



DASU

Pushing Experiments to the Internet's Edge

Mario A. Sánchez

John S. Otto

Zachary S. Bischof

Fabián E. Bustamante

Northwestern U.

David R. Choffnes

U. of Washington

Balachander Krishnamurthy

Walter Willinger

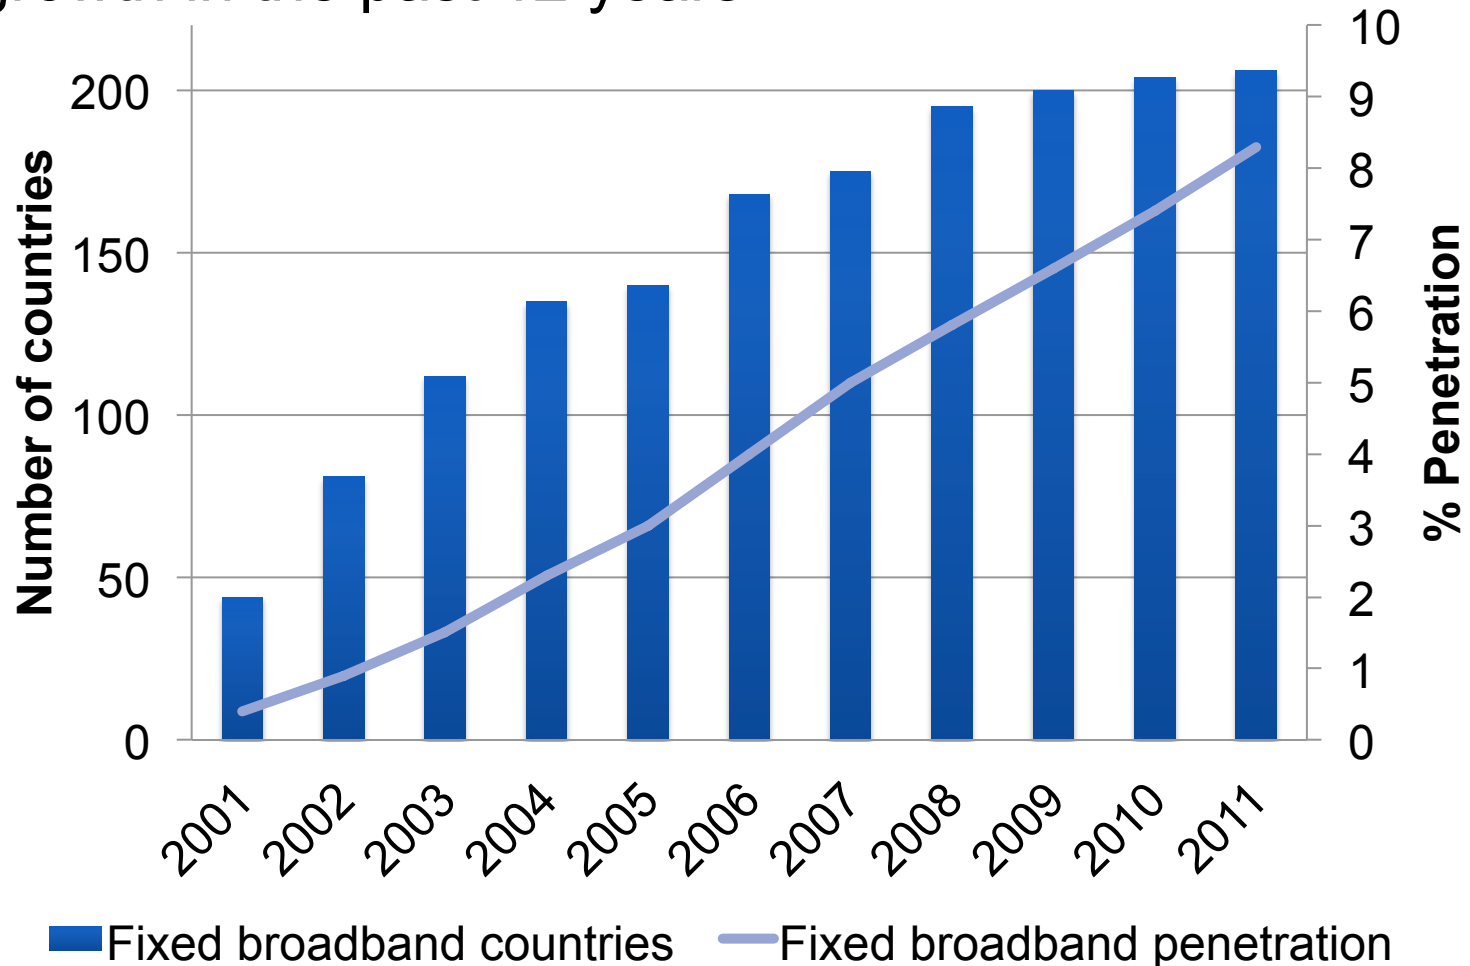
AT&T Labs-Research



The Internet edge is exploding

34% of global population uses the Internet

5x growth in the past 12 years




The view from the edge

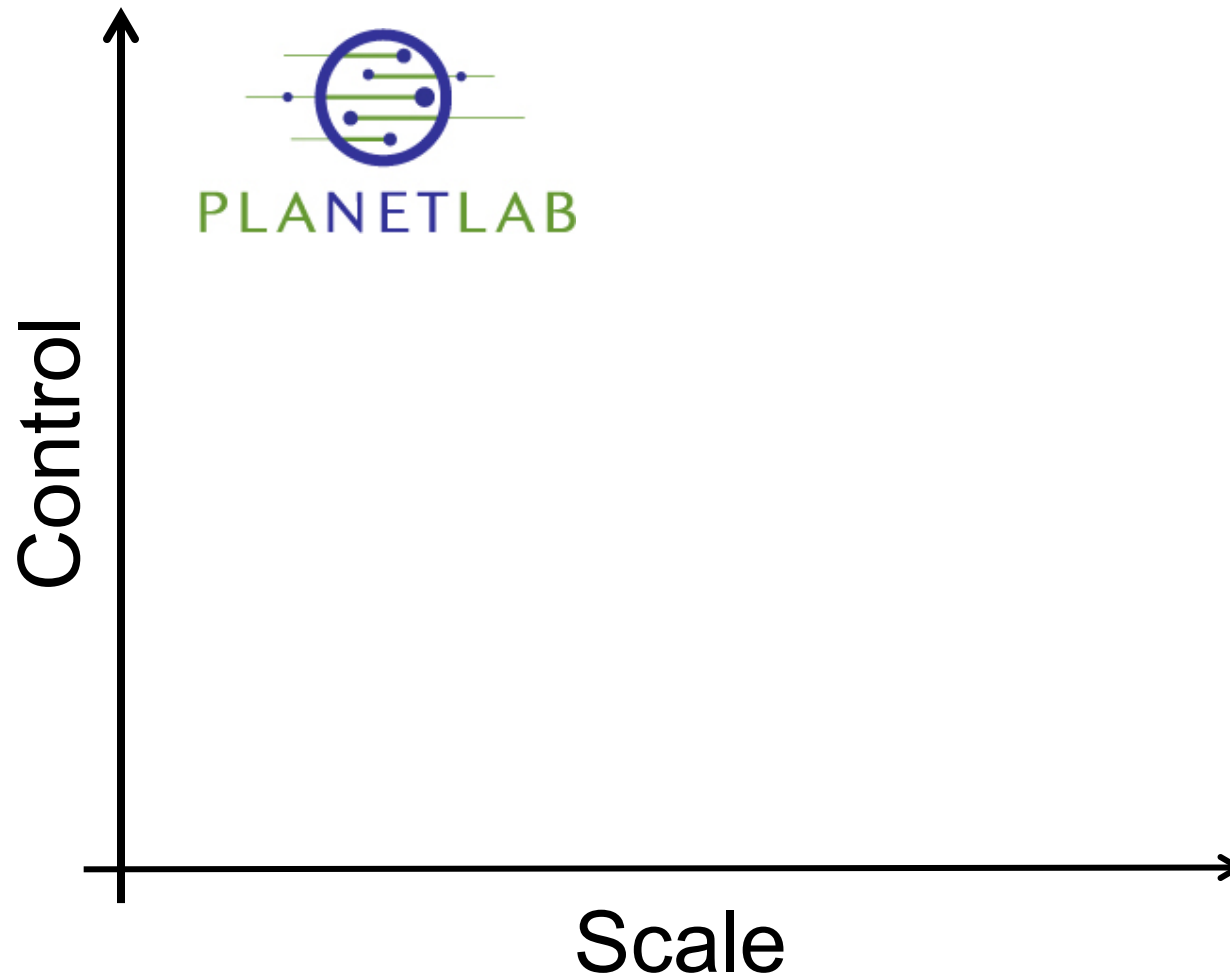
Many of our systems run at the edge

Lack platforms to experiment at or even characterize the edge


The view from the edge



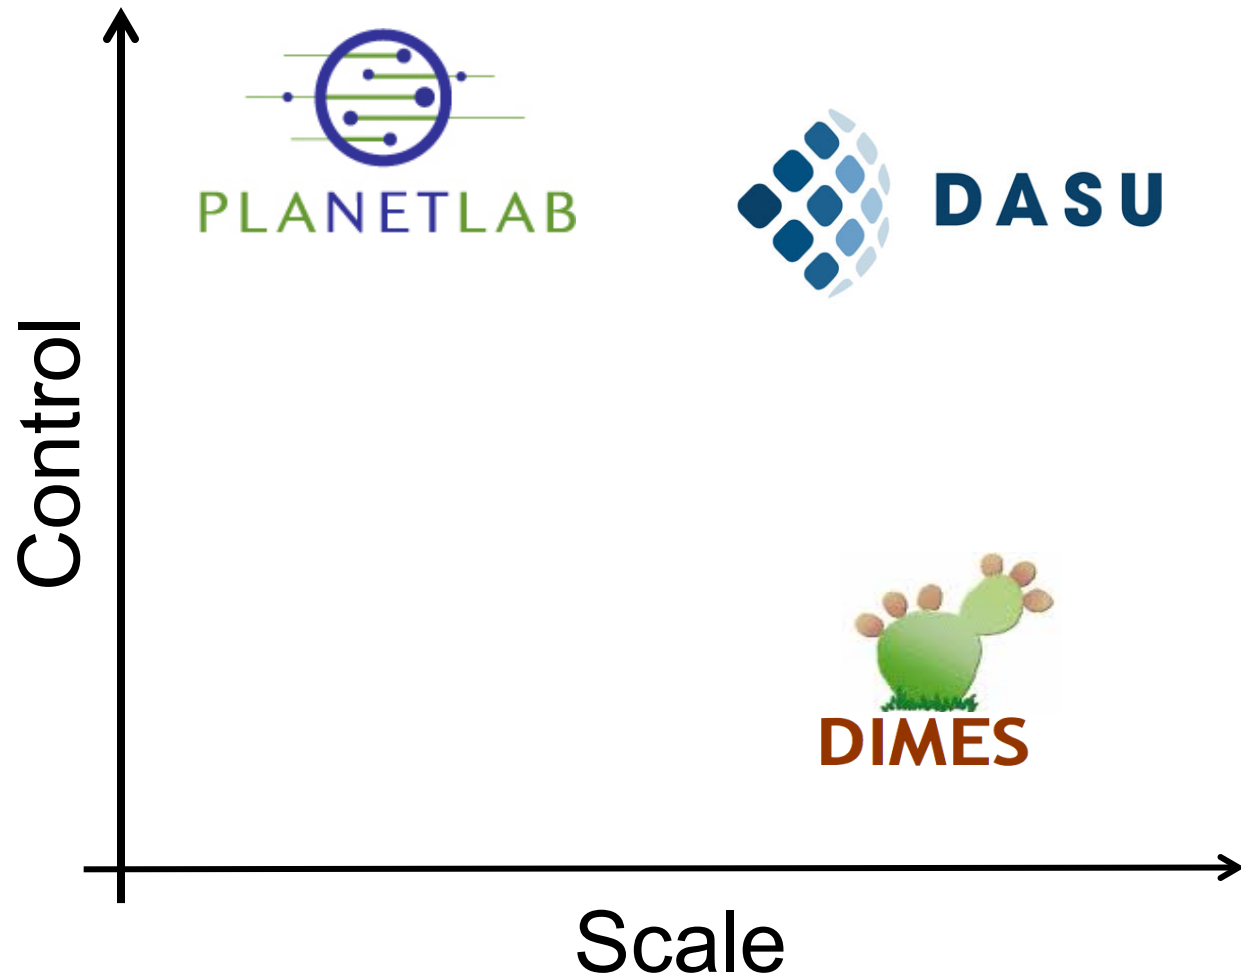
The view from the edge



The view from the edge



The view from the edge



The view from the edge



A flexible experimentation platform
hosted by end users

Our work

Hosted by end users, at scale ...

Guarantees safety of participants

Controls impact of experiments on underlying network and systems

Shares platform resources among concurrent experiments

Dasu

A platform prototype


Software-based (BitTorrent ext., for now)

Hosted at the network's edge

Broadband characterization as incentive for adoption

Dasu

Two purposes by design




- Characterize users' broadband services
- Support experimentation from the edge

Aligns objectives of end-users and experimenters

Dasu

Two purposes by design



- Characterize users' broadband services
- Support experimentation from the edge

Aligns objectives of end-users and experimenters

	End-user	Experimenter
Coverage	✓	✓
Availability	✓	✓
At the edge	✓	✓
Extensibility	✓	✓

End-user based platform – Challenges

No dedicated resources

- We cannot run arbitrary experiments
- Safety of volunteer nodes is key

Large-scale platform → large-scale *impact*

Highly volatile nodes

- We cannot run at arbitrary times
- We cannot run at arbitrary locations

Roadmap

Motivation

Design and implementation

Challenges

A unique perspective

Dasu: Current status

An extension to BitTorrent

- Soon standalone and in a DNS resolver

Could be hosted by other long-running networked app. (e.g. Skype, IPTV)

Over 90,000 users (>150 countries)
with almost no advertisement

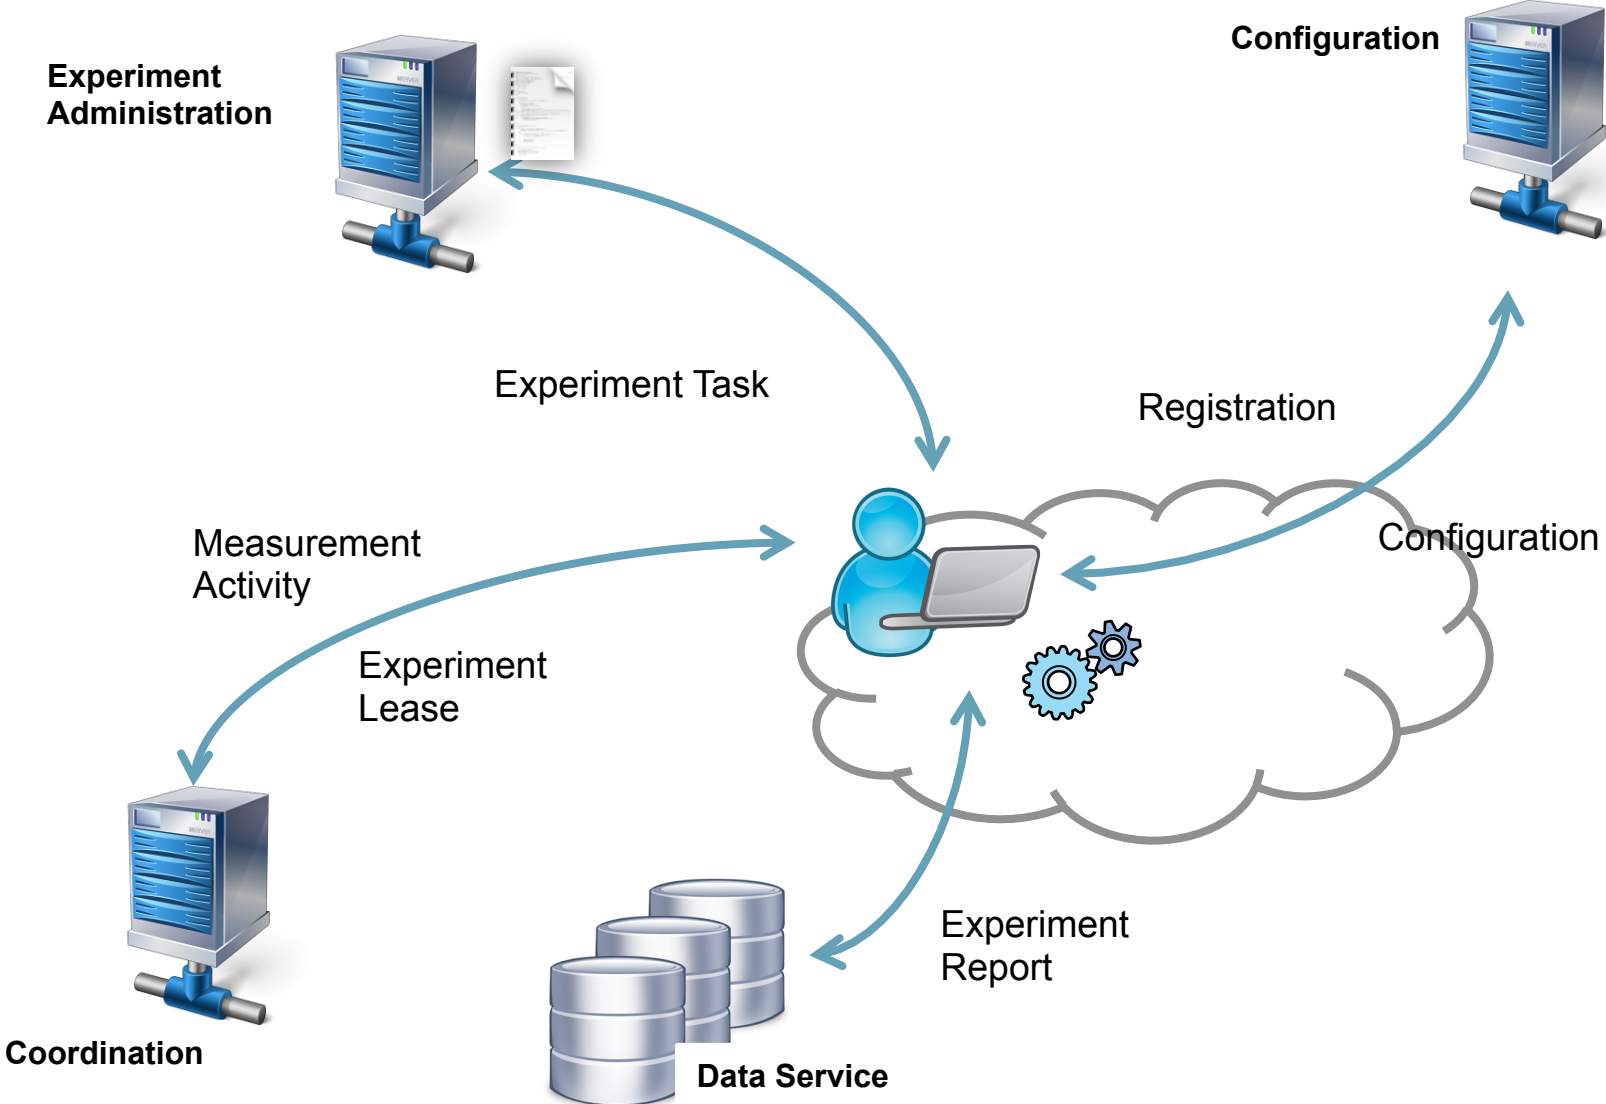
Guiding goals...

Programmable platform

- To run general measurement experiments

Reduce effort needed to develop/deploy large-scale measurements

Design: System components



Design: Experiment specification

package package-name

Imports

variable declaration

- global

- local

functions

rules

```
dialect "java"
package edu.northwestern.dasu.drools

import edu.northwestern.dasu.drools.*;

global FactHalt halt;

declare TestResult
    taskId : String
    todoList : List
    pingList : List
end

function void ping(String ip, int taskId)
{

    AMCoordinator.addProbeTask(
        taskId, ip, ProbeType.PING);

}
```

Design: Experiment specification

package package-name

Imports

variable declaration

- **global**

- **local**

functions

rules

```
dialect "java"  
package edu.northwestern.dasu.drools  
  
import edu.northwestern.dasu.drools.*;  
  
global FactHalt halt;
```

```
declare TestResult  
    taskId : String  
    todoList : List  
    pingList : List  
end
```

```
function void ping(String ip, int taskId)  
{  
  
    AMCoordinator.addProbeTask(  
        taskId, ip, ProbeType.PING);  
  
}
```

Design: Experiment specification

package package-name

Imports

variable declaration

- global

- local

functions

rules

```
dialect "java"
package edu.northwestern.dasu.drools

import edu.northwestern.dasu.drools.*;

global FactHalt halt;

declare TestResult
    taskId : String
    todoList : List
    pingList : List
end

function void ping(String ip, int taskId)
{
    AMCoordinator.addProbeTask(
        taskId, ip, ProbeType.PING);
}
```

Design: Experiment specification

package package-name

Imports

variable declaration

- global

- local

functions

rule "name"

when

...

then

...

end

```
rule "Process Probes"
```

```
when
```

```
  $tr : TestResult();
```

```
  $ping : FactProbePingResult();
```

```
then
```

```
  $tr.getPingList().add($ping);
```

```
  retract($ping);
```

```
  if($tr.getTodoList().size() == 0) {
```

```
    AMCoordinator.commitResult(
```

```
      "Experiment Name", $tr);
```

```
    halt.setHalt(true);
```

```
  } else {
```

```
    ping($tr.getTodoList().remove(0),
```

```
      $tr.getTaskId());
```

```
  }
```

```
end
```

Design: Rule-based programming

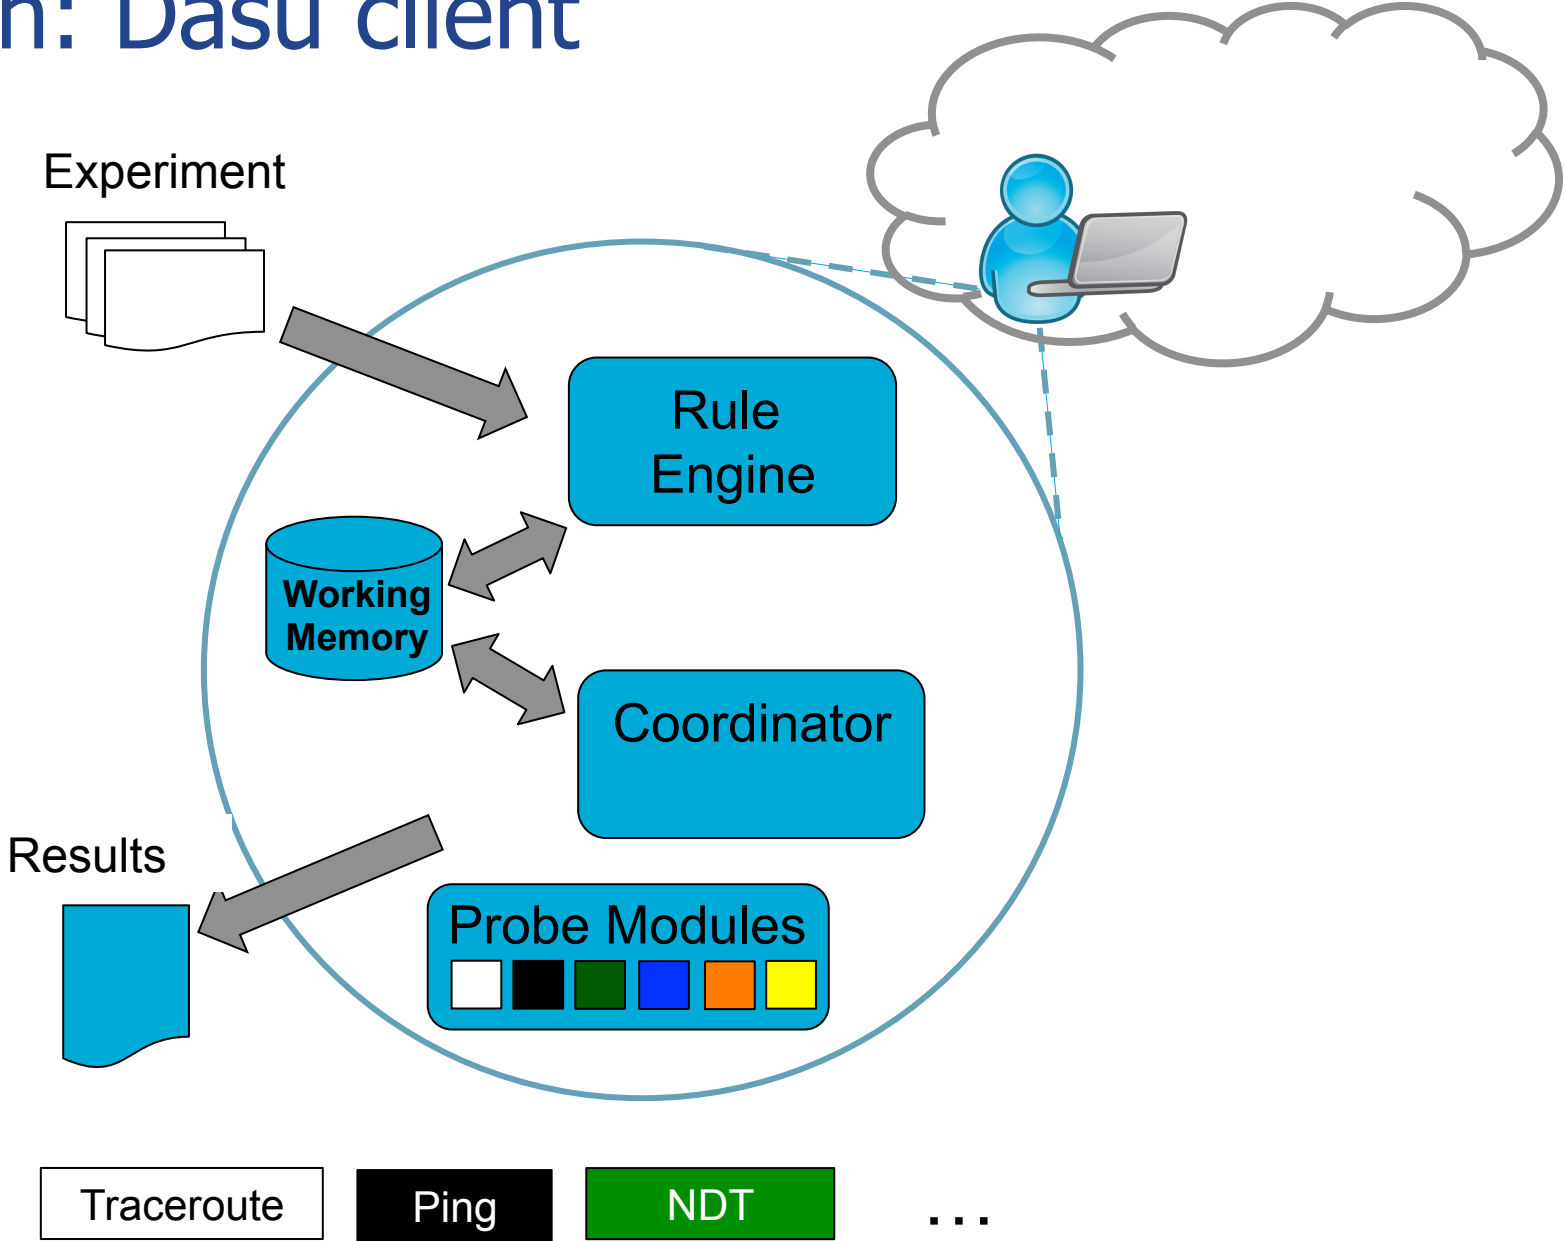
Experiment as a combination of rules

Rules drive execution and flow of experiments

Benefits of a declarative model

- Clear and concise
- Easy to check
- Easy to extend

Design: Dasu client



Challenges: No dedicated resources

Protecting volunteer nodes

- Sandboxed
 - Import restrictions avoid arbitrary code execution
 - Resource profiler / watchdog timer
- Secure communication of configuration and experiments

Challenges: No dedicated resources

Limit resource consumption

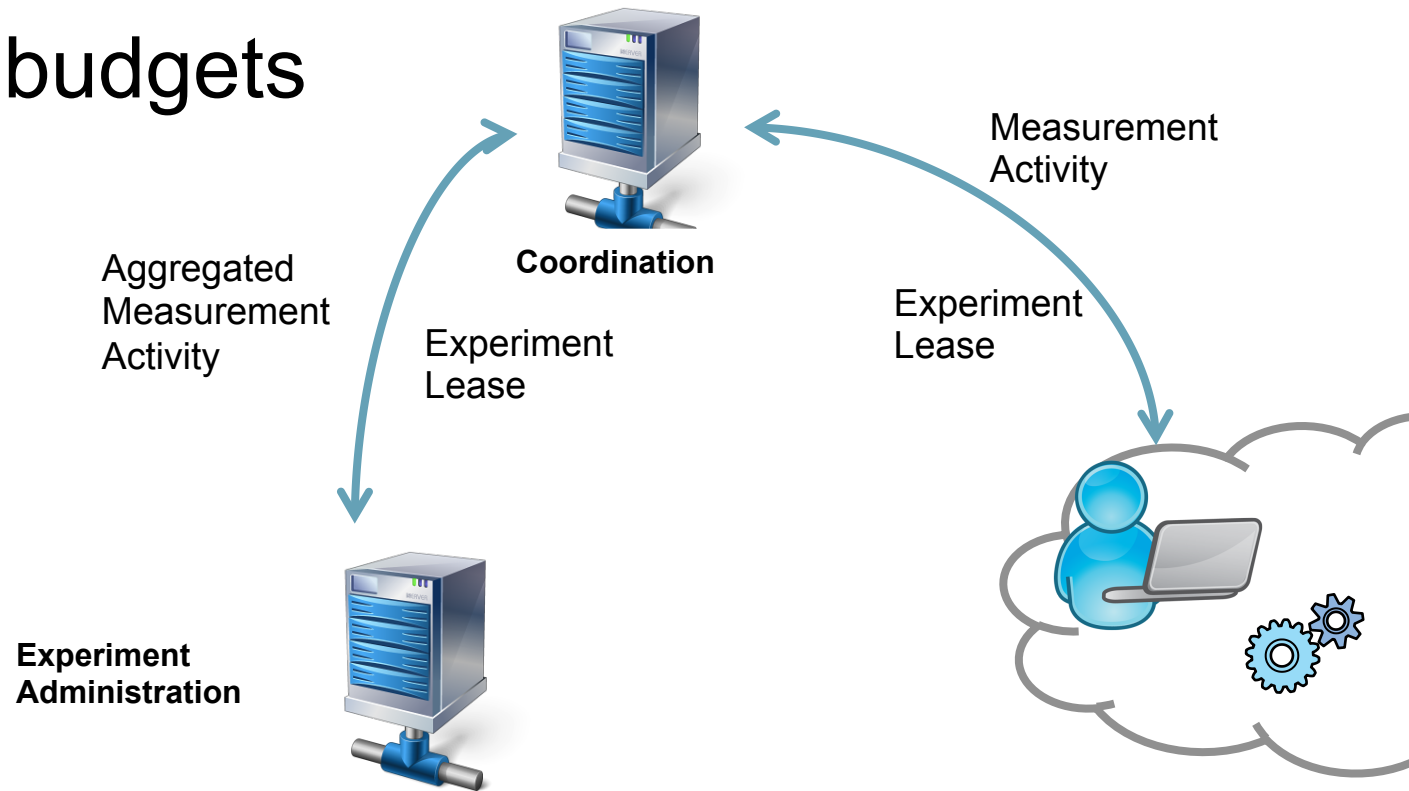
Minimize impact on user's performance:

- Pre-defined probe rates
- Restrict aggregate bandwidth consumption
- Limit probes to low-utilization periods


Challenges: Impact

Scaling measurement coordination across participants


- Experiment leases
- Elastic budgets



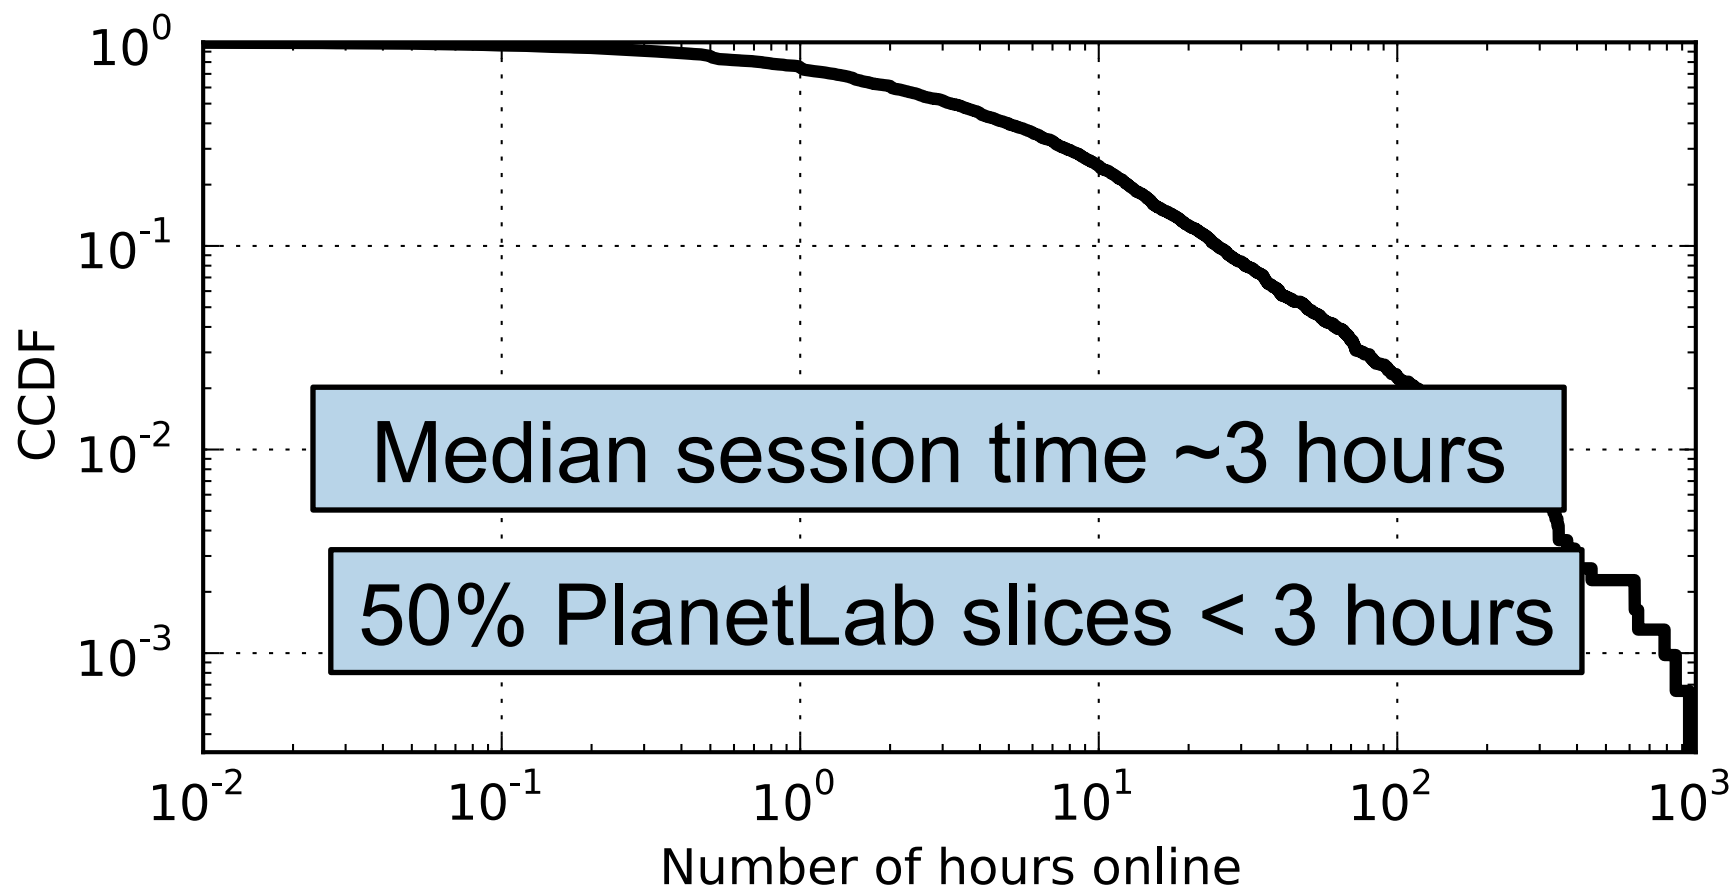
Challenges: Availability and churn



Challenges: Availability and churn




Challenges: Availability and churn



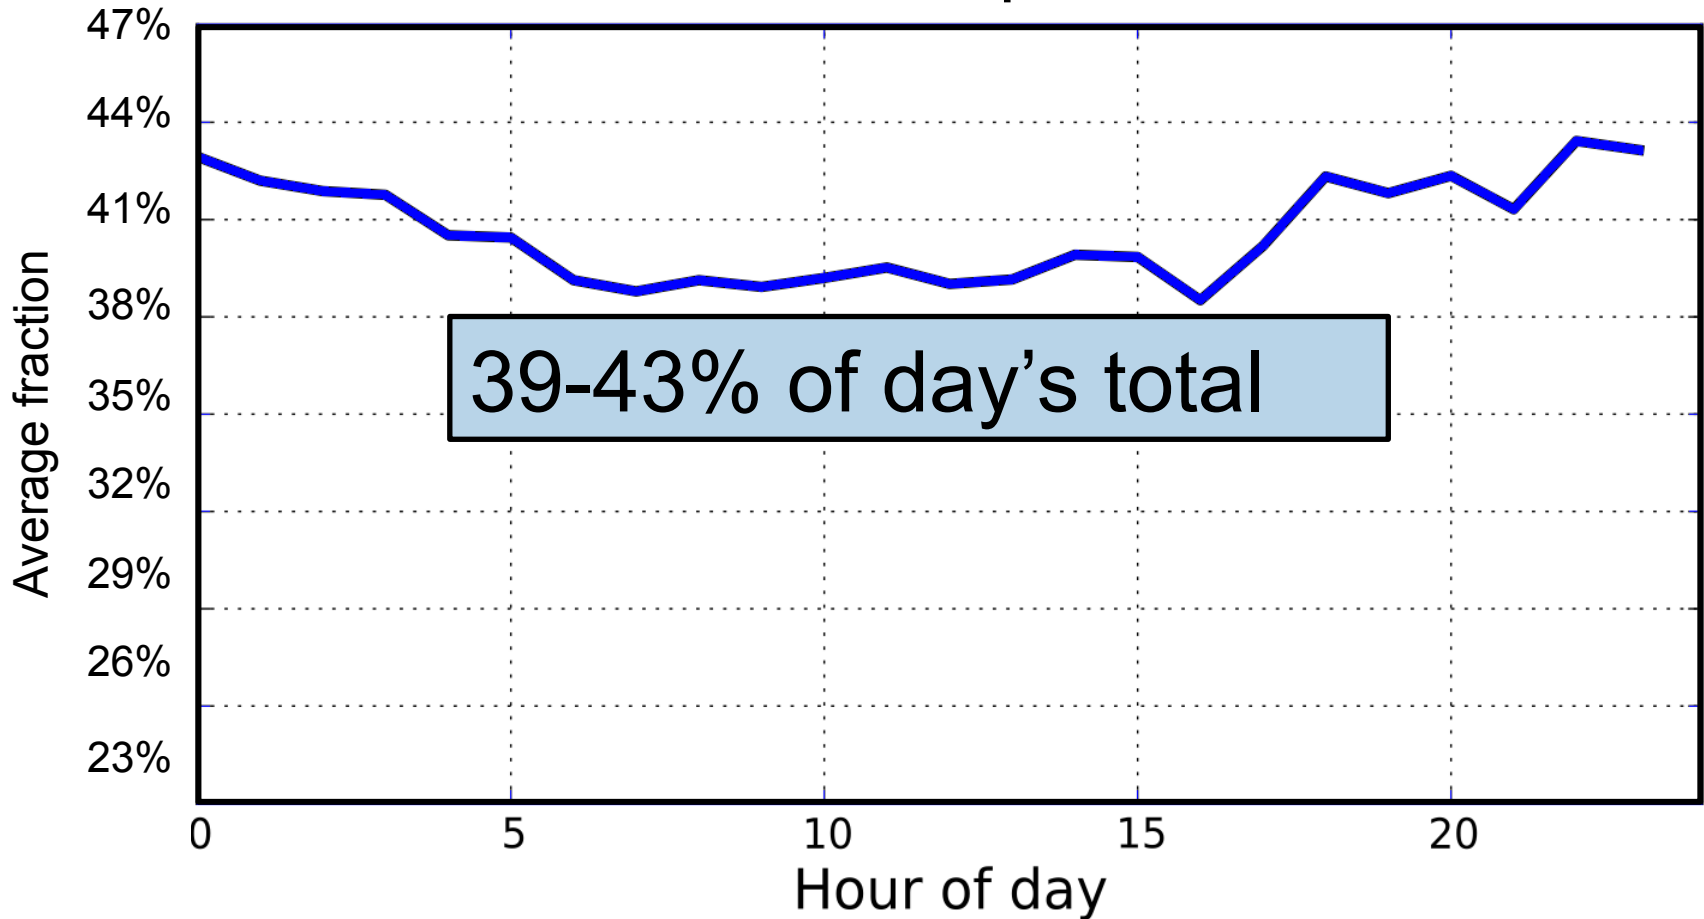
Challenges: Availability and churn

Peers online per hour




Challenges: Availability and churn

Peers online per hour



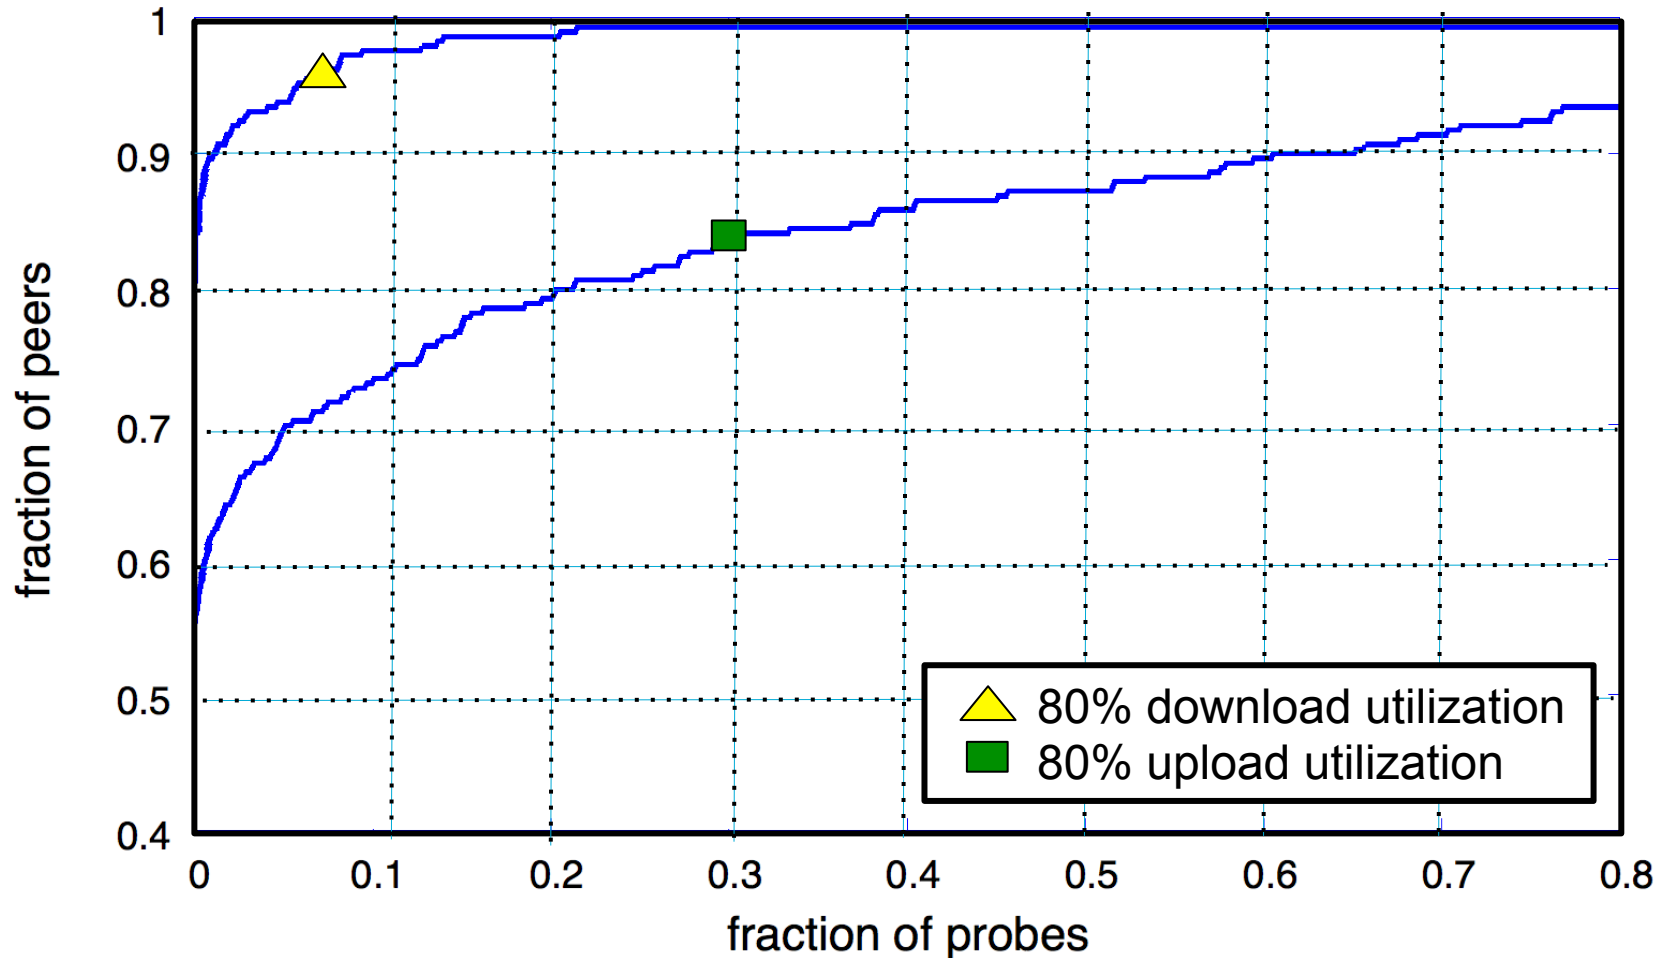
Challenges: Availability and churn

Peers online per hour



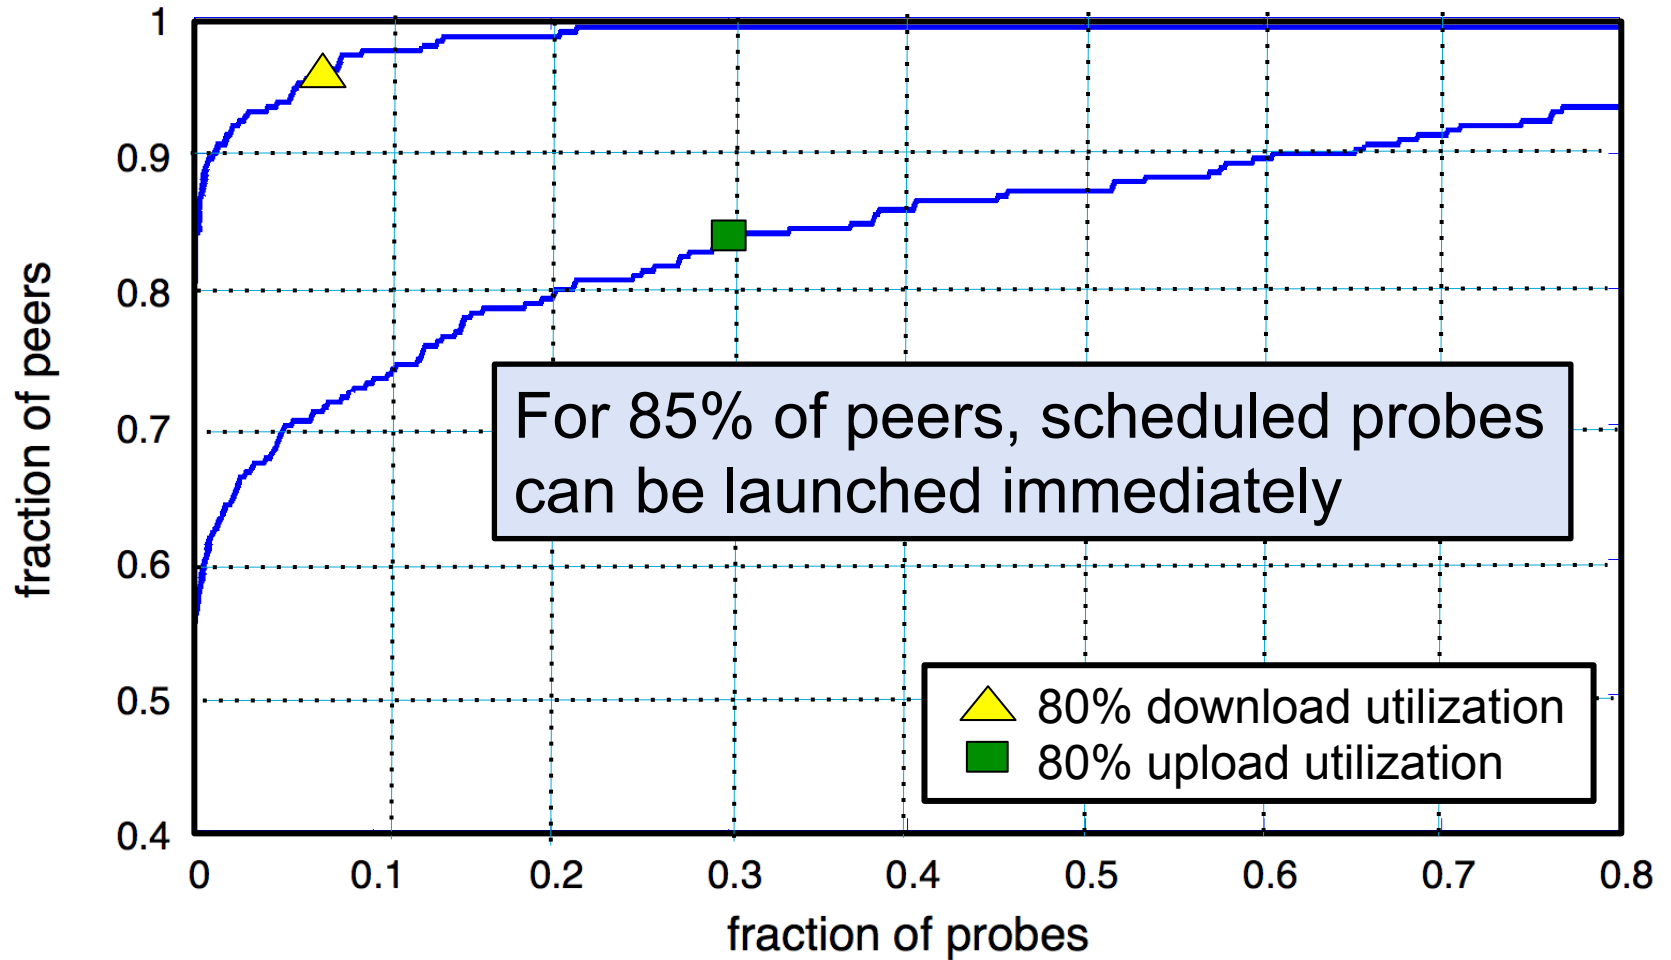
Challenges: Load-control on experiments

Delayed probes per peer



Challenges: Load-control on experiments

Delayed probes per peer




Illustrating its unique perspective

Revisiting past experiments: Routing asymmetry, King


Questioning old assumptions: AS-level connectivity

**Performing novel experiments:
Evaluating a recently-proposed DNS
extension**

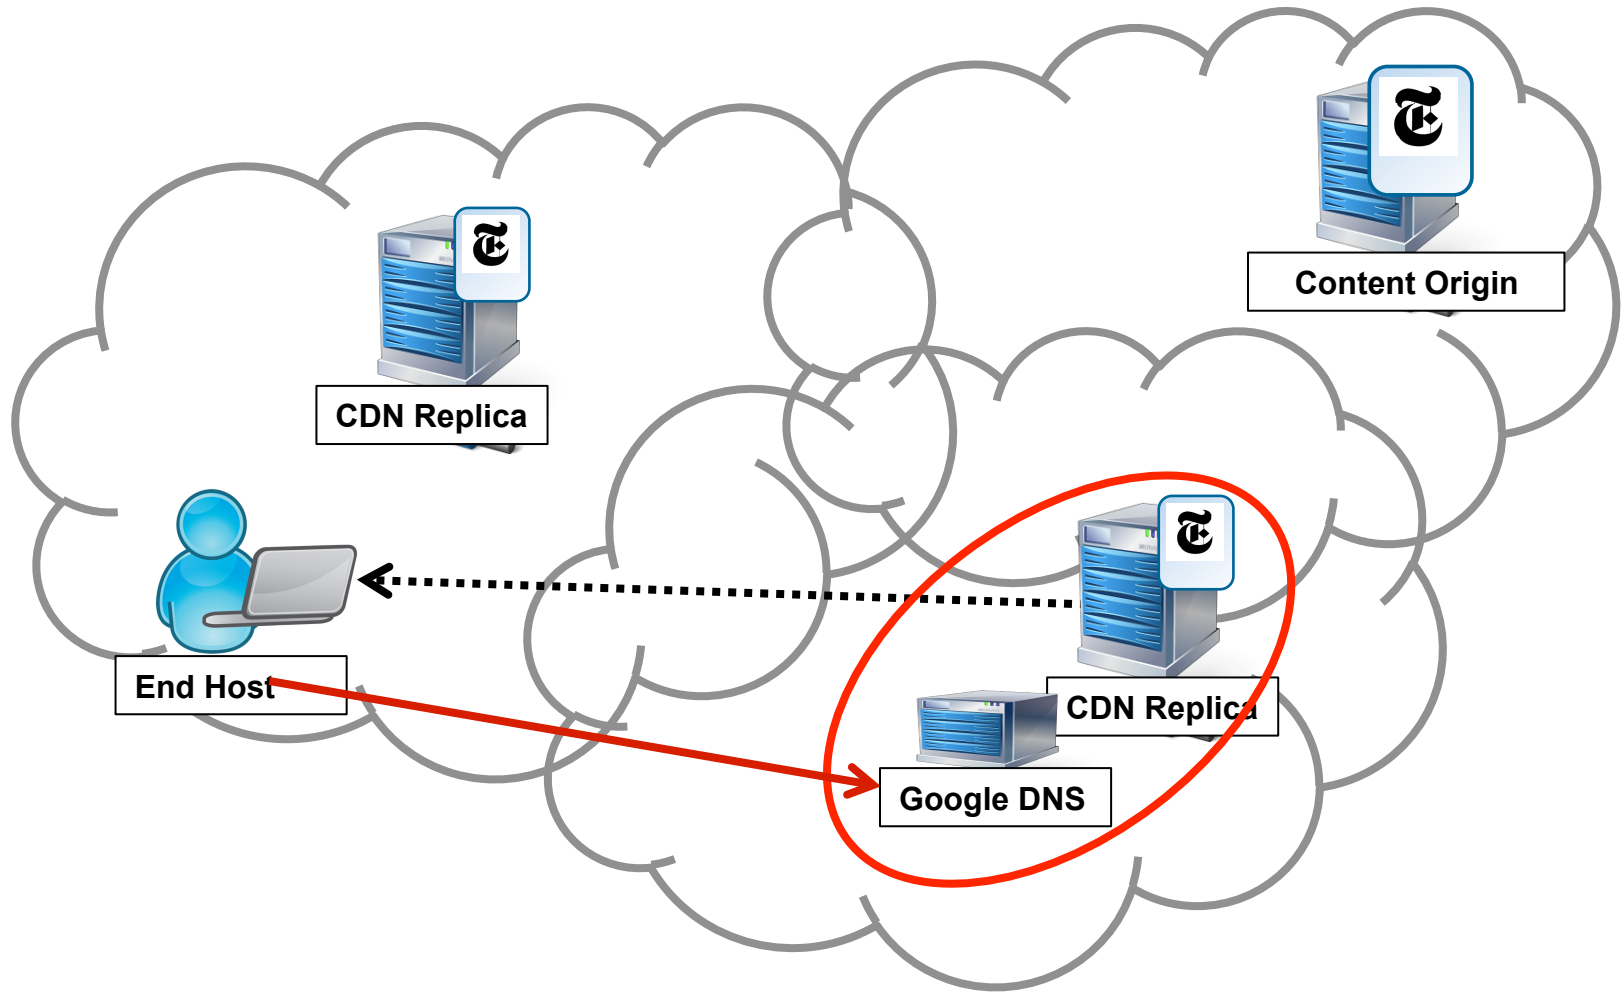
CDNs and DNS-based localization



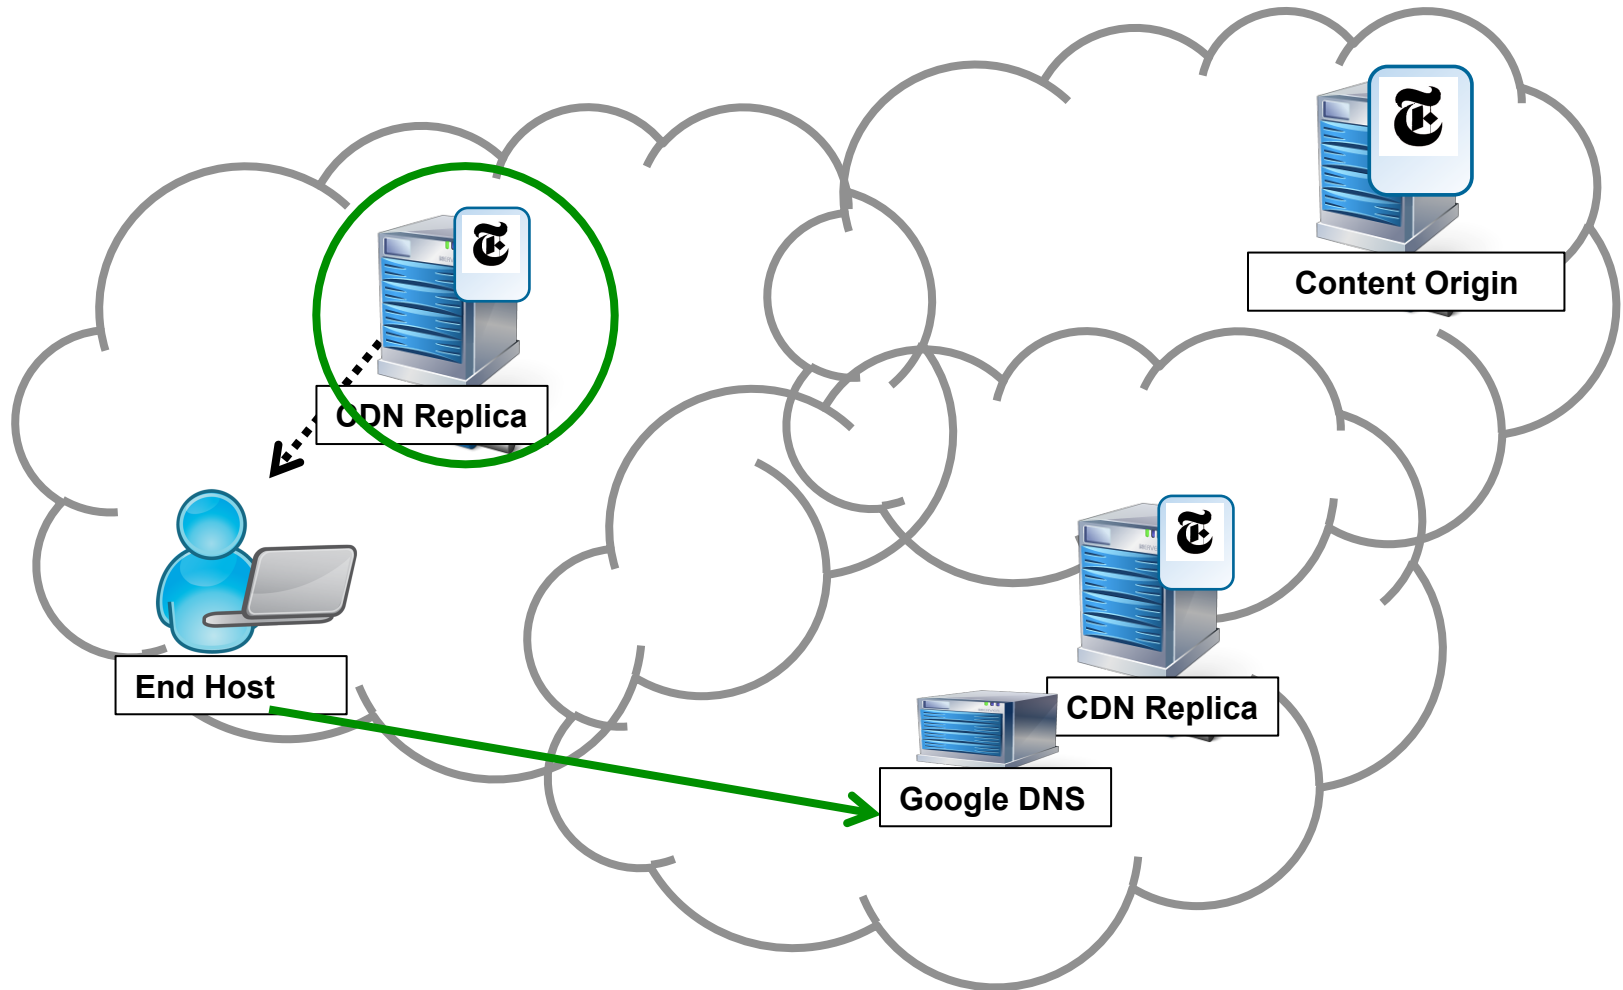
CDNs and DNS-based localization



CDNs and DNS-based localization



CDNs and DNS-based localization



DNS extension approach

Idea: Avoid impact of remote DNS usage by directly providing *client location*

A EDNS0 extension “*edns-client-subnet*”

Resolver adds client’s IP prefix to request

Is it effective?

Understanding performance benefits

Access to large set of vantage points

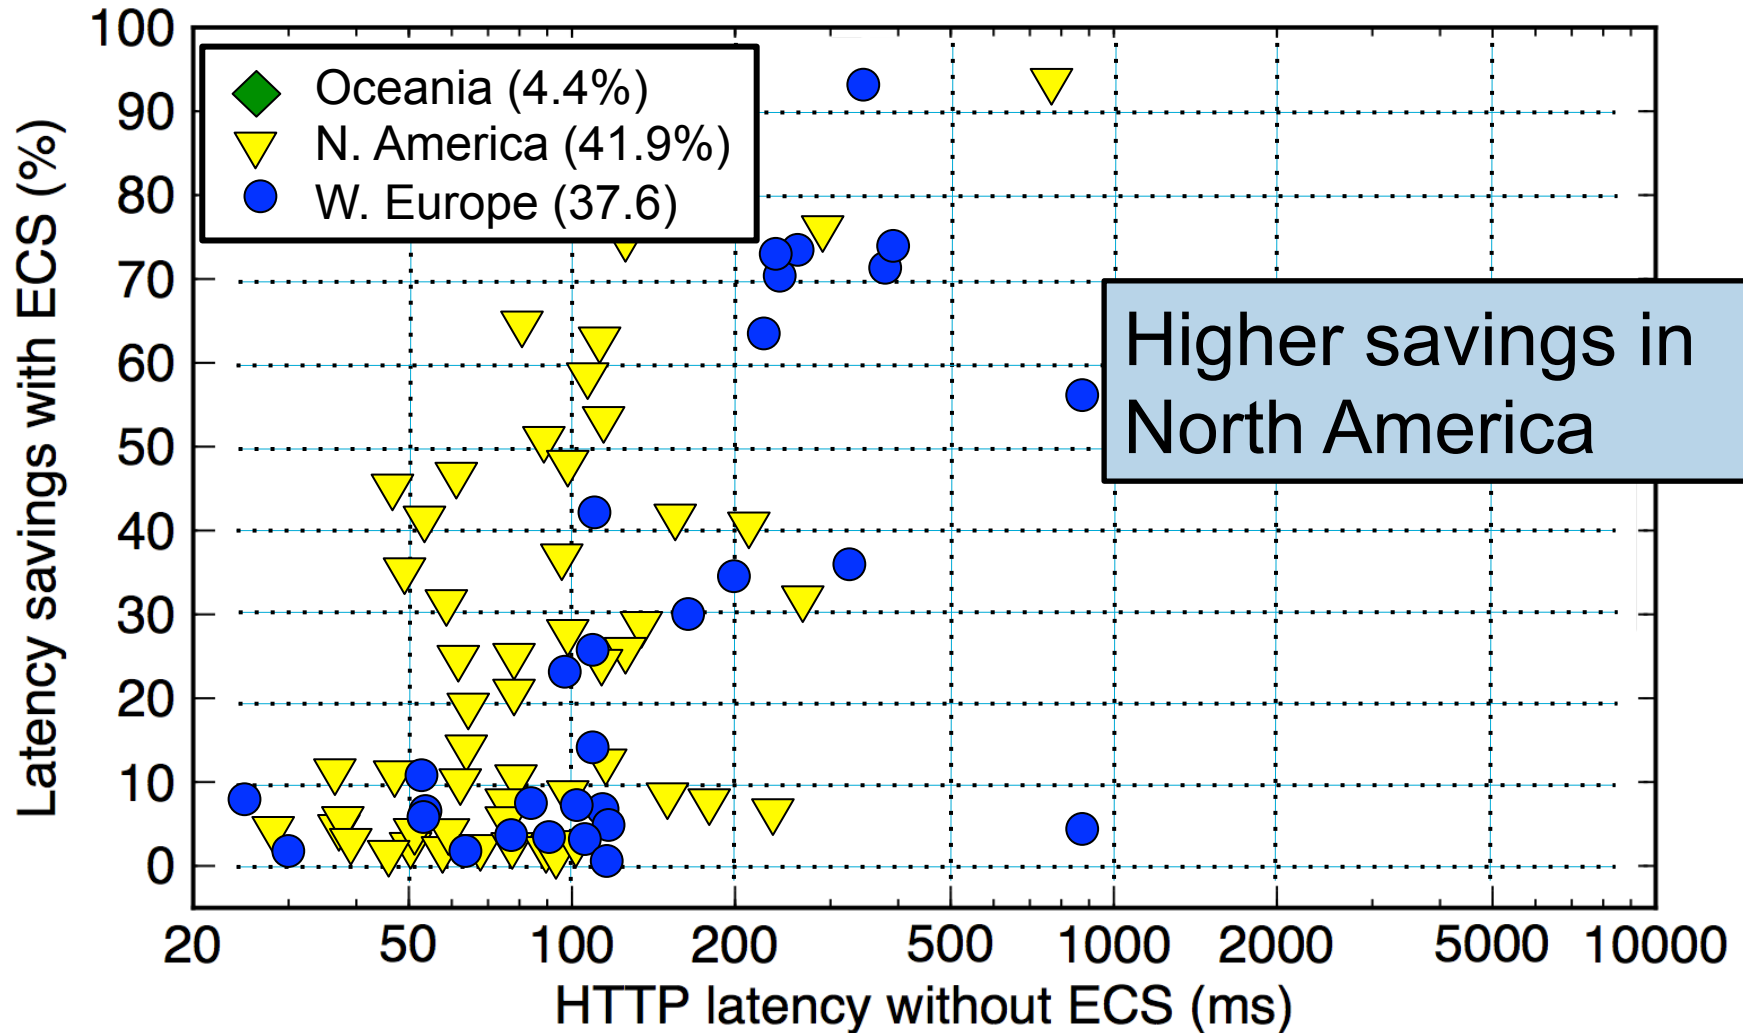
Vantage points located around the world

Interrelated measurement probes

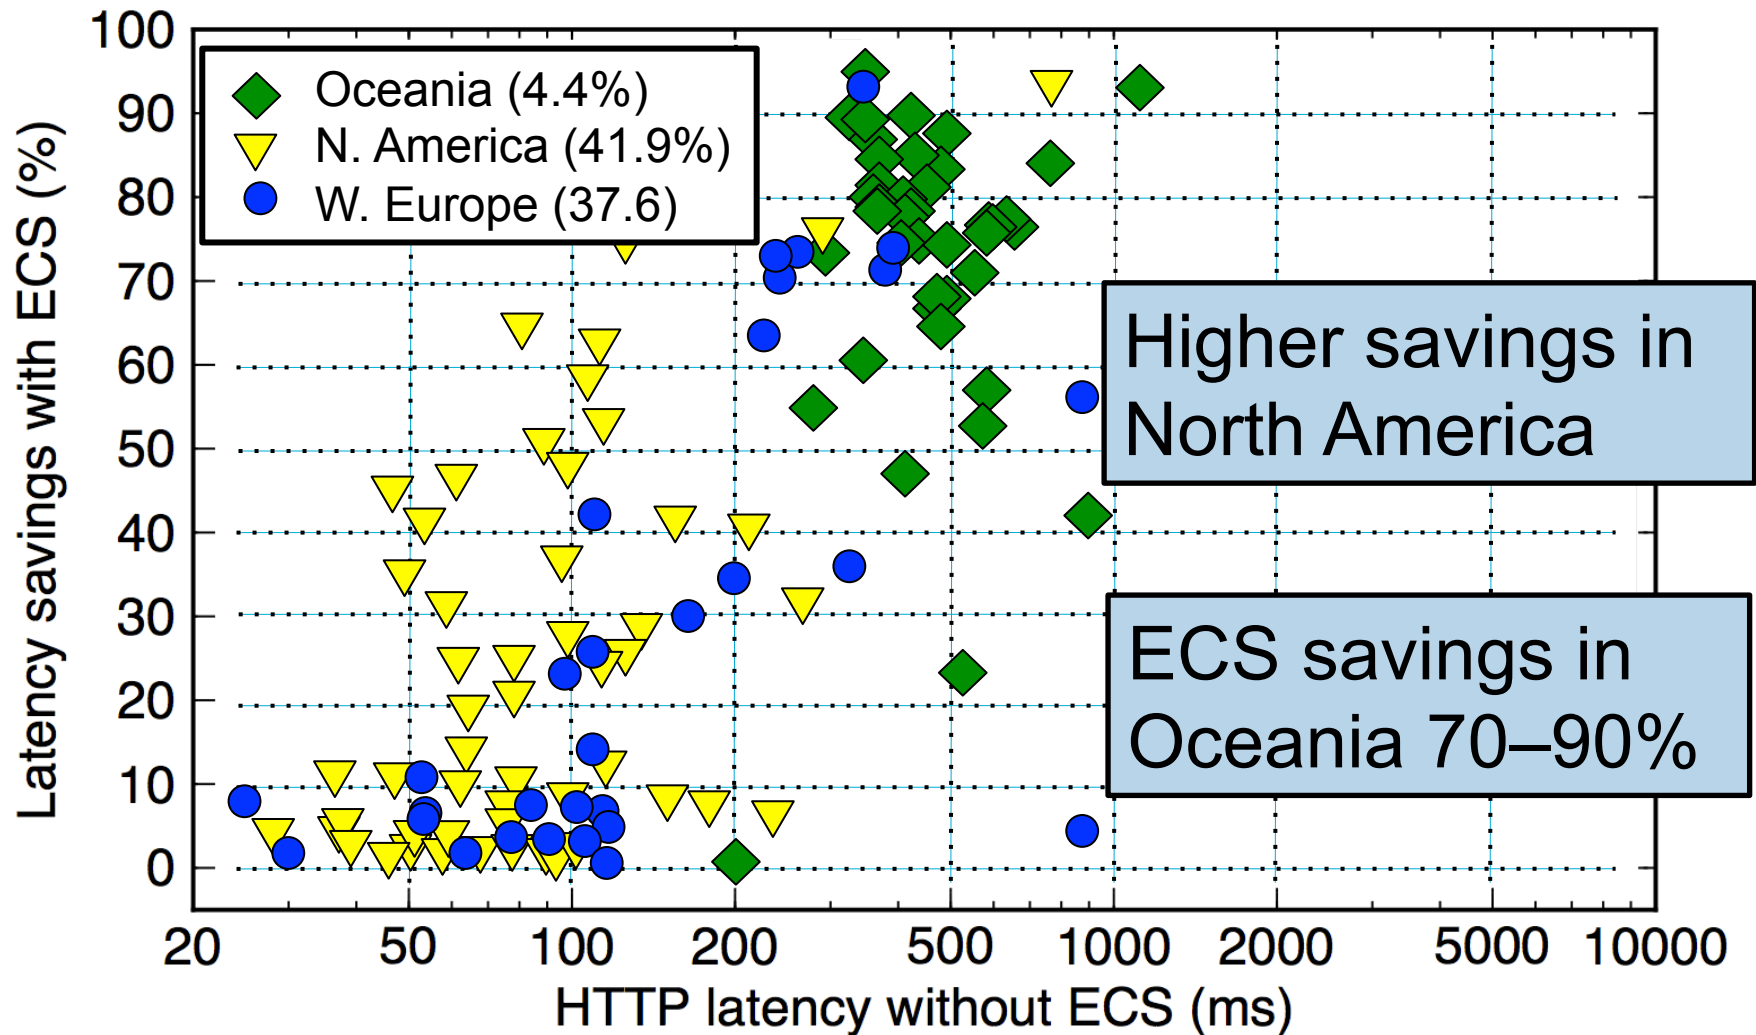
Flexible



Evaluating EDNS



Evaluating EDNS





... more than *just* better CDN performance

- Faster lookups with proactive caching
- Automatic, personalized server selection
- Graceful handling of DNS outages ...

*As of this morning – 52,800 users
over 147 countries!*

<http://aqualab.cs.northwestern.edu/projects/namehelp>

Summary

Experimental platform for the network edge

Relies on a direct incentive model for adoption

- Rather than relying on the kindness of strangers

Participants collectively offer

- Broad network coverage
- High availability
- Fine-grained control

*Contact us to
play with it*