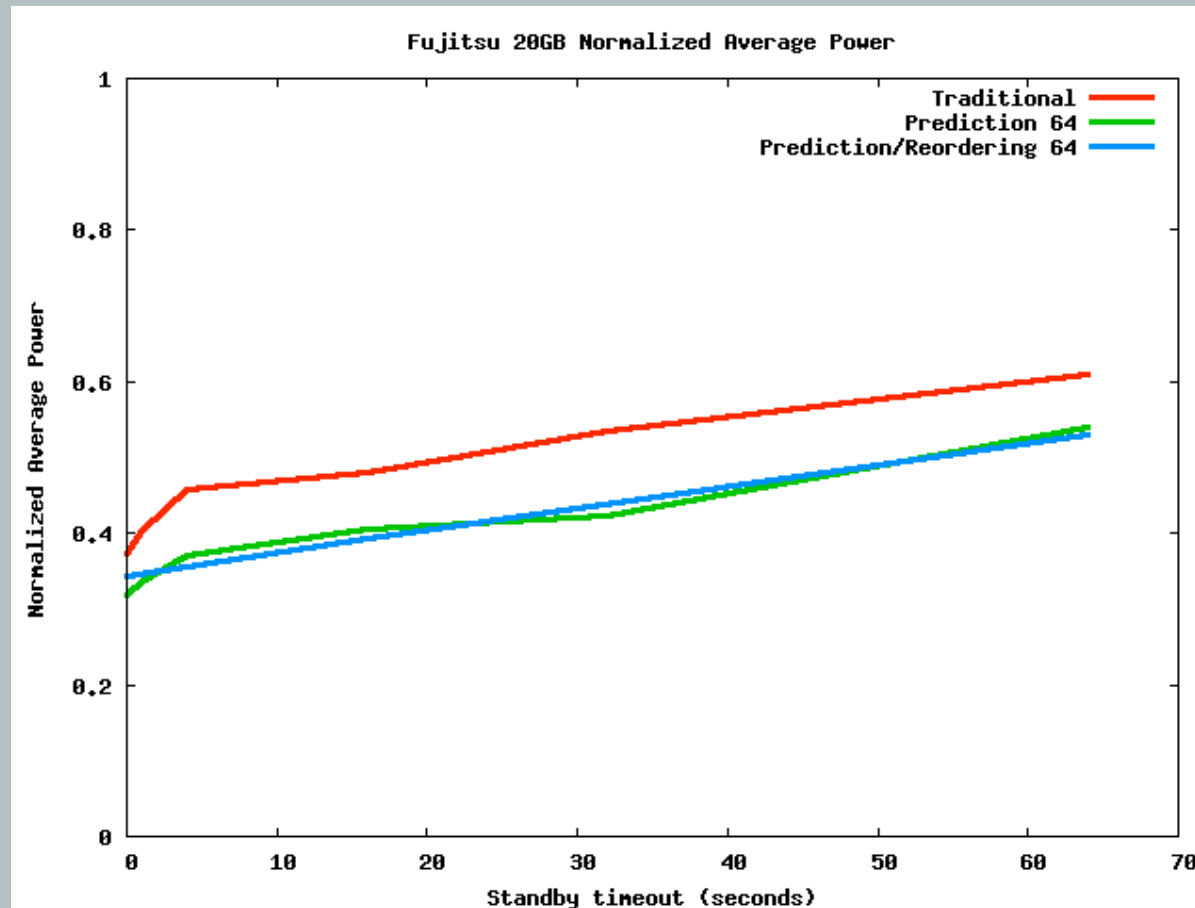
- ▶ *We can reduce arm movements by simply reordering our accesses*
- ▶ *Reordering the access requests simulates effective layout changes*



Prediction Effectiveness

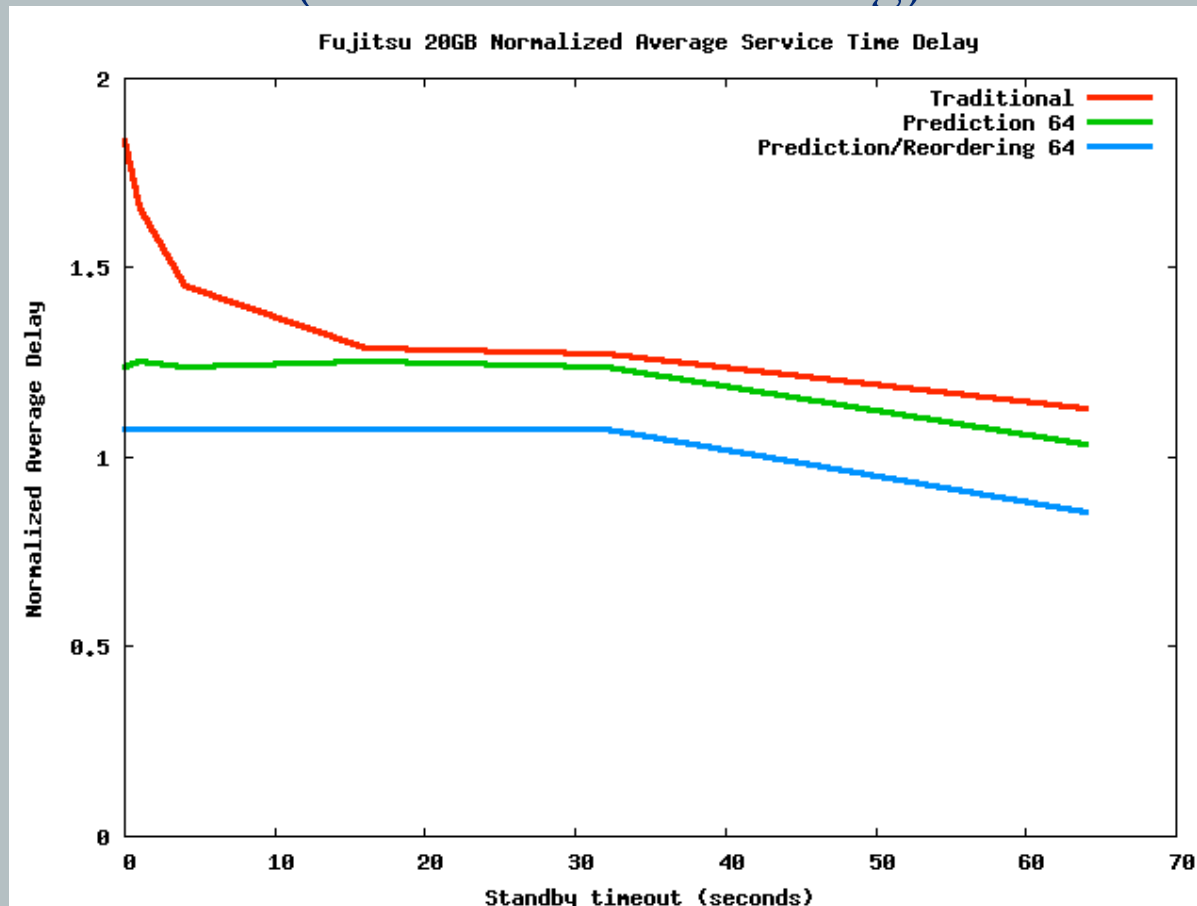


Prediction+Reordering



Latency Results

(Prediction+Reordering)



Summary

- ▶ *Power reduction and access latency need not be conflicting goals*
- ▶ *Prediction yields great power savings, and good access latencies*
- ▶ *Combining Prediction with Reordering further improves power consumption but really shines when we consider access latency, which it not only preserves, but improves*

