Fence: Protecting Device Availability With Uniform Resource Control

Tao Li†, Albert Rafetseder†, Rodrigo Fonseca*,
Justin Cappos†
†New York University *Brown University

Motivation

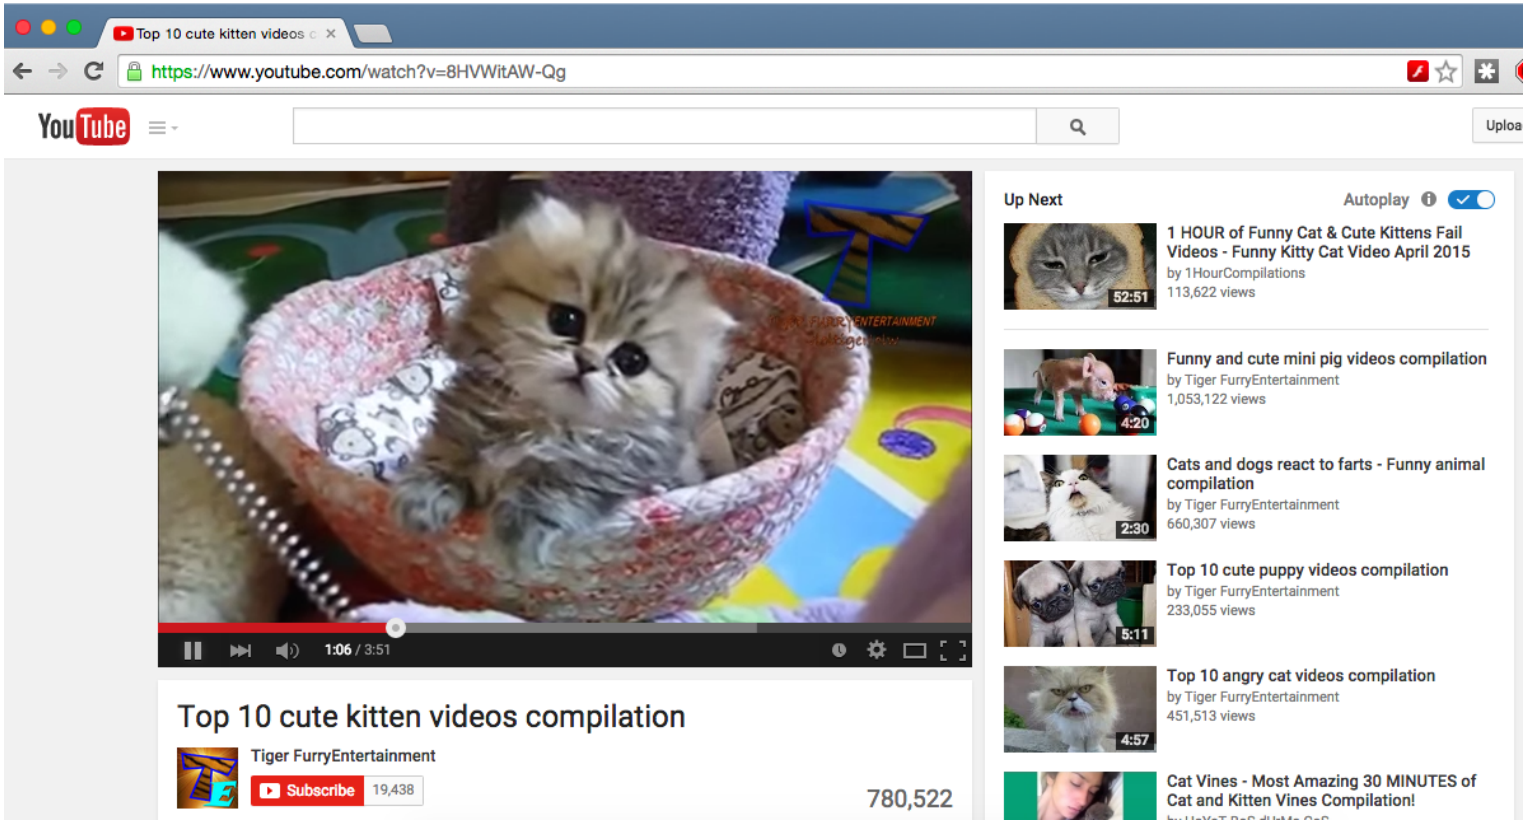
Top 10 cute kitten videos
https://www.youtube.com/watch?v=8HVWitAW-Qg
YouTube
Upload
1:06 / 3:51
Top 10 cute kitten videos compilation
Tiger FurryEntertainment
Subscribe 19,438
780,522
Up Next
Autoplay
1 HOUR of Funny Cat & Cute Kittens Fail Videos - Funny Kitty Cat Video April 2015
by 1HourCompilations
113,622 views
52:51
Funny and cute mini pig videos compilation
by Tiger FurryEntertainment
1,053,122 views
4:20
Cats and dogs react to farts - Funny animal compilation
by Tiger FurryEntertainment
660,307 views
2:30
Top 10 cute puppy videos compilation
by Tiger FurryEntertainment
233,055 views
5:11
Top 10 angry cat videos compilation
by Tiger FurryEntertainment
451,513 views
4:57
Cat Vines - Most Amazing 30 MINUTES of Cat and Kitten Vines Compilation!

Motivation

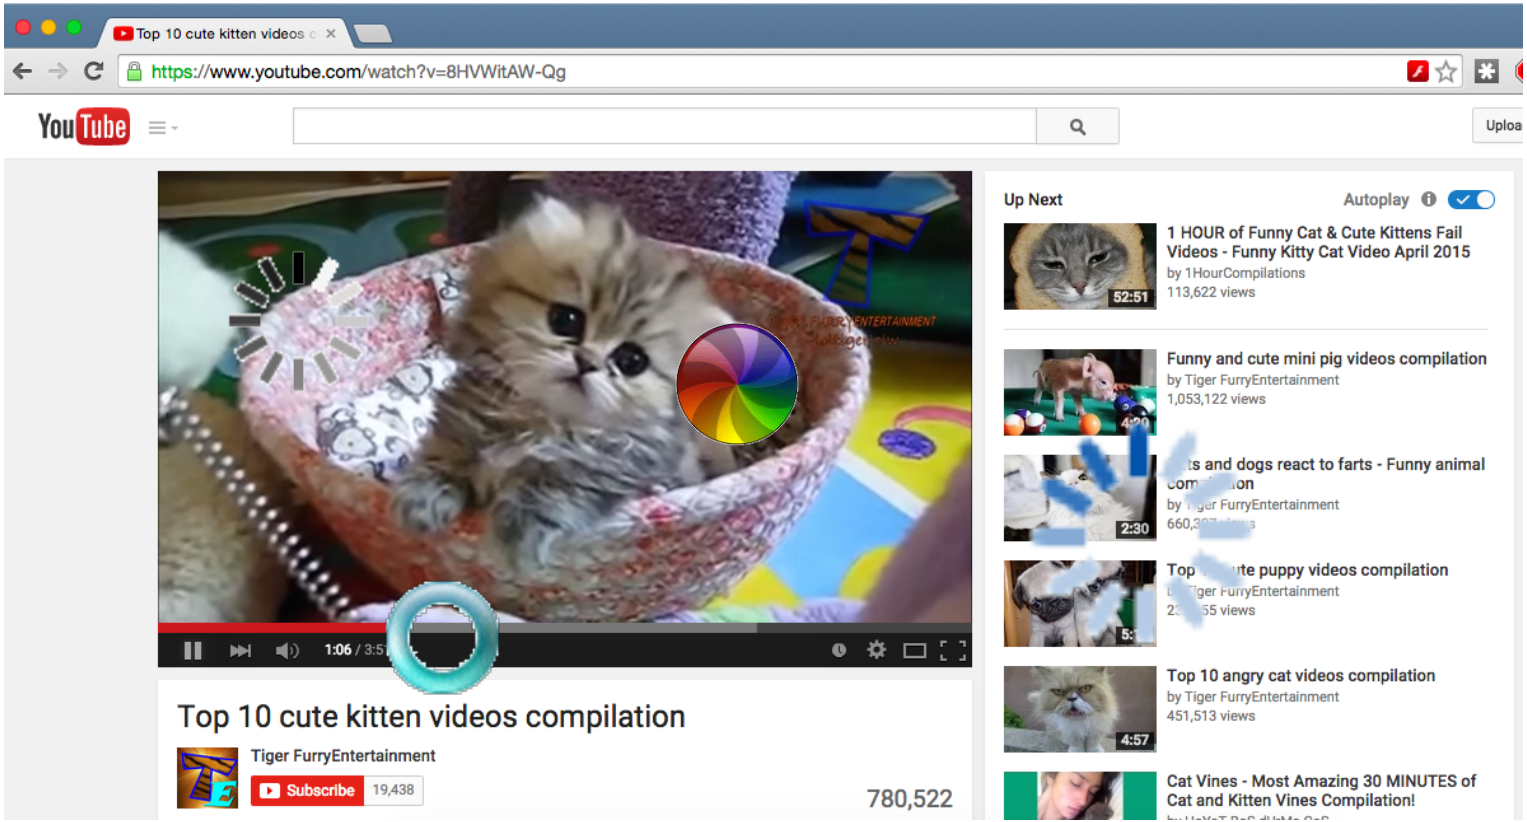
Top 10 cute kitten videos
https://www.youtube.com/watch?v=8HVWitAW-Qg
YouTube
Up Next
Autoplay
1 HOUR of Funny Cat & Cute Kittens Fail Videos - Funny Kitty Cat Video April 2015
by 1HourCompilations
113,622 views
52:51
Funny and cute mini pig videos compilation
by Tiger FurryEntertainment
1,053,122 views
2:30
Top 10 angry cat videos compilation
by Tiger FurryEntertainment
451,513 views
4:57
Cat Vines - Most Amazing 30 MINUTES of Cat and Kitten Vines Compilation!
1:06 / 3:5
Top 10 cute kitten videos compilation
Tiger FurryEntertainment
Subscribe
19,438
780,522

Typical Causes

- DropBox sync
- Browser tabs
- Virus scan
- Software updates
- …

Goal: Performance Isolation

- Control performance, battery, heat, etc.
  - Do not kill -- useful-but-gluttonous apps
- Do not require OS / hardware changes
  - Applicable to sandboxes, browsers, etc.
  - Run everywhere (Linux, Windows, Mac, Android, OpenWrt, etc.)
- Focus on mechanism
  - Necessary for policies to function

Why is performance isolation hard?

- Multiple contended resources
  - Separate mechanisms / policies
  - Creating an overarching policy is difficult
- Some controls are missing
  - Gaps in enforcement -> lack of isolation
- Legacy systems tend to be work-preserving

Key Idea: Uniform Resource Control

- Unifying resource abstraction
  - Two axes per resource: fungible/renewable
    - Fully defines mechanism
  - Easy to cover new resources
  - Easy to implement policies

Resource Abstraction Questions

- Fungible:
  - Are items of this type interchangeable?
    - Yes (disk space) vs No (TCP port)

- Renewable:
  - Are items replenished over time?
    - Yes (Network bandwidth) vs No (RAM)

Resource Controls

	Not Fungible	Fungible

Not Renewable	`is_item_allowed()`
	"Check if permitted"
	e.g. UDP port
		`tattle_add_item()`
		`tattle_remove_item()`
		"Restrict total used"
		e.g. File Descriptors

Renewable	`tattle_quantity()` "Rate limit" e.g. Network b/w	`tattle_quantity()` "Rate limit" e.g. CPU

Enforcement mechanism

- Polling
  - Find value, stop / rate limit if over
    - e.g. CPU uses job control interface (`SIGSTOP` / `SIGCONT`)
- Interposition
  - API code changes to add interposition
- Which depends on implementation

Example Implementation Changes

```
def sendmessage(destip,destport,msg,localip,localport):  # 117 lines
  ...


  # get the OS's UDP socket
  sock = _get_udp_socket(localip, localport)


  # Send this UDP datagram
  bytessent = sock.sendto(msg, (destip, destport))




  ...
```

Example Implementation Changes

```
def sendmessage(destip,destport,msg,localip,localport):  # 117 lines + 10 lines
  ...
  # check that we are permitted to use this port...
  if not fence.is_item_allowed('UDPport',localport):
    raise ResourceAccessDenied('...')
  # get the OS's UDP socket
  sock = _get_udp_socket(localip, localport)


  # Send this UDP datagram
  bytessent = sock.sendto(msg, (destip, destport))




  ...
```

UDP port: Non-fungible, non-renewable

Example Implementation Changes

```
def sendmessage(destip,destport,msg,localip,localport):  # 117 lines + 10 lines
  ...
  # check that we are permitted to use this port...
  if not fence.is_item_allowed('UDPport',localport):
    raise ResourceAccessDenied('...')
  # get the OS's UDP socket
  sock = _get_udp_socket(localip, localport)
  # Register this socket descriptor with fence
  fence.tattle_add_item('outsocketsopened', id(sock))
  # Send this UDP datagram
  bytessent = sock.sendto(msg, (destip, destport))

  ...
```

socket: Fungible, non-renewable

Example Implementation Changes

```
def sendmessage(destip,destport,msg,localip,localport):  # 117 lines + 10 lines
  ...
  # check that we are permitted to use this port...
  if not fence.is_item_allowed('UDPport',localport):
    raise ResourceAccessDenied('...')
  # get the OS's UDP socket
  sock = _get_udp_socket(localip, localport)
  # Register this socket descriptor with fence
  fence.tattle_add_item('outsocketsopened', id(sock))
  # Send this UDP datagram
  bytessent = sock.sendto(msg, (destip, destport))
  # Account for the network bandwidth utilized
  if _is_loopback_ipaddr(destip):
    fence.tattle_quantity('loopbacksend', bytessent + 64)
  else:
    fence.tattle_quantity('internetsend', bytessent + 64)
  ...
```

Network b/w: Fungible, Renewable

Uses of Fence

- Seattle Testbed's Repy sandbox
  - Seattle ~= Peer-to-peer PlanetLab
  - Tens of thousands of diverse devices
- Lind
  - NaCl / POSIX sandbox
- Sensibility Testbed
  - Privacy preserving sensing on Android

Limitations

- Resource consumption must be visible
  - HW / OS hide info
- Minimizes performance impact
  - “Worst case” limits
- Scope of policies
  - Unclear how complete Fence is
  - Worked for us in practice

Evaluation

- How well does Fence work vs legacy controls?
- How well does Fence work across platforms?
- How much overhead does Fence incur?
- Can realistic policies be expressed in Fence?
- How diverse of resources can be metered?
- How hard is it to add resources to Fence?

Evaluation

- How well does Fence work vs legacy controls?
- How well does Fence work across platforms?
- How much overhead does Fence incur?
- Can realistic policies be expressed in Fence?
- How diverse of resources can be metered?
- How hard is it to add resources to Fence?

Fence vs Legacy Controls

- Video on disk (Dell Inspirion 630m w/ Ubuntu 10.04)
- “hog” everything
- worst setting for hog
- best setting for video

https://www.youtube.com/watch?v=OGe8QpPbtz4

Fence vs Legacy Heat / Battery

- Heat / battery
- “hog” everything

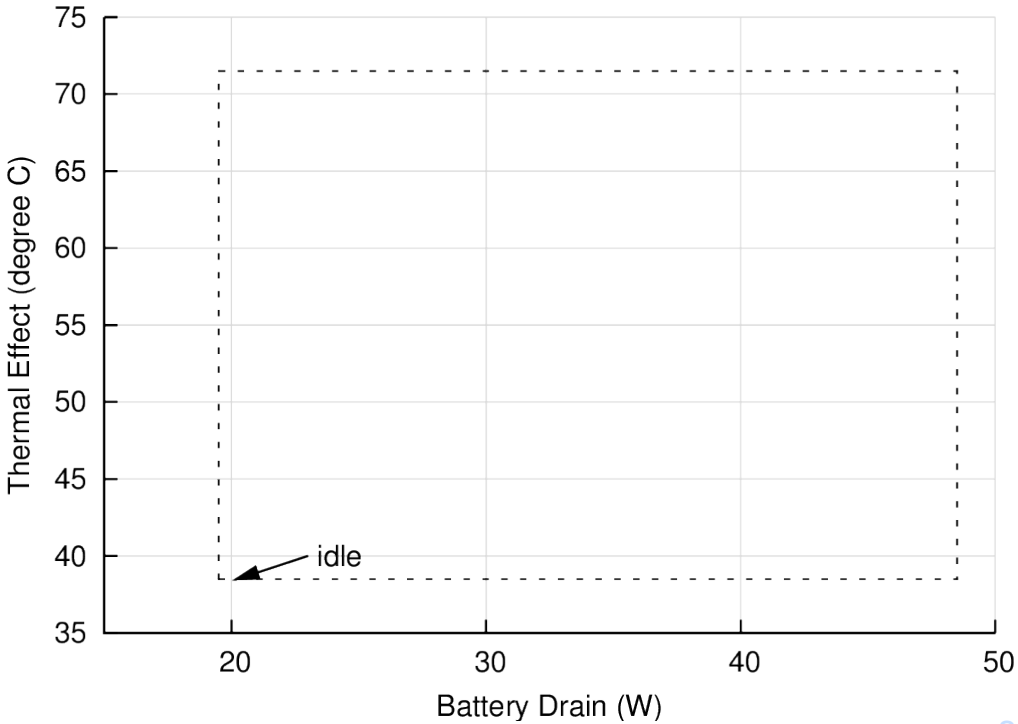

Fence vs Legacy Heat / Battery

- Heat / battery
- “hog” everything

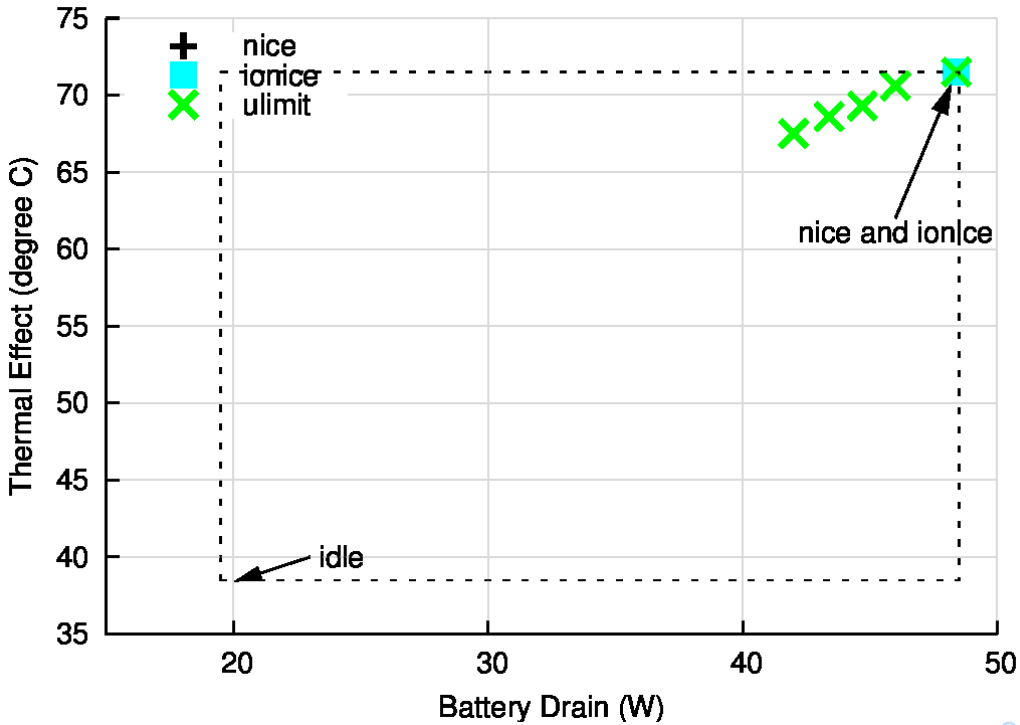

Fence vs Legacy Heat / Battery

- Heat / battery
- “hog” everything

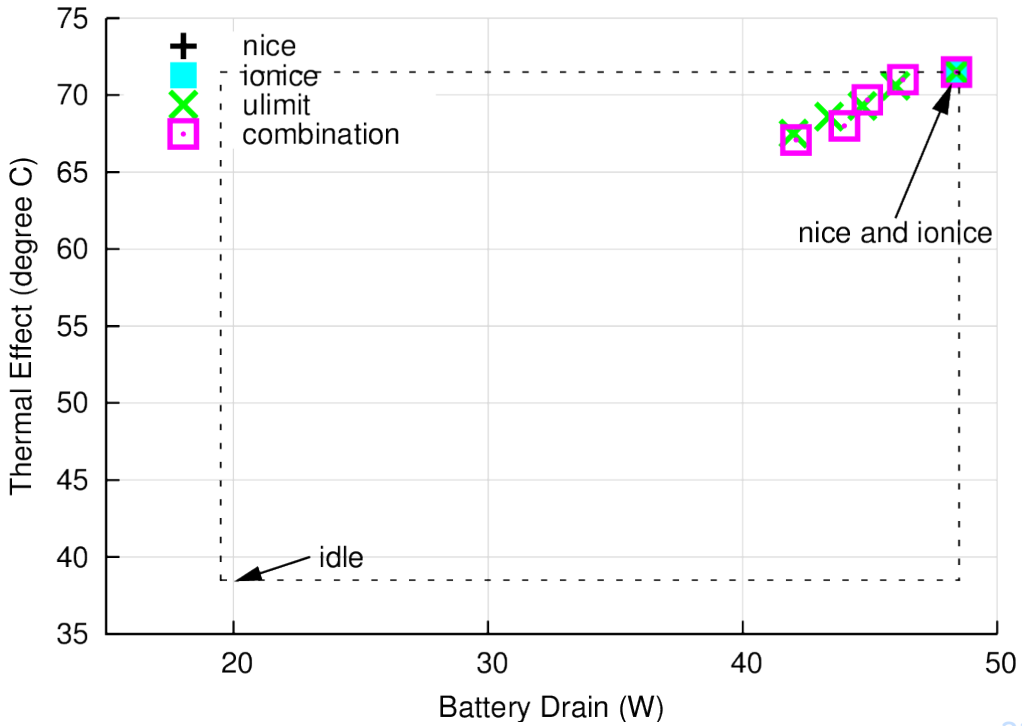

Fence vs Legacy Heat / Battery

- Heat / battery
- “hog” everything

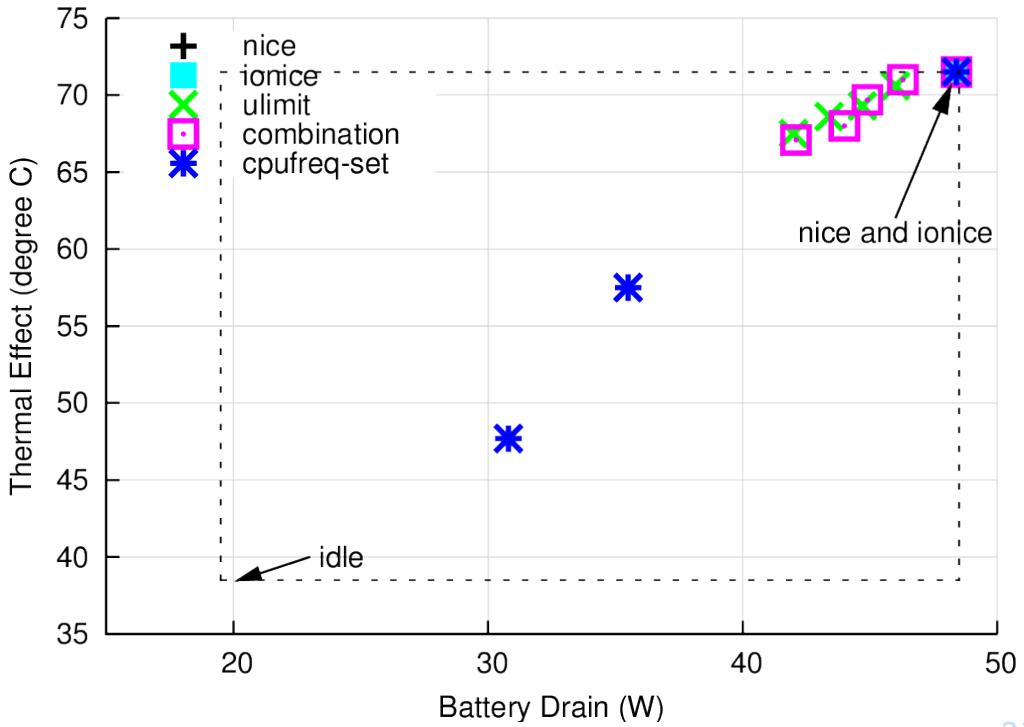

Fence vs Legacy Heat / Battery

- Heat / battery
- “hog” everything

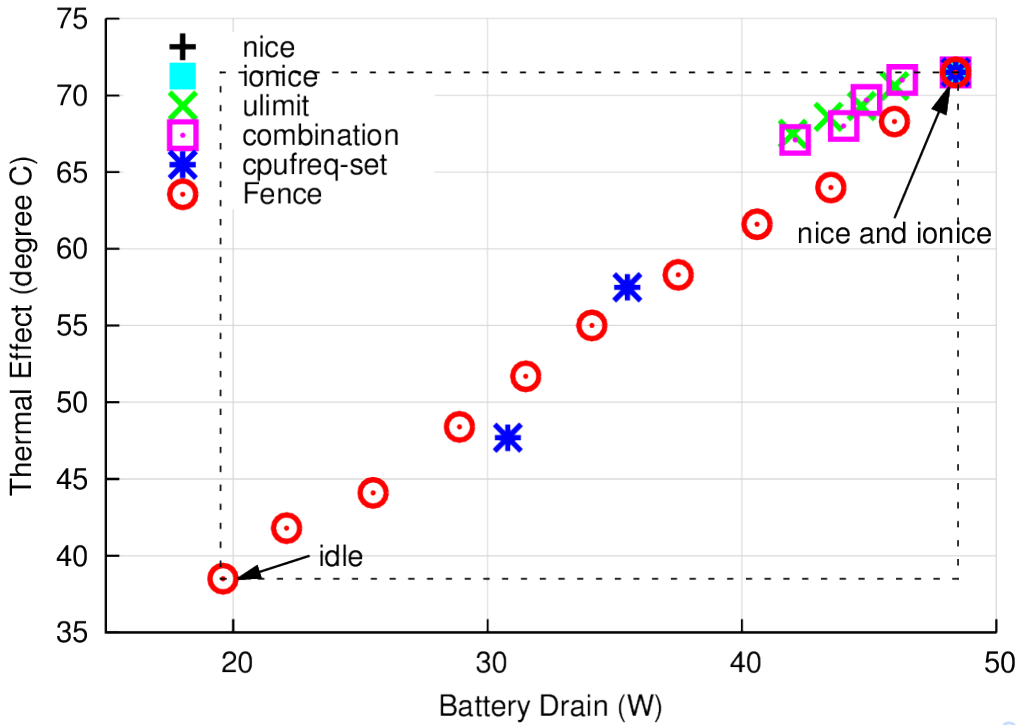

Evaluation

- ~~How well does Fence work vs legacy controls?~~
- How well does Fence work across platforms?
- How much overhead does Fence incur?
- Can realistic policies be expressed in Fence?
- How diverse of resources can be metered?
- How hard is it to add resources to Fence?

Expressing Policies: Cinder

Power draw policy from Cinder [Roy Eurosys 2011]

- Stores energy w/ a tap (token bucket)
- Polling using ACPI (updates every 15 seconds)

Program: Richards benchmark in a run / sleep cycle

Expressing Policies: Cinder (cont)

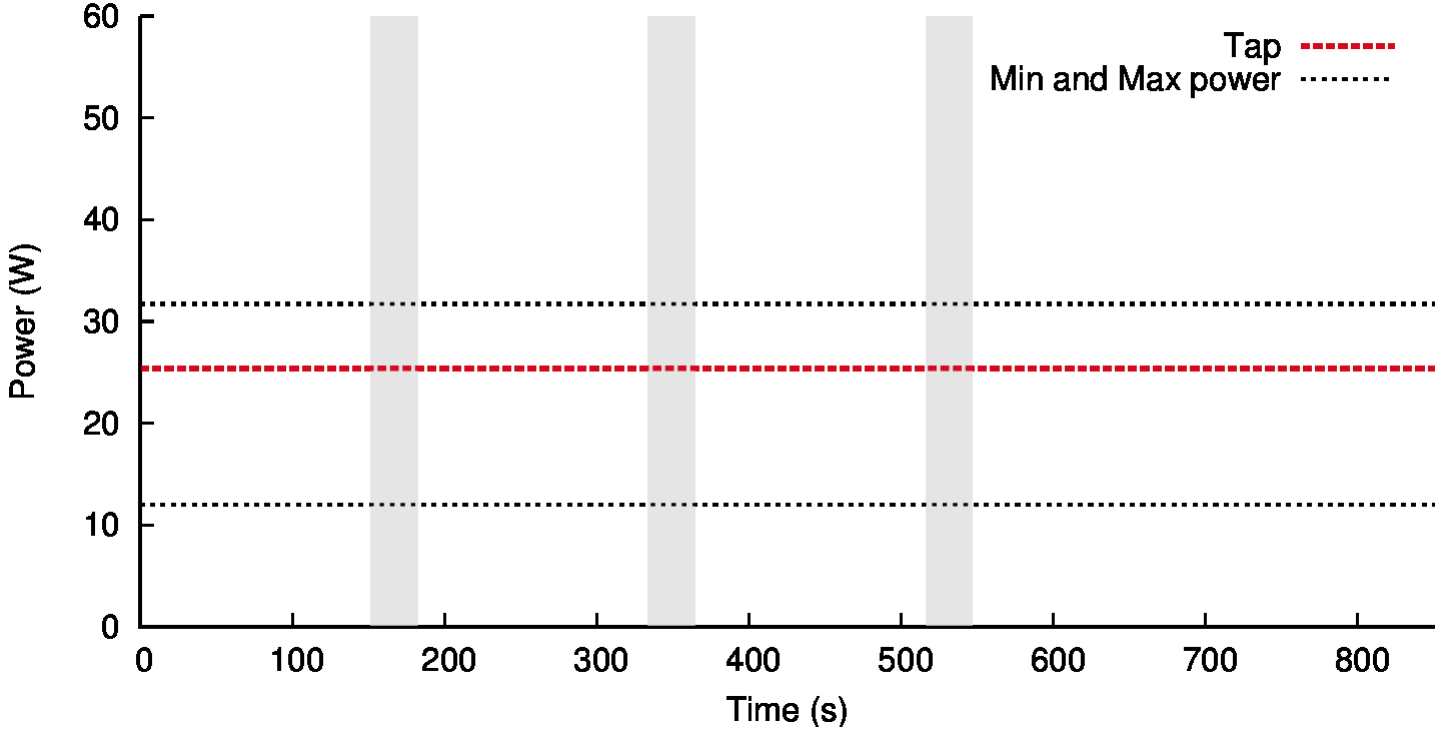

Expressing Policies: Cinder (cont)

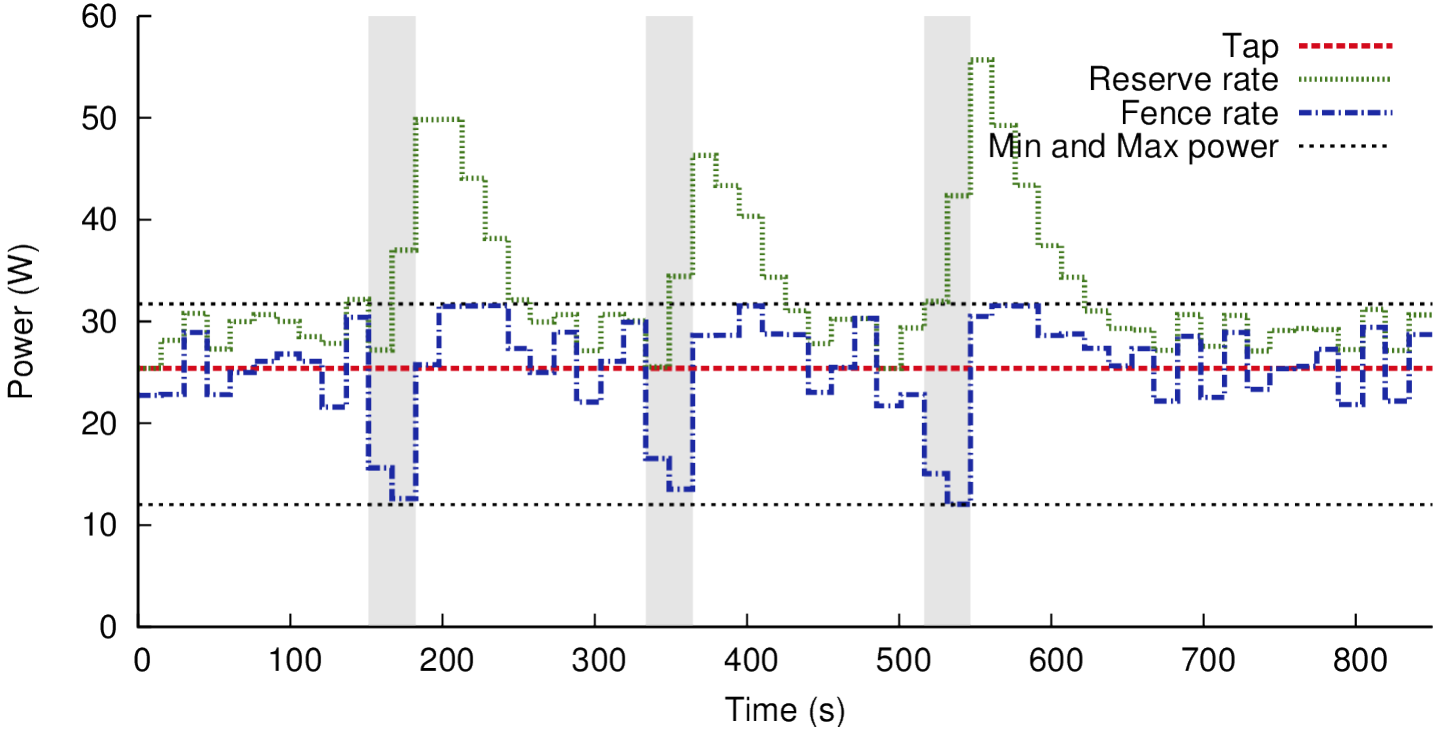

Expressing Policies: Cinder (cont)

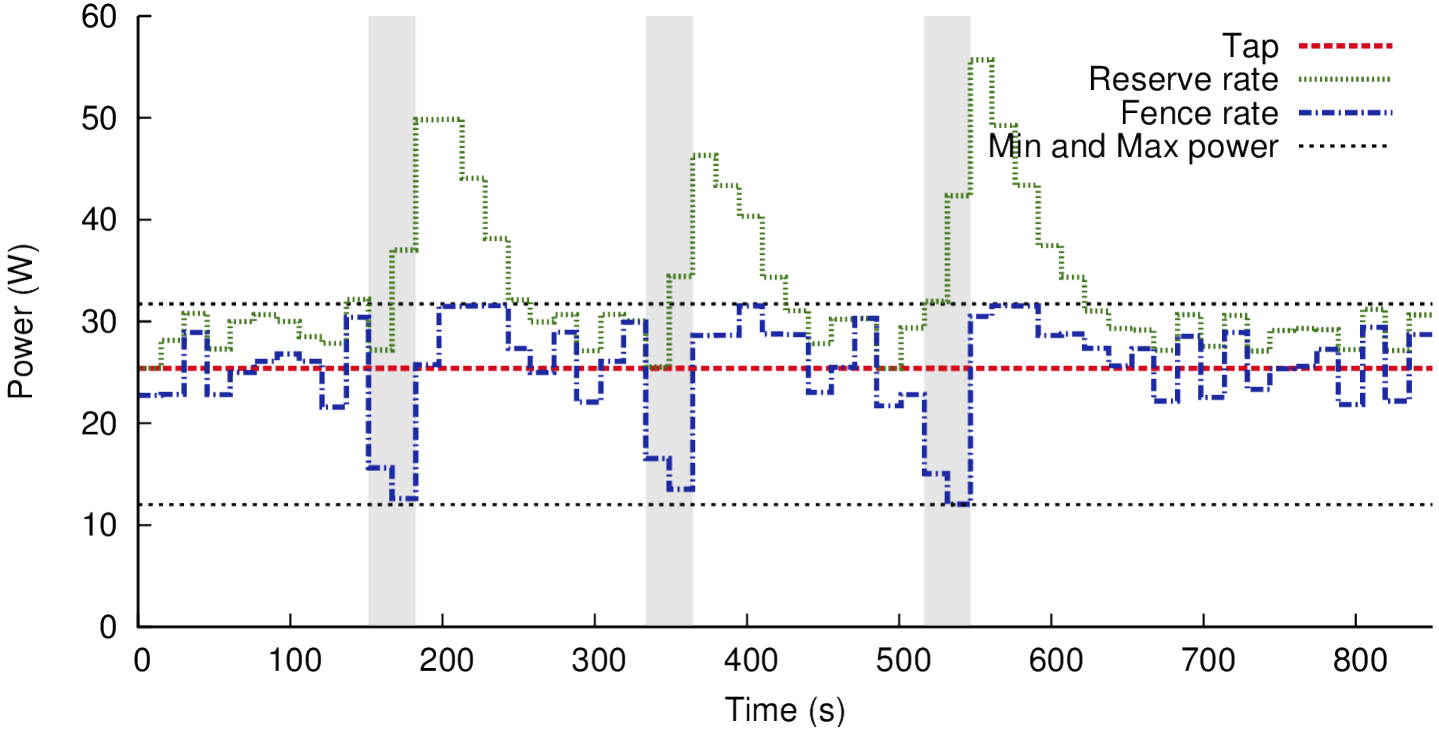

150 LOC!

Conclusion

- Performance isolation is still a challenge

- Uniform Resource Control
  - Same simple reasoning for all resources
  - Fungible / Renewable
  - Easy to implement / use
  - Effective in practice

NYU is Hiring!

- Secure software distribution
  - Adoption by Python, Ruby, Docker, LEAP, CoreOS, Go, Rust, Haskell, OCaml, etc.
  - Plausible standard for many new domains
- Hiring Post Doc / Research Professor / Dev

Questions?

Fence vs Legacy Controls (cont)

<20% frames for each legacy tool
(Combining tools, only gives 22% of frames)

99% of frames for Fence

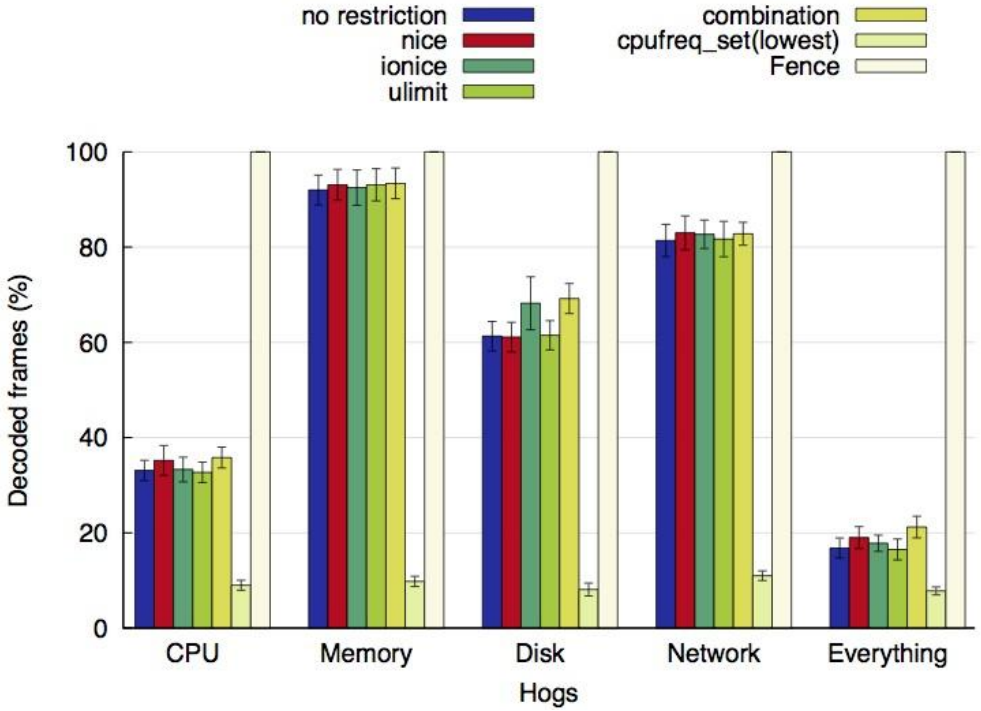

Example Resource Categorization

	Not Fungible	Fungible
Not Renewable	UDP ports TCP ports	Threads Memory (RAM) Storage Space Open Sockets Open Files
Renewable	Network read / write	CPU File read / write HW random