Enabling Multi-hop ISP-Hypergiant Collaboration

Cristian Munteanu*, Oliver Gasser*, Ingmar Poese^,
Georgios Smaragdakis~, Anja Feldmann*

*Max Planck Institute for Informatics, ^Benocs, ~TU Delft

Hypergiants and ISPs

Hypergiants and ISPs

Hypergiants and ISPs

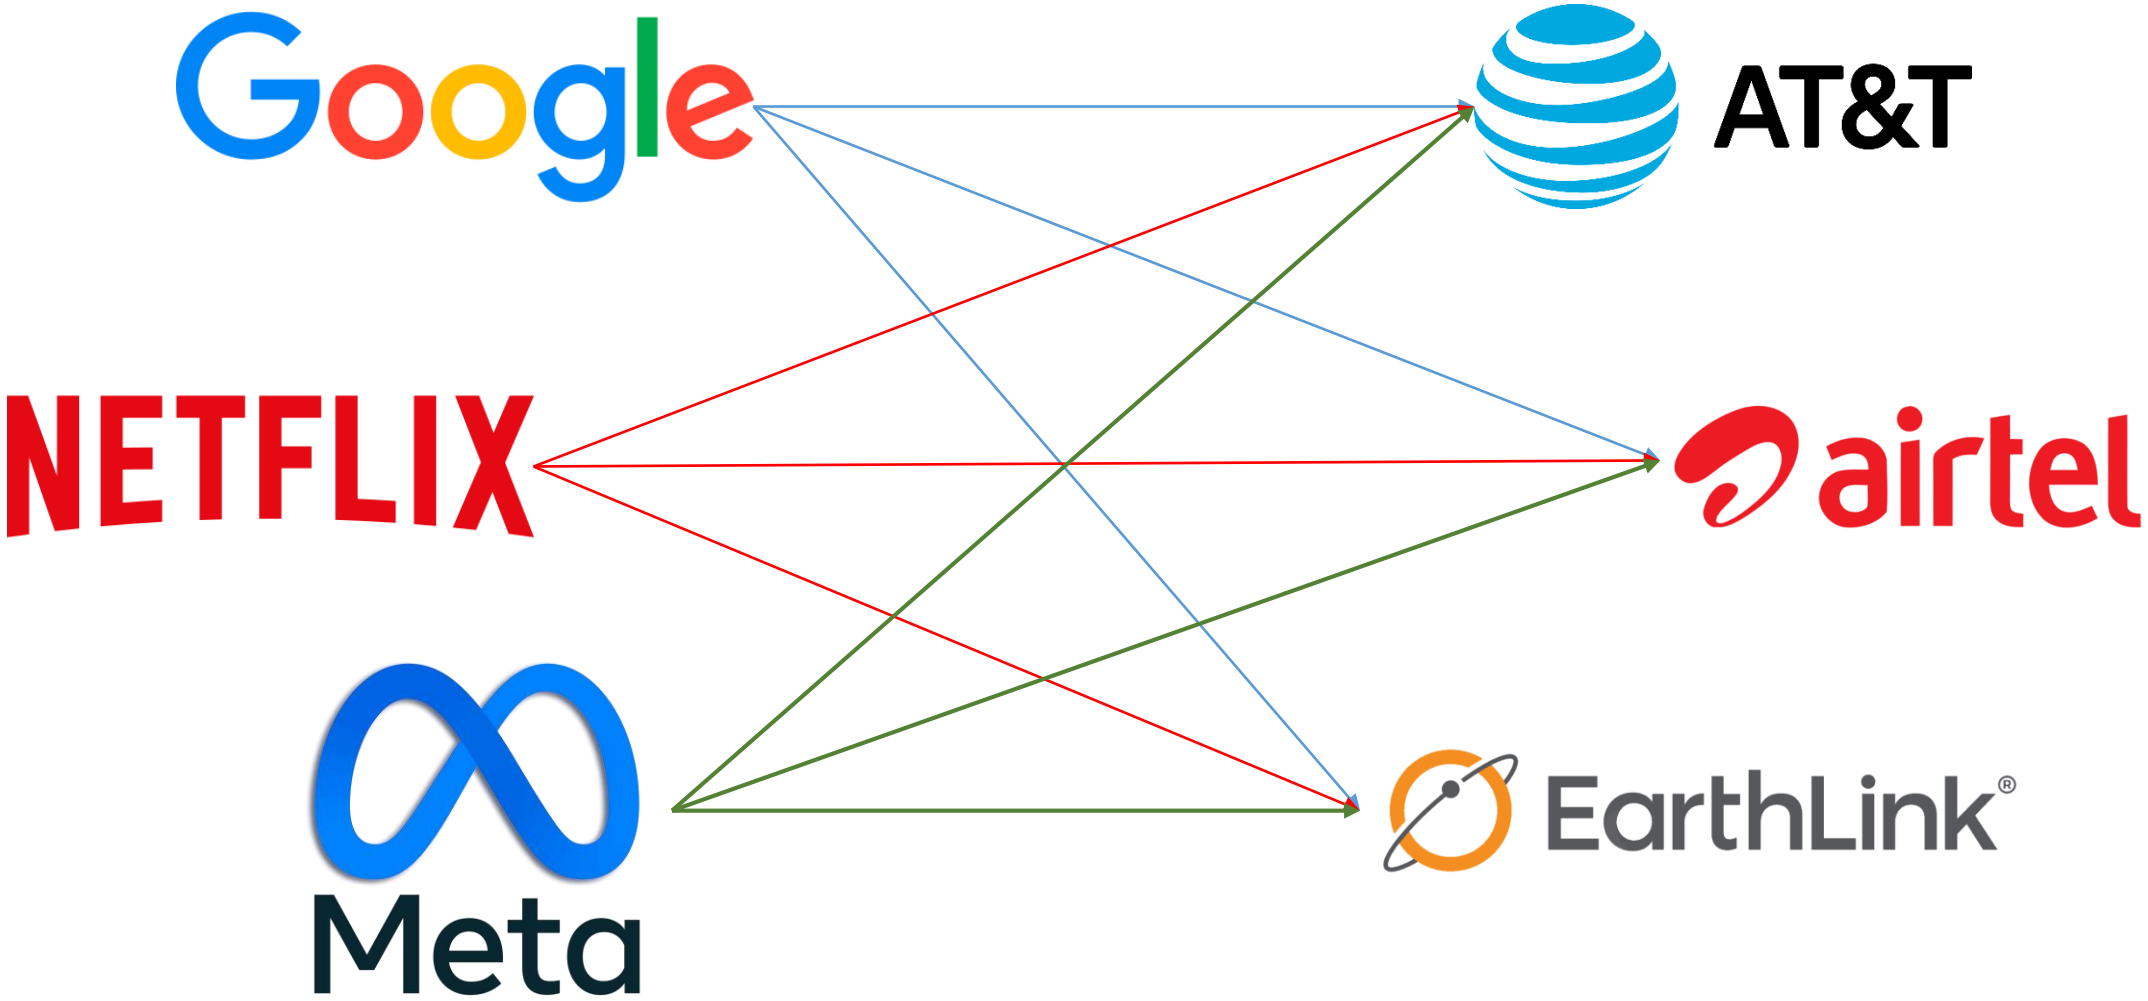

Hypergiants and ISPs

Large Hypergiants peer with more than **10K networks**

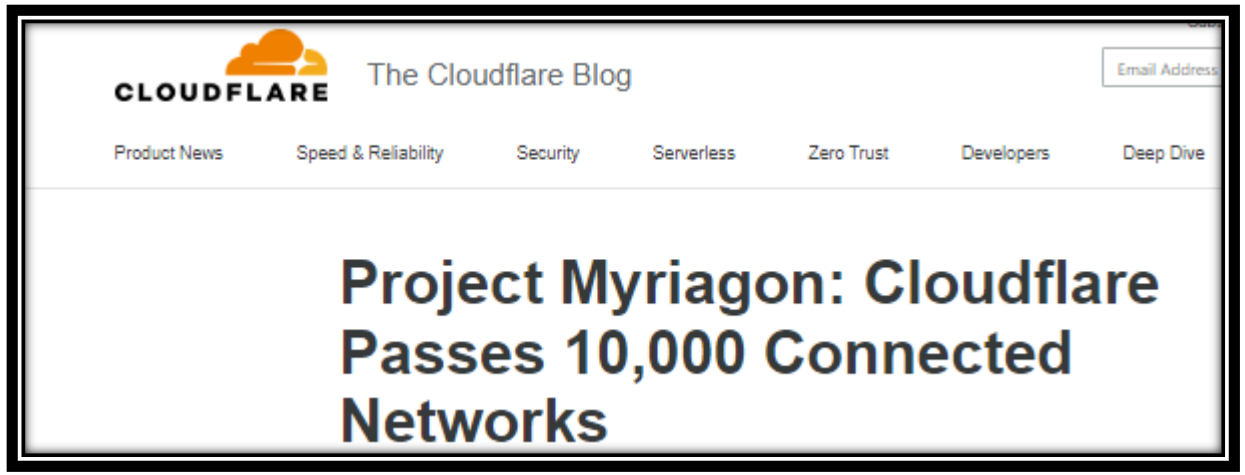

Google >12K

>10K

Server selection

Hypergiants need to select the optimal server

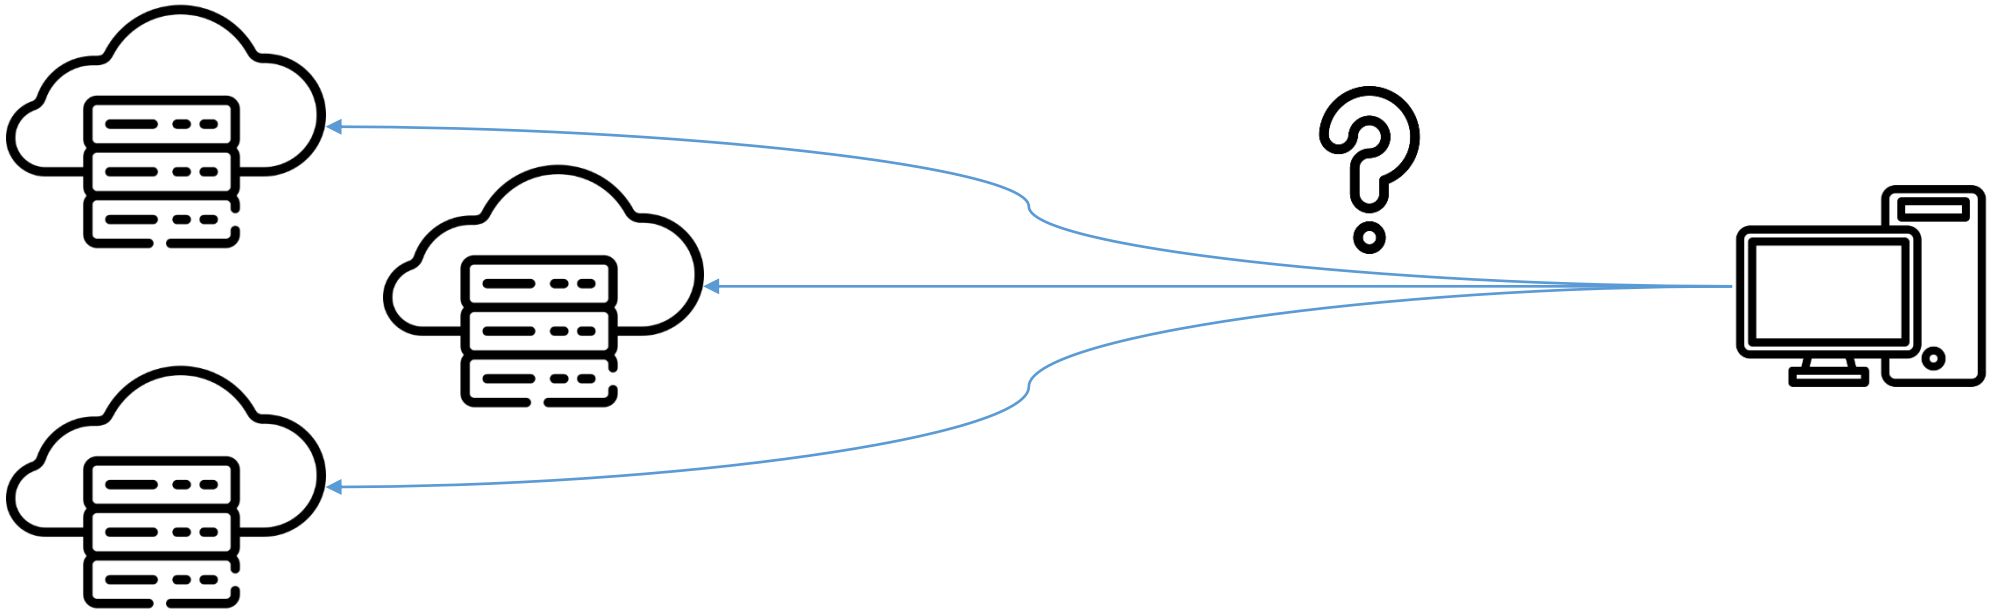

Server selection

Hypergiants need to select the optimal server

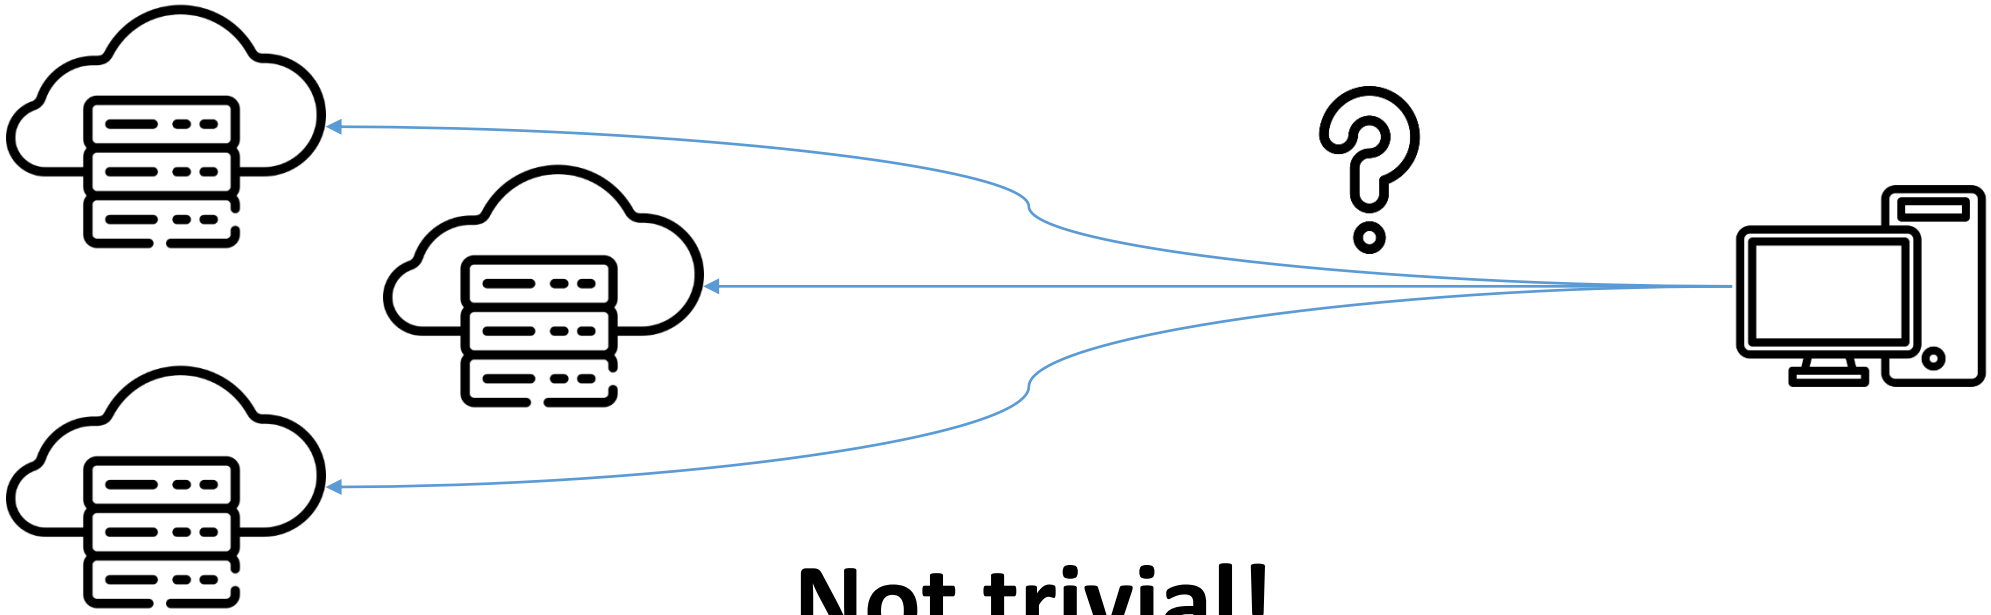

Not trivial!

Previous work

*Pujol et. al. designed a system that help the Hypergiants to improve their server selection for the clients of "*neighbor*" ISPs.

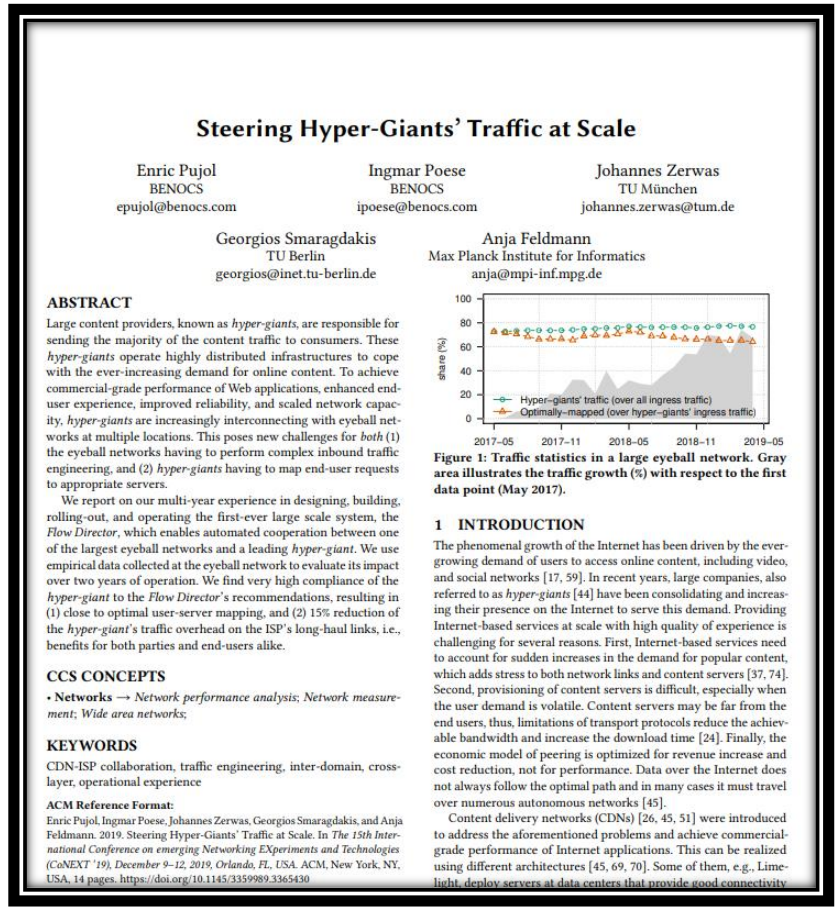

Steering Hyper-Giants' Traffic at Scale

Enric Pujol
BENOCS
epujol@benocs.com

Ingmar Poese
BENOCS
ipoese@benocs.com

Johannes Zerwas
TU München
johannes.zerwas@tum.de

Georgios Smaragdakis
TU Berlin
georgios@inet.tu-berlin.de

Anja Feldmann
Max Planck Institute for Informatics
anja@mpi-inf.mpg.de

ABSTRACT

Large content providers, known as *hyper-giants*, are responsible for sending the majority of the content traffic to consumers. These *hyper-giants* operate highly distributed infrastructures to cope with the ever-increasing demand for online content. To achieve commercial-grade performance of Web applications, enhanced end-user experience, improved reliability, and scaled network capacity, *hyper-giants* are increasingly interconnecting with eyeball networks at multiple locations. This poses new challenges for *both* (1) the eyeball networks having to perform complex inbound traffic engineering, and (2) *hyper-giants* having to map end-user requests to appropriate servers.

We report on our multi-year experience in designing, building, rolling-out, and operating the first-ever large scale system, the *Flow Director*, which enables automated cooperation between one of the largest eyeball networks and a leading *hyper-giant*. We use empirical data collected at the eyeball network to evaluate its impact over two years of operation. We find very high compliance of the *hyper-giant* to the *Flow Director*'s recommendations, resulting in (1) close to optimal user-server mapping, and (2) 15% reduction of the *hyper-giant*'s traffic overhead on the ISP's long-haul links, i.e., benefits for both parties and end-users alike.

CCS CONCEPTS

• **Networks** → *Network performance analysis; Network measurement; Wide area networks;*

KEYWORDS

CDN-ISP collaboration, traffic engineering, inter-domain, cross-layer, operational experience

ACM Reference Format:

Enric Pujol, Ingmar Poese, Johannes Zerwas, Georgios Smaragdakis, and Anja Feldmann. 2019. Steering Hyper-Giants' Traffic at Scale. In *The 15th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '19), December 9–12, 2019, Orlando, FL, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3359989.3365430

Figure 1: Traffic statistics in a large eyeball network. Gray area illustrates the traffic growth (%) with respect to the first data point (May 2017).

1 INTRODUCTION

The phenomenal growth of the Internet has been driven by the ever-growing demand of users to access online content, including video, and social networks [17, 59]. In recent years, large companies, also referred to as *hyper-giants* [44] have been consolidating and increasing their presence on the Internet to serve this demand. Providing Internet-based services at scale with high quality of experience is challenging for several reasons. First, Internet-based services need to account for sudden increases in the demand for popular content, which adds stress to both network links and content servers [37, 74]. Second, provisioning of content servers is difficult, especially when the user demand is volatile. Content servers may be far from the end users, thus, limitations of transport protocols reduce the achievable bandwidth and increase the download time [24]. Finally, the economic model of peering is optimized for revenue increase and cost reduction, not for performance. Data over the Internet does not always follow the optimal path and in many cases it must travel over numerous autonomous networks [45].

Content delivery networks (CDNs) [26, 45, 51] were introduced to address the aforementioned problems and achieve commercial-grade performance of Internet applications. This can be realized using different architectures [45, 69, 70]. Some of them, e.g., Limelight, deploy servers at data centers that provide good connectivity

**Enric Pujol, Ingmar Poese, Johannes Zerwas, Georgios Smaragdakis, and Anja Feldmann. "Steering hyper-giants' traffic at scale". CoNEXT 2019.*

How about the networks that do not peer with the Hypergiants?

There are around 40K networks that do not peer with a Hypergiant!

Collaboration with a Large European Transit provider

A large number of ISPs that do not peer with majority of Hypergiants and rely on their transit provider!

Small ISP - Hypergiant Traffic

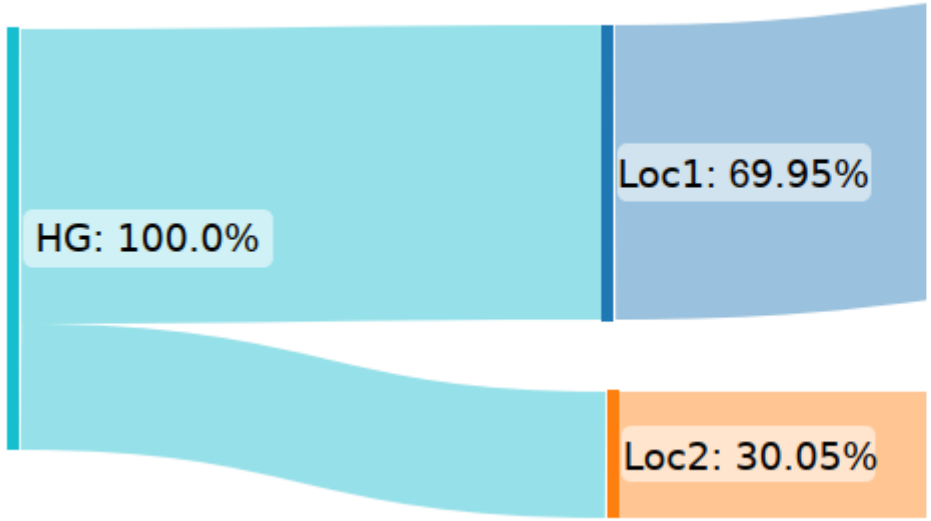

Hypergiant Transit AS

Small ISP - Hypergiant Traffic

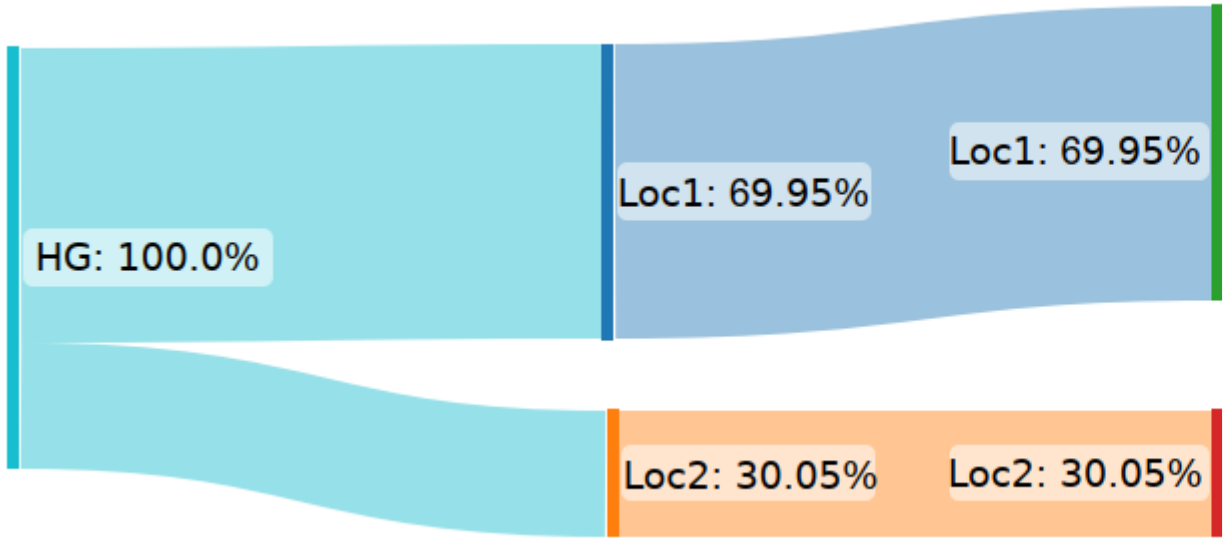

Hypergiant Transit AS Small European ISP

Small ISP - Hypergiant Traffic

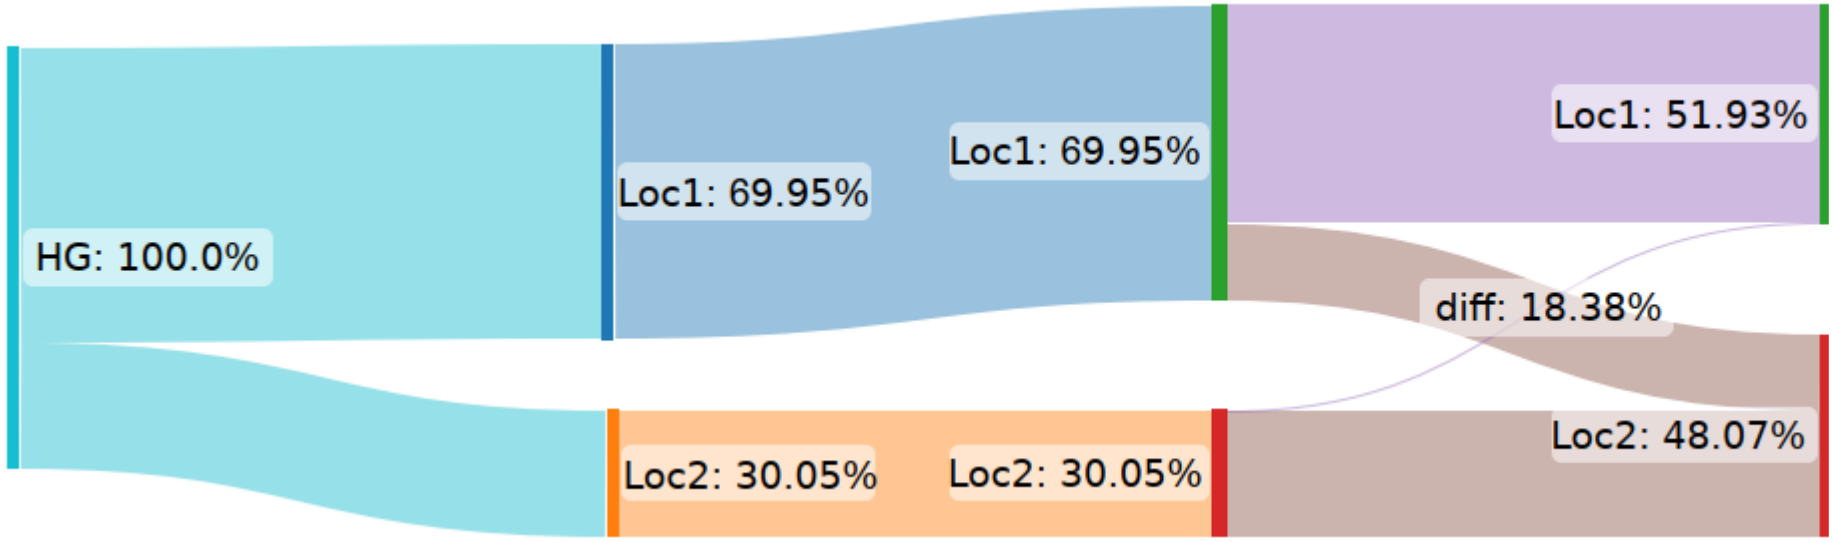

Hypergiant Transit AS Small European ISP

Small ISP - Hypergiant Traffic

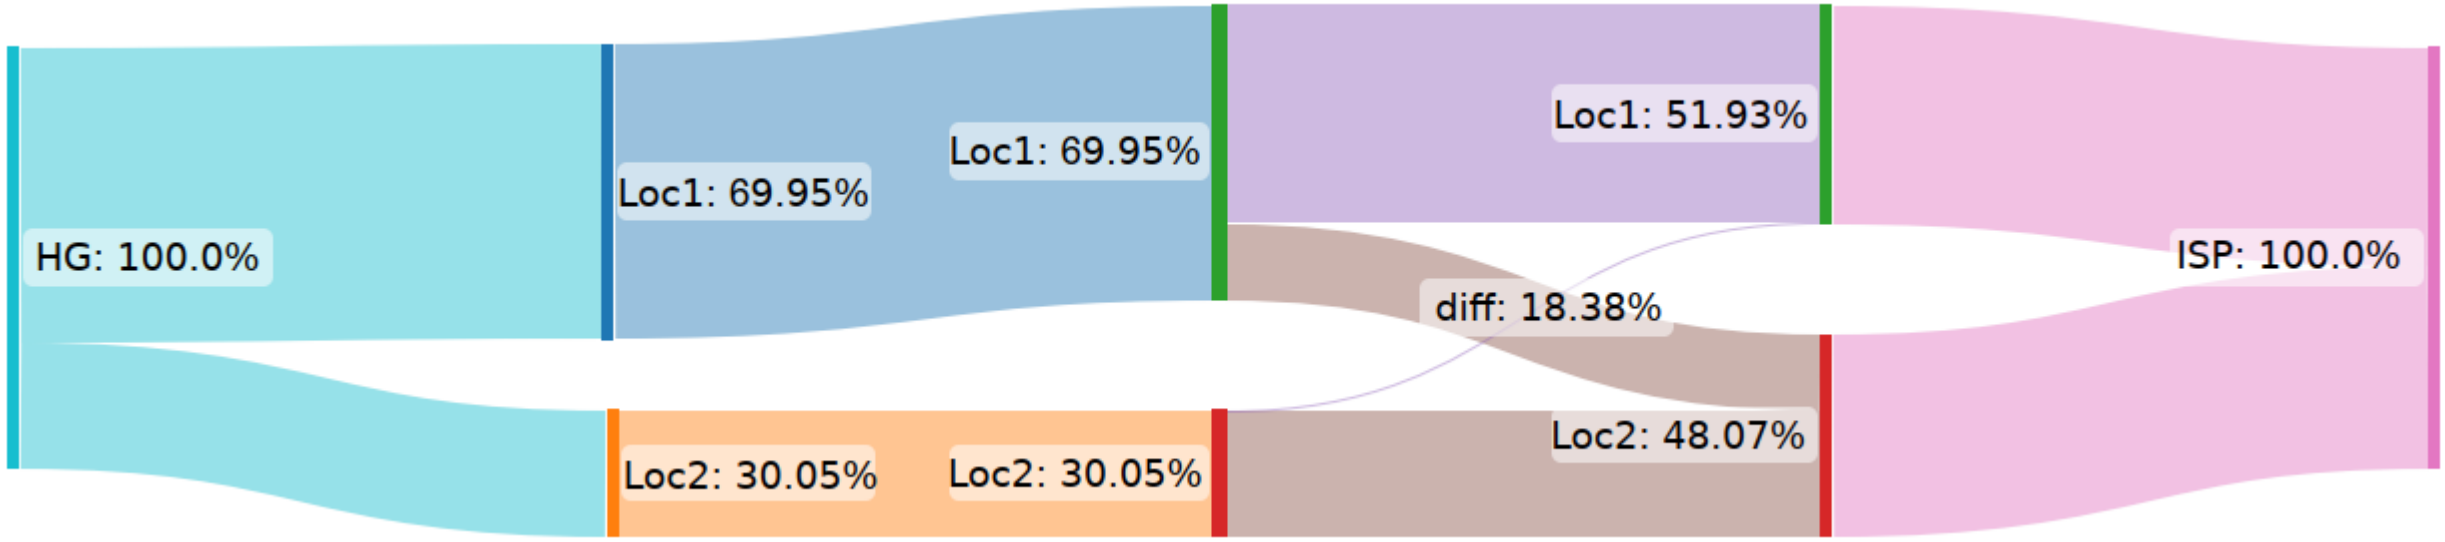

Hypergiant Transit AS Small European ISP

Small ISP - Hypergiant Traffic

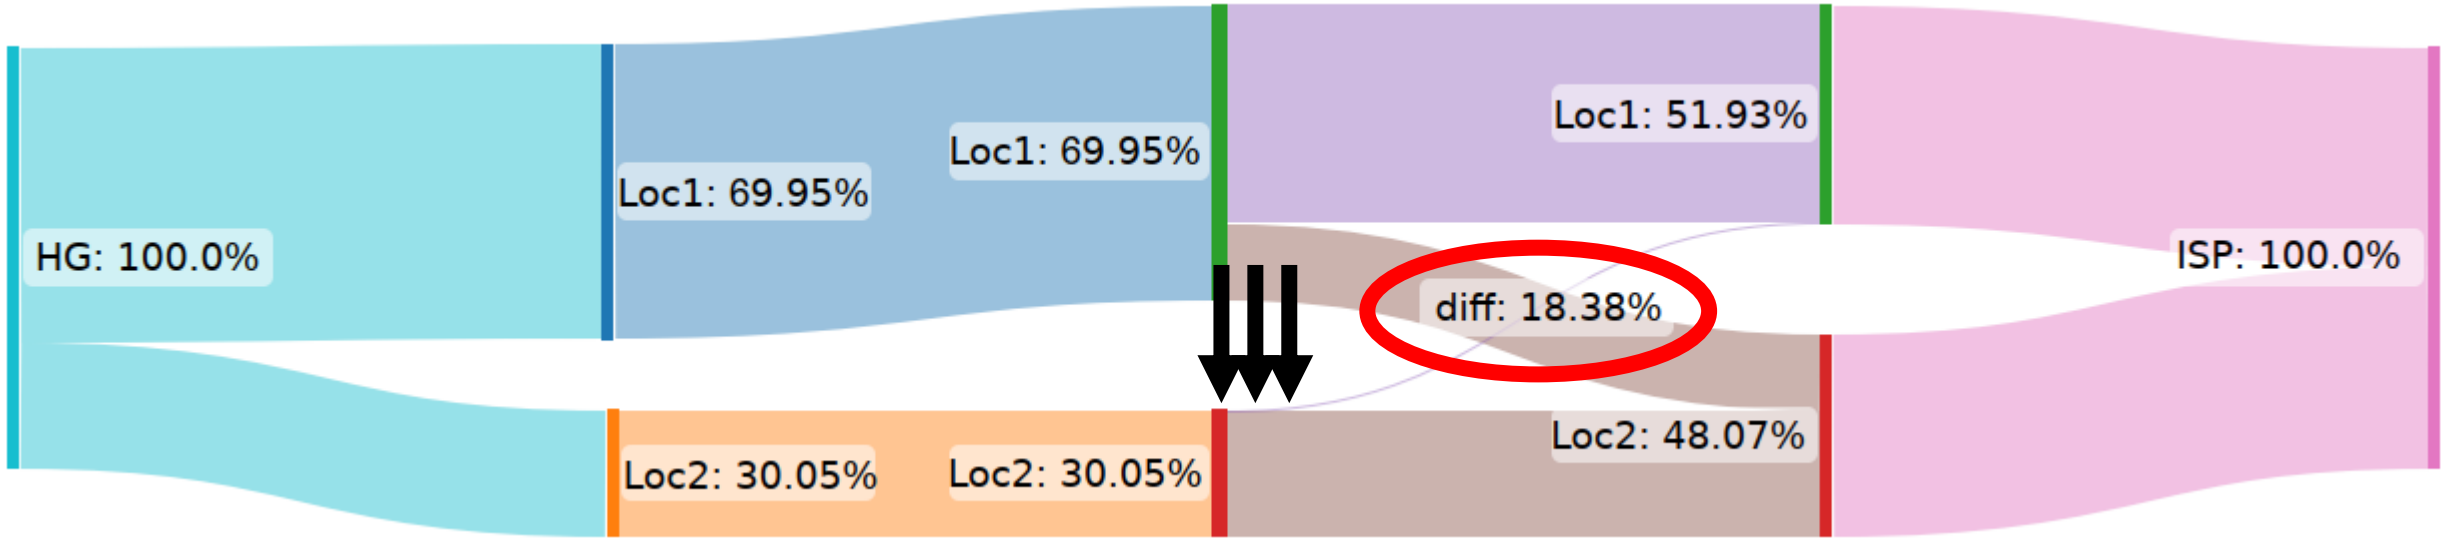

Hypergiant Transit AS Small European ISP

Small ISP - Hypergiant Traffic

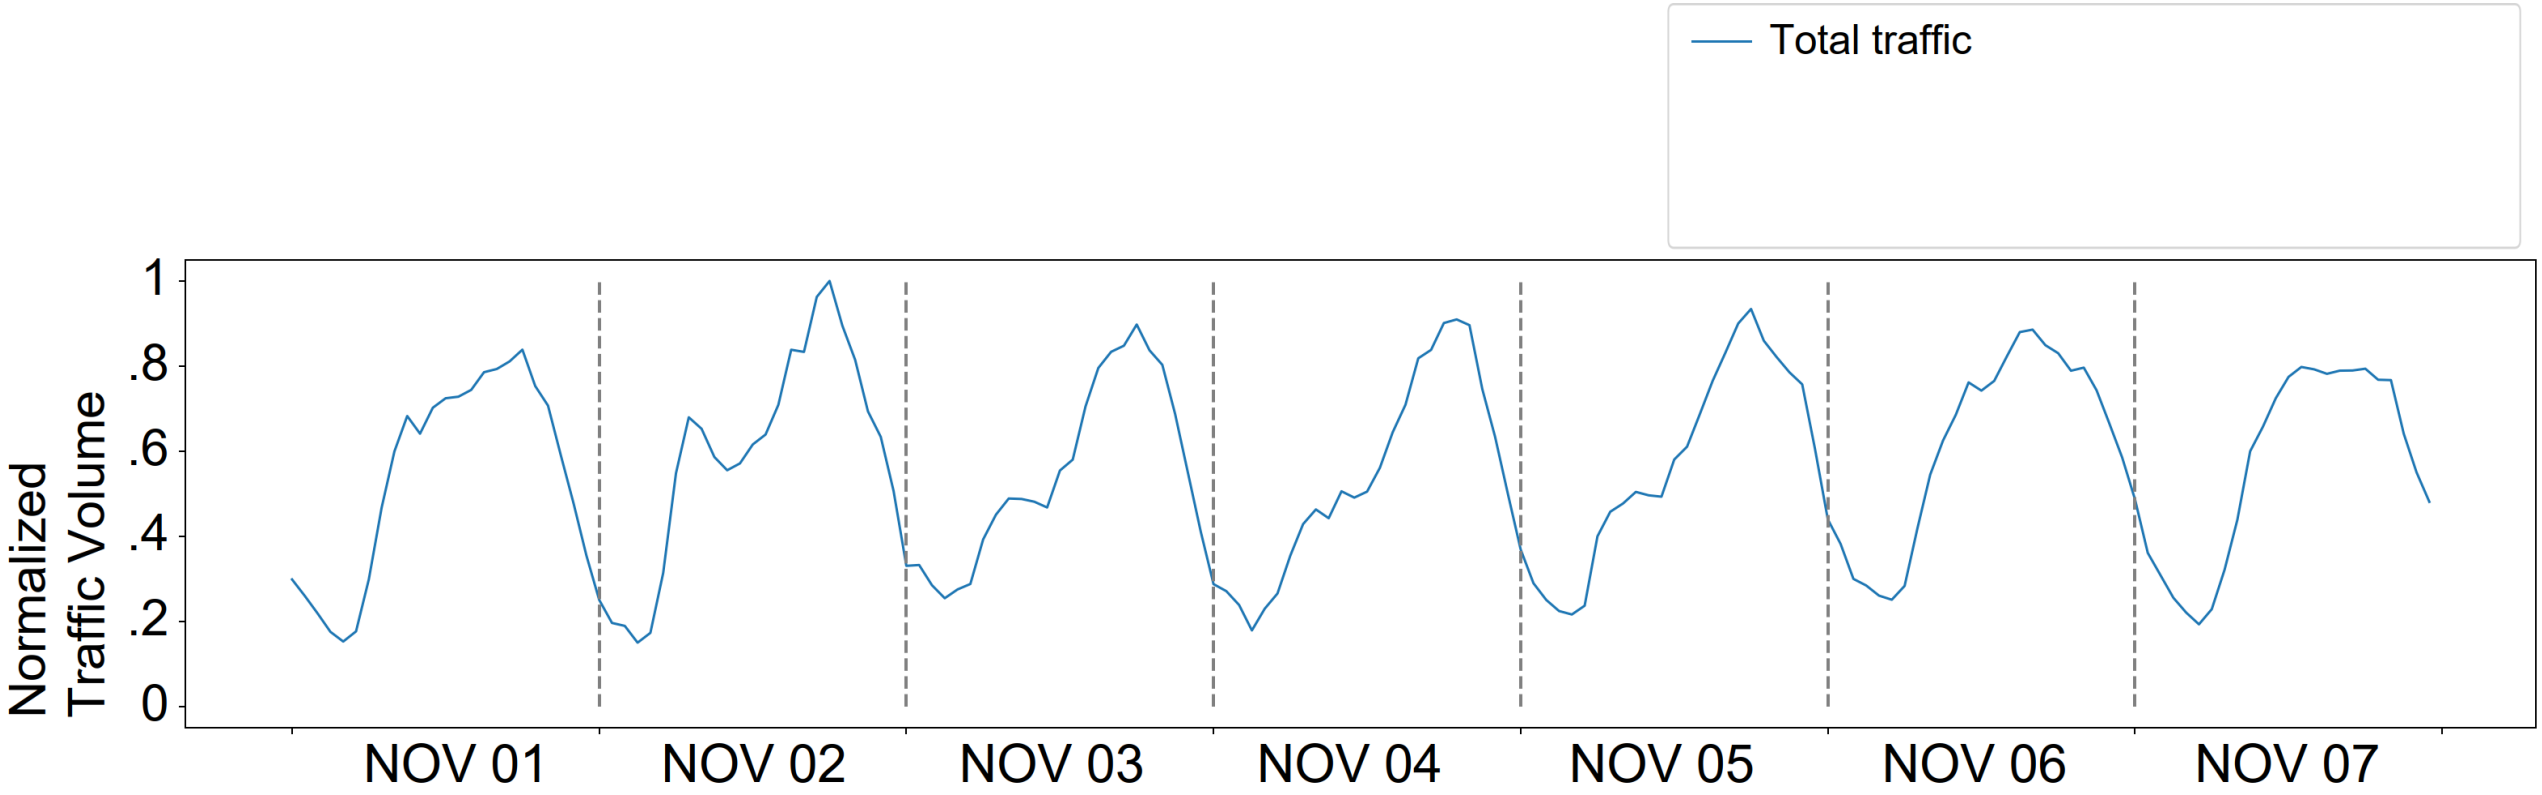

Small ISP - Hypergiant Traffic

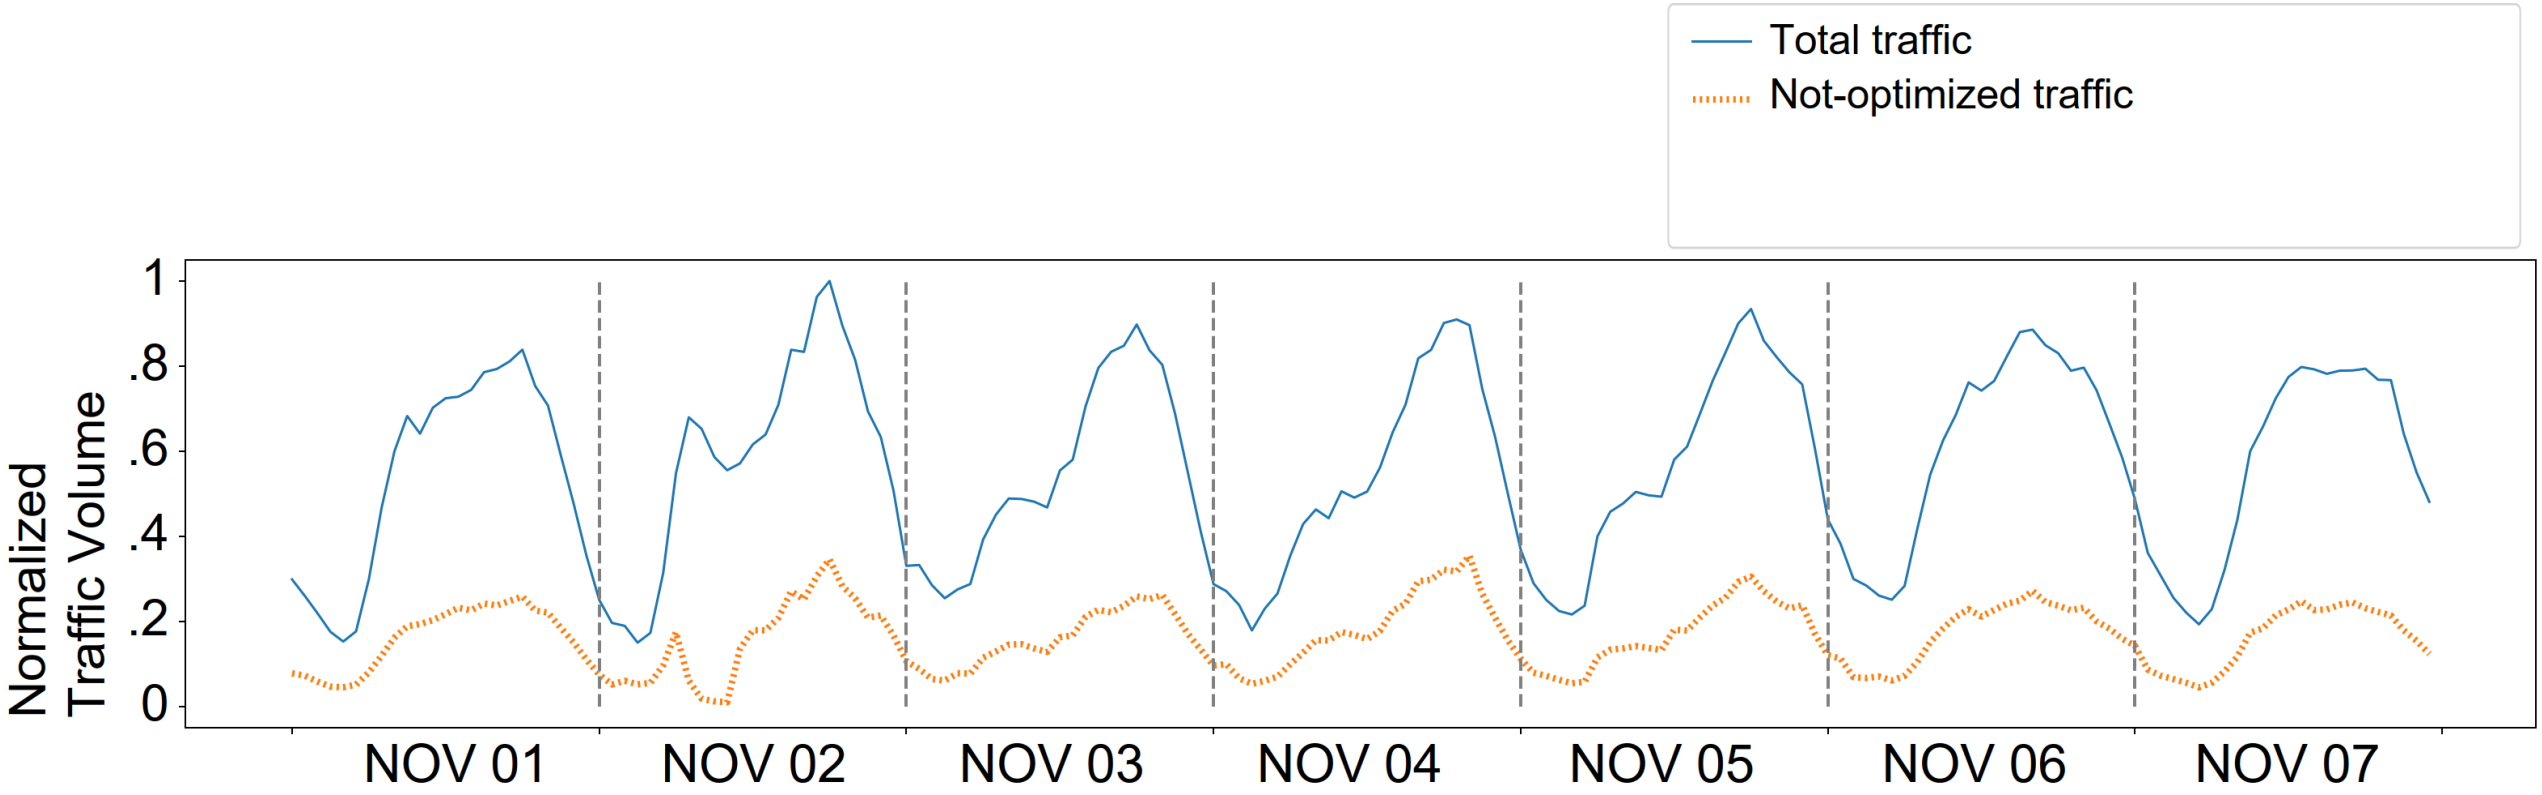

Small ISP - Hypergiant Traffic

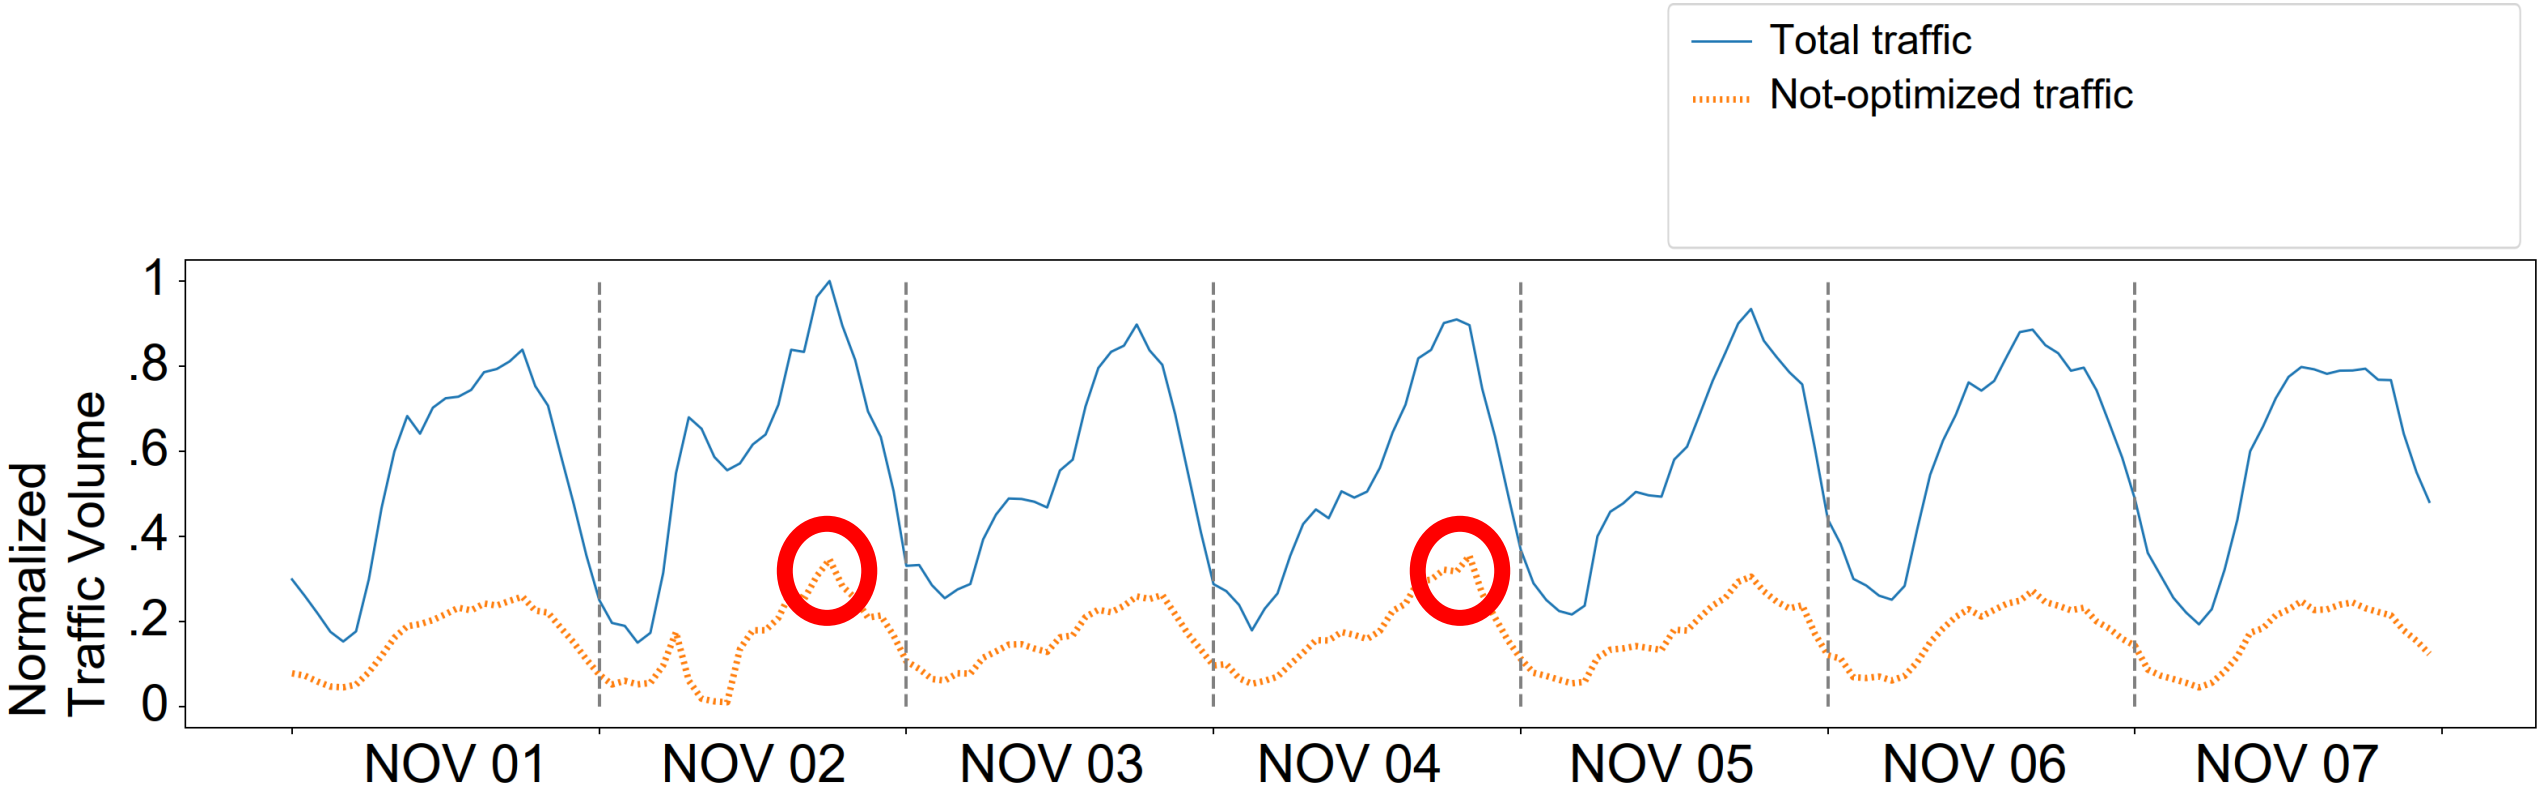

More than 20 European ISPs encounter similar problems!

Can we help the Hypergiants improve the server selection for not directly connected ISPs?

Can we reduce the 18% ?

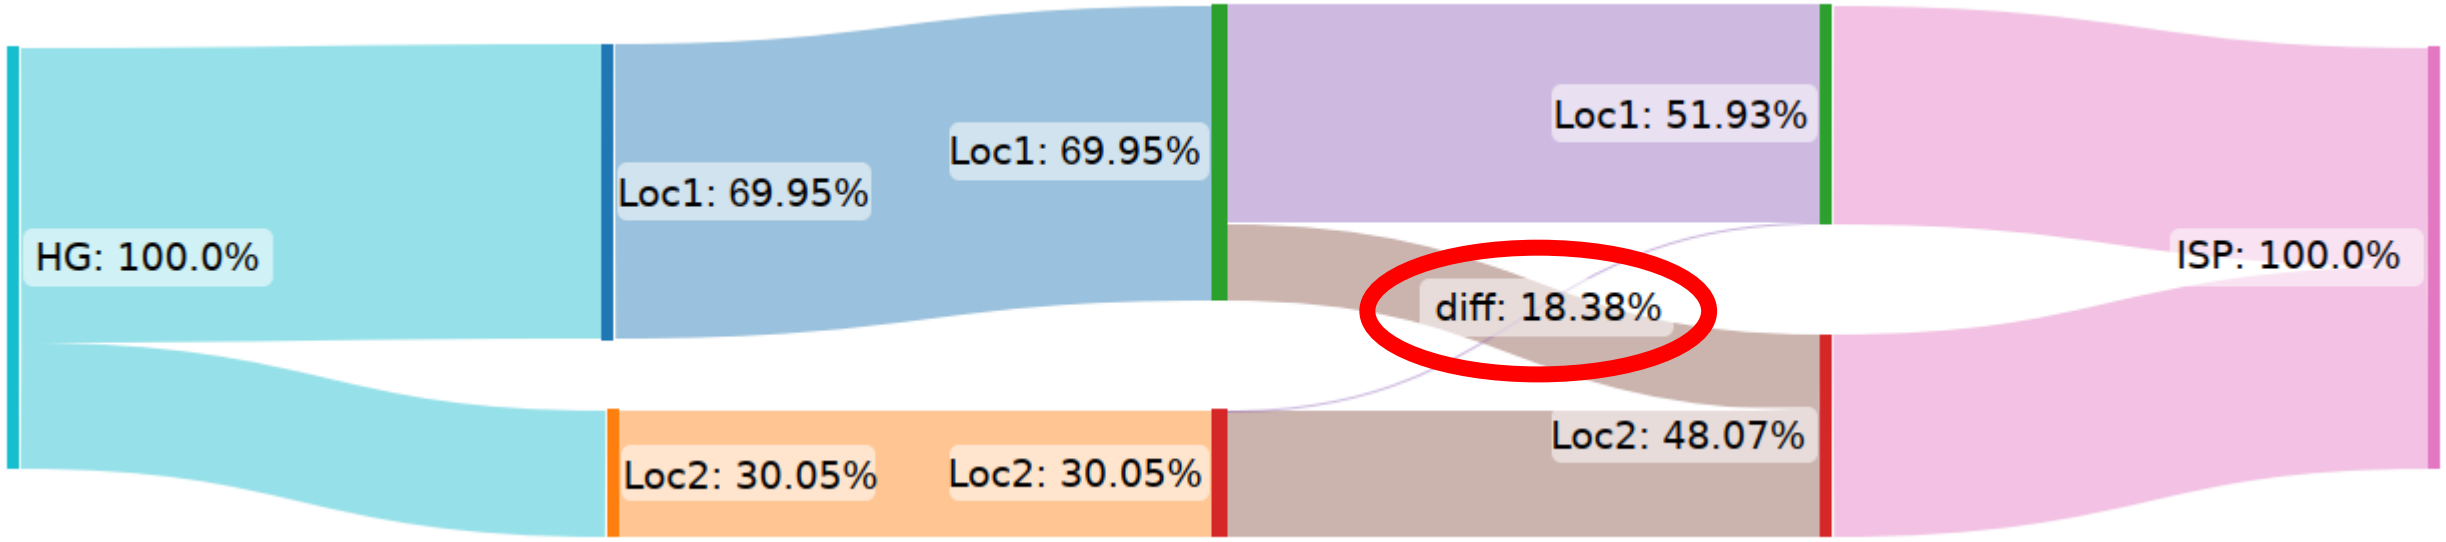

Hypergiant Transit AS Small European ISP

Can we reduce the 18% ?

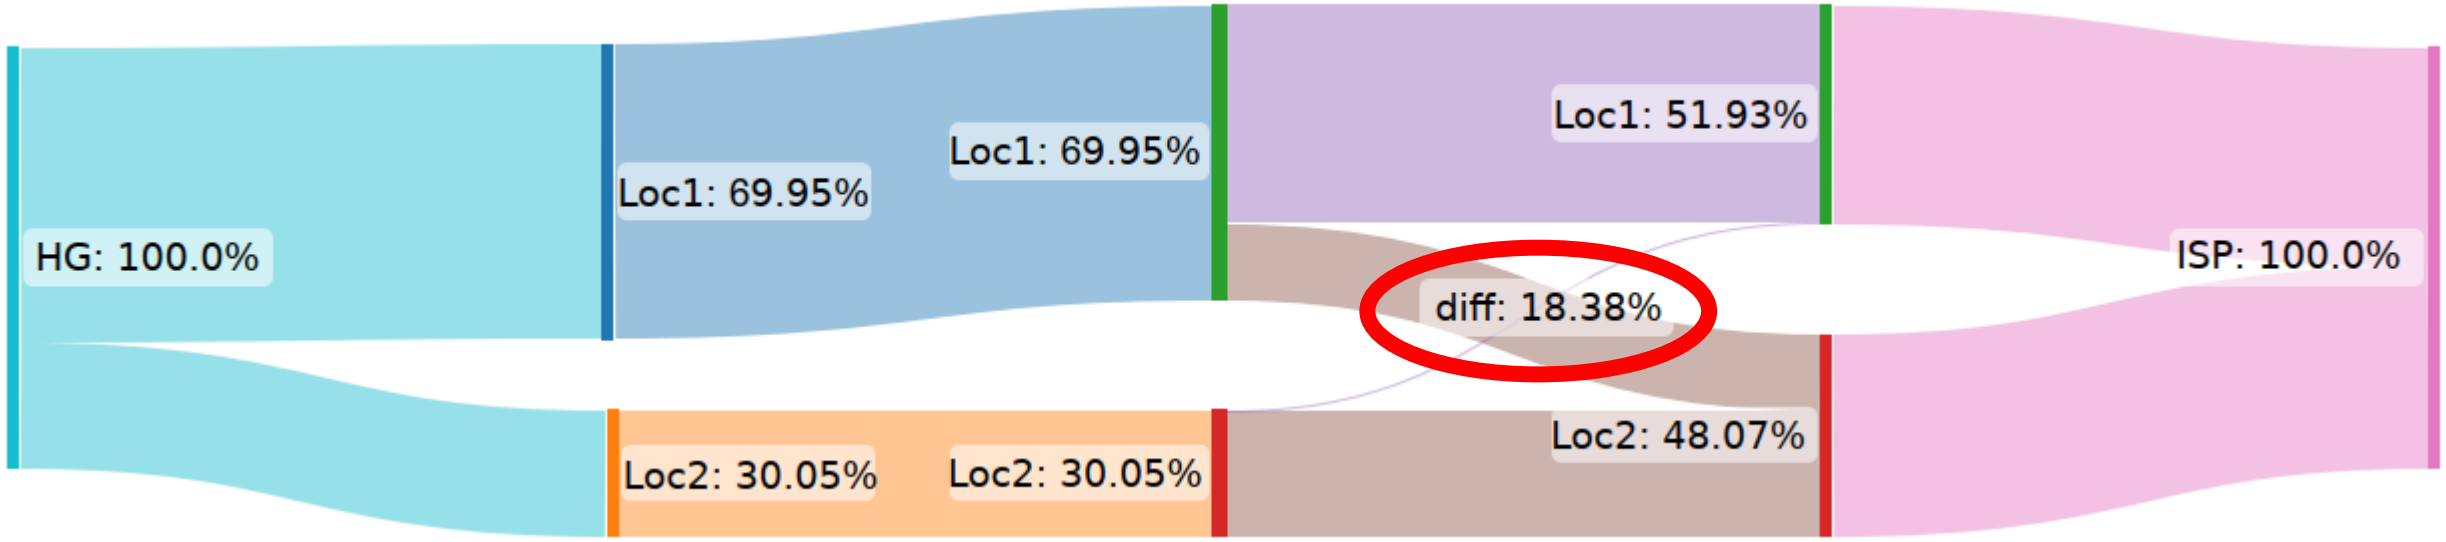

Hypergiant Transit AS Small European ISP

YES!

ISP-Hypergiant Collaboration

ISP sends additional information to the Hypergiant to improve server selection.

Collaboration

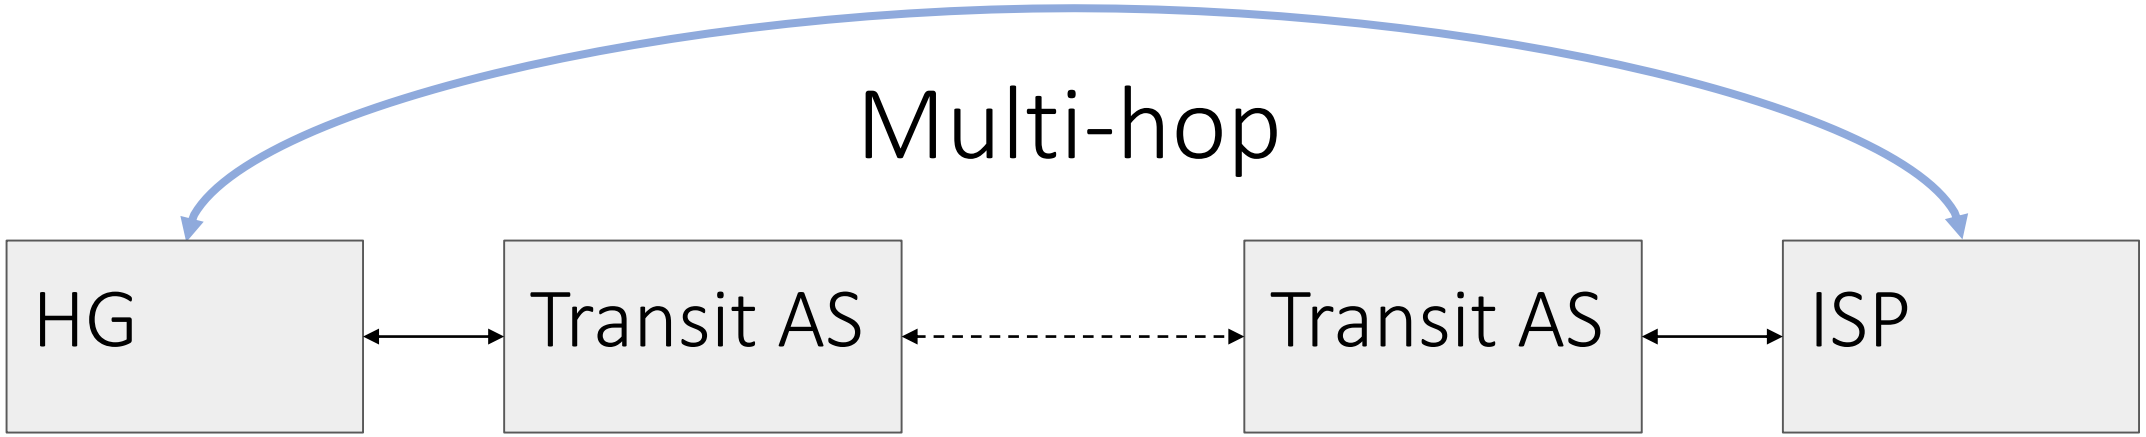

Collaboration

One+-hop

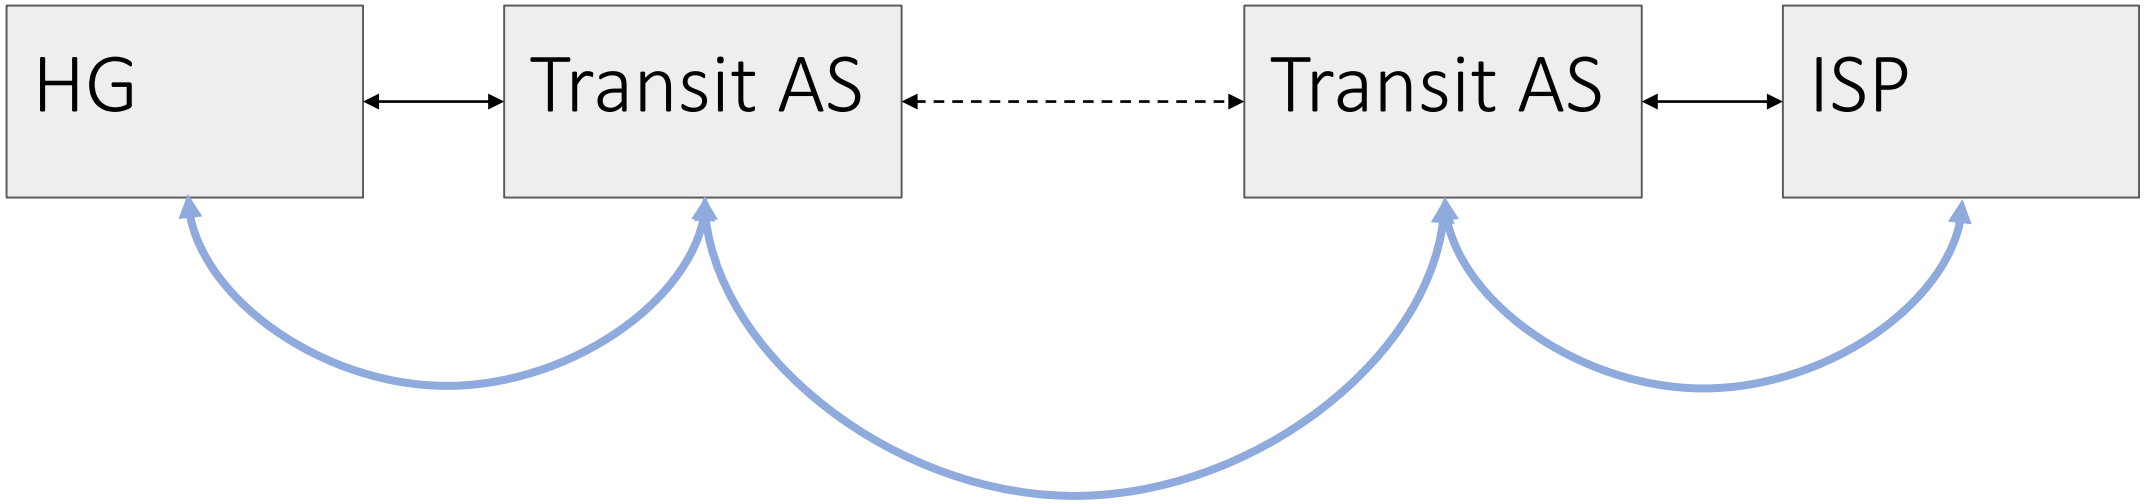

Multi-hop Collaboration

ISP sends a set of *key*:*value* pairs to the HG

Multi-hop Collaboration

ISP sends a set of *key*:*value* pairs to the HG

- "*key*" : IP Prefix
- "*value*" : [list of similar IP Prefixes]

Multi-hop Collaboration

ISP sends a set of *key*:*value* pairs to the HG

“*key*” : IP Prefix

“*value*” : [list of similar IP Prefixes]

Example: “IP Prefix A”:[“IP Prefix B, IP Prefix C”]

Select prefixes

- BGP announced prefixes
  (spoiler – not efficient)
- ISP DNS-Resolver working prefixes (*DNS-default*)
- /24 disaggregation - *If DNS ECS possible*

Benefits: One Hypergiant

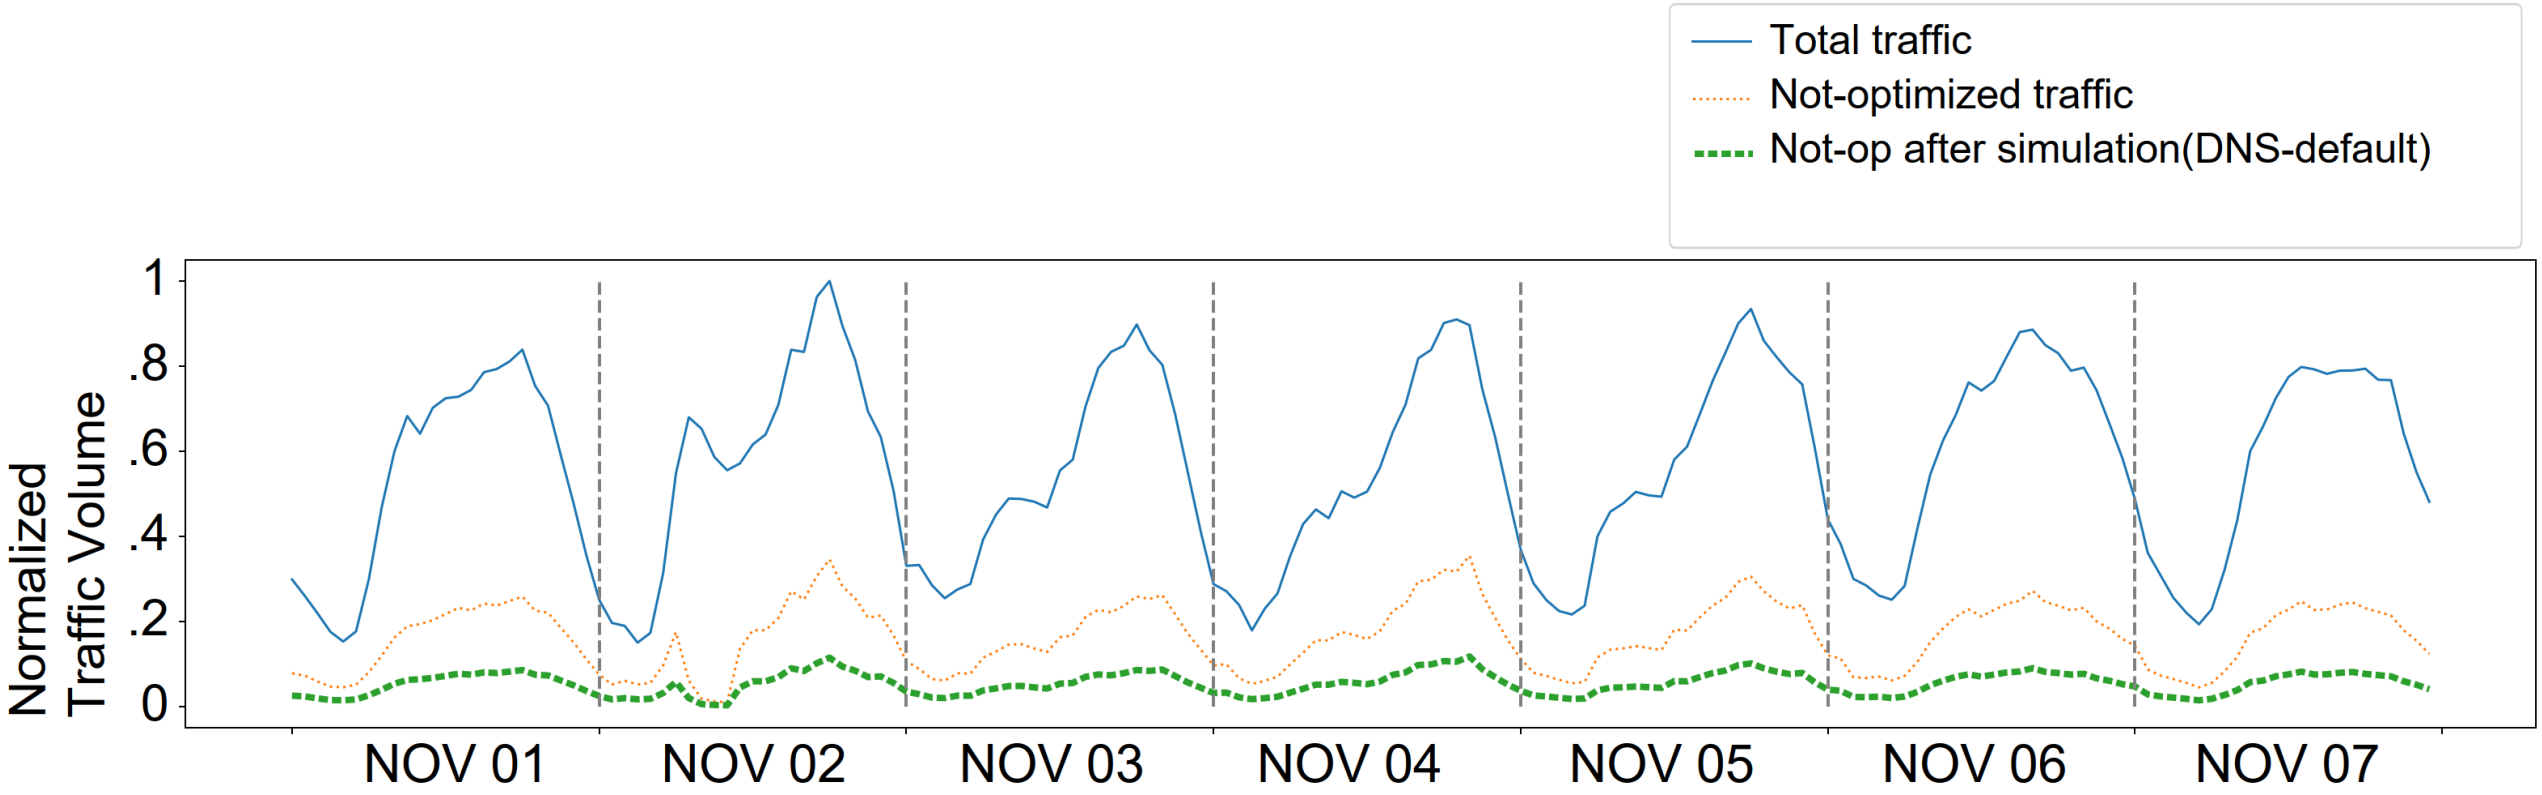

Benefits: One Hypergiant

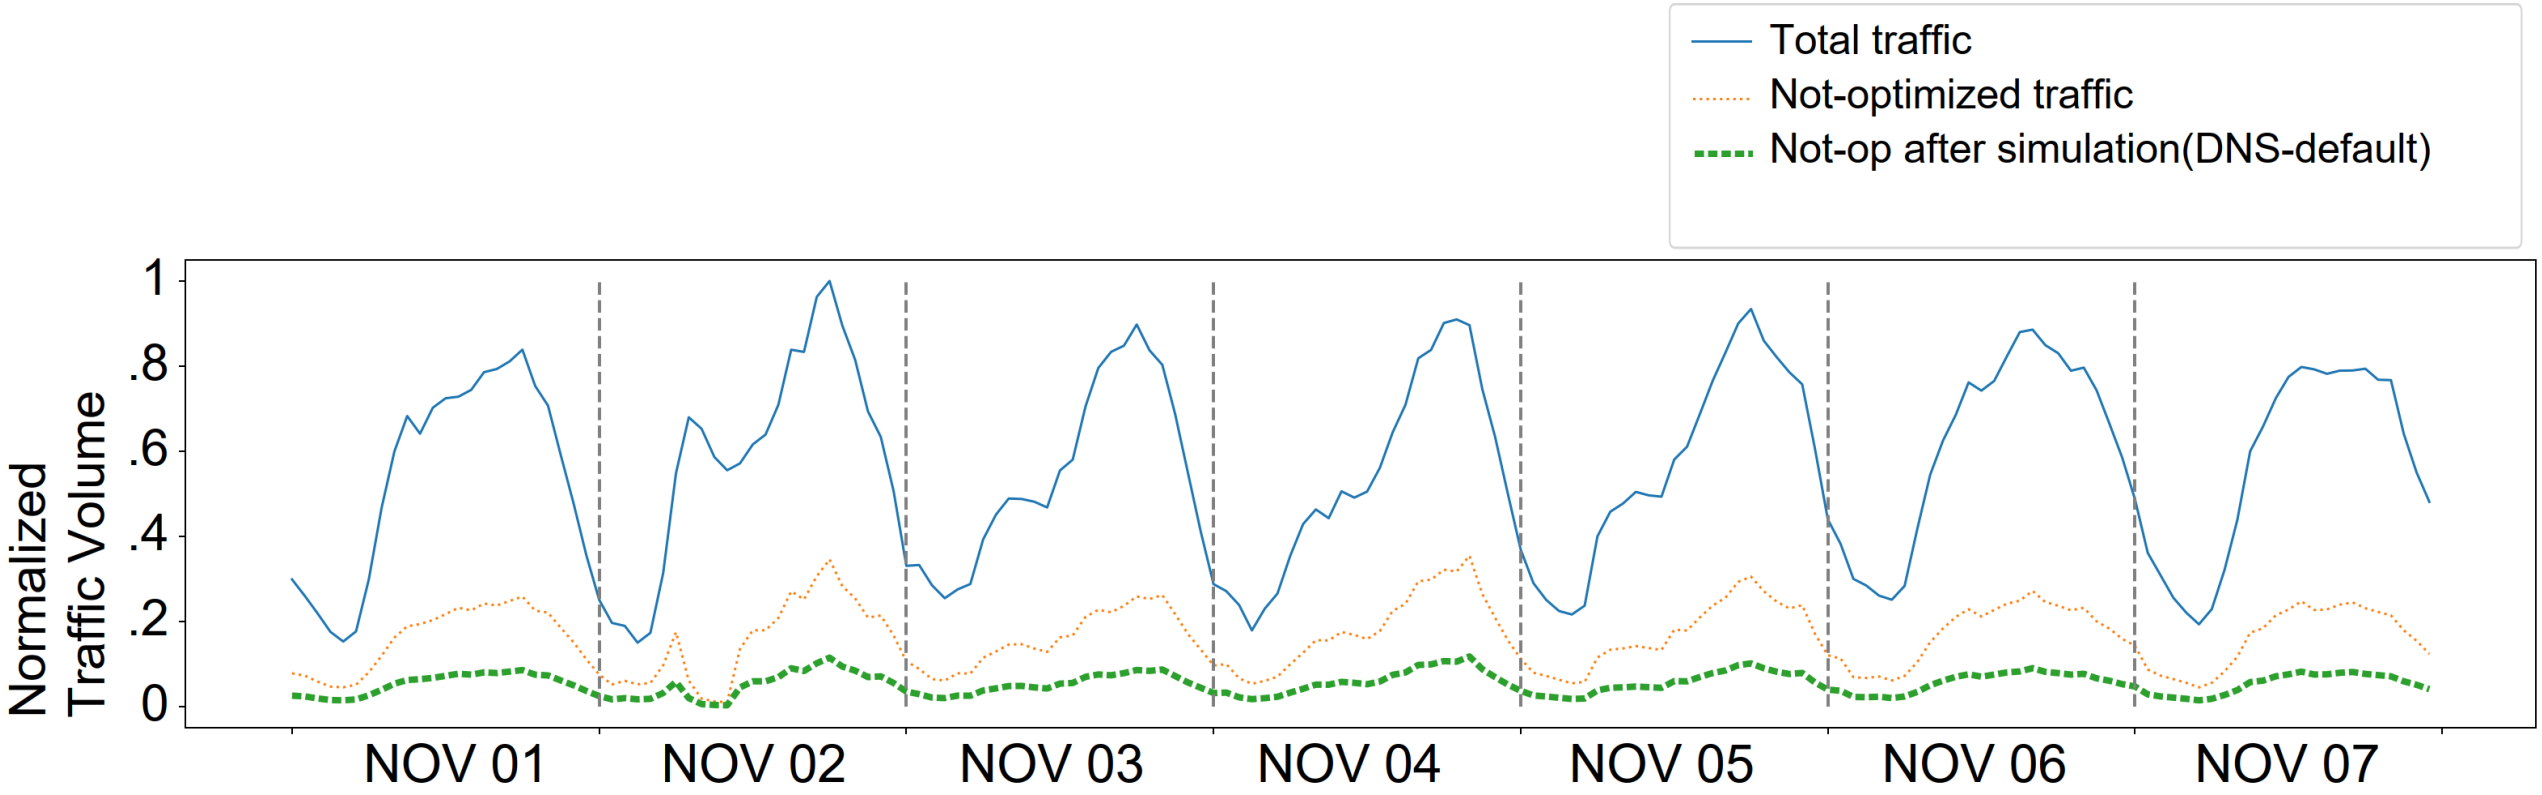

DNS-default ->from 18% to 1.3%

Benefits: One Hypergiant

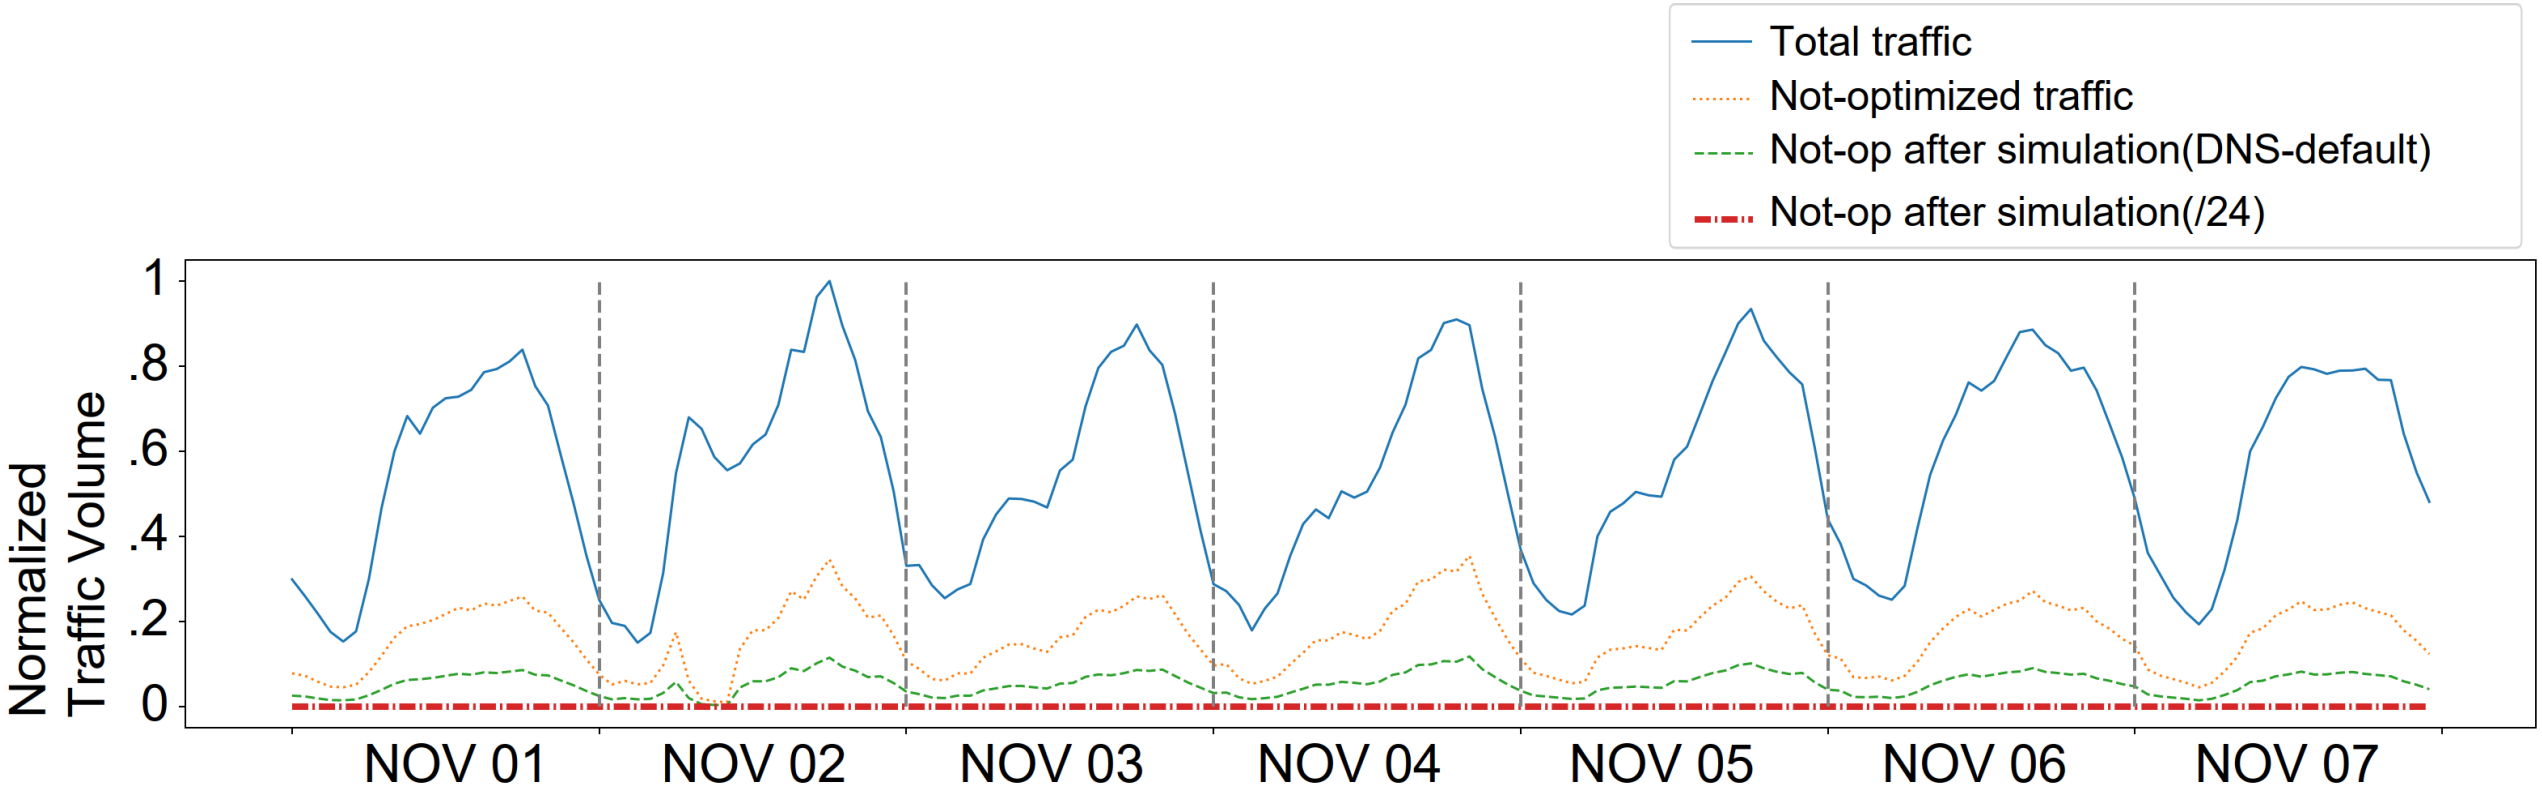

/24 -> Optimal traffic!

Benefits: Multiple Hypergiants

Hypergiant	Traffic %	Not-optimized %	Not-optimized % per own traffic share
HG1	31.93%	0.59%	1.86%
HG2	16.17%	2.97%	18.38%
HG3	8.15%	1.78%	21.90%
HG4	6.96%	3.21%	46.15%
HG5 *	4.46%	1.70%	38.10%
HG6	3.09%	1.07%	34.62%
HG7	2.62%	0.06%	2.27%
HG8	2.26%	0.24%	10.53%
HG9	2.26%	0.78%	34.21%
HG10 *	2.08%	0.75%	36.00%
HG11 *	2.08%	0.76%	37.00%
Others	17.95%	—	—
Total	**100%**	**13.91%**	

Benefits: Multiple Hypergiants

Hypergiant	Traffic %	Not-optimized %	Not-optimized % per own traffic share
HG1	31.93%	0.59%	1.86%
HG2	16.17%	2.97%	18.38%
HG3	8.15%	1.78%	21.90%
HG4	6.96%	3.21%	46.15%
HG5 *	4.46%	1.70%	38.10%
HG6	3.09%	1.07%	34.62%
HG7	2.62%	0.06%	2.27%
HG8	2.26%	0.24%	10.53%
HG9	2.26%	0.78%	34.21%
HG10 *	2.08%	0.75%	36.00%
HG11 *	2.08%	0.76%	37.00%
Others	17.95%	—	—
Total	**100%**	**13.91%**	

Benefits: Multiple Hypergiants

Hypergiant	Traffic %	Not-optimized %	Not-optimized % per own traffic share
HG1	31.93%	0.59%	1.86%
HG2	16.17%	2.97%	18.38%
HG3	8.15%	1.78%	21.90%
HG4	6.96%	3.21%	46.15%
HG5 *	4.46%	1.70%	38.10%
HG6	3.09%	1.07%	34.62%
HG7	2.62%	0.06%	2.27%
HG8	2.26%	0.24%	10.53%
HG9	2.26%	0.78%	34.21%
HG10 *	2.08%	0.75%	36.00%
HG11 *	2.08%	0.76%	37.00%
Others	17.95%	—	—
Total	**100%**	**13.91%**	

Benefits: Multiple Hypergiants

Hypergiant	Traffic %	Not-optimized %	Not-optimized % per own traffic share
HG1	31.93%	0.59%	1.86%
HG2	16.17%	2.97%	18.38%
HG3	8.15%	1.78%	21.90%
HG4	6.96%	3.21%	46.15%
HG5 *	4.46%	1.70%	38.10%
HG6	3.09%	1.07%	34.62%
HG7	2.62%	0.06%	2.27%
HG8	2.26%	0.24%	10.53%
HG9	2.26%	0.78%	34.21%
HG10 *	2.08%	0.75%	36.00%
HG11 *	2.08%	0.76%	37.00%
Others	17.95%	—	—
Total	**100%**	**13.91%**	

Benefits: Multiple Hypergiants

Hypergiant	Traffic %	Not-optimized %	Not-optimized % per own traffic share
HG1	31.93%	0.59%	1.86%
HG2	16.17%	2.97%	18.38%
HG3	8.15%	1.78%	21.90%
HG4	6.96%	3.21%	46.15%
HG5 *	4.46%	1.70%	38.10%
HG6	3.09%	1.07%	34.62%
HG7	2.62%	0.06%	2.27%
HG8	2.26%	0.24%	10.53%
HG9	2.26%	0.78%	34.21%
HG10 *	2.08%	0.75%	36.00%
HG11 *	2.08%	0.76%	37.00%
Others	17.95%	—	—
Total	**100%**	**13.91%**	

Benefits: Multiple Hypergiants

Hypergiant	Traffic %	Not-optimized %	Not-optimized % per own traffic share
HG1	31.93%	0.59%	1.86%
HG2	16.17%	2.97%	18.38%
HG3	8.15%	1.78%	21.90%
HG4	6.96%	3.21%	46.15%
HG5 *	4.46%	1.70%	38.10%
HG6	3.09%	1.07%	34.62%
HG7	2.62%	0.06%	2.27%
HG8	2.26%	0.24%	10.53%
HG9	2.26%	0.78%	34.21%
HG10 *	2.08%	0.75%	36.00%
HG11 *	2.08%	0.76%	37.00%
Others	17.95%	—	—
Total	**100%**	**13.91%**	

Benefits: Multiple Hypergiants

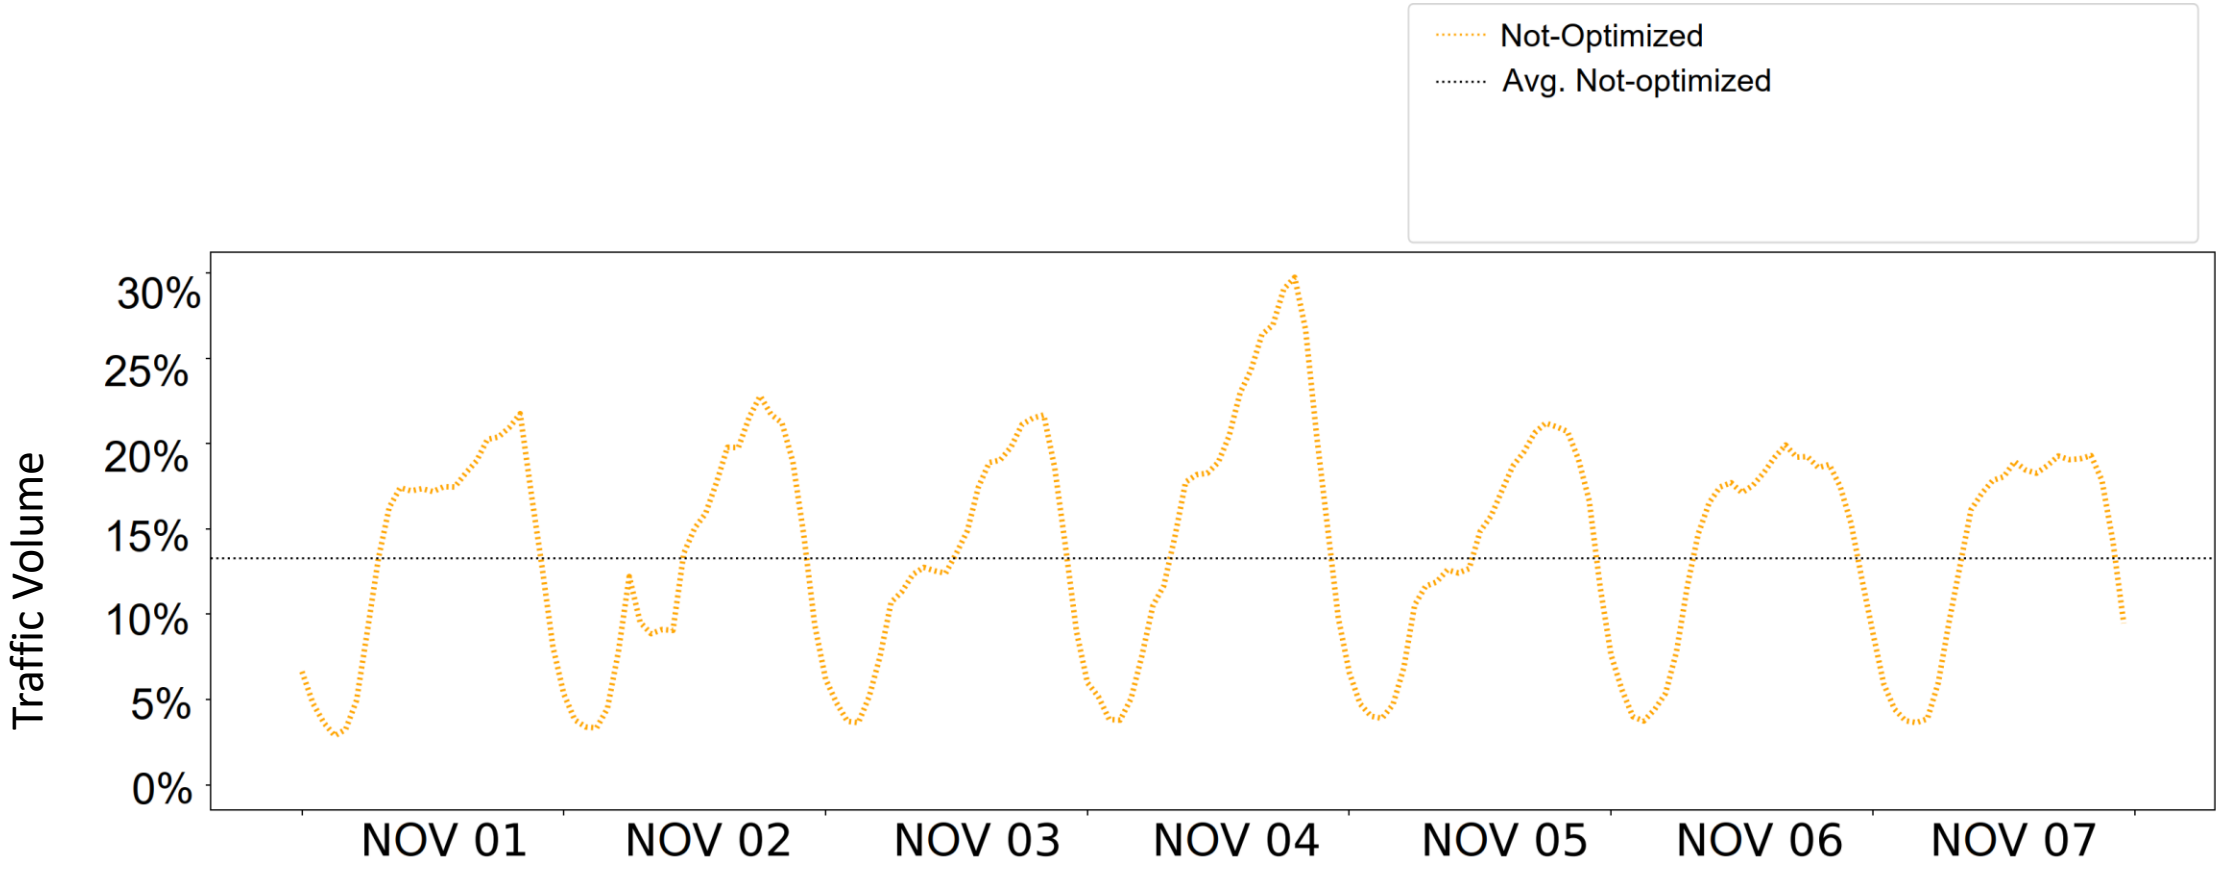

14% not-optimized traffic

Benefits: Multiple Hypergiants

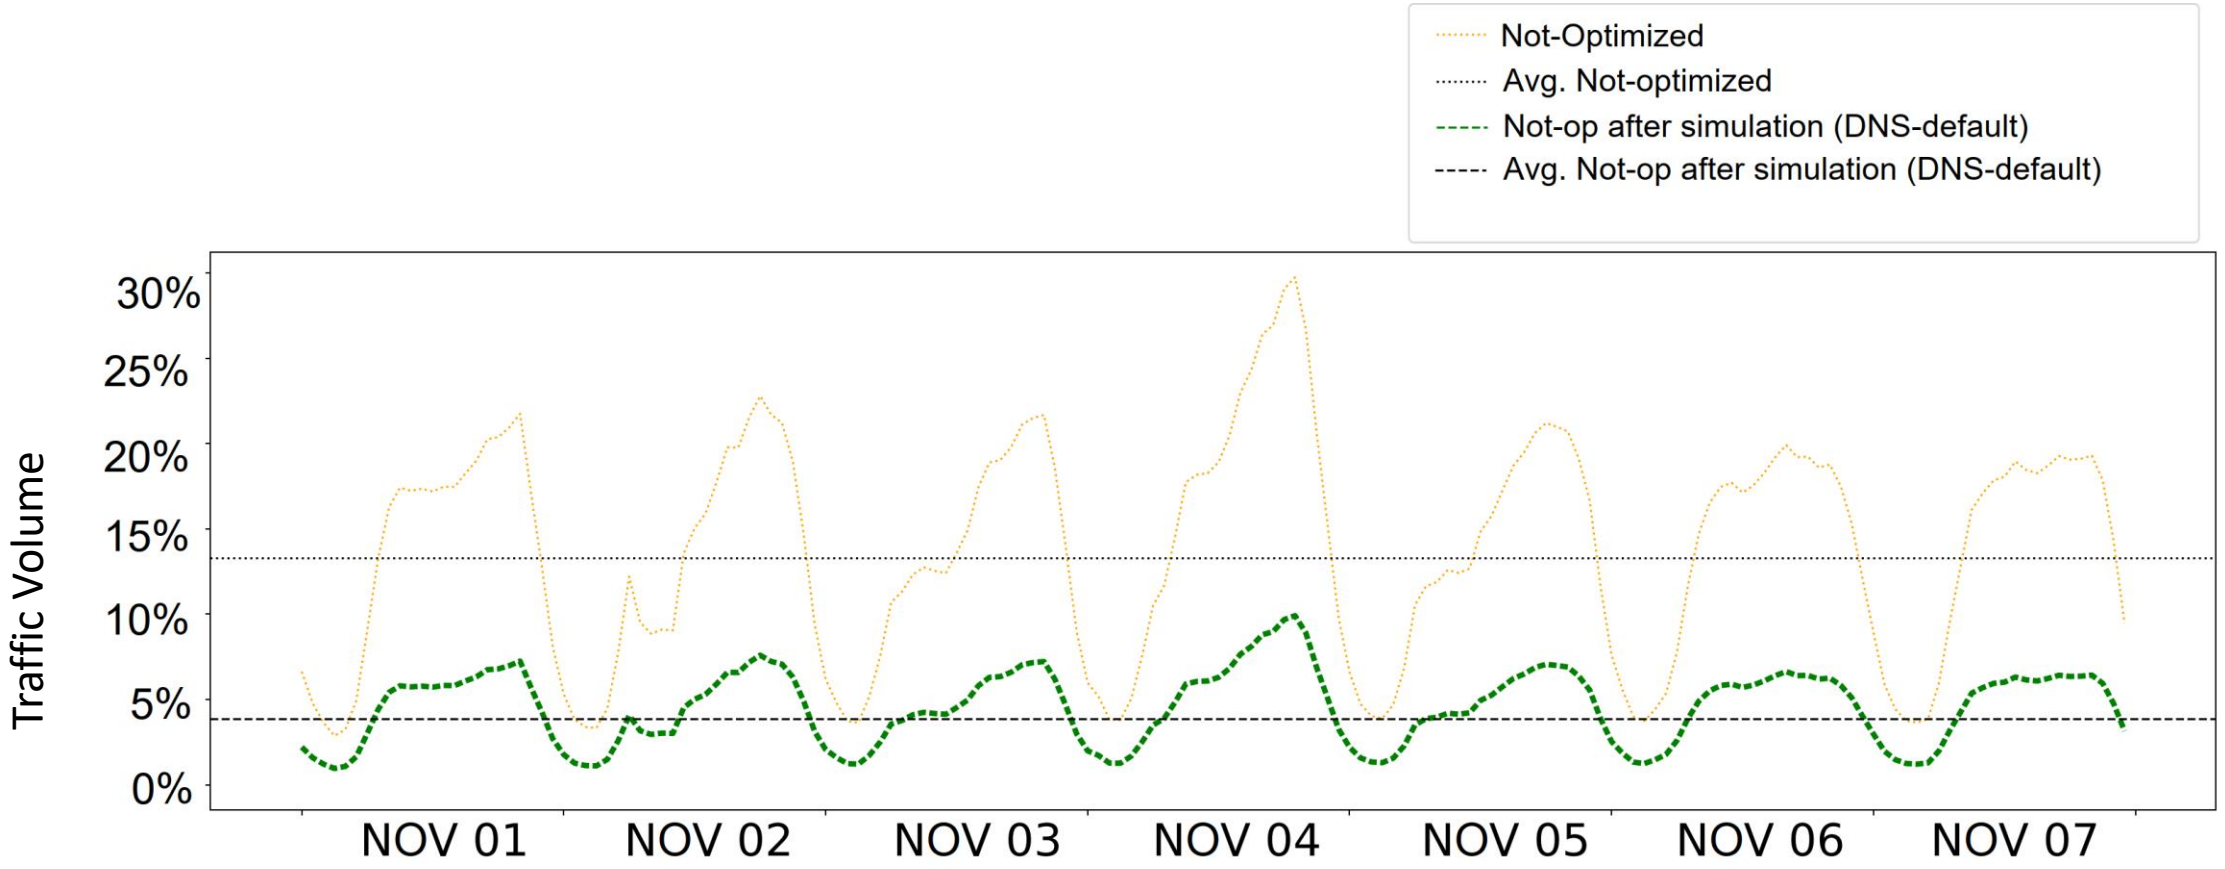

DNS-default -> 4% not-optimized traffic

Benefits: Multiple Hypergiants

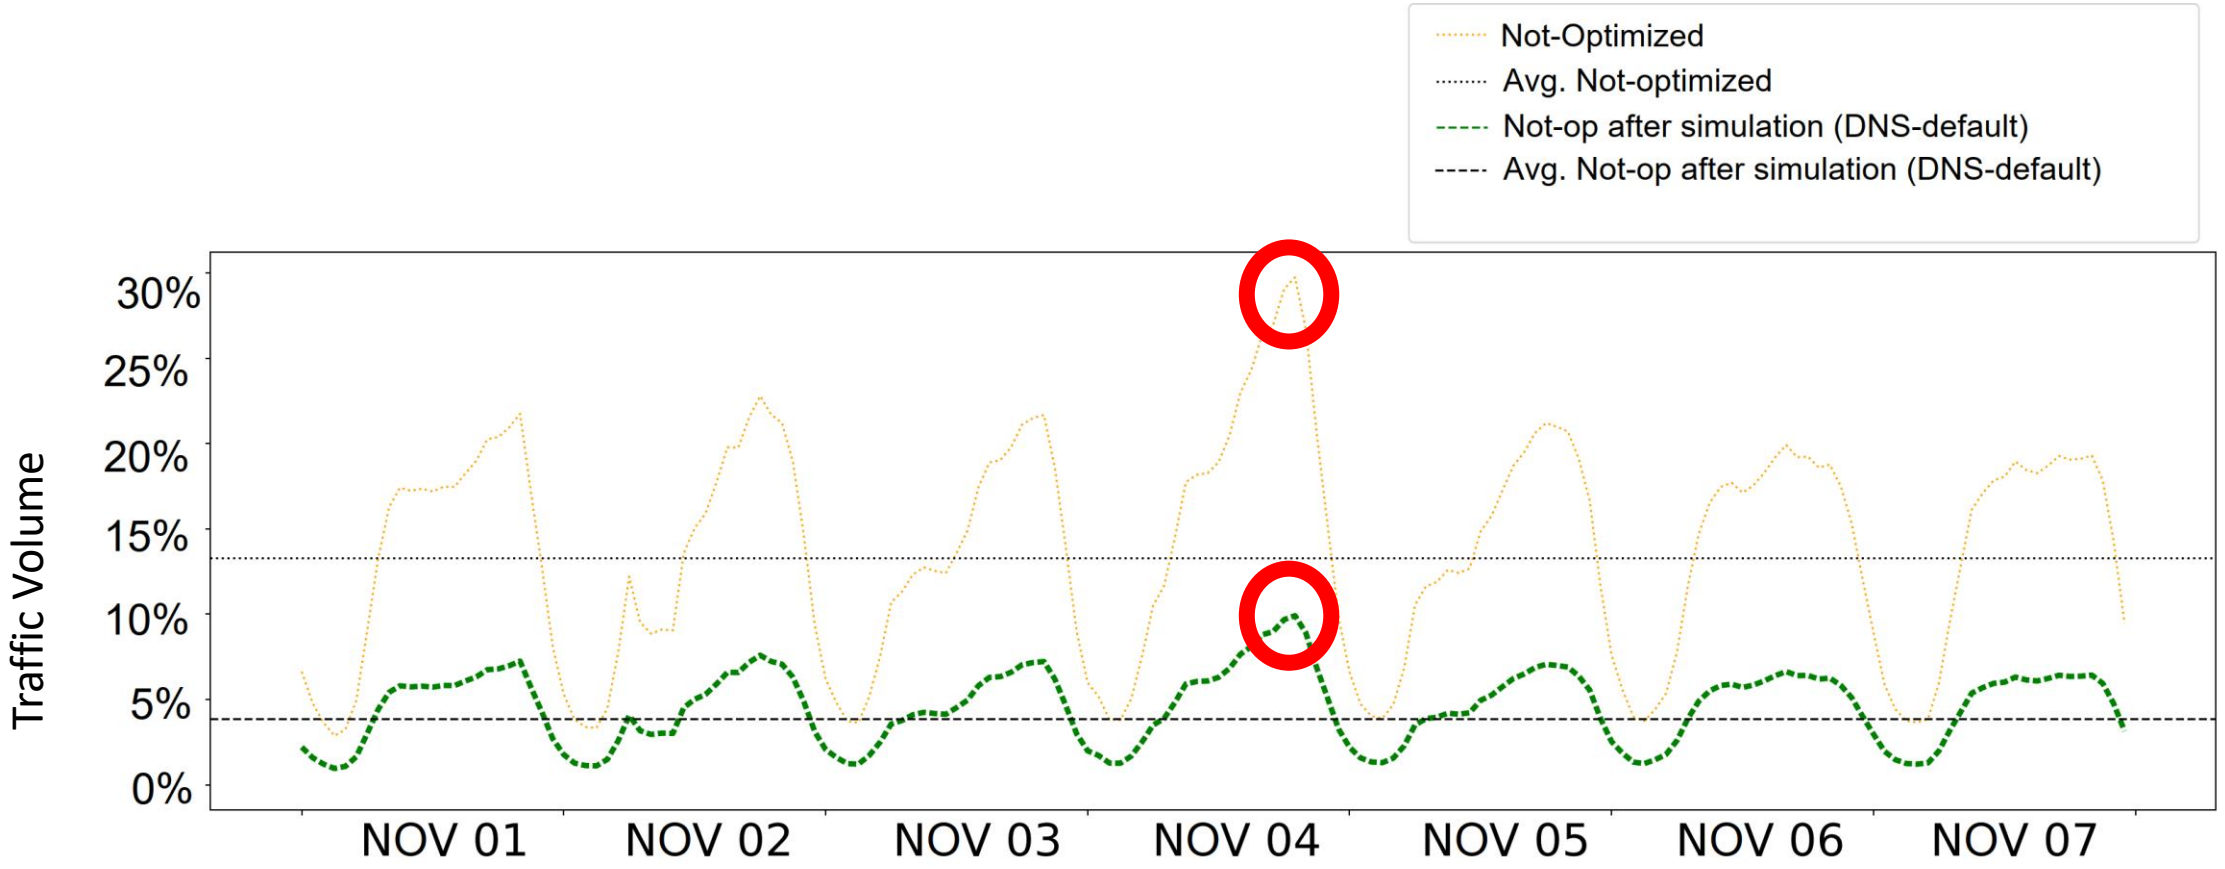

DNS-default -> 4% not-optimized traffic

Benefits: Multiple Hypergiants

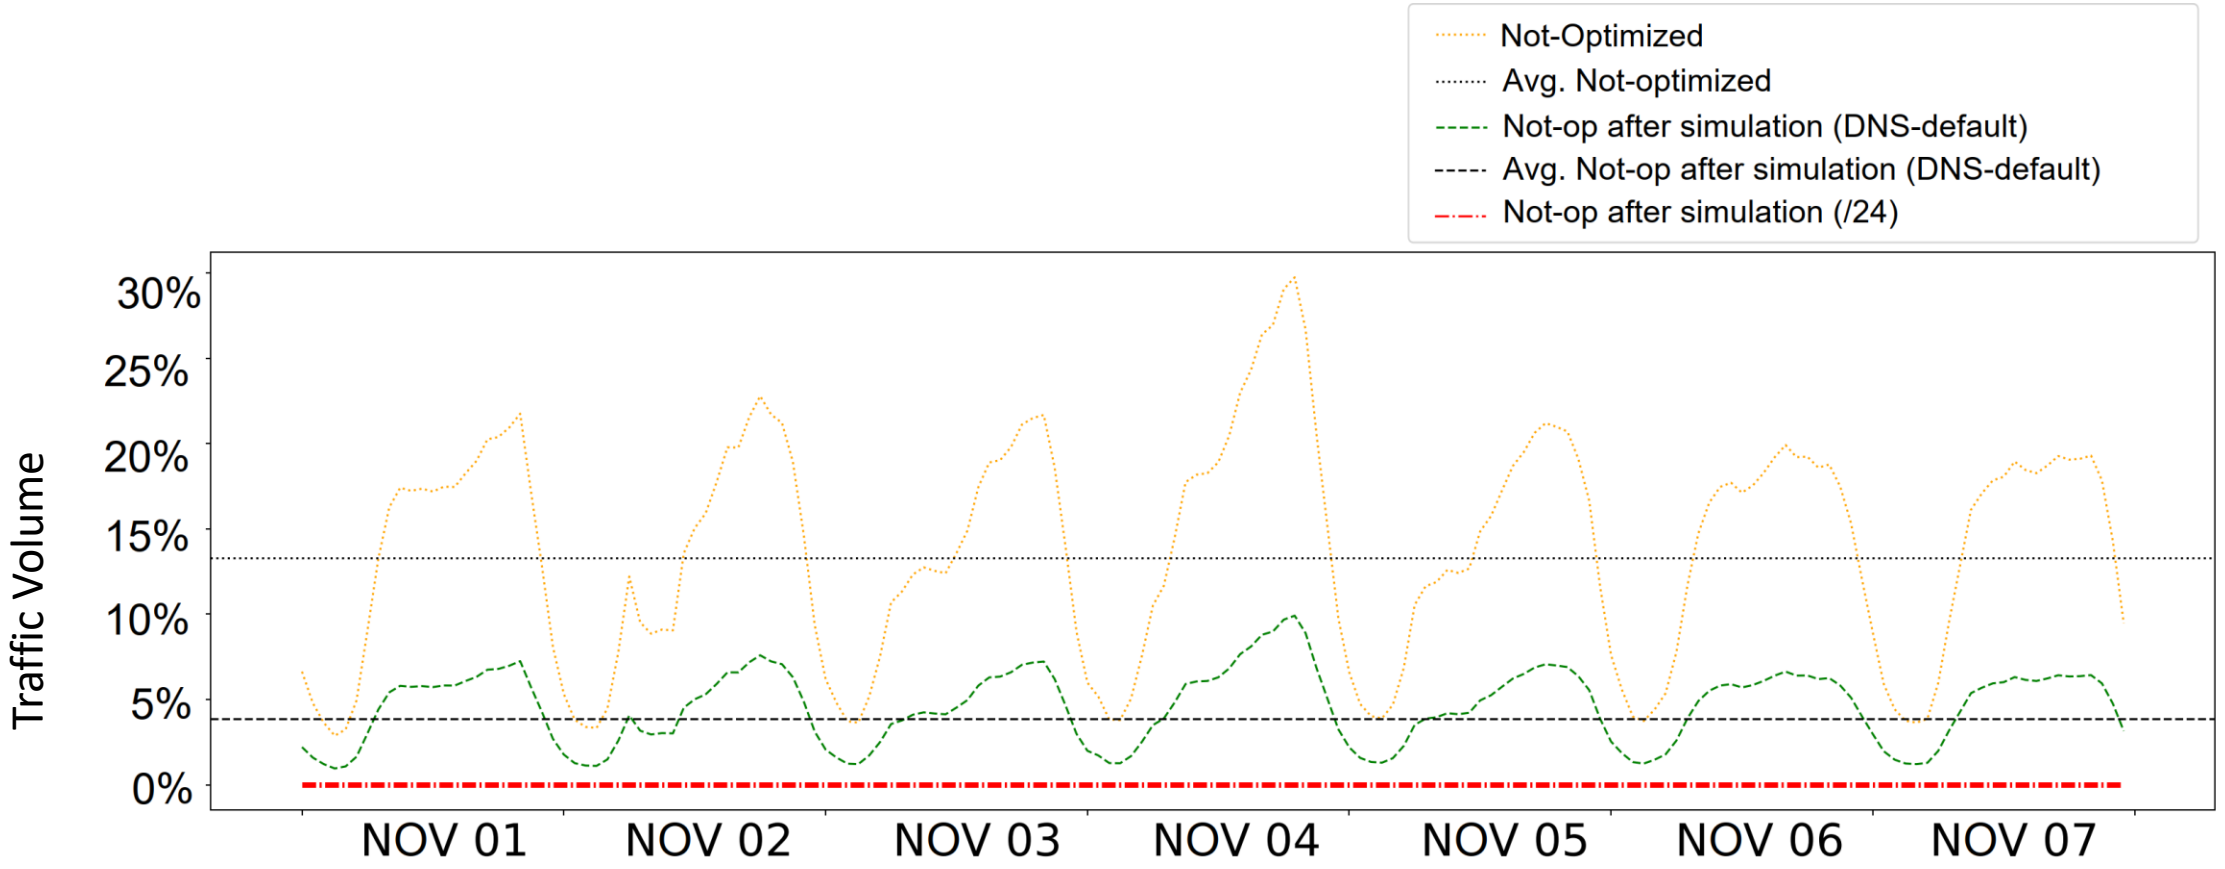

/24 -> 0% not-optimized traffic

Conclusion

- It is possible to improve server selection even if there is no direct peering between ISP and Hypergiant.
- We show, using real ISP data, that the system can improve non-optimized traffic up to **10%**.
- Results also show that for some Hypergiants, **up to 46%** of their traffic is delivered via non-optimal interconnection.
- More than **40K networks** can potentially benefit.

mpii max planck institut informatik

BENOCS YEARS

TU Delft

Benefits: Multiple Hypergiants

Hypergiant	Original Not-opt	Not-opt after Simulation	
	BGP ann. (#prf.)	'/24' (#prf.)	DNS-default (#prf.)
HG1	1.86% (8)	0% (371)	1.86% (69)
HG2	18.38% (8)	0% (273)	1.37% (70)
HG3	21.90% (8)	0% (268)	11.44% (62)
HG4	42.80% (8)	0% (182)	8.93% (40)
HG6	34.62% (8)	0% (145)	15.44% (28)
HG7	2.27% (8)	0% (144)	2.27% (25)
HG8	10.53% (8)	0% (138)	7.62% (24)
HG9	34.21% (8)	0% (132)	6.21% (24)