



Sara Alspaugh, Laura Keys, Andrew Krioukov

Motivation

- ❖ Energy consumption in clusters is not proportional to load
 - Underlying nodes are not power proportional
- ❖ Cluster may be consuming large fraction of peak power when only lightly loaded

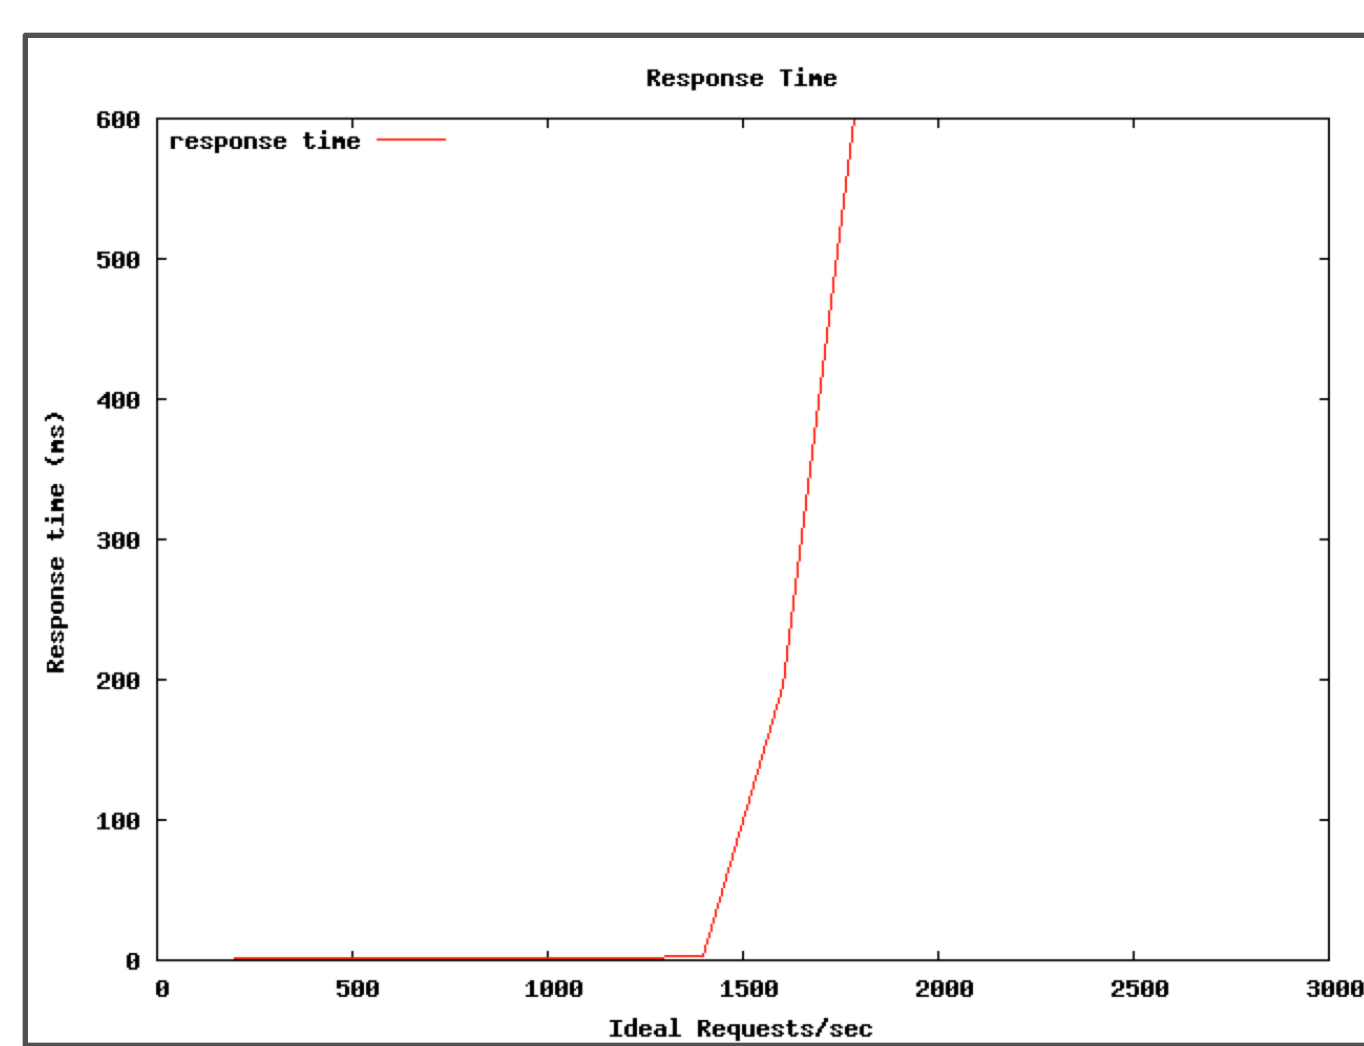
Goals

- ❖ Build a power proportional web server cluster
 - Use energy-efficient nodes
 - Maintain response time guarantees
 - Handle same peak load as original cluster

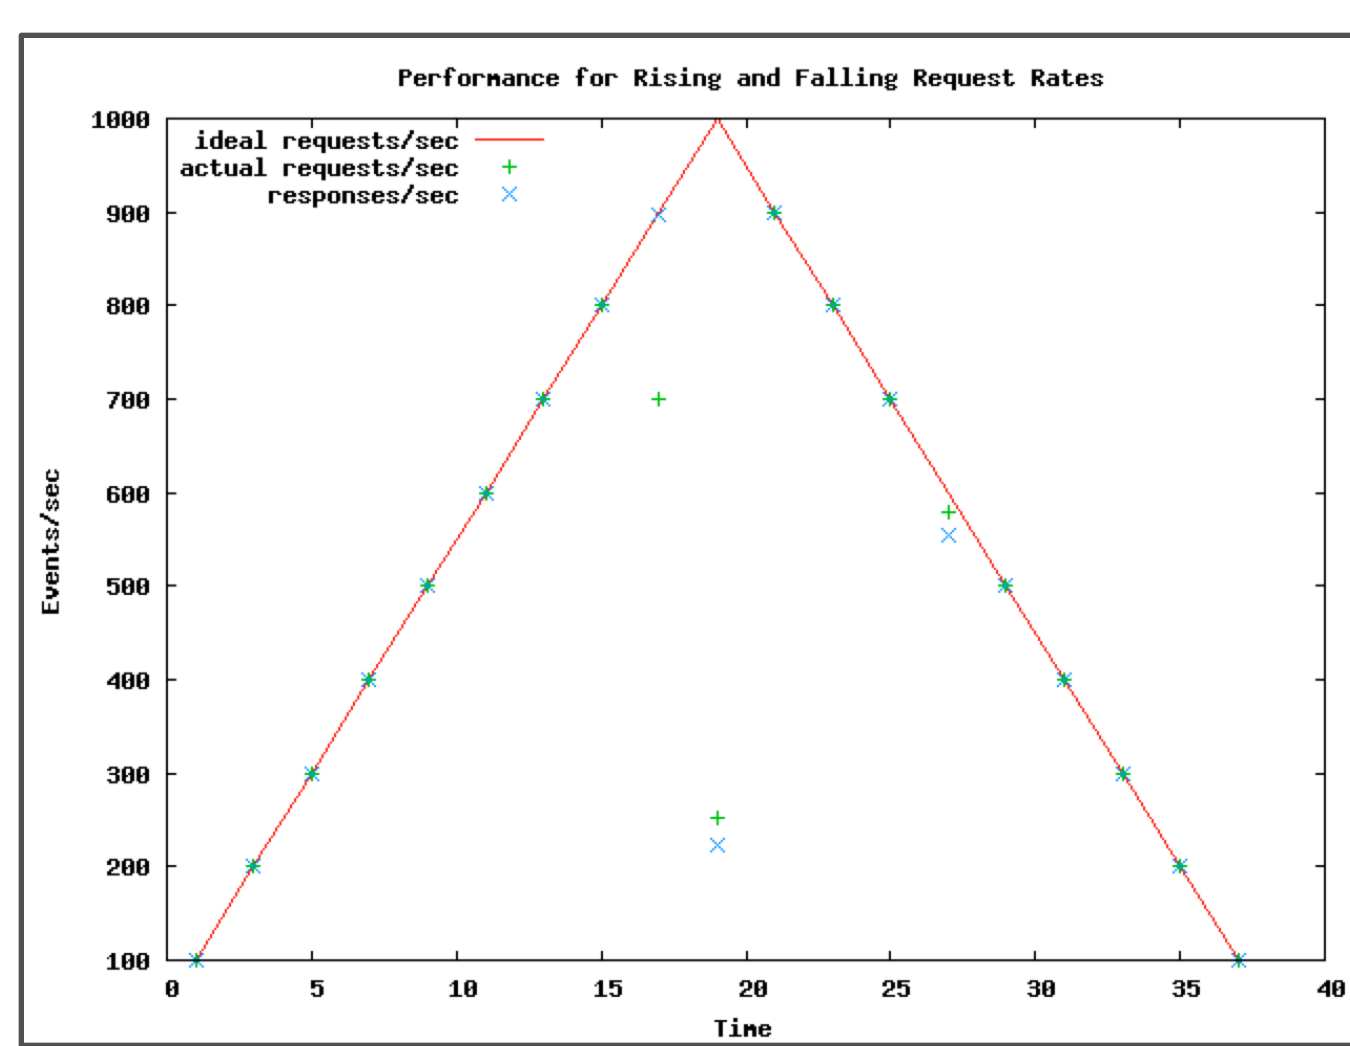
Workloads

- ❖ Synthetic httperf stress tests
 - Homogenous: requests single file repeatedly
 - Heterogeneous: uses actual file distribution from web server trace
- ❖ Real world web server traces

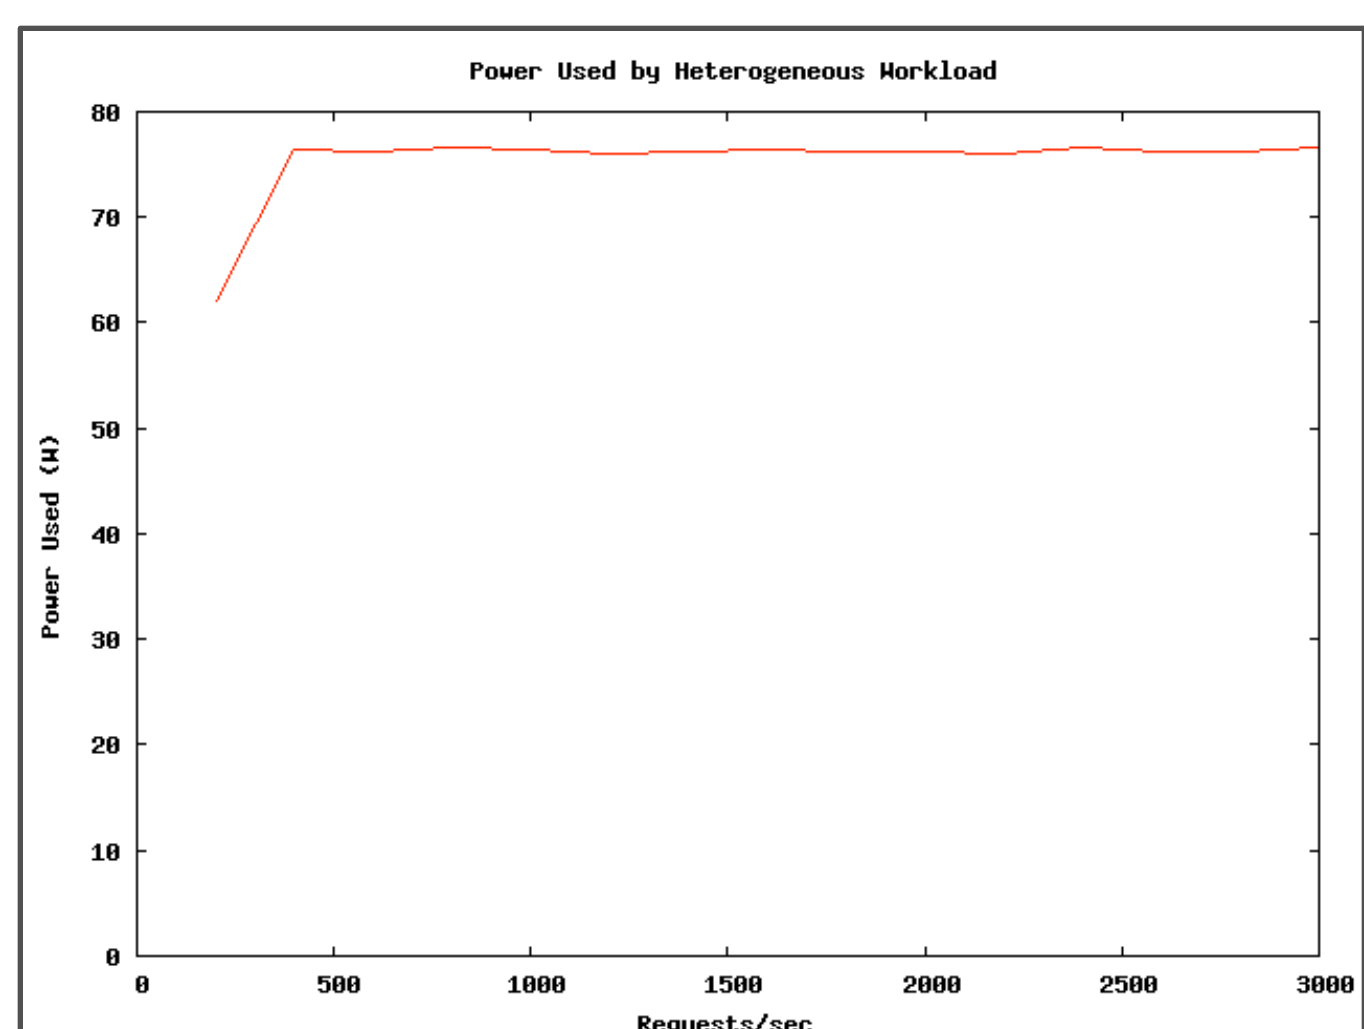
Evaluation



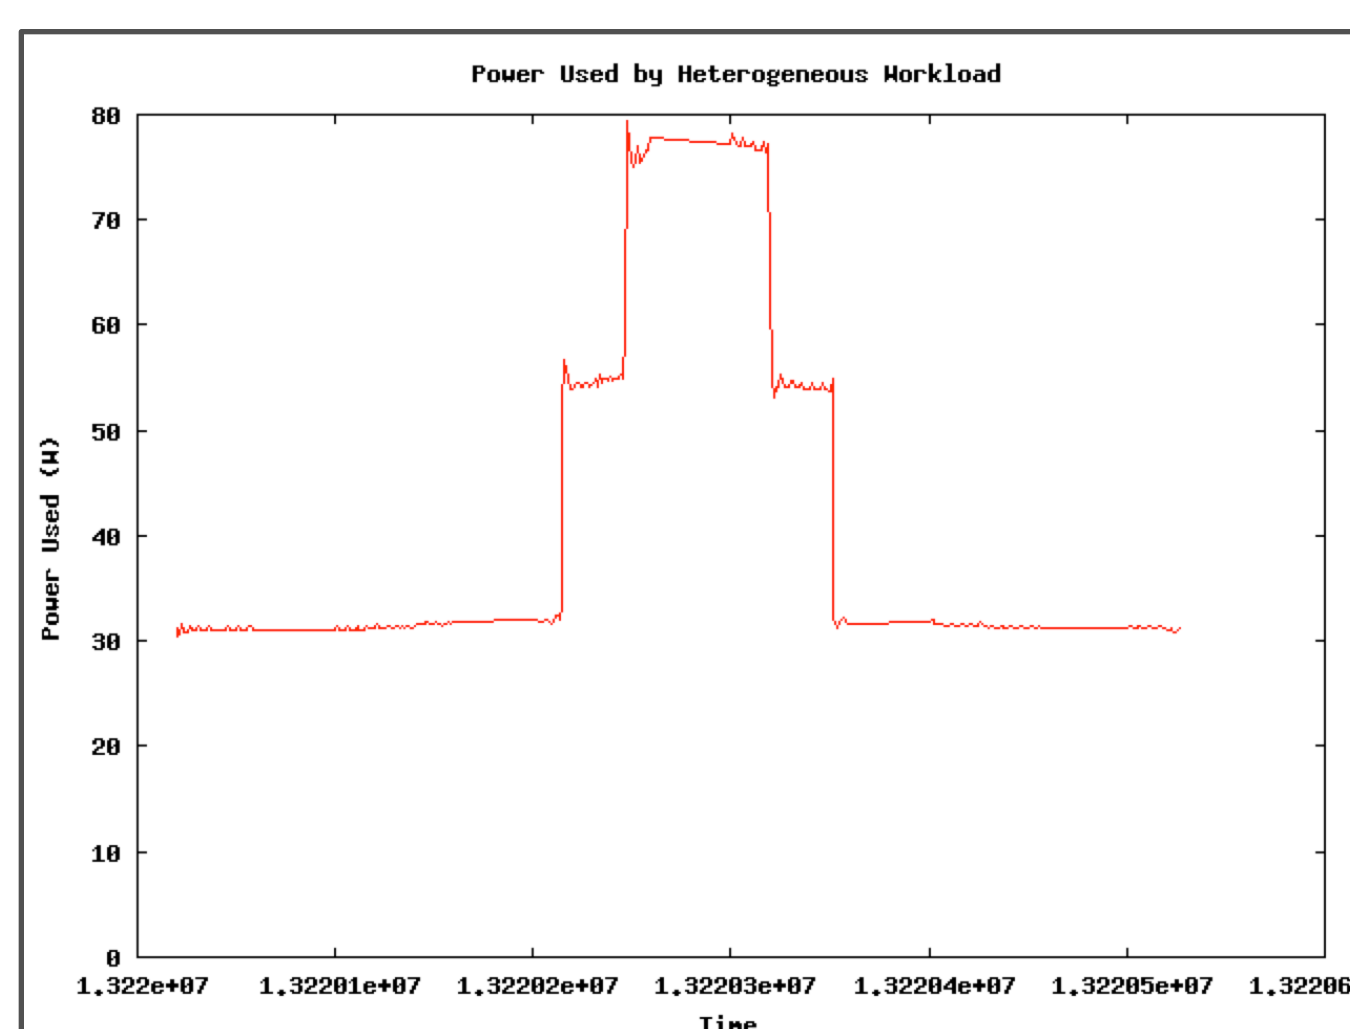
Response time of single Apache node



Offered load

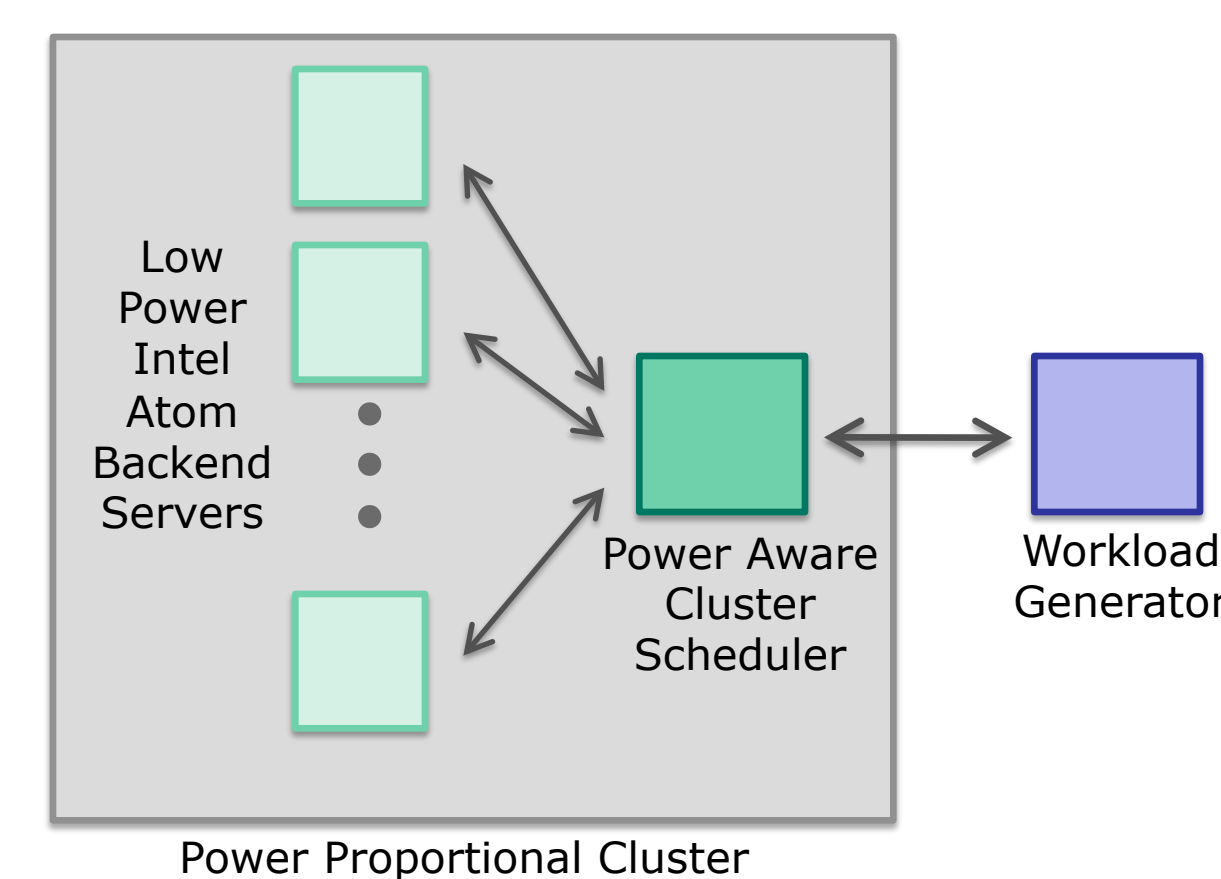


Power usage of simple scheduler



Power usage of energy-aware scheduler

Cluster Design



System architecture

	Idle		Peak		Sleep	
	no USB	with USB	no USB	with USB	no USB	with USB
Beagle *	1.65	3.35	--	3.4	--	--
Atom	22.08		24		1.6	

* without power management additions

Node power usage (W)

	to S3	from S3
Atom	2.6	2.4
Centrino	4.2	4.9
Nehalem	--	4.9

Node transition times (s)

$$\text{desired_servers} := \frac{(\text{incoming req. rate}) + (\text{max req. rate increase})(\text{server wakeup time})}{(\text{max req. rate per server})}$$

Determining the number of servers required to service the current load plus potential spikes

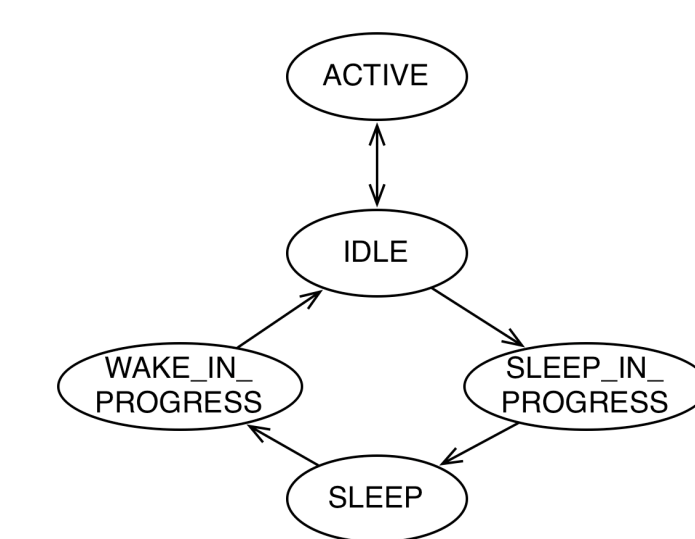
```

on_new_request()
{
  if (active_servers < desired_servers)
    wake_up_server();
}

on_end_request()
{
  if (server.requests == 0 &&
      active_servers > desired_servers)
    sleep_server_after_timeout();
}

```

Waking up or sleeping servers as needed as load levels change



Node states

```

update_max_req_rate(req)
{
  if (req.response_time < SLA)
    rate += A;
  else
    rate /= B;
}

```

Adjusting to changes to workload

Conclusions

- ❖ Energy-aware cluster scheduler
 - reduces power usage during periods of underutilization
 - causes no performance degradation
- ❖ Trade-off between aggressive power savings and responsiveness to load spikes