

Measuring the Effectiveness of Happy Eyeballs

Vaibhav Bajpai and Jürgen Schönwälder
{v.bajpai, j.schoenwaelder}@jacobs-university.de

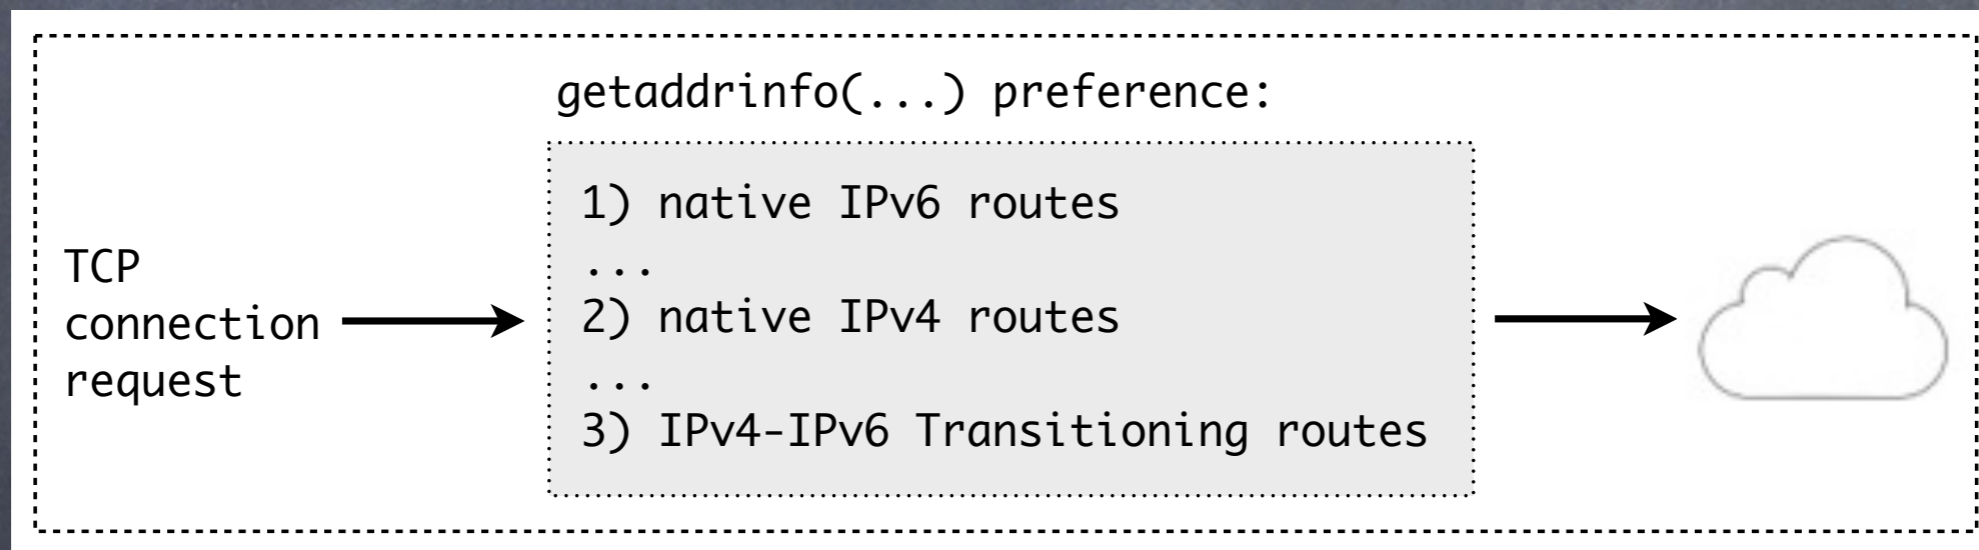
RIPE66, Dublin

Computer Networks and Distributed Systems
Jacobs University Bremen
Bremen, Germany

May 2013

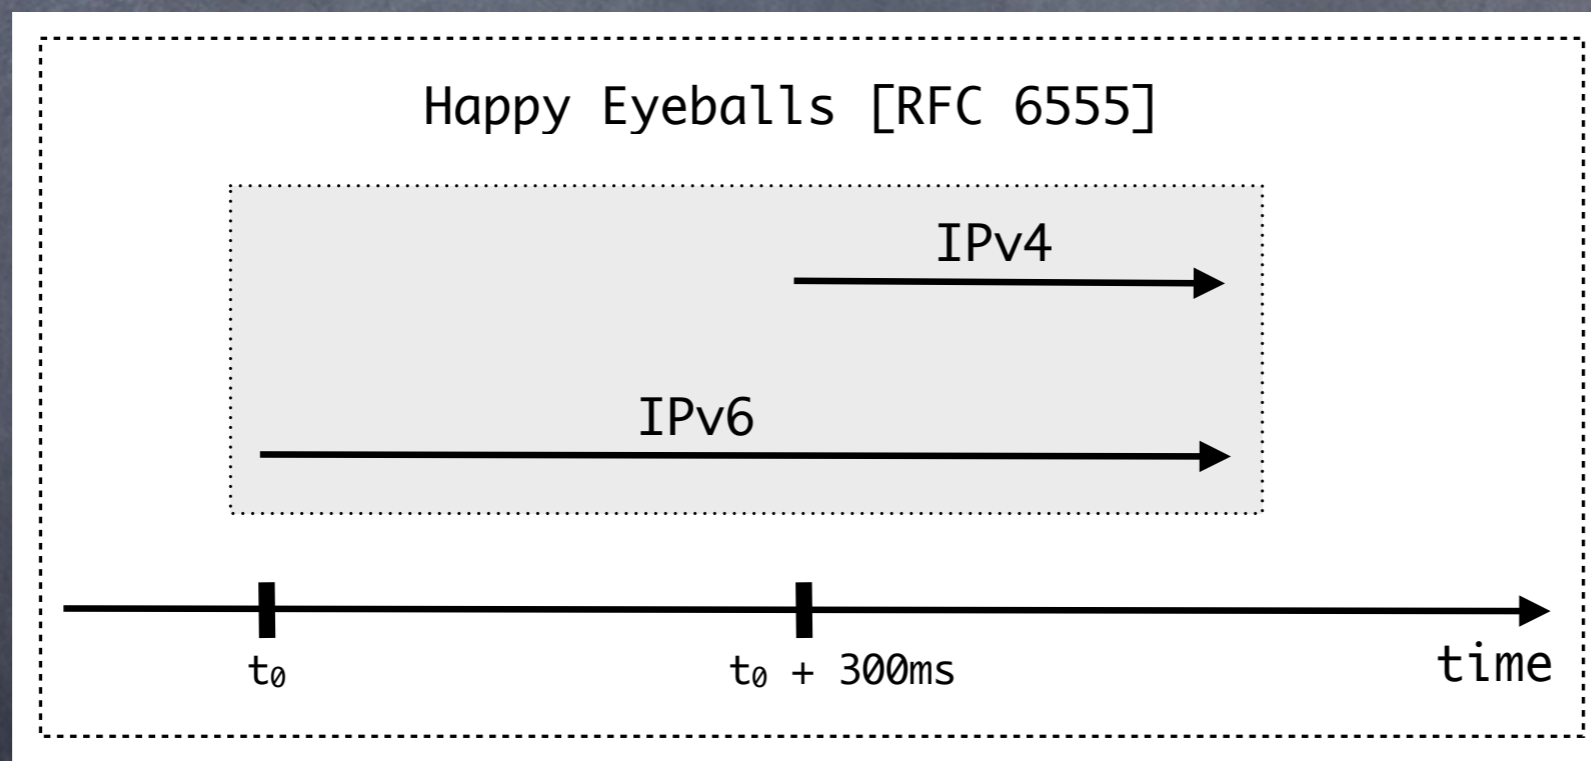
motivation

- `getaddrinfo(...)` behavior:
 - returns list of endpoints in an order that prioritizes IPv6 upgrade path
 - order is dictated by [RFC 6724] and `/etc/gai.conf`
 - if IPv6 is broken, application is unresponsive in order of seconds



motivation

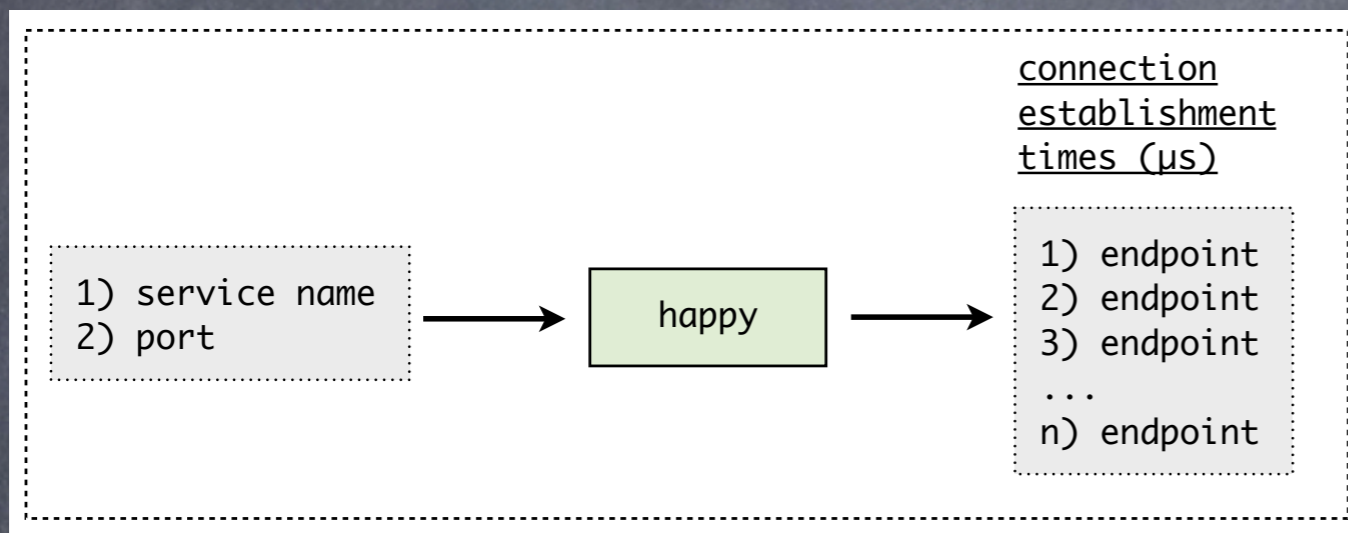
- happy eyeballs algorithm [RFC 6555]:
 - initiate a TCP connect(...) with the first endpoint, give it 300ms
 - switch over with a TCP connect(...) to a different address family otherwise
 - the competition runs fair after 300ms



- does the algorithm help improve the user experience?

metric and implementation

- developed a simple TCP happy eyeballs [RFC 6555] probing tool



- uses `getaddrinfo(...)` to resolve service names to endpoints
- uses non-blocking `connect(...)` to connect to all endpoints of a service
- uses a short-delay between connection attempts to avoid SYN floods
- the service name resolution time is not accounted in the output
- can produce either human-readable or machine-readable output
- file locking capability

```
>> ./happy -q 1 -m www.google.com www.facebook.com
HAPPY.0;1360681039;OK;www.google.com;80;173.194.69.105;8626
HAPPY.0;1360681039;OK;www.google.com;80;2a00:1450:4008:c01::69;8884
HAPPY.0;1360681039;OK;www.facebook.com;80;2a03:2880:10:6f01:face:b00c::8;170855
HAPPY.0;1360681039;OK;www.facebook.com;80;31.13.72.39;26665
```

measurement trials


- dual-stacked web service name list:
 - HE.net maintains a list of top 100 dual-stacked service names
 - they use 1M service names from Alexa Top Sites
 - some domains we expect are missing from the list
 - some services only provide a IPv6 endpoint on prepending a www
 - HE.net does not follow CNAMEs (for e.g. wikipedia.org)
 - amazon has made 1M service name list public
 - we use it and script it ourselves to explicitly follow CNAMEs
- measurement agents:
 - native IPv6, 6in4, Teredo, IPv6 tunnel broker endpoints, native IPv4
 - located at Bremen, Amsterdam, Braunschweig
- measurement cycle length:
 - 1 month

how does IPv6 compare in
performance to IPv4?

to what extent is IPv6 preferred
when connecting to a dual-
stacked service?


IPv6 preference levels

Native IPv6 [Bremen]



- IPv4 connectivity via DFN [AS 680]
- IPv6 connectivity via DFN [AS 680]

Teredo IPv6 [Amsterdam]




- IPv4 connectivity via LambdaNet Communications
- IPv6 connectivity via Teredo

how slow is a happy eyeballed
winner to that of a loser?

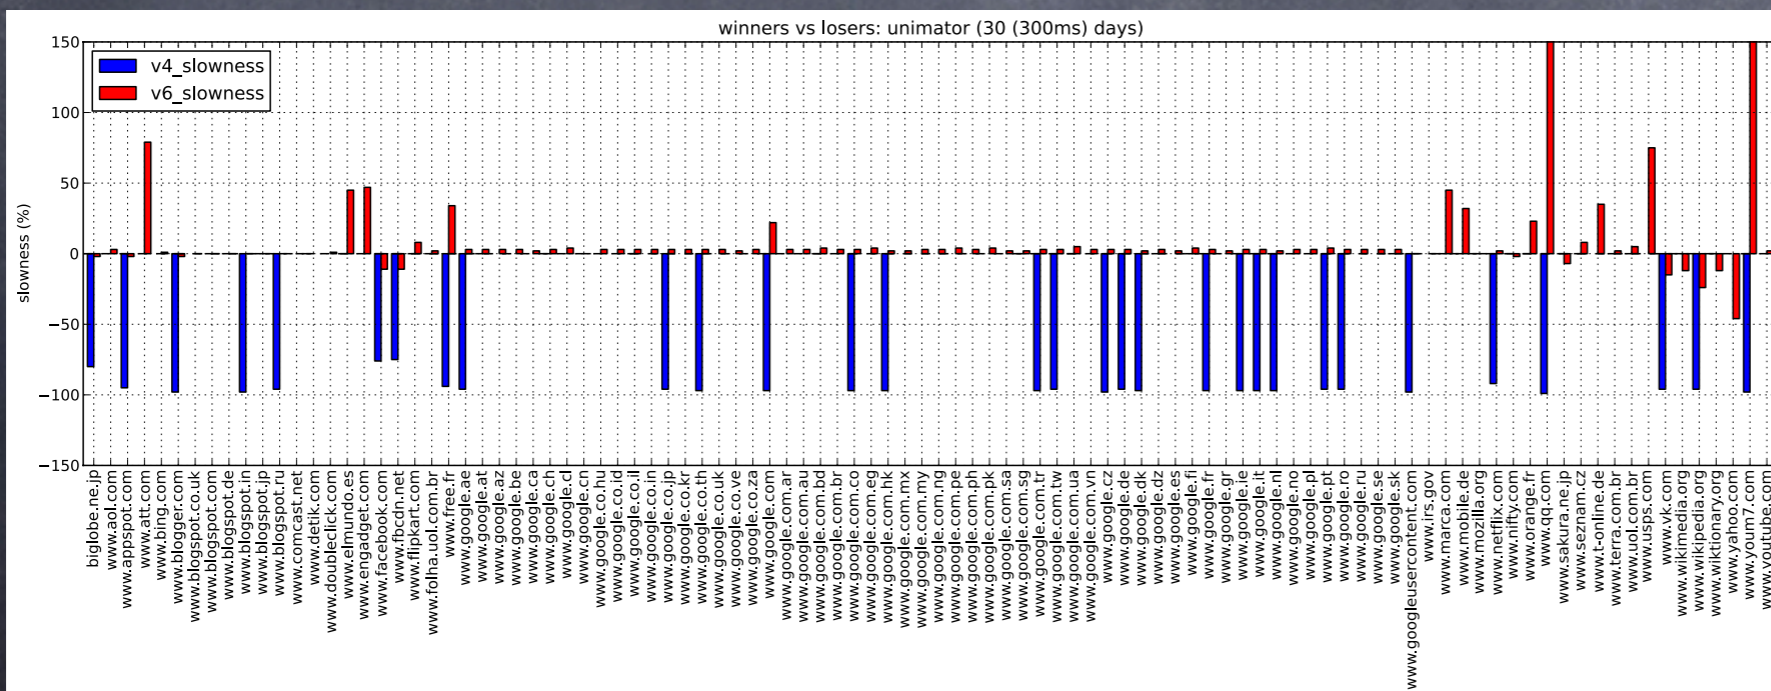
winner slowness to loser

Native IPv6 [Bremen]



- IPv4 connectivity via DFN [AS 680]
- IPv6 connectivity via DFN [AS 680]

Native IPv6 [Braunschweig]




- IPv4 connectivity via DFN [AS 680]
- IPv6 connectivity via DFN [AS 680]

what are repercussions of
reducing the IPv6 advantage from
300ms to 10ms

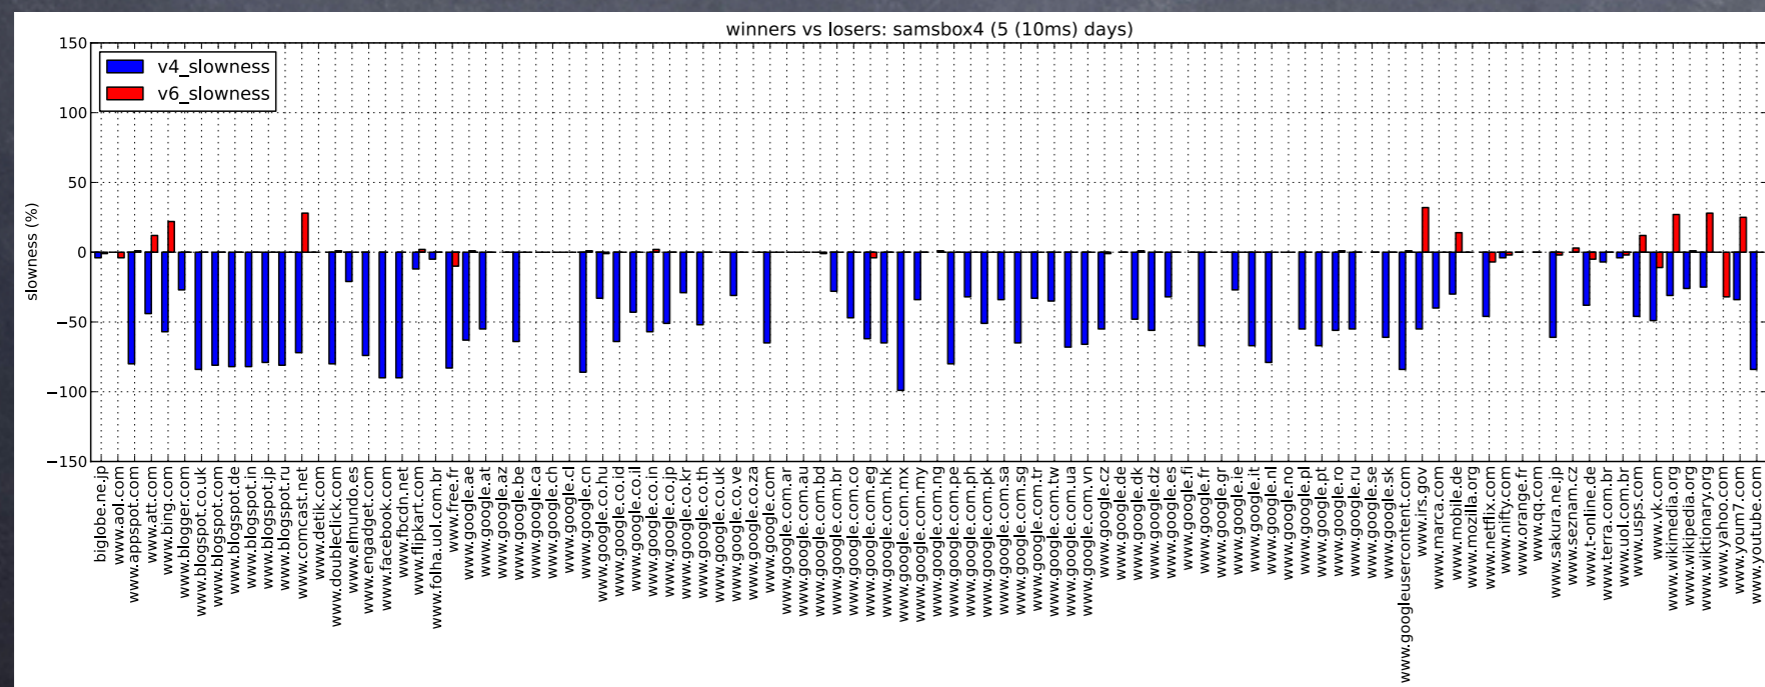
happy eyeballs advantage: 10ms

Native IPv6 [Bremen]



- IPv4 connectivity via Deutsche Telekom AG [AS3320]
- IPv6 connectivity via Deutsche Telekom AG [AS3320]

Native IPv6 [Bremen]



- IPv4 connectivity via Deutsche Telekom AG [AS3320]
- IPv6 connectivity via Deutsche Telekom AG [AS3320]

conclusion

- higher connection times and variations over IPv6
- will never use Teredo IPv6 unless IPv4 connectivity is broken
- 300ms advantage leaves 1% chance to prefer IPv4 (even though faster)
- IPv6 happy eyeballed winner is rarely faster than IPv4 route
- 10ms advantage helps remove outliers where IPv6 connectivity is bad

- request:
 - happy must be run from a wider standpoint to get a more comprehensive picture
 - looking for hosts with native IPv6 connectivity to host our happy test.
 - send me your shipment address*, and we ship you a SamKnows probe.

* v.bajpai@jacobs-university.de