Delivering
Application-Layer Traffic Optimization Services based on Public Routing Data at IXPs

Danny Alex Lachos Perez

Orientador: Prof. Dr. Christian Rodolfo Esteve Rothenberg

INTRIG
INFORMATION & NETWORKING TECHNOLOGIES RESEARCH & INNOVATION GROUP

FEEC
Faculdade de Engenharia Elétrica e de Computação
UNICAMP

Agenda

1. Introduction
2. Background
3. Related work
4. Objectives
5. Design of ALTO-as-a-Service (AaaS)
6. Prototype Implementation
7. Experimental Evaluation
8. Future Work /Conclusions

Agenda

1. Introduction
2. Background
3. Related work
4. Objectives
5. Design of ALTO-as-a-Service (AaaS)
6. Prototype Implementation
7. Experimental Evaluation
8. Future Work / Conclusions

Distributed Applications

File sharing applications, Content Delivery Networks (CDNs), real-time communication, among others, use a significant amount of network resource to connect nodes across the Internet and transfer a large amount of data.

Distributed Applications

- ✓ Topological information used only has a localized view provided by the ISPs.
- ✓ Without knowledge of the underlying network topology.
- ✓ Selection of resources are provided randomly.
- ✓ Impacting both applications and networking infrastructure.

Introduction

MOTIVATION

Application-Layer Traffic Optimization (ALTO)

ISP => Focus of ALTO implementations

Internet eXchange Point (IXP)

Third-parties => ALTO Information <=ALTO Protocol

Agenda

1. Introduction
2. **Background**
3. Related work
4. Objectives
5. Design of ALTO-as-a-Service (AaaS)
6. AaaS Prototype
7. Experimental Evaluation
8. Future Work / Conclusions

Background :
Application-Layer Traffic Optimization (ALTO)

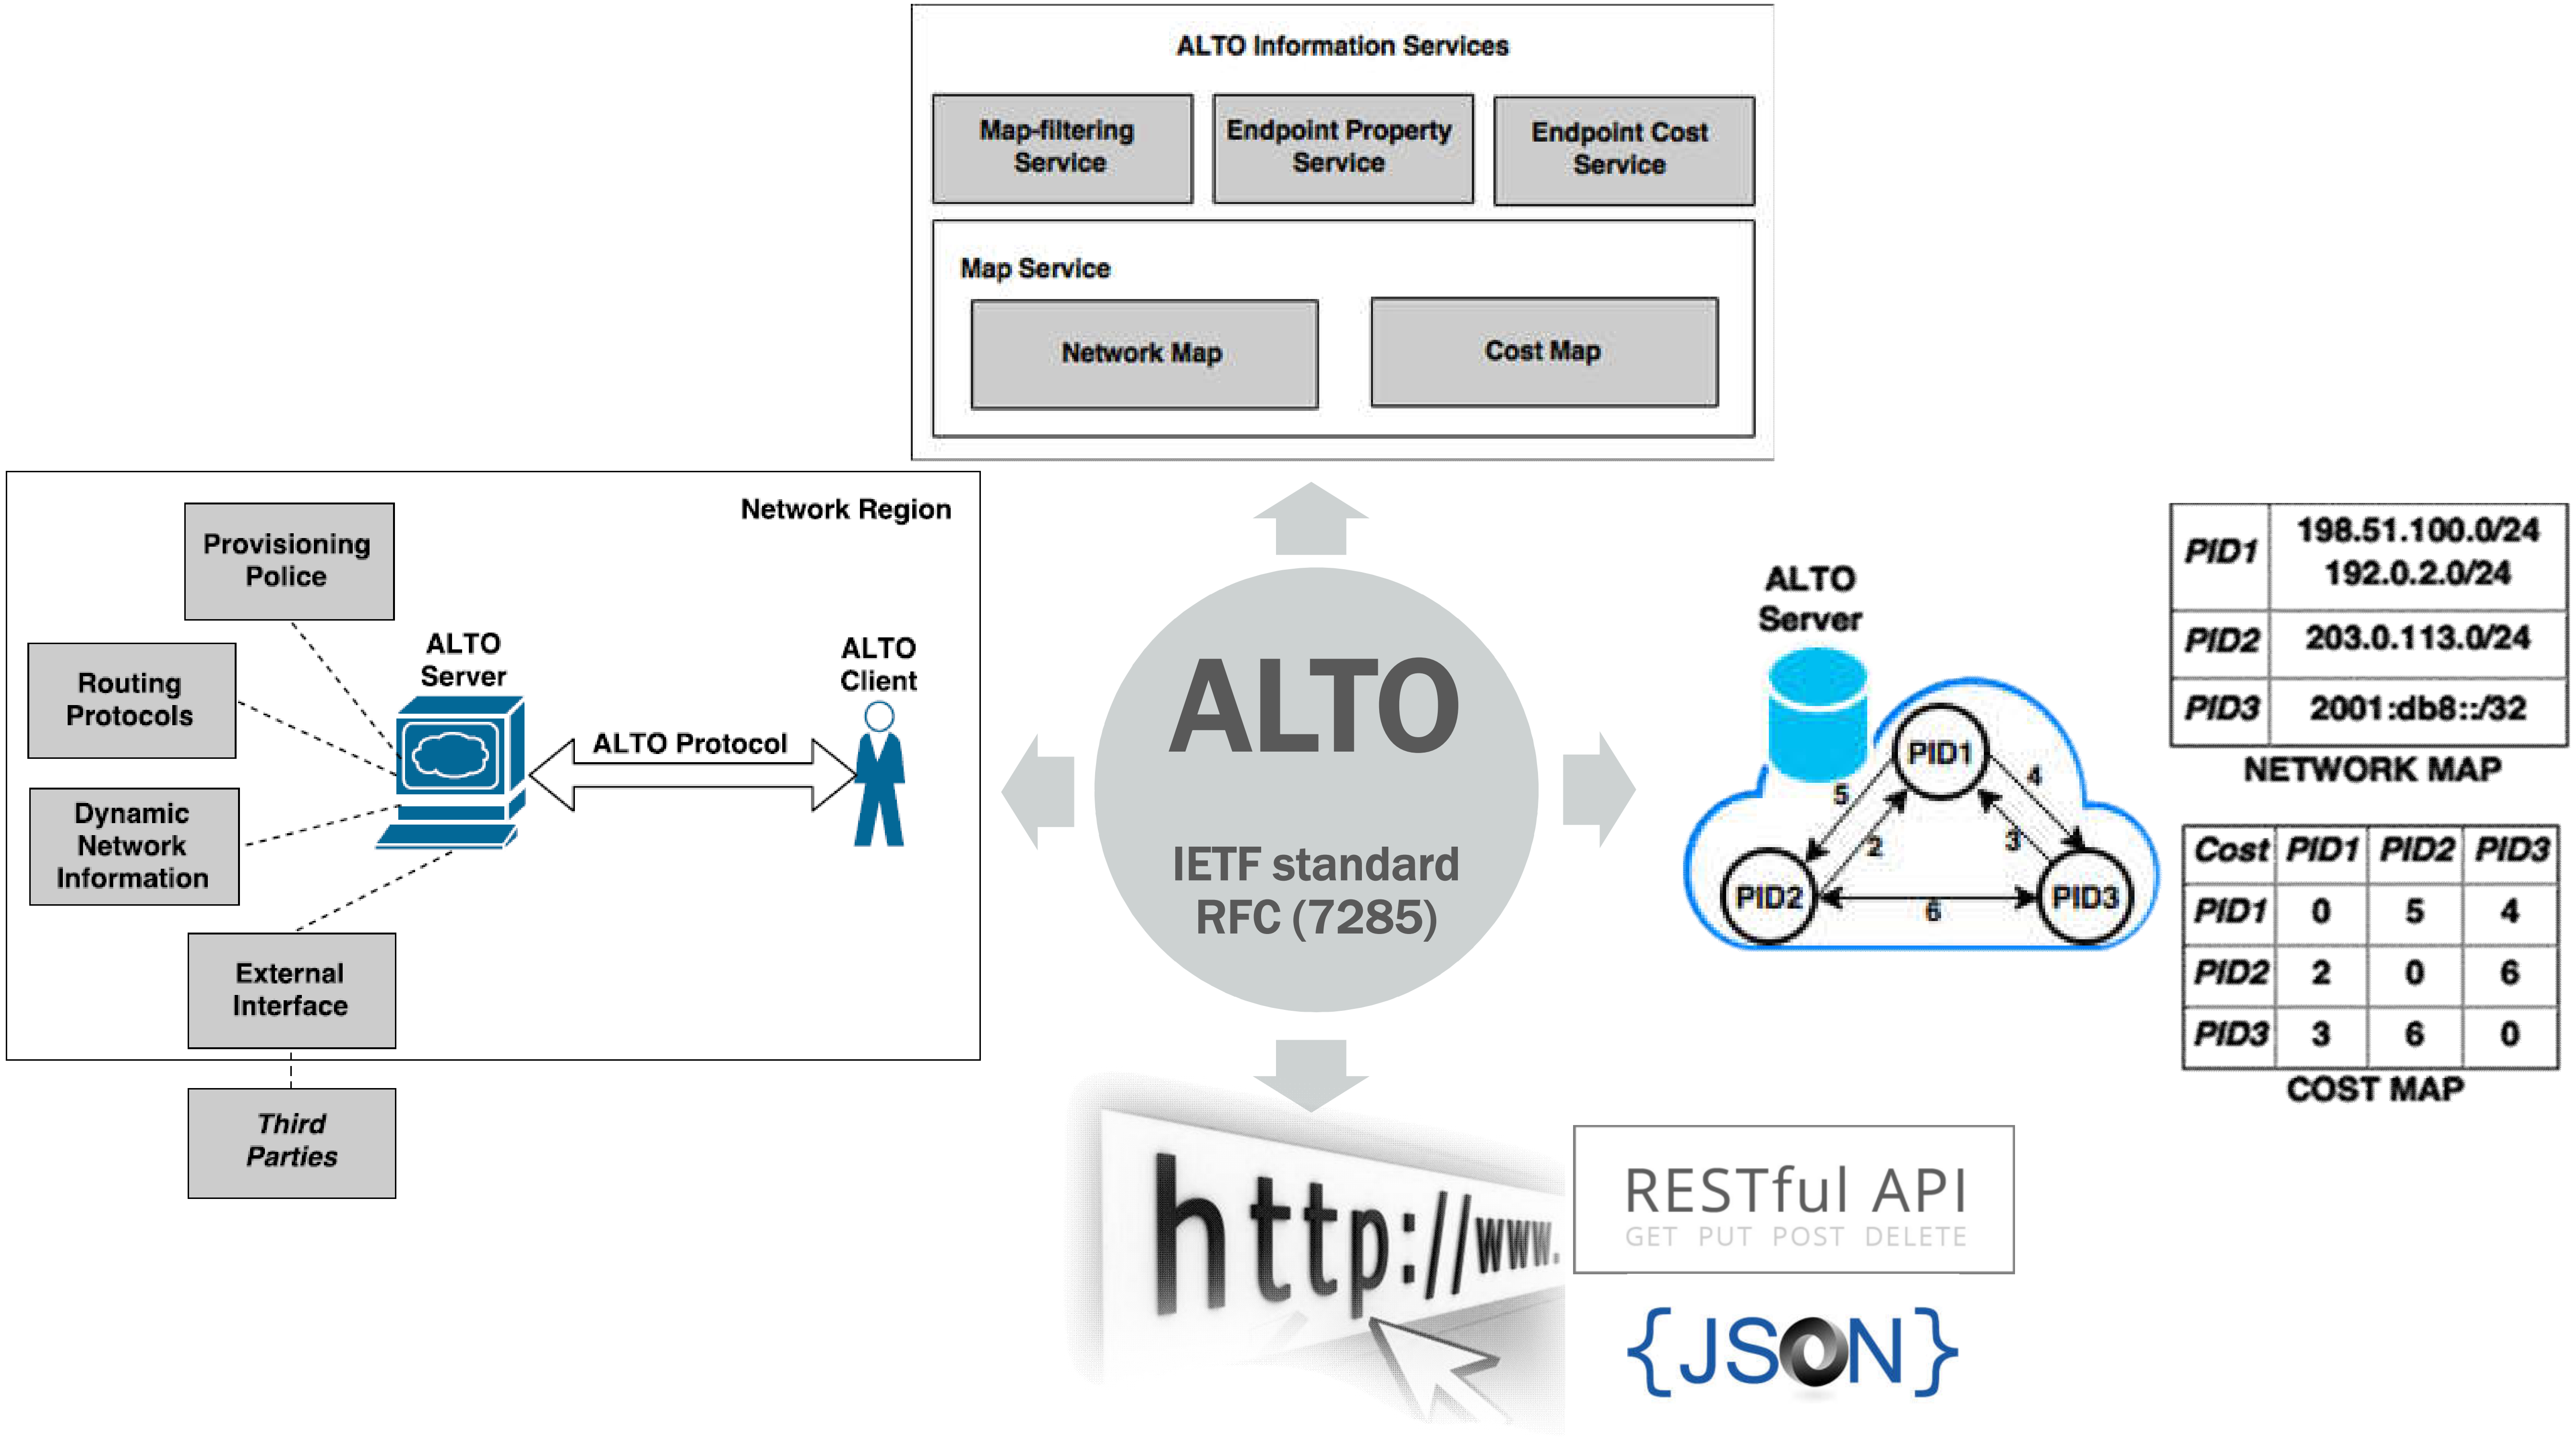

PID1	198.51.100.0/24 192.0.2.0/24
PID2	203.0.113.0/24
PID3	2001:db8::/32

NETWORK MAP

Cost	PID1	PID2	PID3
PID1	0	5	4
PID2	2	0	6
PID3	3	6	0

COST MAP

Background :

Internet eXchange Point (IXP)

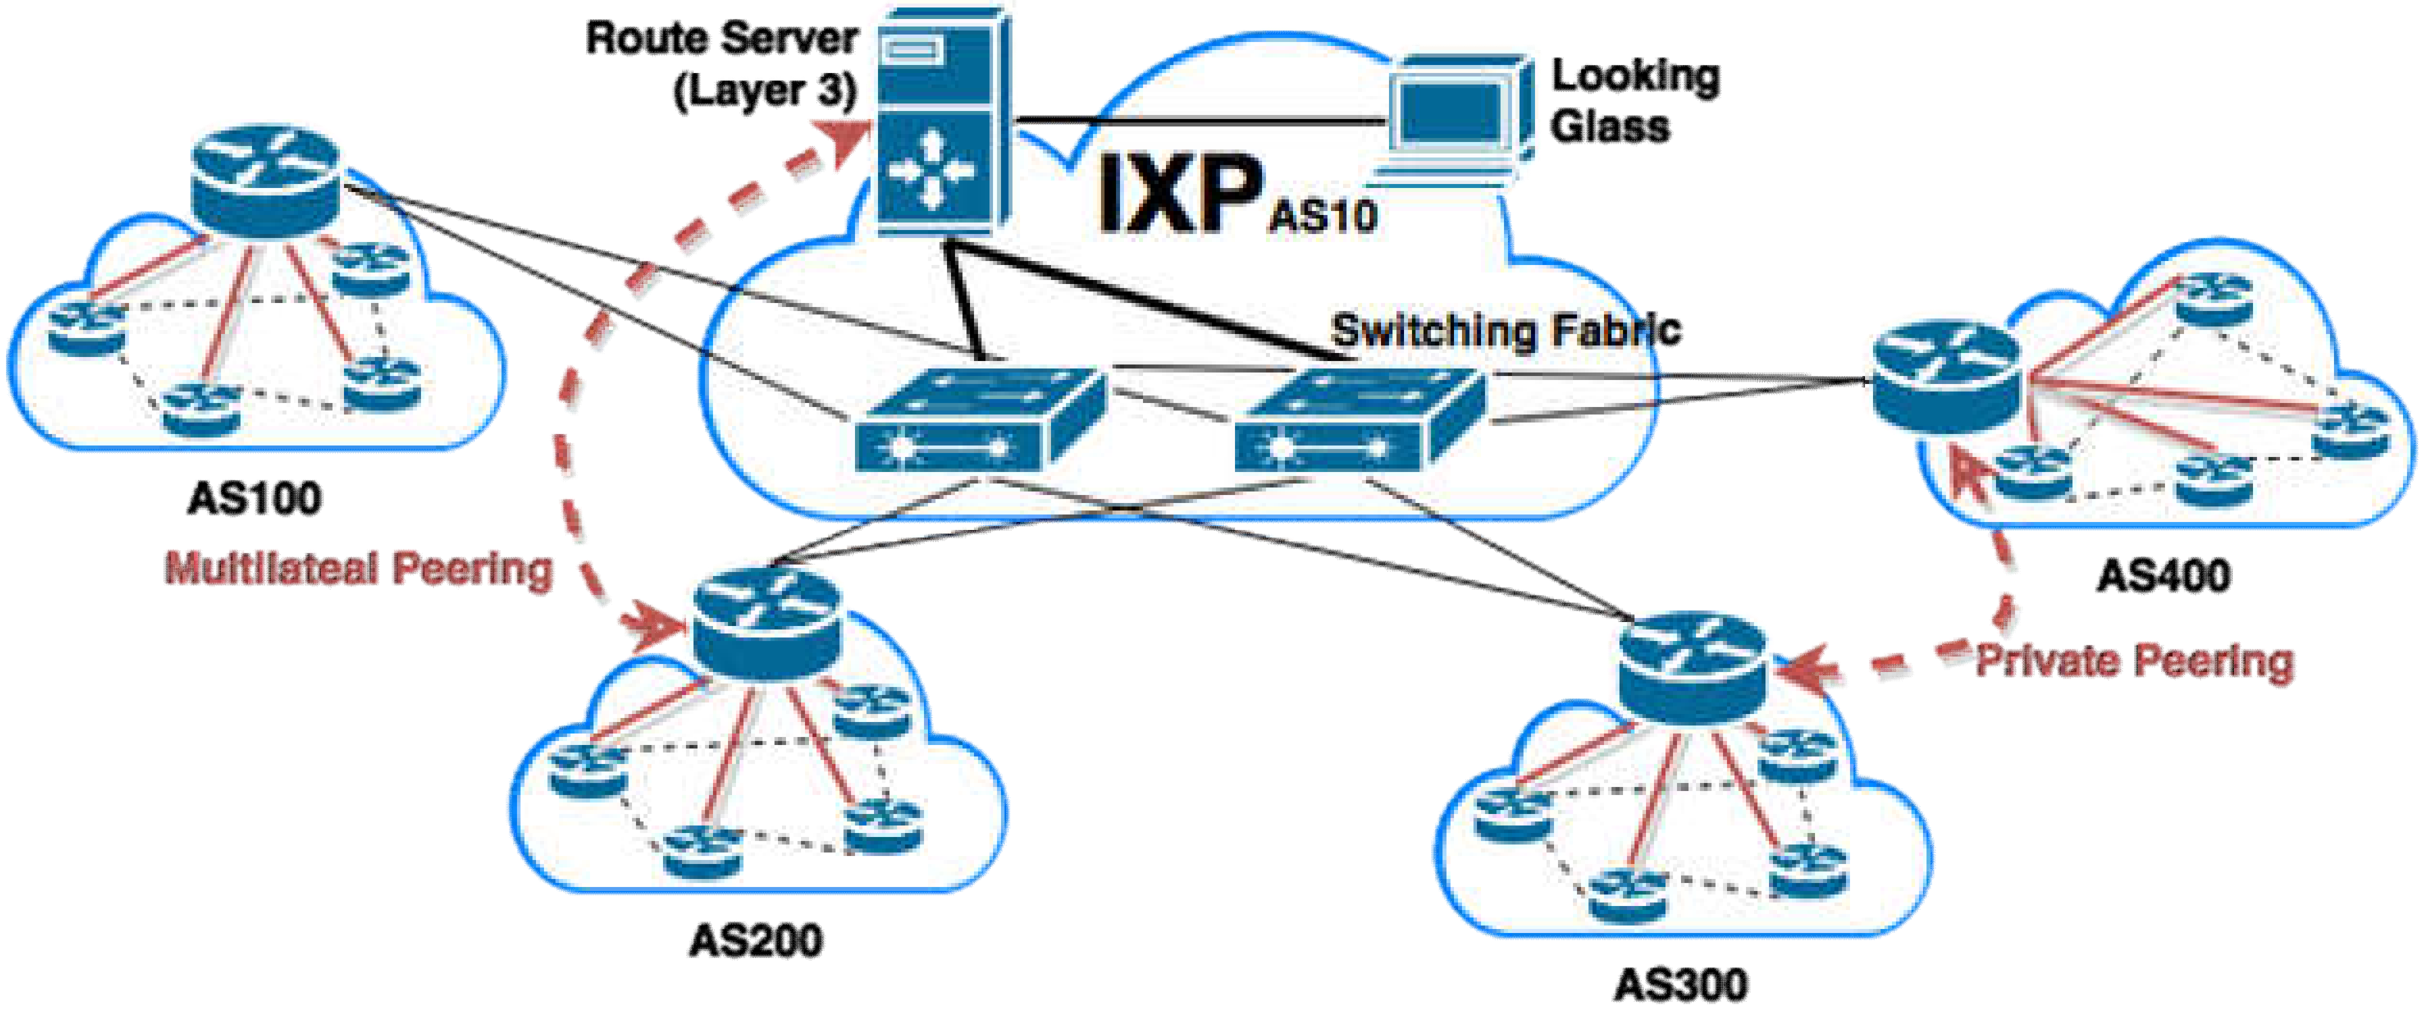

A shared network infrastructure

The Internet performance is optimized

Multilateral peering

Allow bilateral agreements

See IX.br public service approach (CGI)

Agenda

1. Introduction
2. Background
3. **Related work**
4. Objectives
5. Design of ALTO-as-a-Service (AaaS)
6. AaaS Prototype
7. Experimental Evaluation
8. Future Work / Conclusions

Related Work

BGP Information from LG Servers
Brito et al. [1] collected BGP information from public IXPs operating in Brazil. It was one of the first (if not the first) efforts to comprehend the ecosystem operation of Brazilian IXPs. Some result: AS-level connectivity graphs, Peering density, Depth of AS-PATH, etc. The authors in [2] collected raw data from 245 LG servers across 110 countries. They create and analyze the LG server-based AS topology against BGP (IRL [3]), traceroute (Ark [4] and iPlane [5]), and IRR [6] based AS topologies. Among other results, they observed 11 K new ASes links and 686 new ASes.

Related Work

Provide Network Topology Information
Choffnes et al. [7], Bonaventure et al. [8] and Xie et al. [9] provide a P2P mechanism to find reliable information source to apply better-than-random peer selection, but not in the form of ALTO services (i.e. difficult to deploy globally without an open standard).

Related Work

Create and/or Provide ALTO Information
The work in [10] creates network topology and cost maps for ALTO from public sources of information, specifically using the United States Federal Communications Commission (FCC) public database from the Measuring Broadband America (MBA) program. Pinthong et al. [11] and Guanxiu et al. [12] explore the ALTO protocol by mapping IP addresses to an AS, to create a network map and BGP route announcements from ASes to create a cost map. Both maps were created by using simulation platforms (PeerSim and GPS, respectively). SHIBUYA et al. [13] uses a SQL-based database for storing the Prefix List and the Priority Map.

Agenda

1. Introduction
2. Background
3. Related work
4. **Objectives**
5. Design of ALTO-as-a-Service (AaaS)
6. Prototype Implementation
7. Experimental Evaluation
8. Future Work / Conclusions

Objectives

MAIN OBJECTIVE: Create and provide Application-Layer Traffic Optimization (ALTO) services based on public information, more specifically BGP routing information publicly available at Internet eXchange Points (IXPs).

Issue: Random selection of available peers	**Providing a mechanism for giving peers information which enables choosing closest neighbors rather randomly.**
Issue: Cannot directly use traditional traffic optimization techniques	**Applications could have a large impact on the overall traffic generated using services with knowledge of network topology.**
Issue: Limited knowledge of the underlying network topology	**Working at a layer above, with a better understanding of the underlying network would give applications a detailed guide to achieve the desired resources.**

Agenda

1. Introduction
2. Background
3. Related work
4. Objectives
5. **Design of ALTO-as-a-Service (AaaS)**
6. Prototype Implementation
7. Experimental Evaluation
8. Future Work / Conclusions

Design of ALTO-as-a-Service (AaaS) :
AaaS Workflow

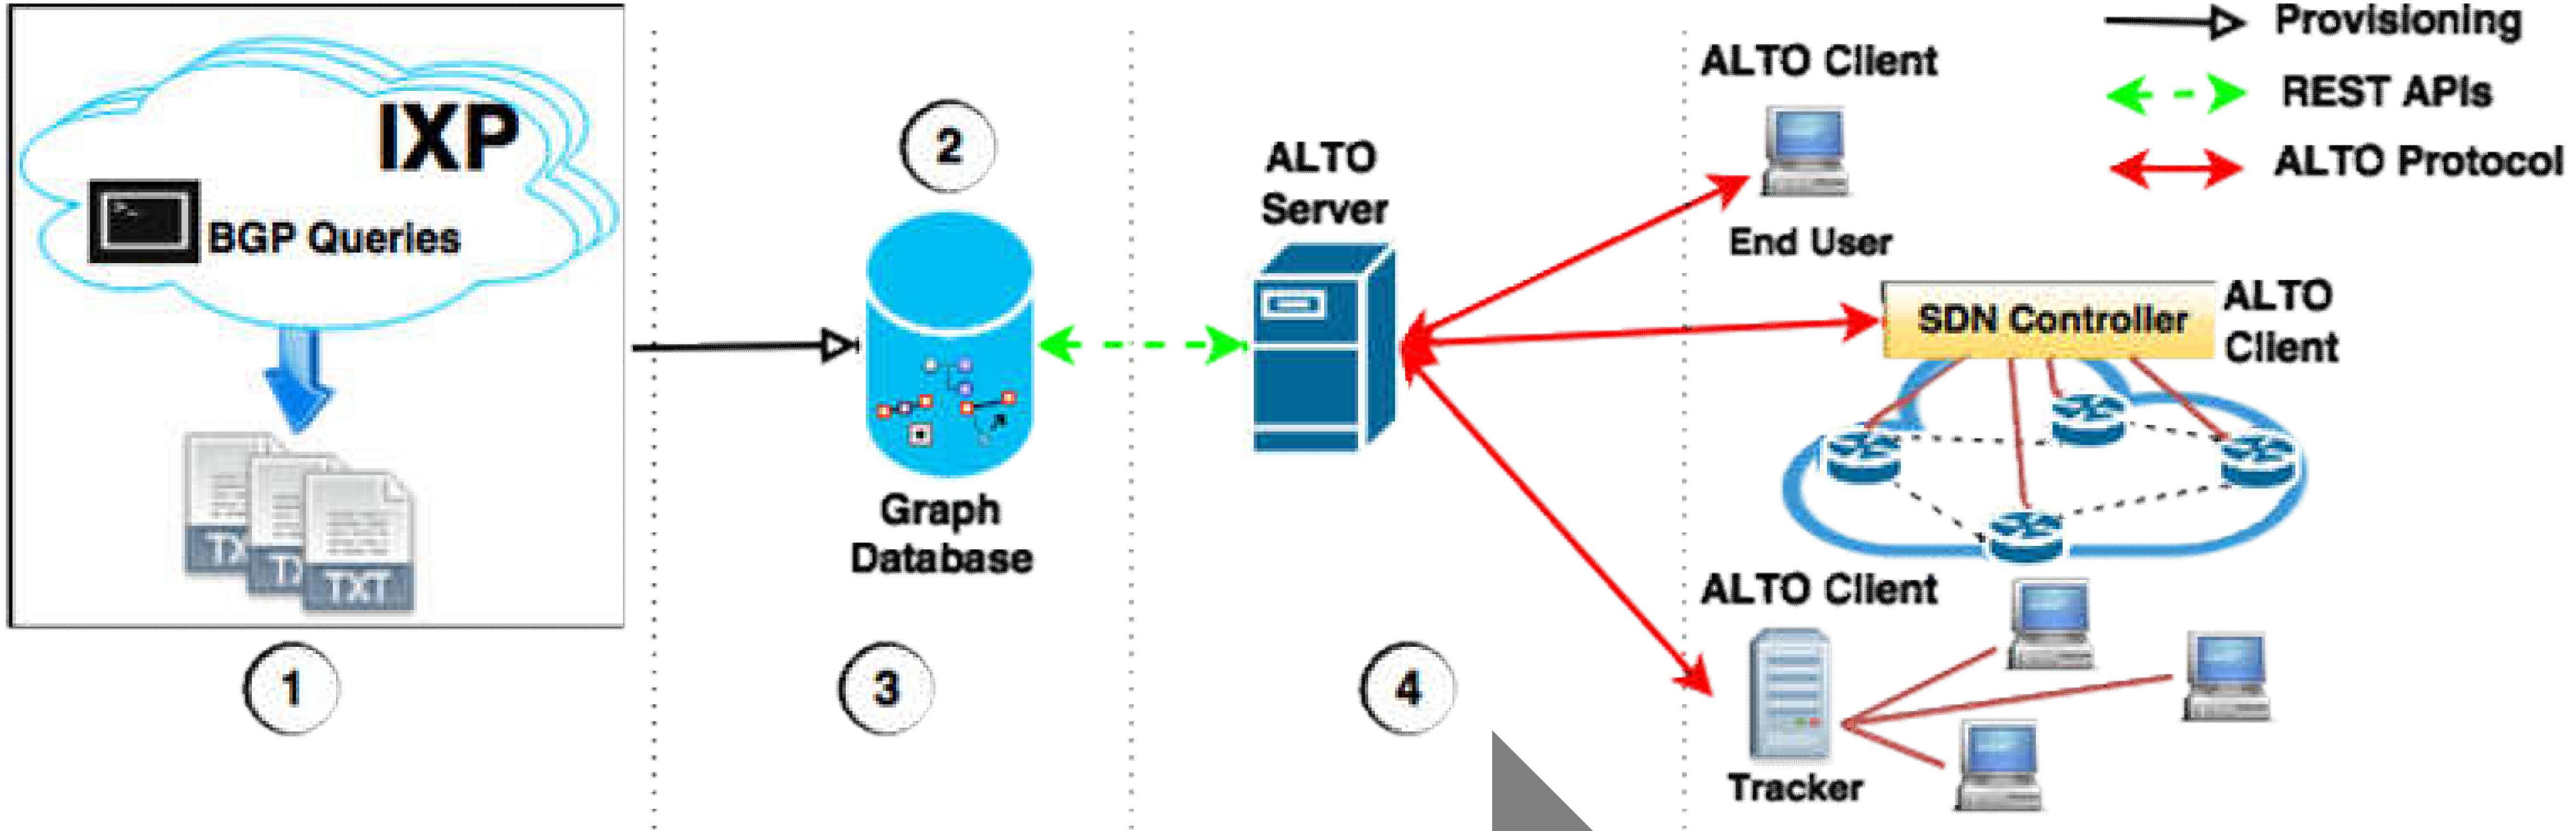

1

Acquiring Input Data

Design of ALTO-as-a-Service (AaaS) :

Acquiring Input Data

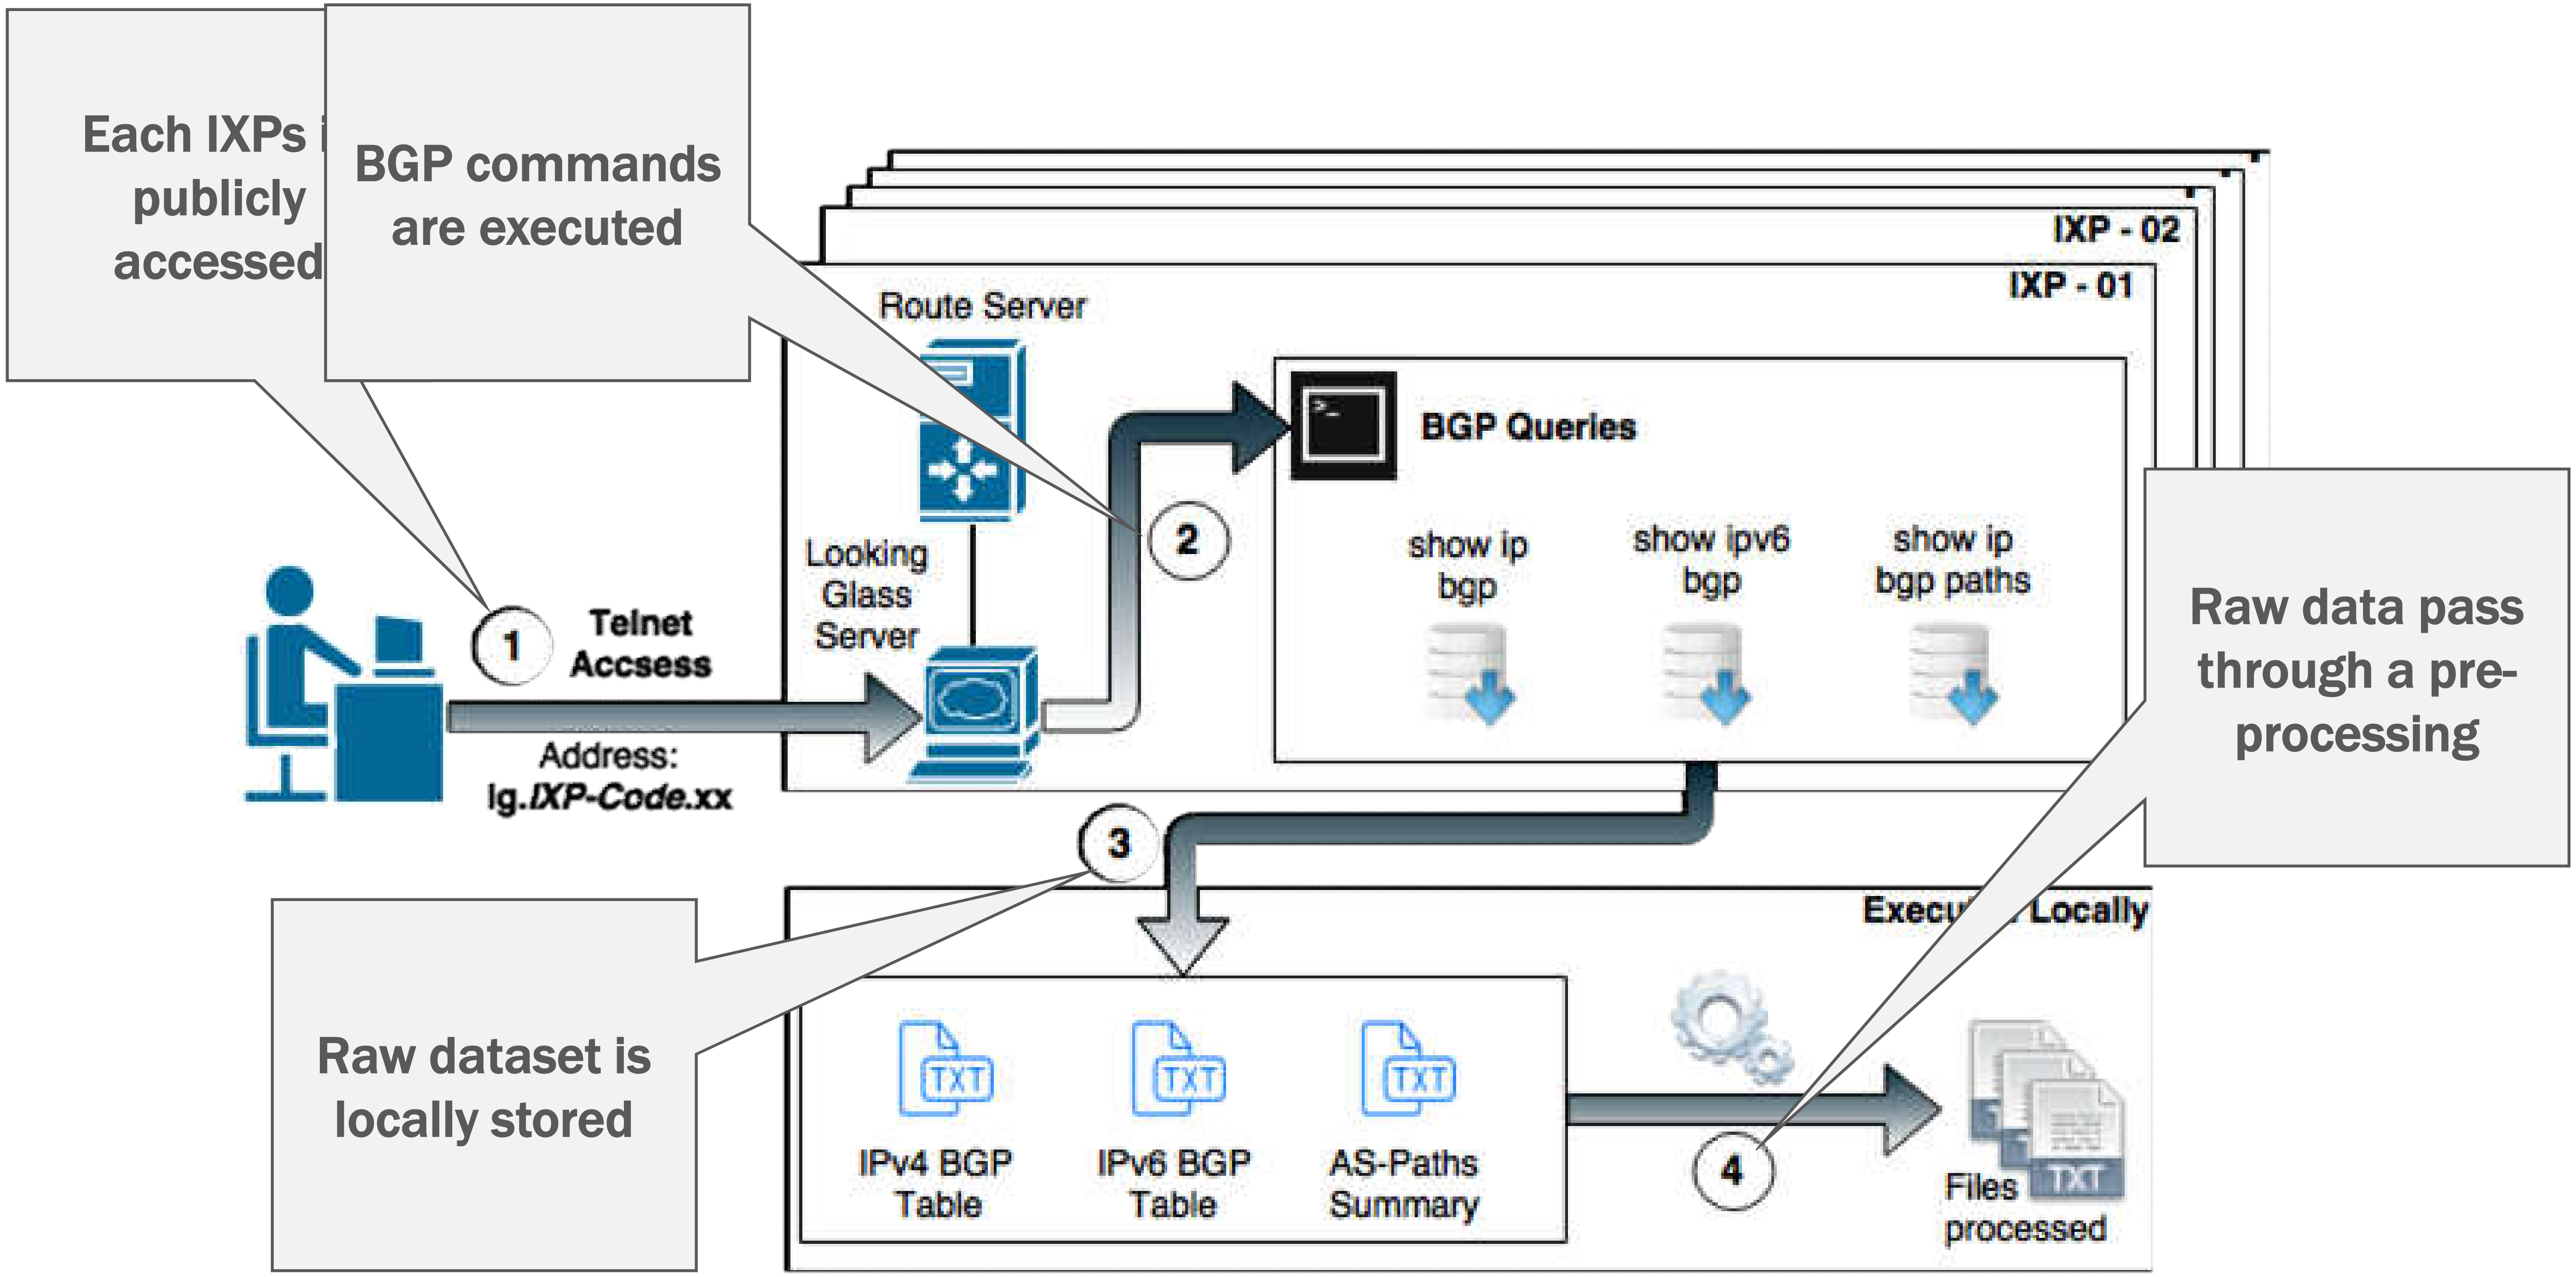

Design of ALTO-as-a-Service (AaaS) :

AaaS Workflow

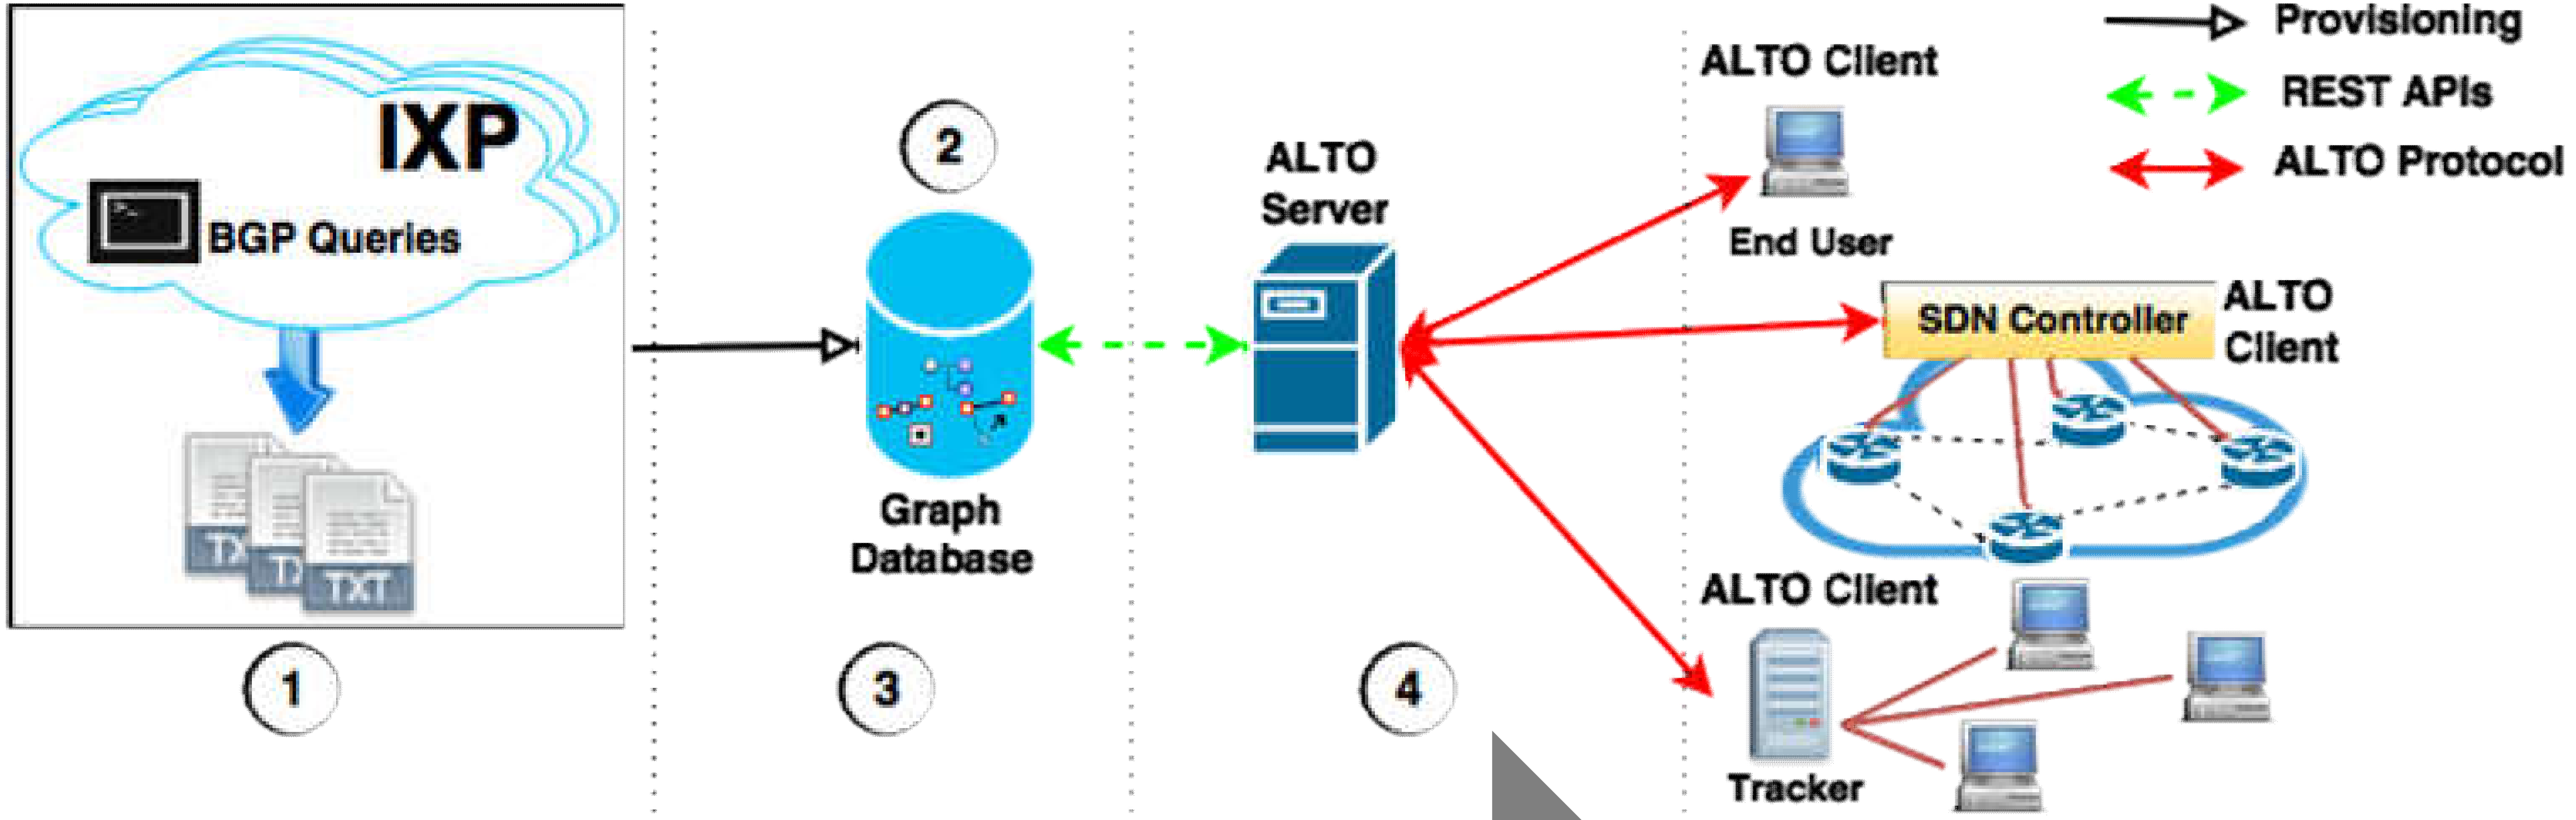

1
Acquiring Input Data

2
Graph Data Modeling

Design of ALTO-as-a-Service (AaaS) :

Building Graph Data Models

RFC 7285 IETF

Nodes:

1. Provider-Defined Identifier
2. Endpoint Address
3. Address Type
4. Version Tag

Relationships:

1. Has_PID
2. Has_EndPoint
3. Type_EndPoint
4. Cost

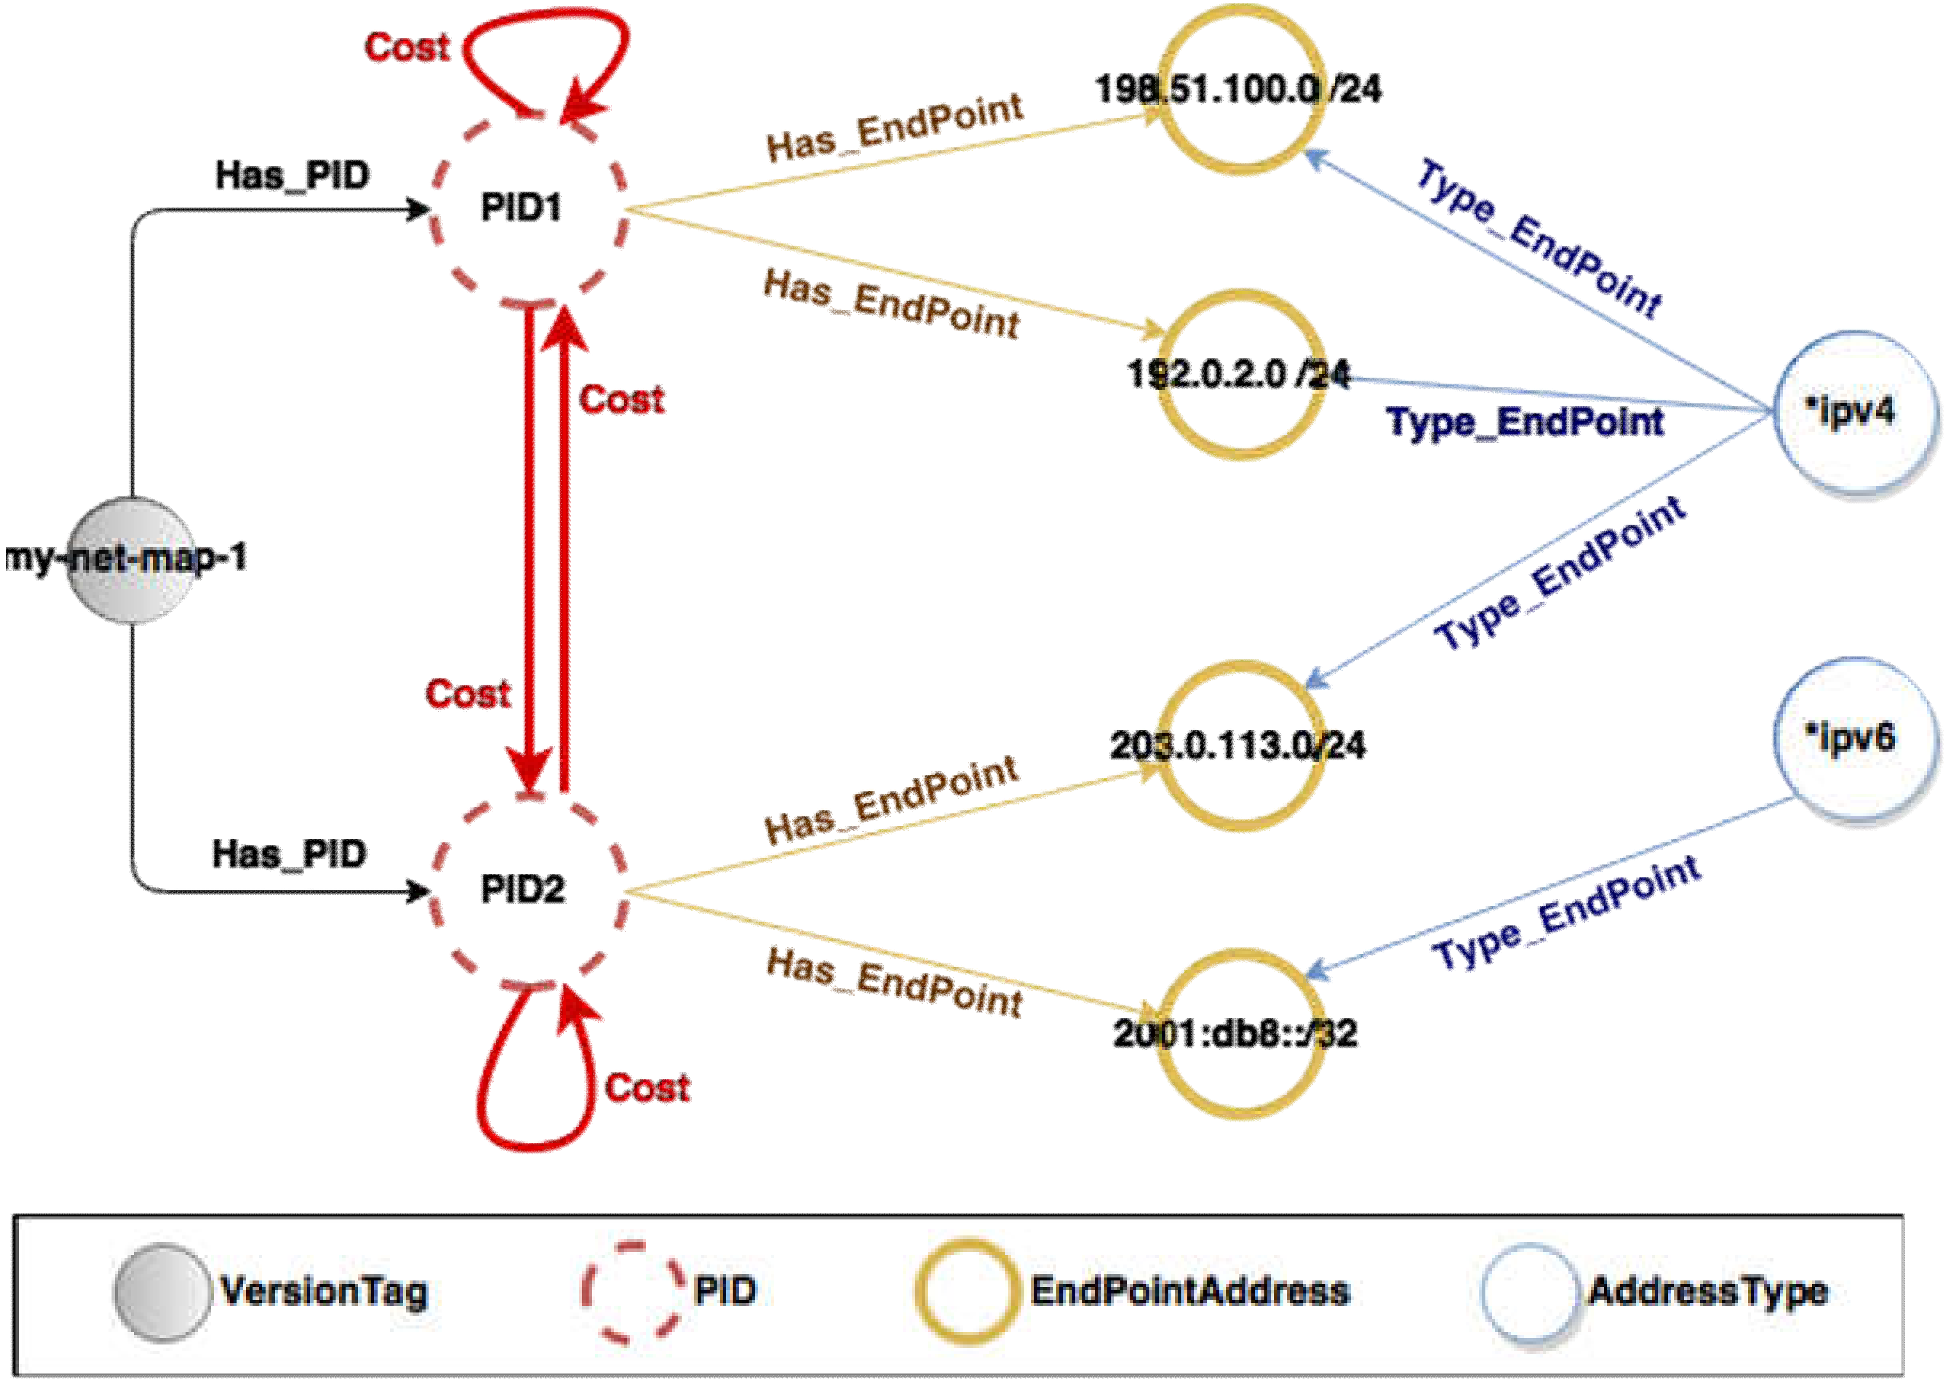

Design of ALTO-as-a-Service (AaaS) :

AaaS Workflow

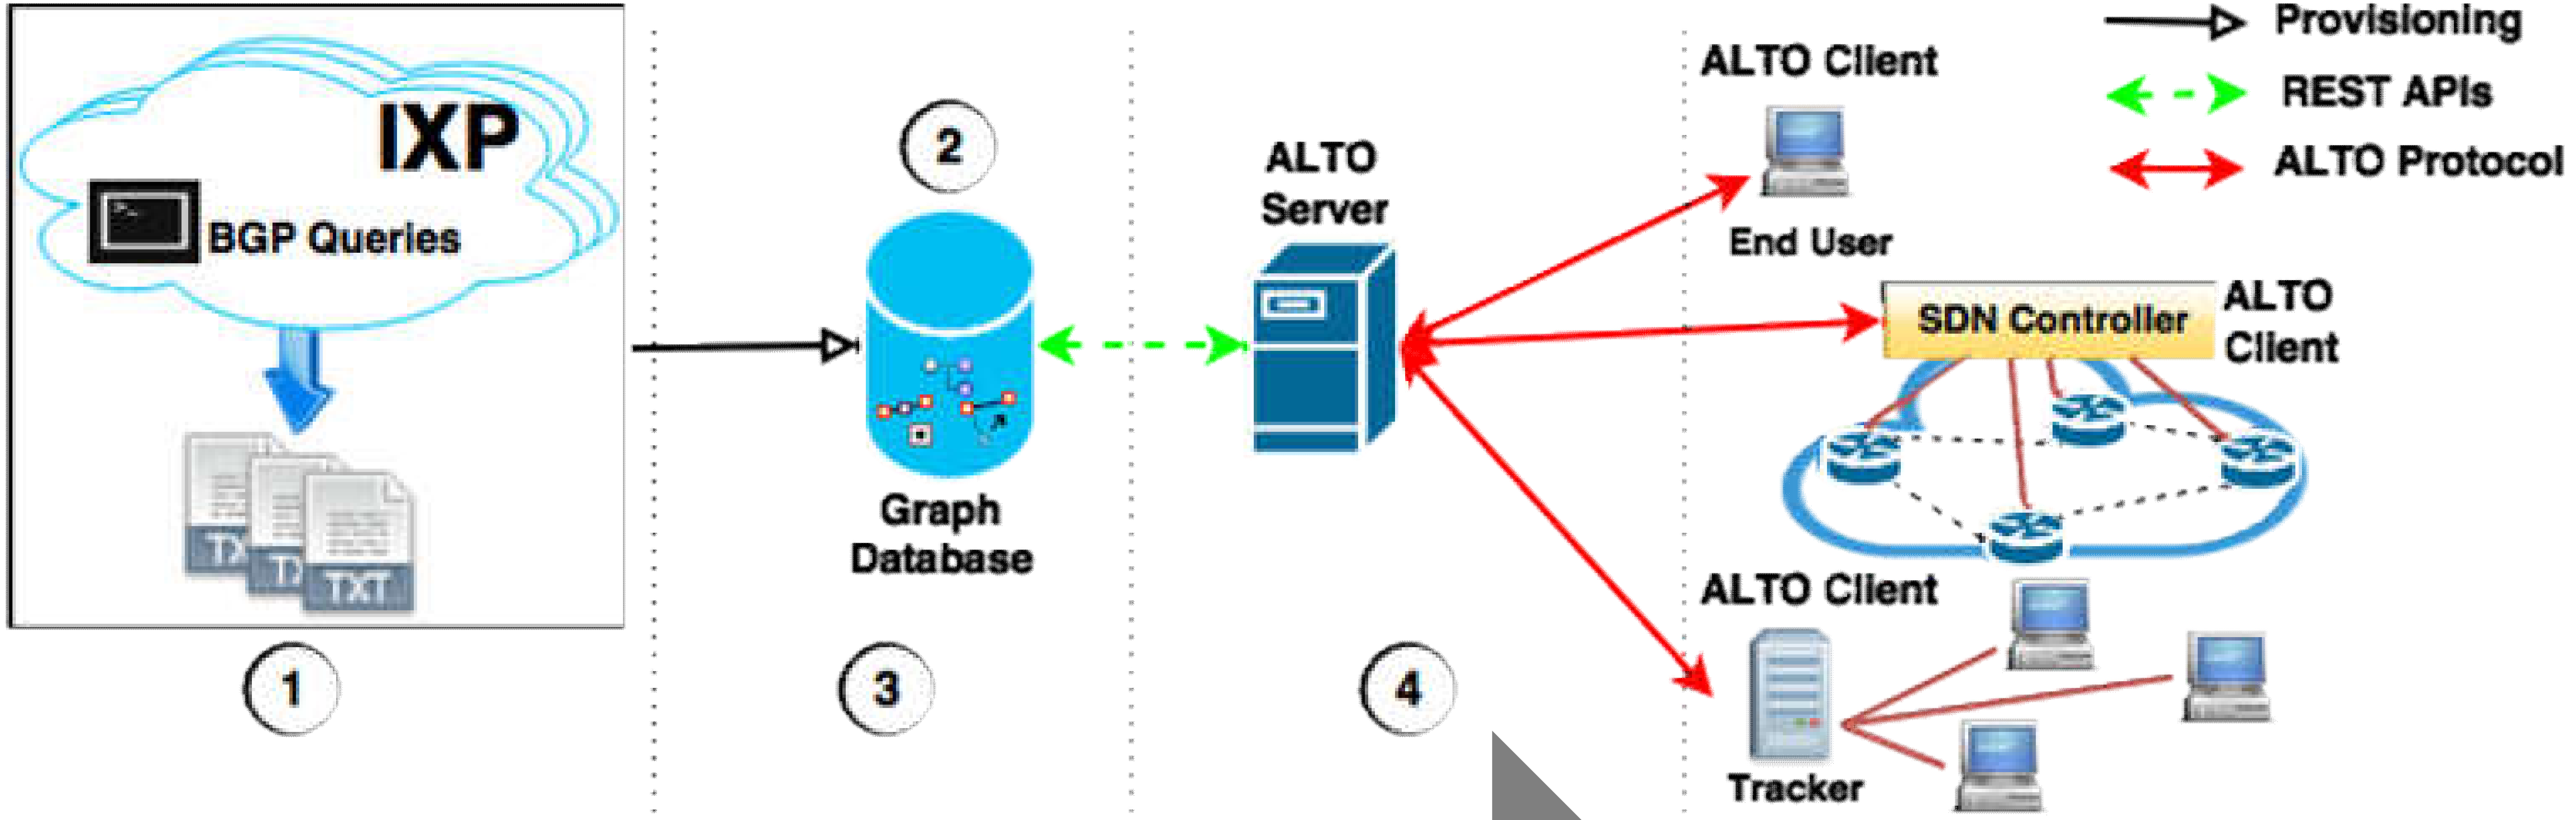

1 Acquiring Input Data

2 Graph Data Modeling

3 Creating ALTO Information

4 Delivering ALTO Services

Agenda

1. Introduction
2. Background
3. Related work
4. Objectives
5. Design of ALTO-as-a-Service (AaaS)
6. **Prototype Implementation**
7. Experimental Evaluation
8. Future Work / Conclusions

Prototype Implementation:

Input: IX.br BGP Data Set

25 IXPs are part of PTTMetro project[1], created to promote the necessary infrastructure for direct interconnection between networks of Brazilian metropolitan regions.

1 The largest IXP ecosystem in Latin America and one of the world's top ten.

1374 Members[2].

+1.0 Tbps average throughput [3]

+2.5 GB of real BGP data.

[1]http://ix.br

[2]http://ix.br/particip, Accessed: September, 2015

[3]http://ix.br/trafego, Accessed: May, 2015

Table 4.1: Public IXPs Operating in Brazil

#	City	State	Looking Glass	Gbps	Members	
01	Belem	PA	lg.bel.ptt.br	0.44	14	
02	Belo Horizonte	MG	lg.mg.ptt.br	2.07	33	
03	Brasilia	DF	lg.df.ptt.br	2.98	30	
04	Campina Grande	PB	lg.cpv.ptt.br	0.69	10	
05	Campinas	SP	lg.cas.ptt.br	3.57	35	(*)
06	Cuiaba	MT	lg.cgb.ptt.br	0.00	9	(*)
07	Caxias do Sul	RS	lg.cxj.ptt.br	0.08	5	(*)
08	Curitiba	PR	lg.pr.ptt.br	16.10	68	(1)
19	Florianopolis	SC	lg.sc.ptt.br	1.28	34	(*)
10	Fortaleza	CE	lg.ce.ptt.br	2.72	29	
11	Goiania	GO	lg.gyn.ptt.br	1.06	24	
12	Lajeado	RS	lg.laj.ptt.br	0.01	8	(*)
13	Londrina	PR	lg.lda.ptt.br	1.62	32	
14	Manaus	AM	lg.mao.ptt.br	0.02	8	(*)
15	Maringa	PR	lg.mgf.ptt.br	0.28	21	(*)
16	Natal	RN	lg.nat.ptt.br	0.26	13	(*)
17	Porto Alegre	RS	lg.rs.ptt.br	20.85	117	
18	Recife	PE	lg.pe.ptt.br	0.69	16	
19	Rio de Janeiro	RJ	lg.rj.ptt.br	39.22	68	
20	Salvador	BA	lg.ba.ptt.br	1.47	47	(*)
21	Sao Carlos	SP	lg.sca.ptt.br	0.00	3	(*)
22	Sao Jose dos Campos	SP	lg.sjc.ptt.br	0.47	13	
23	Sao Jose do Rio Preto	SP	lg.sjp.ptt.br	0.62	11	(*)
24	Sao Paulo	SP	lg.sp.ptt.br	429.45	667	(1)
25	Vitoria	ES	lg.vix.ptt.br	0.80	22	

(1) There are filters in LG compromising the BGP table.

(*) Data provided by NIC.br, since publicly access was denied.

Prototype Implementation :

ALTO Server Backend: Neo4j

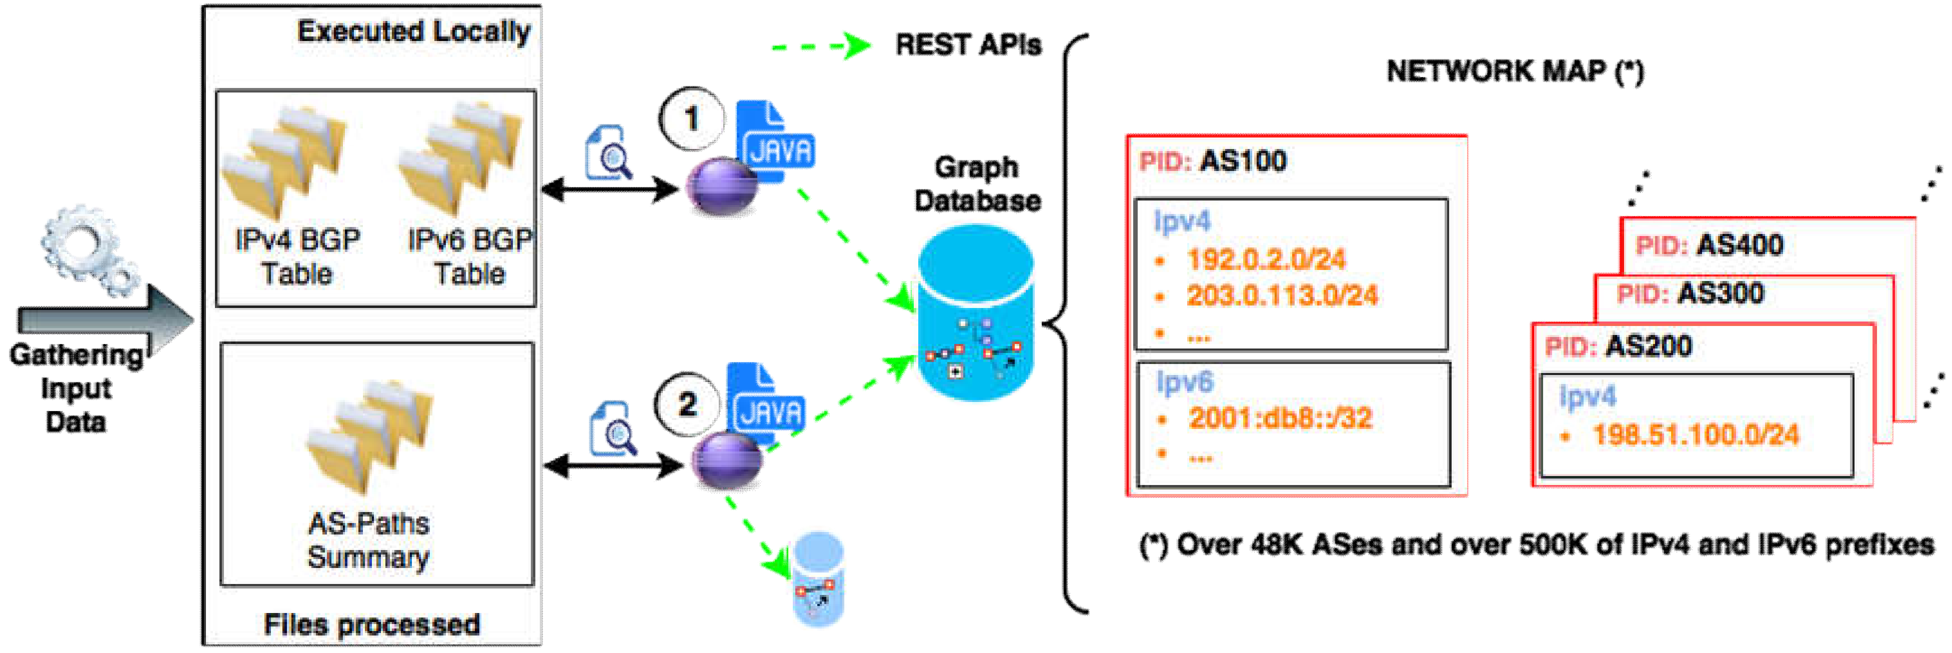

COST MAP (+)

AS hops	AS100	AS200	AS300	...
AS100	0	5	4	...
AS200	2	0	6	...
AS300	3	3	0	...
...	...	...	...	...

(+) Over 2.3 billions of pah costs

1. Network Map

✓ The ASNs serve as our grouping, with each ASN corr

✓ Each prefix (be it IPv4 and IPv6) will be associate w

2. Cost Map

✓ Path costs = The topological distance, expressed in the n of traversed AS.

✓ Two Cost Maps: Absolute distance and Relative distance.

Hops between ASes present in the same IXPs are zeroed to favor intra-IXP traffic.

Prototype Implementation:

ALTO Server Front-End: OpenDaylight

- ✓ ALTO in ODL is focused on implementing basic ALTO services as RESTful web services (Northbound APIs) for ALTO client/server communications.
- ✓ ALTO Northbound APIs generate ALTO services from data stored in the MD-SAL data store (an ODL core component)
  - It was necessary to modify the Northbound APIs to generate ALTO services from the data stored in the Neo4j GDB (instead of the MD-SAL topology).

Agenda

1. Introduction
2. Background
3. Related work
4. Objectives
5. Design of ALTO-as-a-Service (AaaS)
6. Prototype Implementation
7. **Experimental Evaluation**
8. Future Work / Conclusions

Functional Evaluation

RFC 7285 IETF

Postman - REST Client

{JSON}

URI : "http://intrig.dca.fee.unicamp.br/controller/nb/v2/alto/filtered/networkmap/my-default-network-map"
HTTP Method : "POST"
Content-Type : "application/alto-networkmapfilter+json"
Input Parameters :

```
{
    "pids" : ["AS100","AS200"]
}
```

HTTP Response :

```
{
  "meta": {
    "vtag": { "resource-id": "my-default-network-map",
      "tag": "da65eca2eb7a10ce" }
  },
  "network-map": {
    "AS100": { "ipv4": [ "192.0.2.0/24","203.0.113.0/24" ],
               "ipv6": [ "2001:db8::/32" ],
    "AS200": { "ipv4": [ "198.51.100.0/24" ] }
    }
  }
}
```

URI : "http://intrig.dca.fee.unicamp.br/controller/nb/v2/alto/filtered/costmap/my-default-network-map"
HTTP Method : "POST"
Content-Type : "application/alto-costmapfilter+json"
Input Parameters :

```
{  "cost-type" :{"cost-mode": "Numerical",
   "cost-metric": "HopsNumber" },
   "pids" : { "srcs" : [ "AS100" ],
   "dsts" : [ "AS100","AS200","AS300" ] }
}
```

HTTP Response :

```
{
 "meta": {
  "dependent-vtags": [
          { "resource-id": "my-default-network-map",
            "tag": "da65eca2eb7a10ce" } ],
  "cost-type": { "cost-mode": "Numerical",
          "cost-metric": "HopsNumber" }},
  "cost-map": { "AS100":
          { "AS100": 0, "AS300": 4, "AS200": 5 } }
}
```

(a) Filtered Network Map

(b) Filtered Cost Map (abs. distance)

Experimental Evaluation :

System Performance Profiling

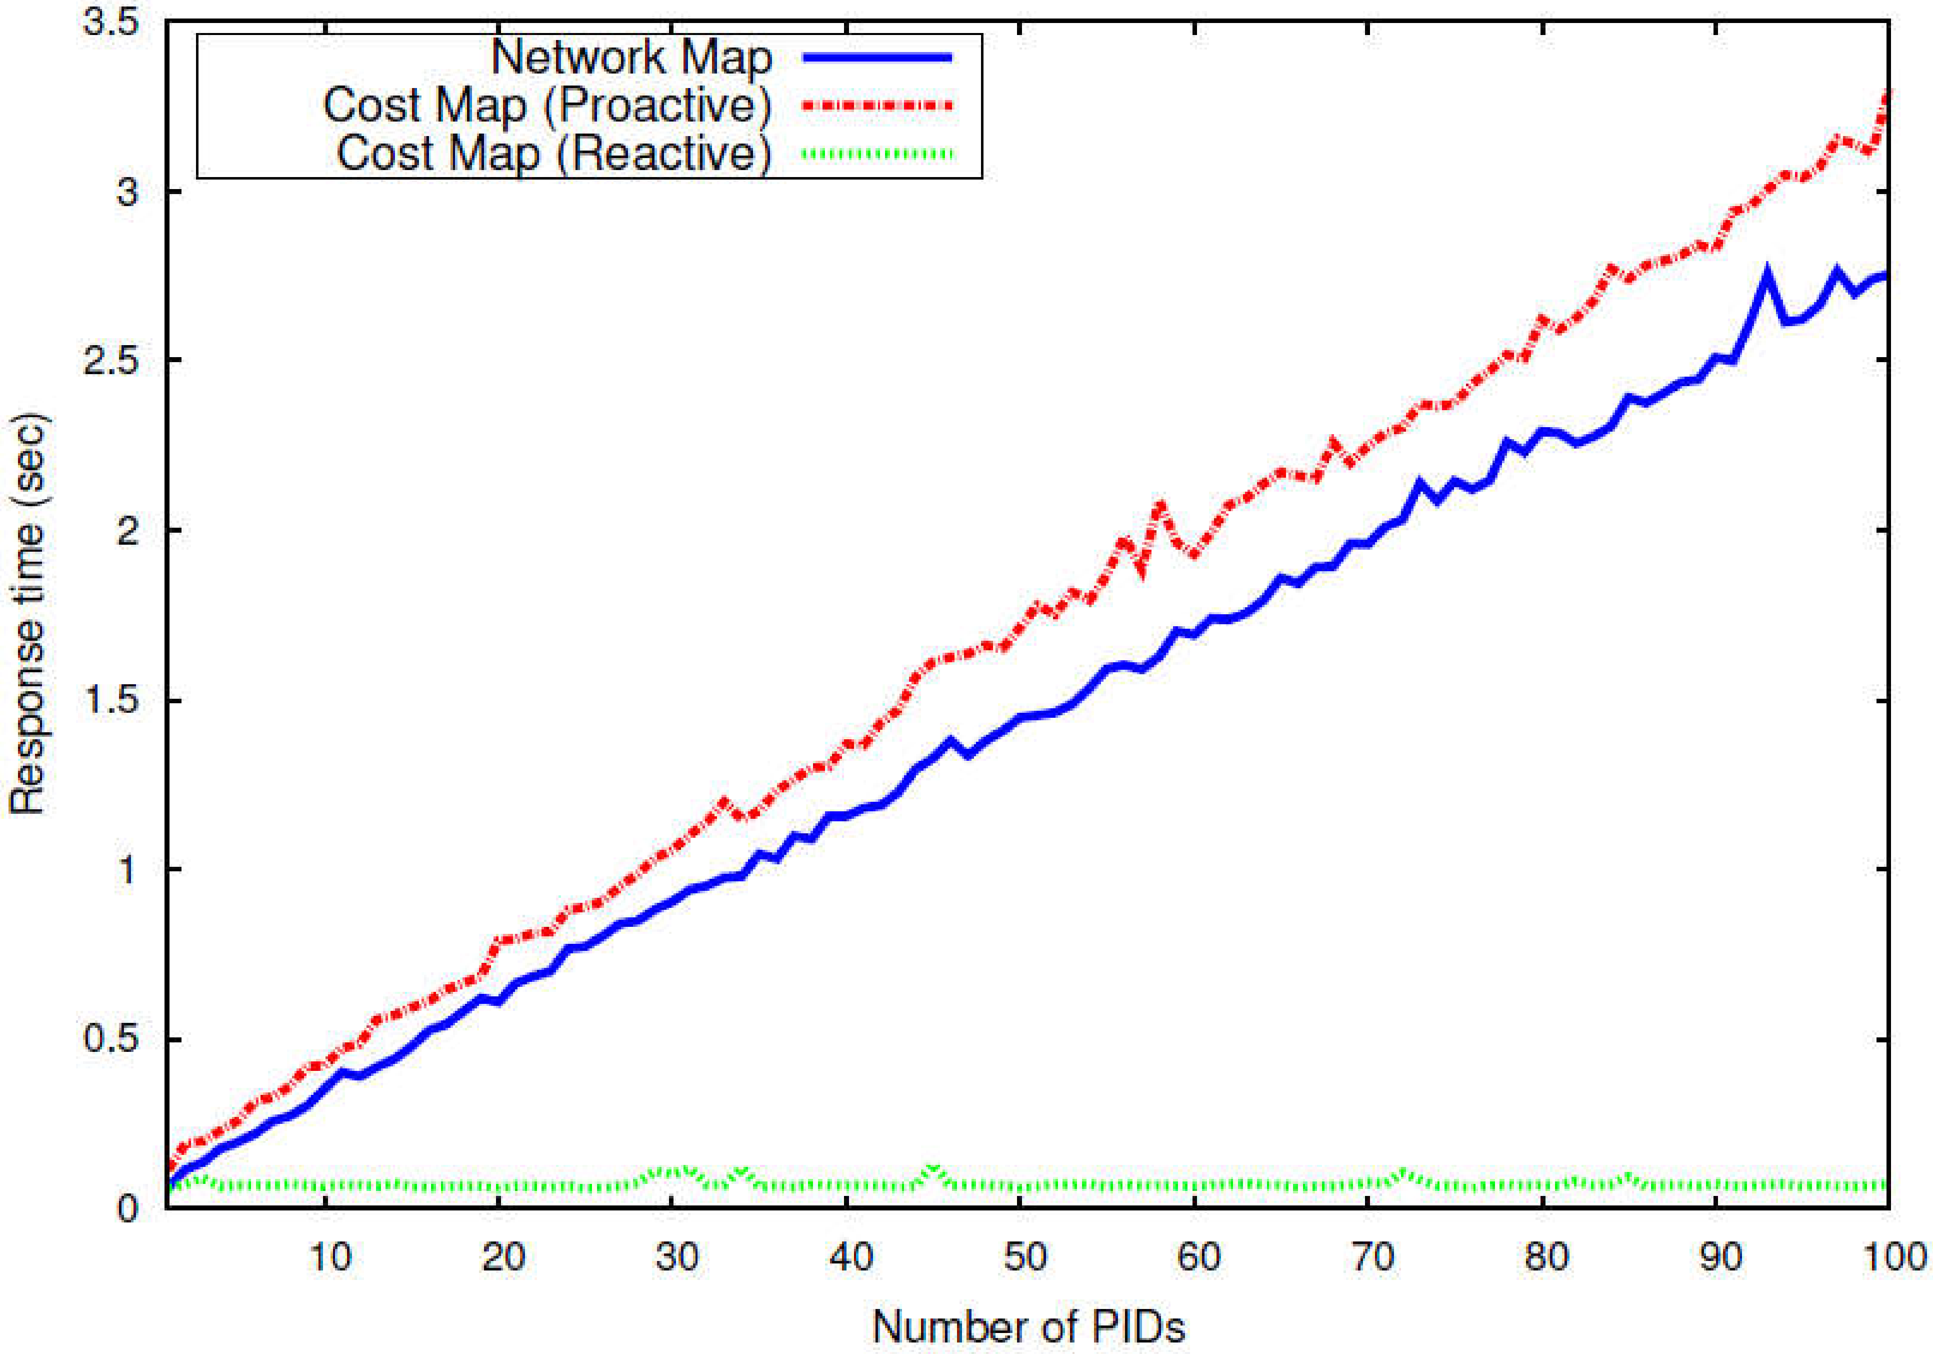

Response processing time for the Network and Cost Map (Absolute Distance) services

Experimental Evaluation :

1) IP Endpoint Selection with AaaS (*strawman*)

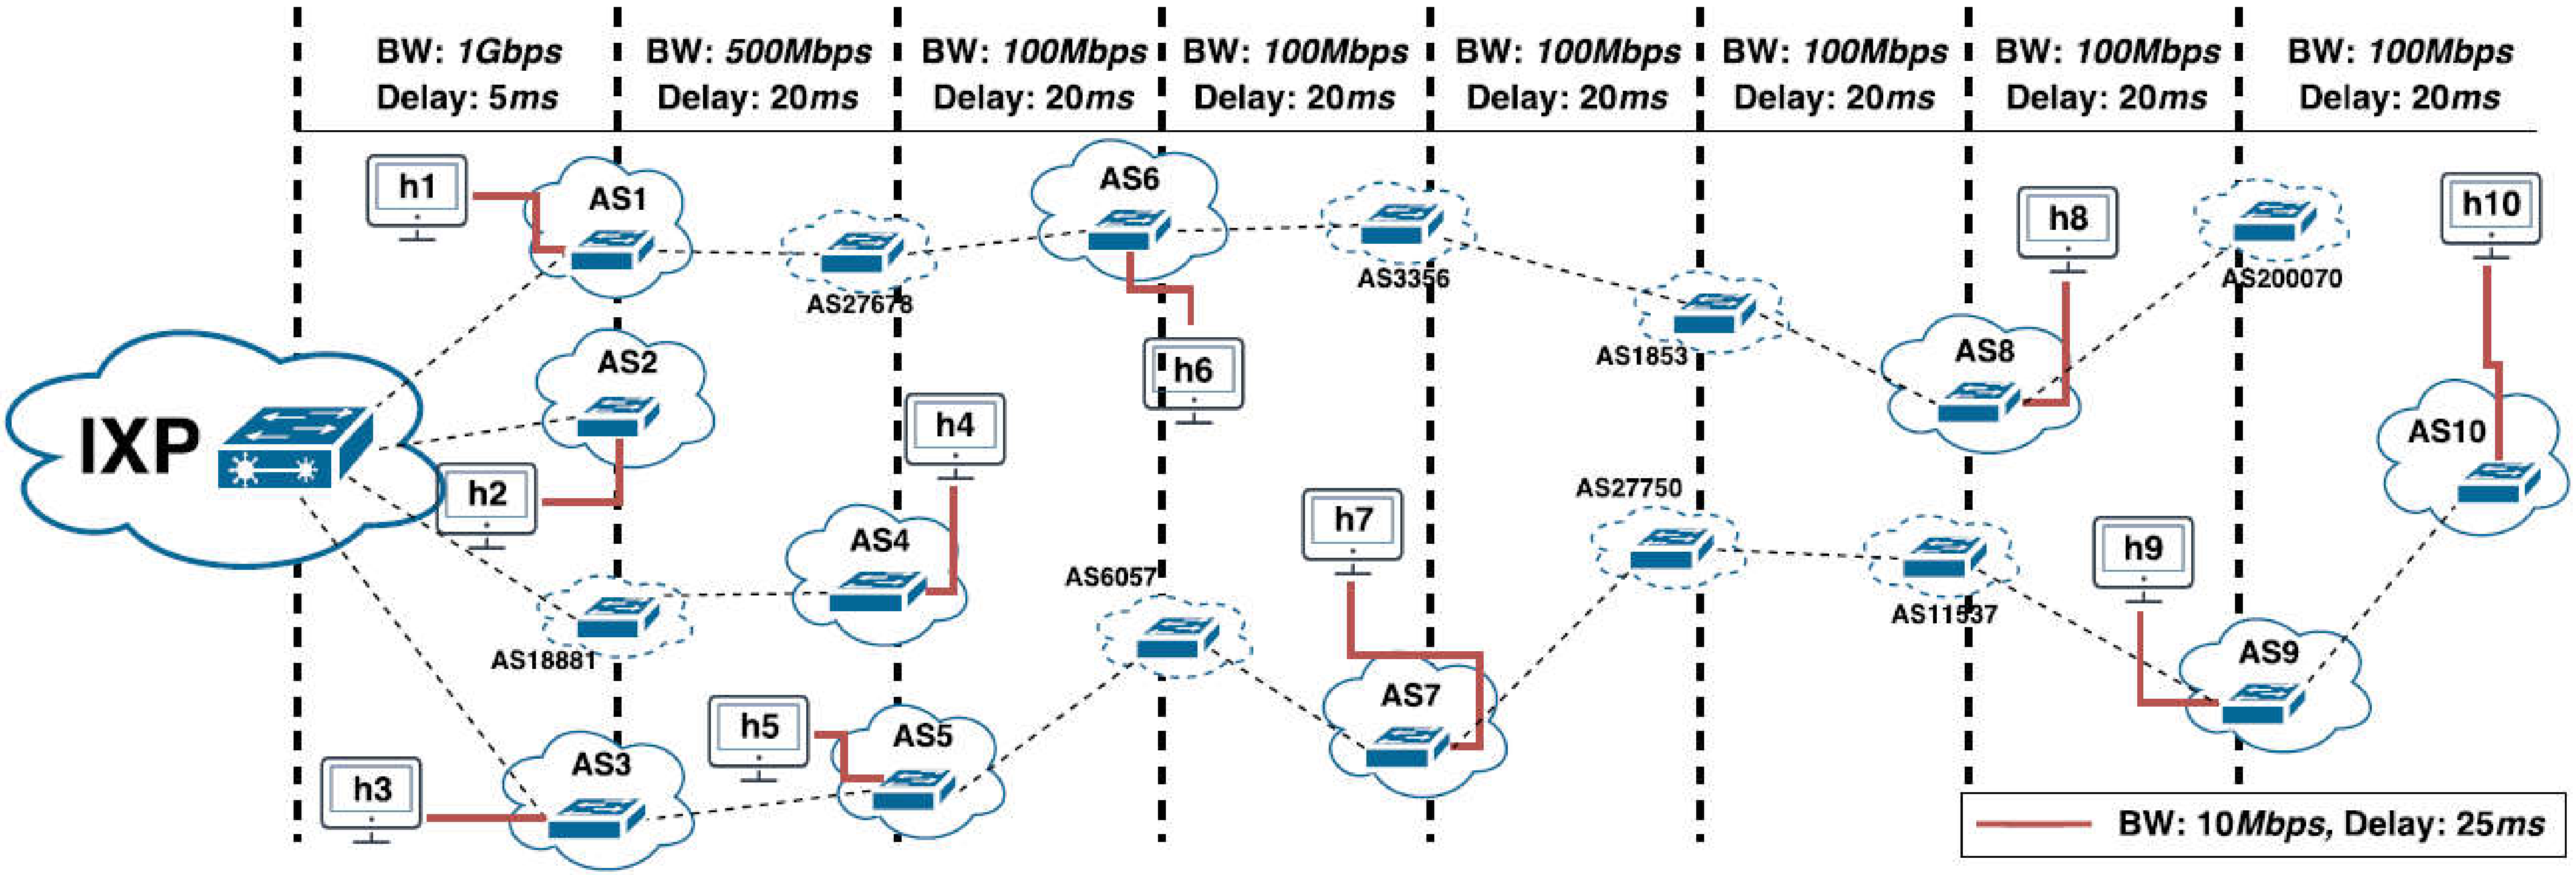

- ✓ 19 ASes, each represented by a switch abstraction in the Mininet emulator.
- ✓ The large AS switch represents the IXP.
- ✓ Links between ASes were set with larger bandwidth and lower delay when closer to the IXP.
- ✓ 10 communicating peers are represented as Mininet hosts attached to the switches.

Experimental Evaluation :

1) IP Endpoint Selection with AaaS (*strawman*)

ALTO information

Absolute distance (HopsNumber)

	AS1	AS2	AS3	AS4	AS5	AS6	AS7	AS8	AS9	AS10
AS1	0	2	2	3	3	2	5	5	8	9
AS2	2	0	2	3	3	4	5	7	8	9
AS3	2	2	0	3	1	4	3	7	6	7

Relative distance (HopsNumberPTT)

AS	AS1	AS2	AS3	AS4	AS5	AS6	AS7	AS8	AS9	AS10
AS1	0	0	0	1	1	2	3	5	6	7
AS2	0	0	0	1	1	2	3	5	6	7
AS3	0	0	0	1	1	2	3	5	6	7

ALTO Clients

- ✓ Hosts whose ASes are directly connected at IXP (ie., h1, h2, h3).

Workload and Metrics

- ✓ With a background traffic using the D-ITG traffic sending TCP traffic (512 byte packet size, 1,000 pps rate).
- ✓ We run end-to-end round-trip time measurements and available bandwidth with the remaining nine hosts using ping and iperf tools, respectively.

Experimental Evaluation :

Network Latency and Throughput

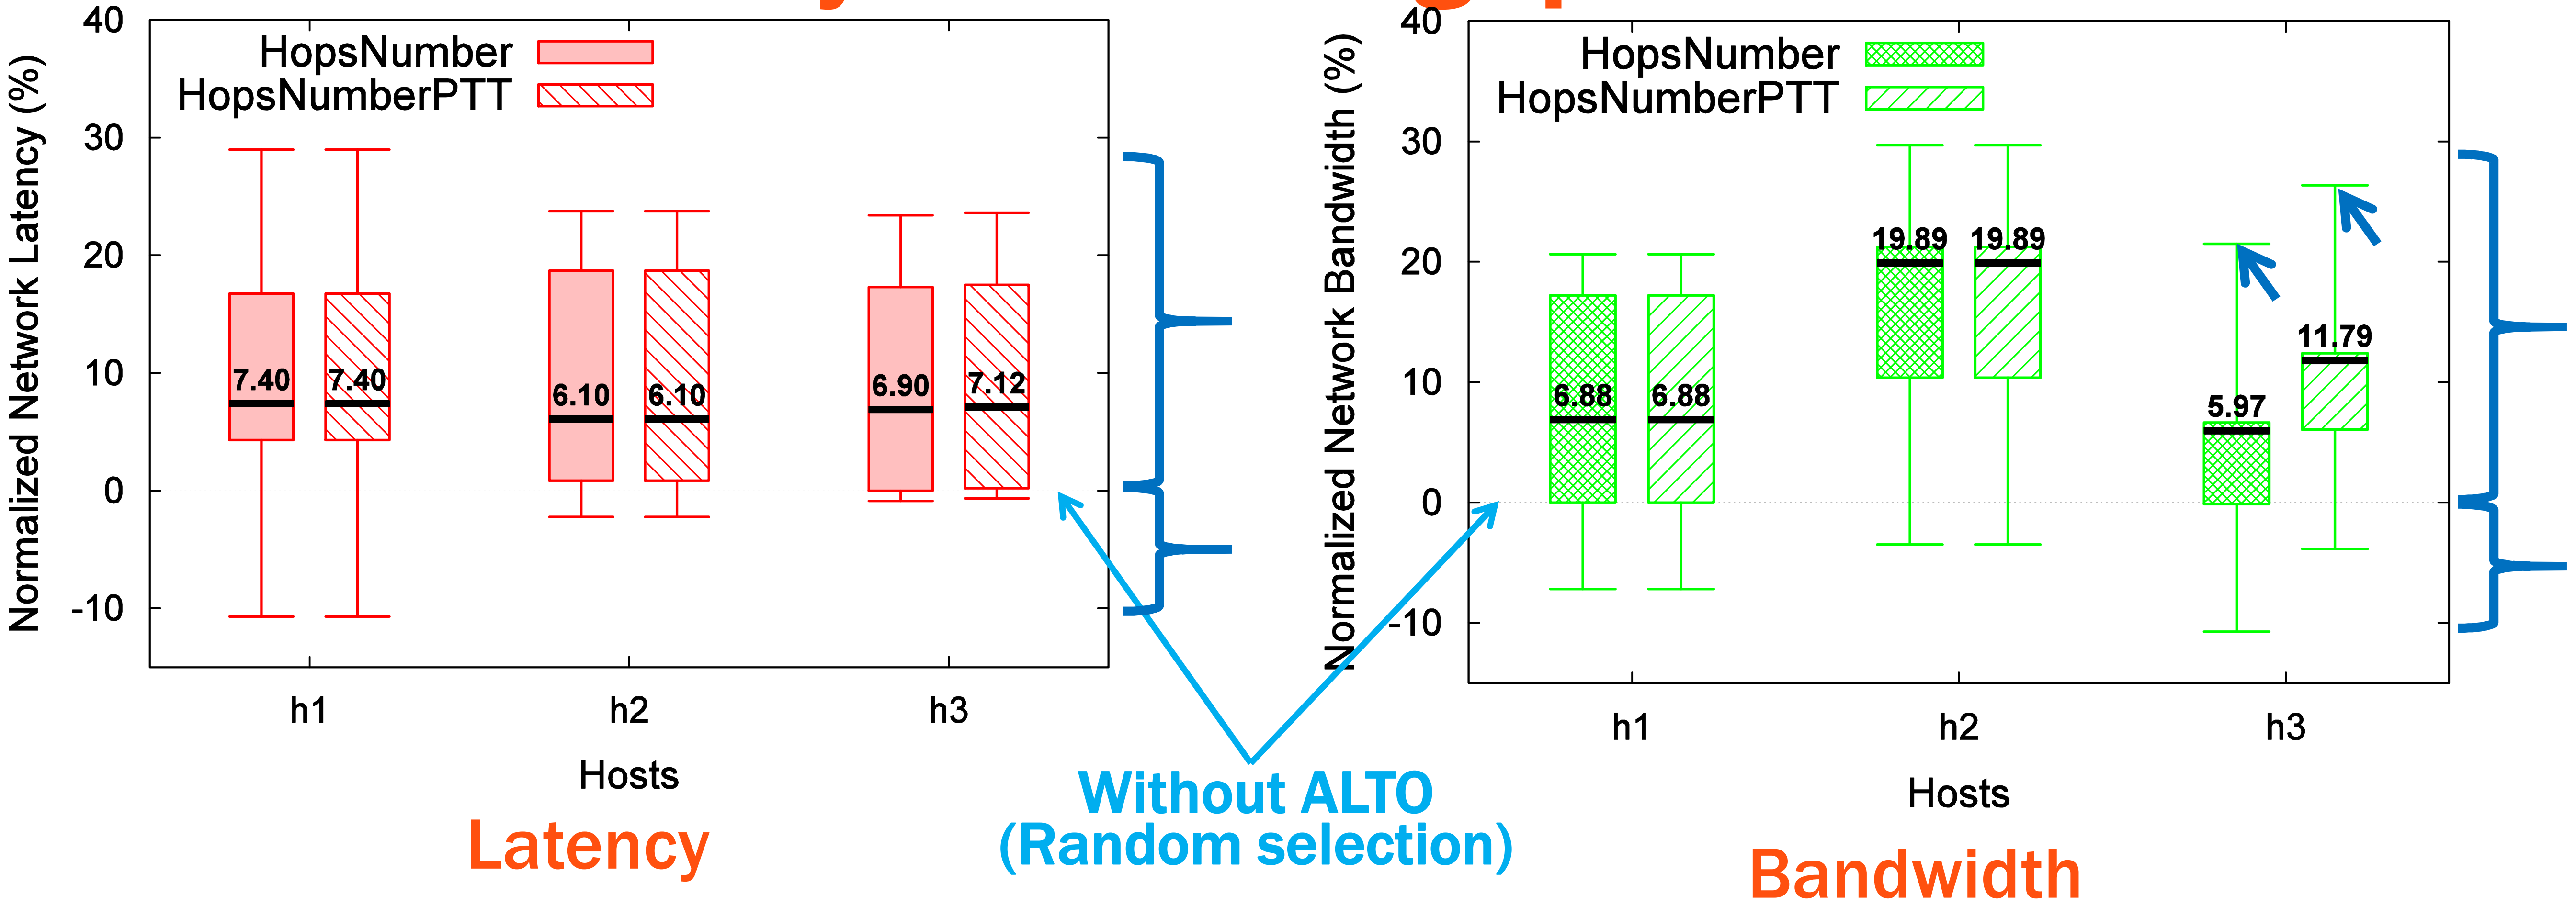

Results Analysis

- ✓ Improvements in latency and throughput of up to 29%.
- ✓ In a few cases, peers selected through AaaS ended with less bandwidth or higher latency (between 1% and 11%) : Shortest paths not necessary the best performing ones.
- ✓ When h3 uses the IXP infrastructure to select a peer (h1), as HopsNumberPTT suggests, further throughput improvements (up to 26%).

Experimental Evaluation :

2)Inter-Domain Traffic Reduction

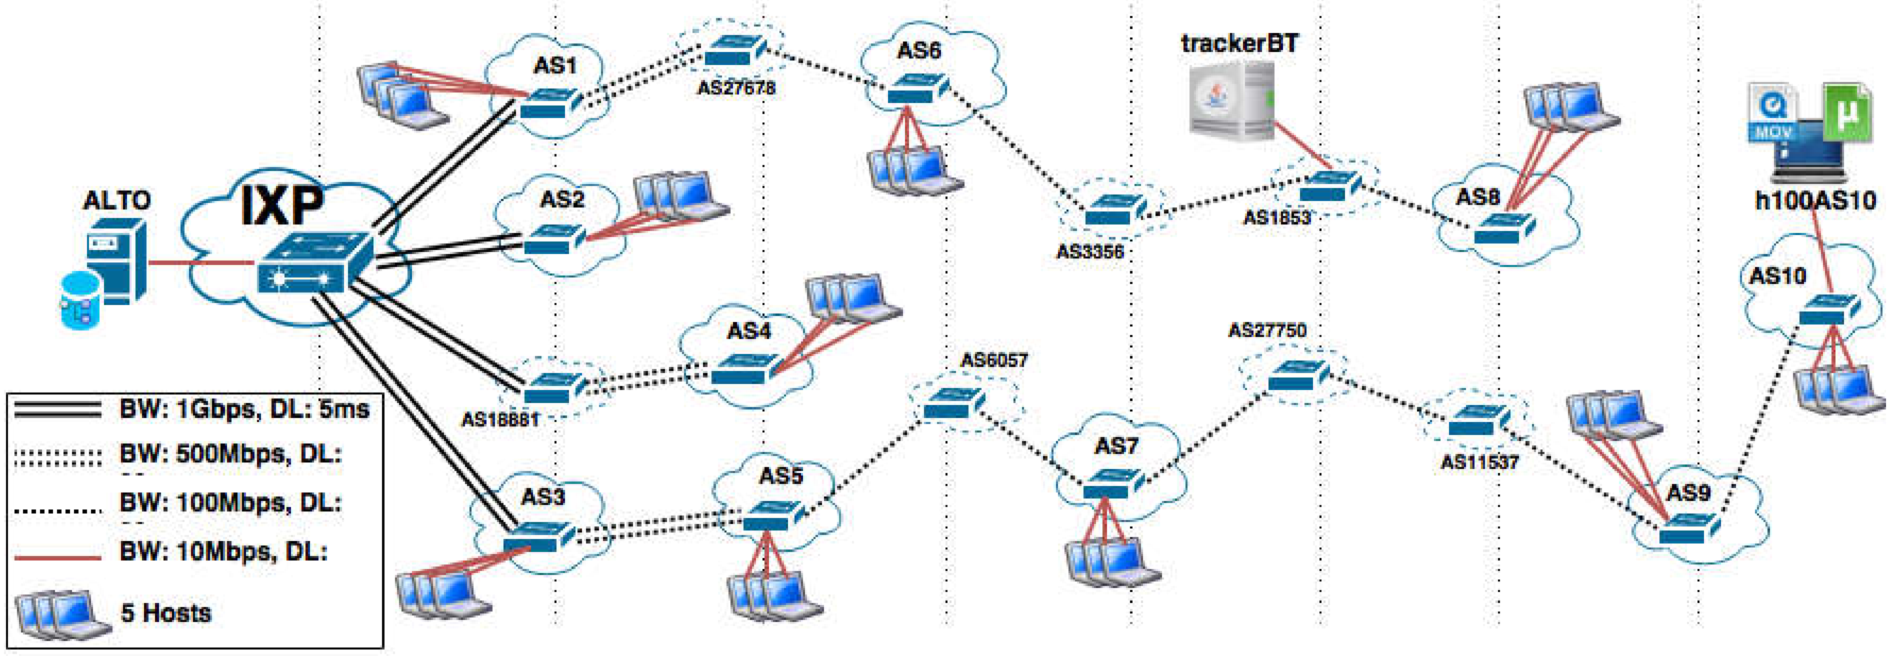

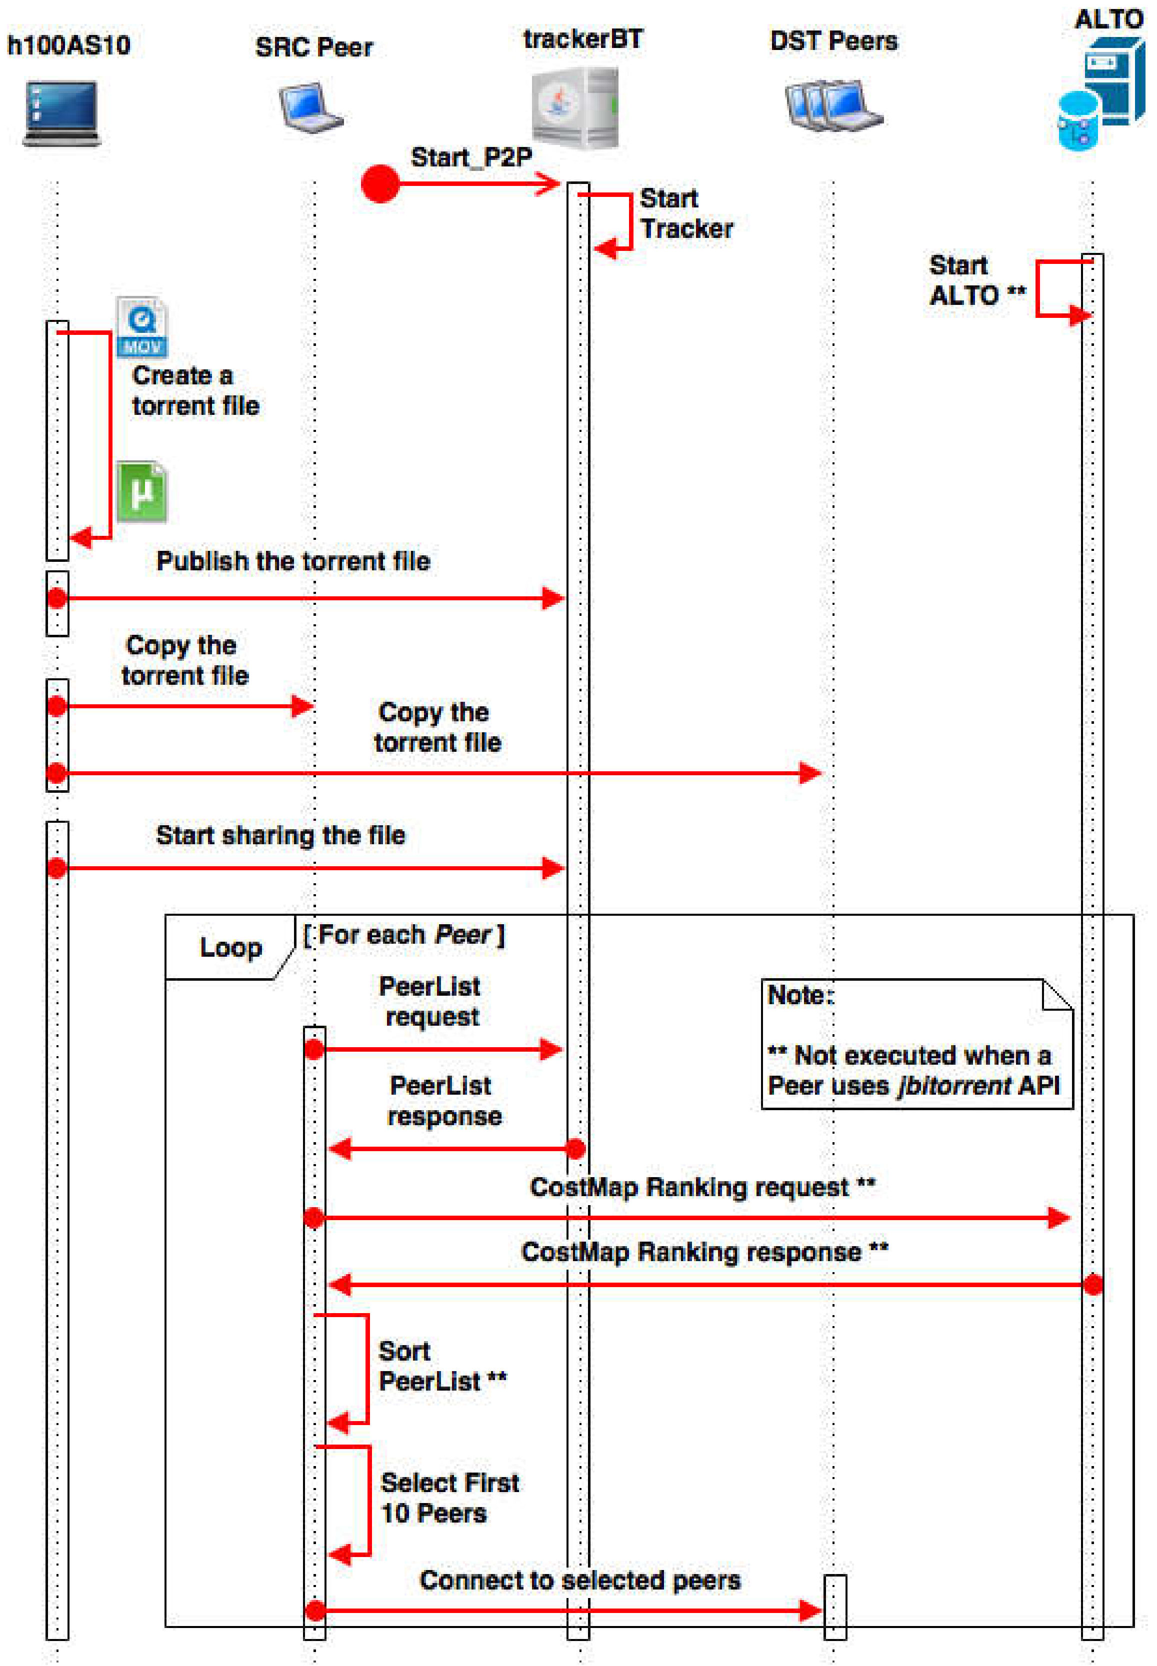

- ✓ we re-use a Java implementation of the BitTorrent protocol[1]
  - **trackerBT:** Single and easy-to-use Tracker.
  - **jbittorrent**: Allows the creation of torrent files and the download and sharing of files between peers (JBIT).
- ✓ A new package based on the jbittorrent API was generated to consume the ALTO services (JBIT_AaaS).

[1]https://github.com/cloudspaces/jbittorrent

Experimental Evaluation :

2) Inter-Domain Traffic Reduction

ALTO information

Absolute distance (HopsNumber)

	AS1	AS2	AS3	AS4	AS5	AS6	AS7	AS8	AS9	AS10
AS1	0	2	2	3	3	2	5	5	8	9
AS2	2	0	2	3	3	4	5	7	8	9
AS3	2	2	0	3	1	4	3	7	6	7

ALTO Clients

✓ Hosts whose ASes are directly connected at IXP (AS1, AS2 and AS3).

Workload and Metrics

✓ Two scenarios:

1. Better-than-random peer selection using an ALTO server (JBIT_AaaS).
2. Peer selection following the standard workflow of the BitTorrent protocol (JBIT).

✓ Each SRC peer generates a sample with the number of pieces extracted from the DST peers.

- Since each peer is associated with an ASN, we compute the amount of AS-hops to reach DST peers.
- Finally, we calculate the number of pieces retrieved in a given number of AS-hops from a SRC peer.

34

Experimental Evaluation :

AS-Hops Distribution

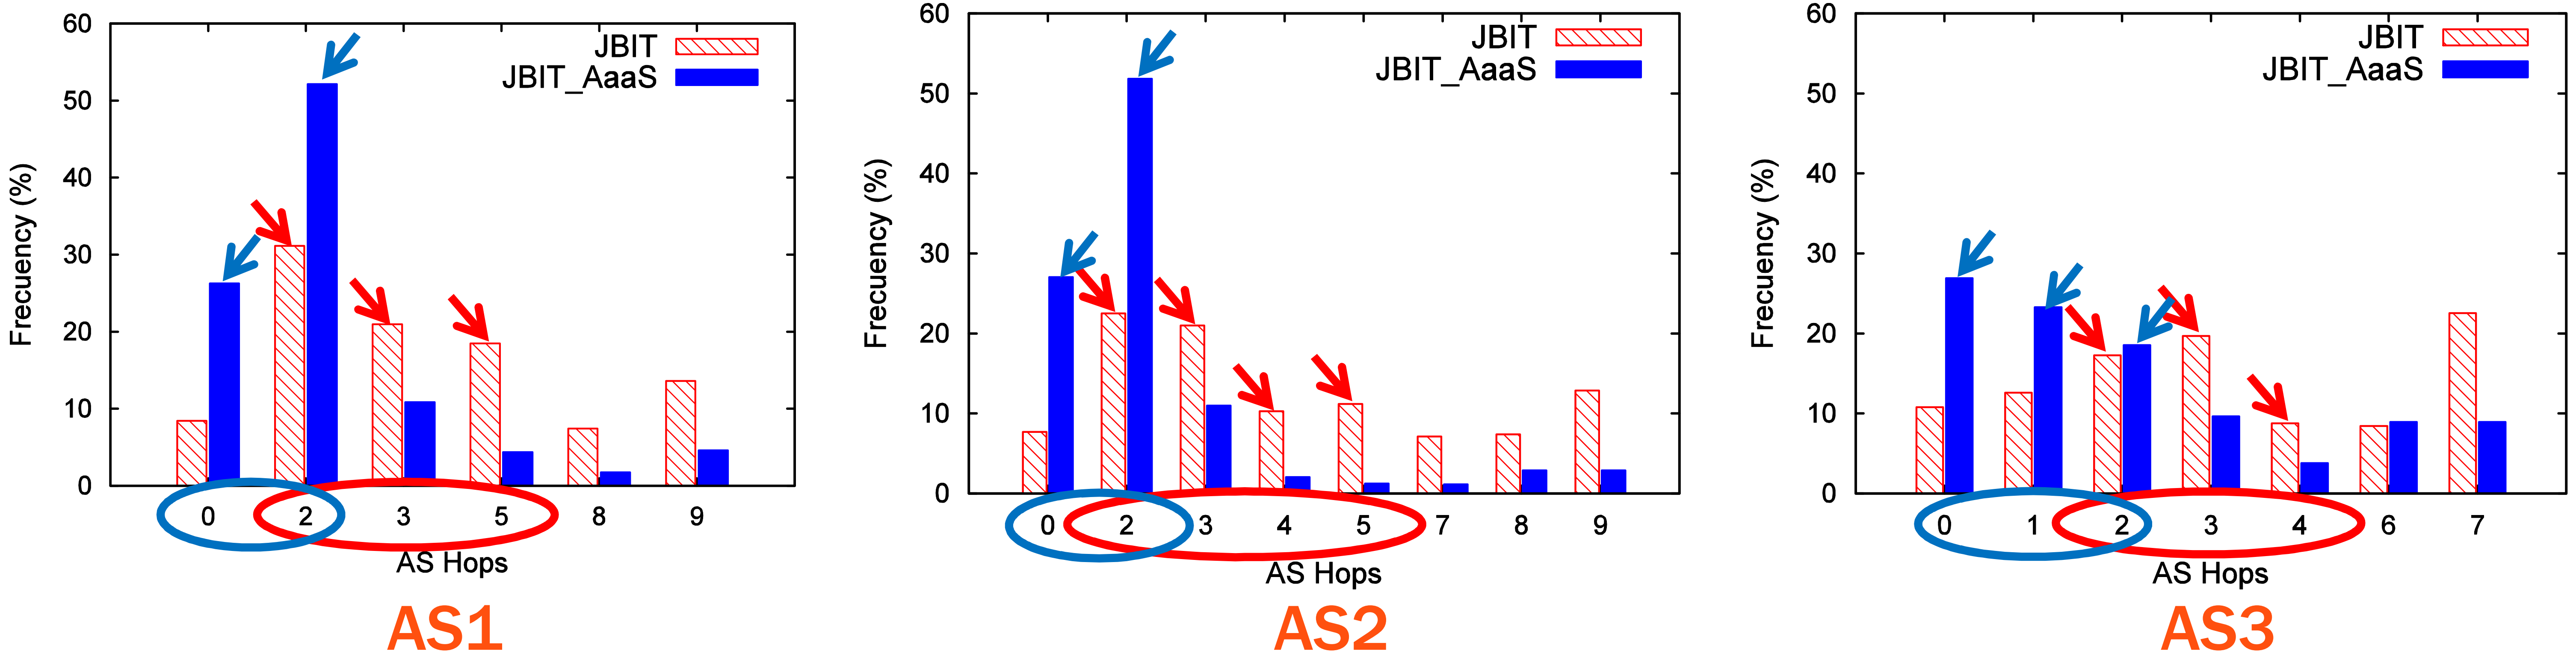

Results Analysis

✓ The higher concentration of pieces downloaded is at a depth between 2 to 5 AS-hops.

- The average number of AS-level hops varies between 3.5 and 4.3 hops.

✓ Selecting DST peers using JBIT_AaaS has a better performance. The higher concentration of pieces downloaded is at a depth between 0 and 2 AS-hops.

- The average number of AS-hops is between 1.9 and 2.1 AS-hops.

Experimental Evaluation :

Accumulated AS-Level Hops

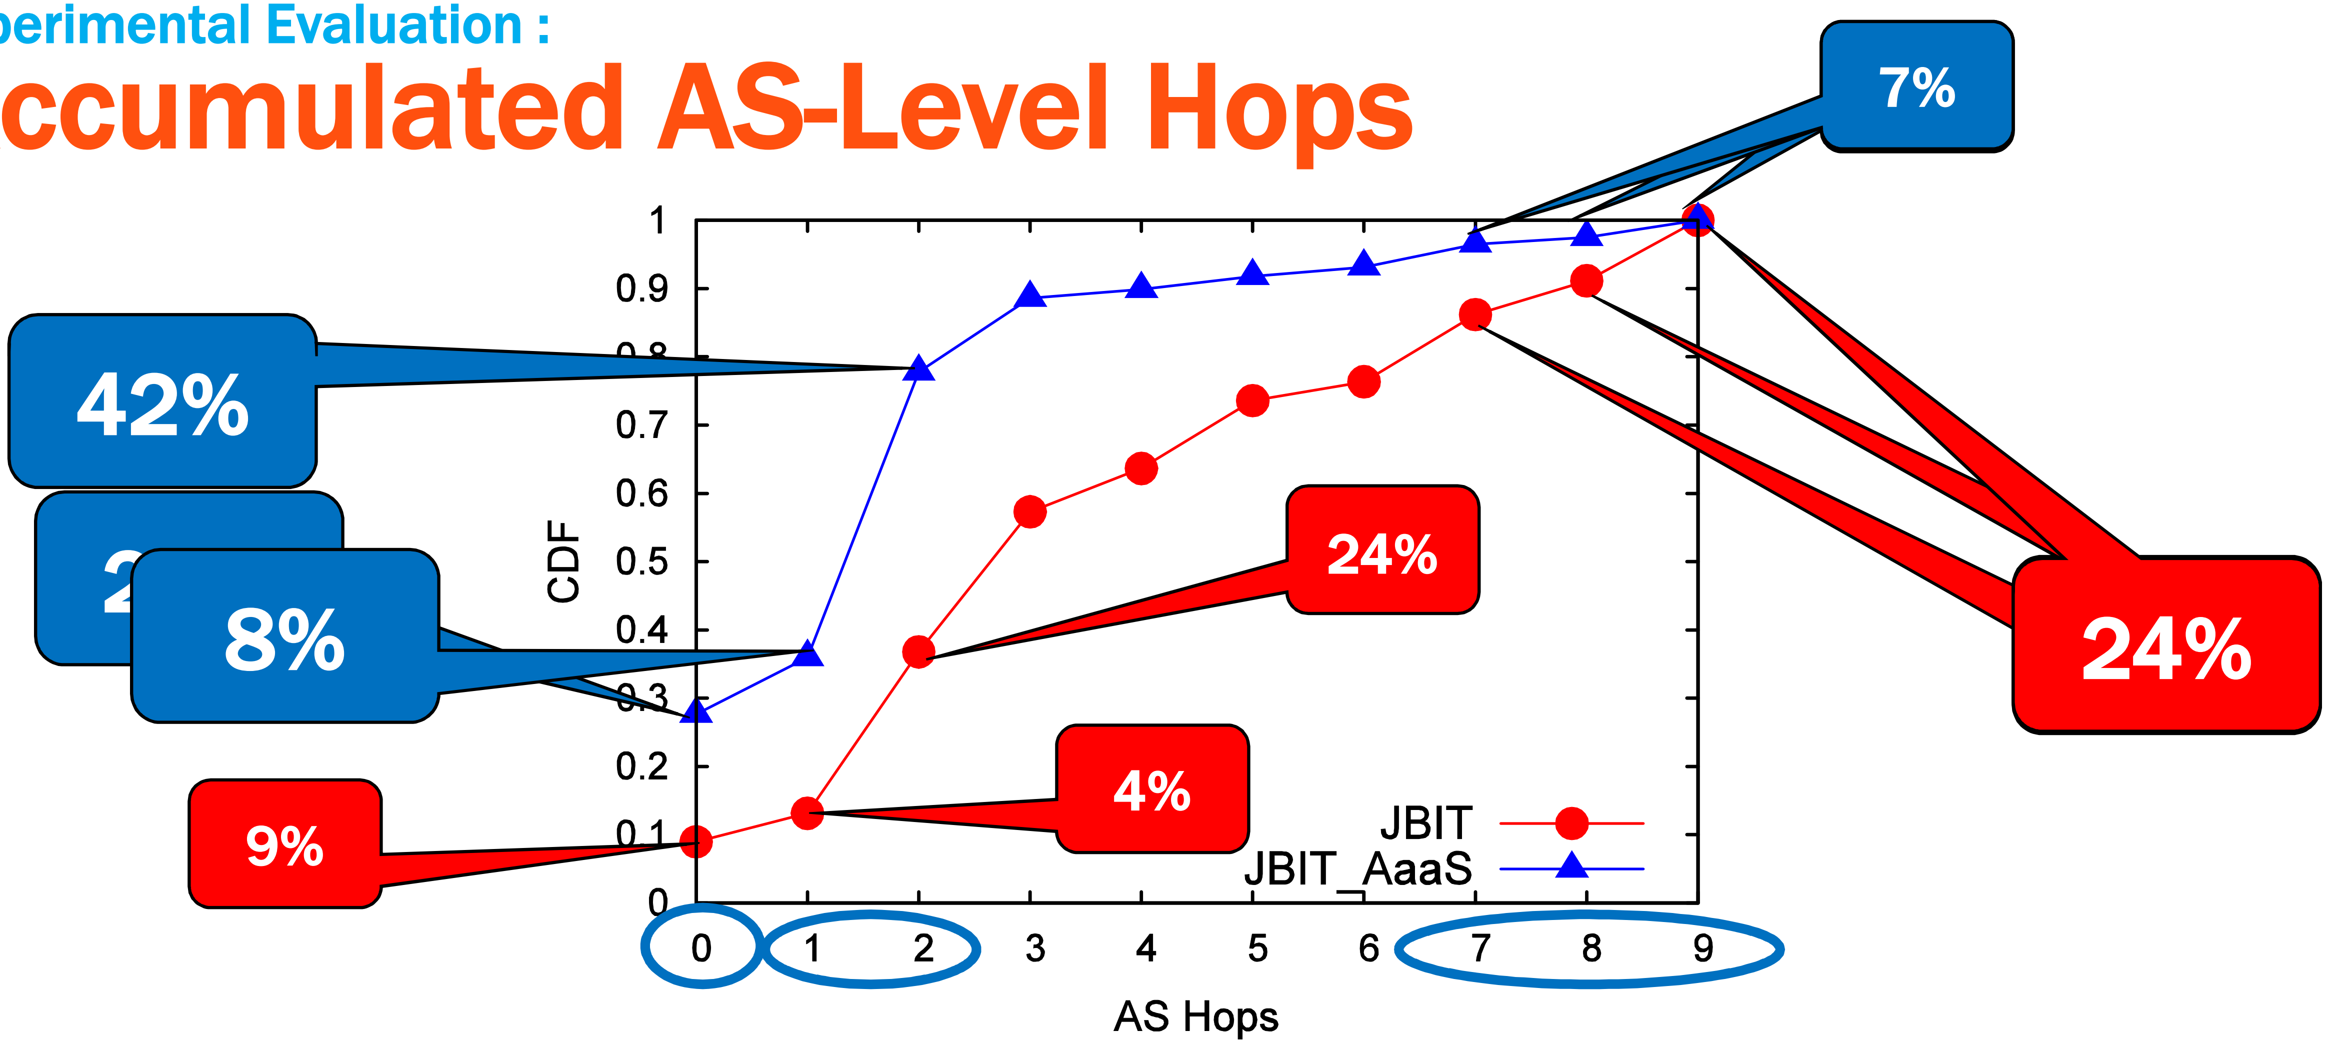

Results Analysis

- ✓ Over 28% of the times, peers selected by JBIT_AaaS do not leave the source AS, this is more than three times the performance achievable through JBIT (9%) at random.
- ✓ 50% of the total pieces are retrieved at a depth of 1 and 2 AS-hops, while JBIT-recommended peers in the same depth can only download 28% of the blocks for a file.
- ✓ The sum of pieces retrieved from AS8, AS9 and AS10 through SRC peers running the JBIT_AaaS API (7%) is less than three times the sum of pieces retrieved at random by JBIT (24%).

Experimental Evaluation :

3) Performance Improvement of a Real P2P App.

Table 10 – Torrent Files description (April 19th, 2016)

Name	# Seeder	# Leecher	Piece Length	File Name	File Size
Torrent1.torrent	1446	42	512 KB	ubuntu-15.10-desktop-i386.iso	1.14 GB
Torrent2.torrent	183	6	512 KB	debian-8.4.0-amd64-CD-1.iso	630 MB
Torrent3.torrent	558	22	256 KB	LibreOffice_5.1.2_Win_x86.msi	211.36 MB
Torrent4.torrent	304	35	512 KB	ubuntu-14.10-desktop-i386.iso	1.11 GB
Torrent5.torrent	3576	68	512 KB	ubuntu-15.10-desktop-amd64.iso	1.10 GB

- ✓ We use a single SRC peer located on University of Campinas (ASN: 53187).
  - Our Java implementation supports only HTTP GET requests to the TCP-based Trackers.
- ✓ Max number of DST peer list to receive from a Tracker is 100 (only IPv4 addresses) and half of peers are selected for the final DST peer list.
- ✓ JBIT_AaaS: A new input parameter is added in order to obtain the absolute o relative distance between SRC peer and DST peers.

Experimental Evaluation :

Global Download Rate

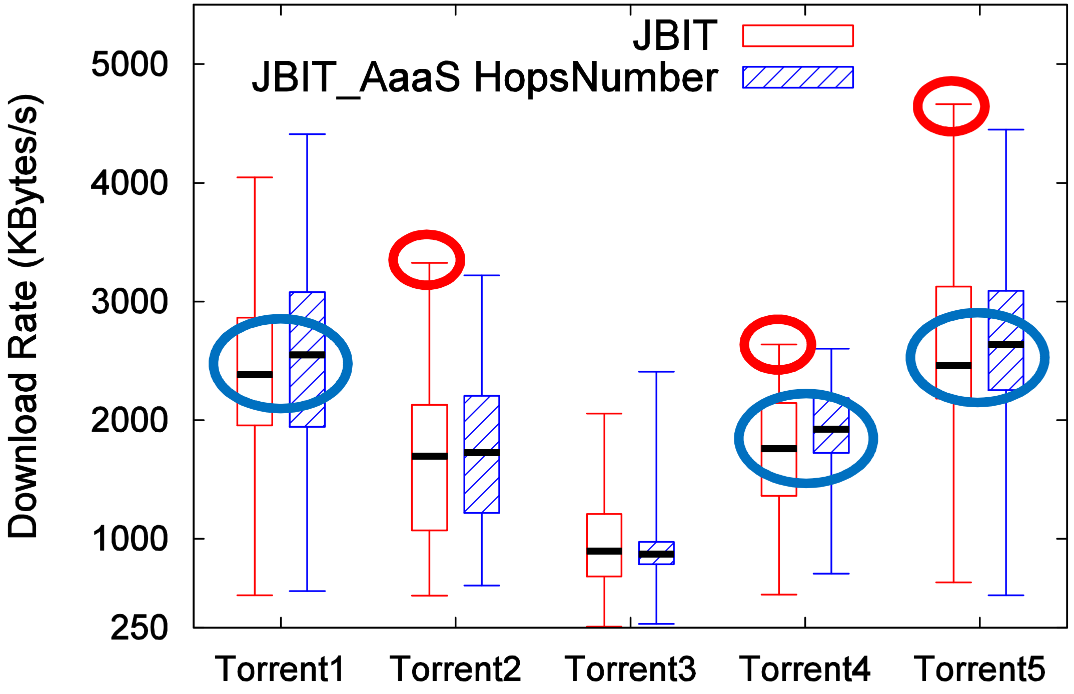

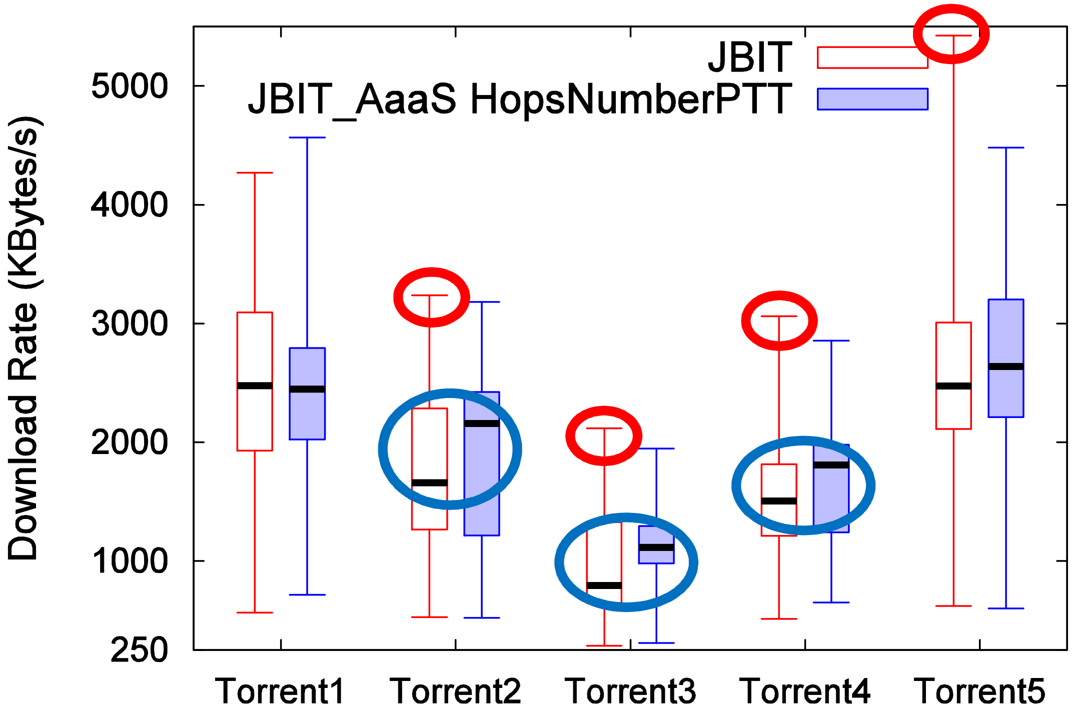

Results Analysis

- ✓ The difference of median download rates for JBIT_AaaS and JBIT is between 166 and 178 KBytes/s. Exceptions: torrent2 (30 KBytes/s) and torrent3 (-26 KBytes/s).
- ✓ JBIT provides the highest peaks of download rate (torrent2, torrent4 and torrent5).

- ✓ The difference in median download rate is above 300 KBytes/s (torrent3 and torrent4), reaching around 500 KBytes/s in the case of torrent2. Exceptions: torrent5 (163 Kbytes/s) and torrent1 (-29 KBytes/s).
- ✓ JBIT achieves the higher peak download rate of all our experiments (5426 KBytes/s).

Experimental Evaluation :

Mean Global Download Rate (95% CL)

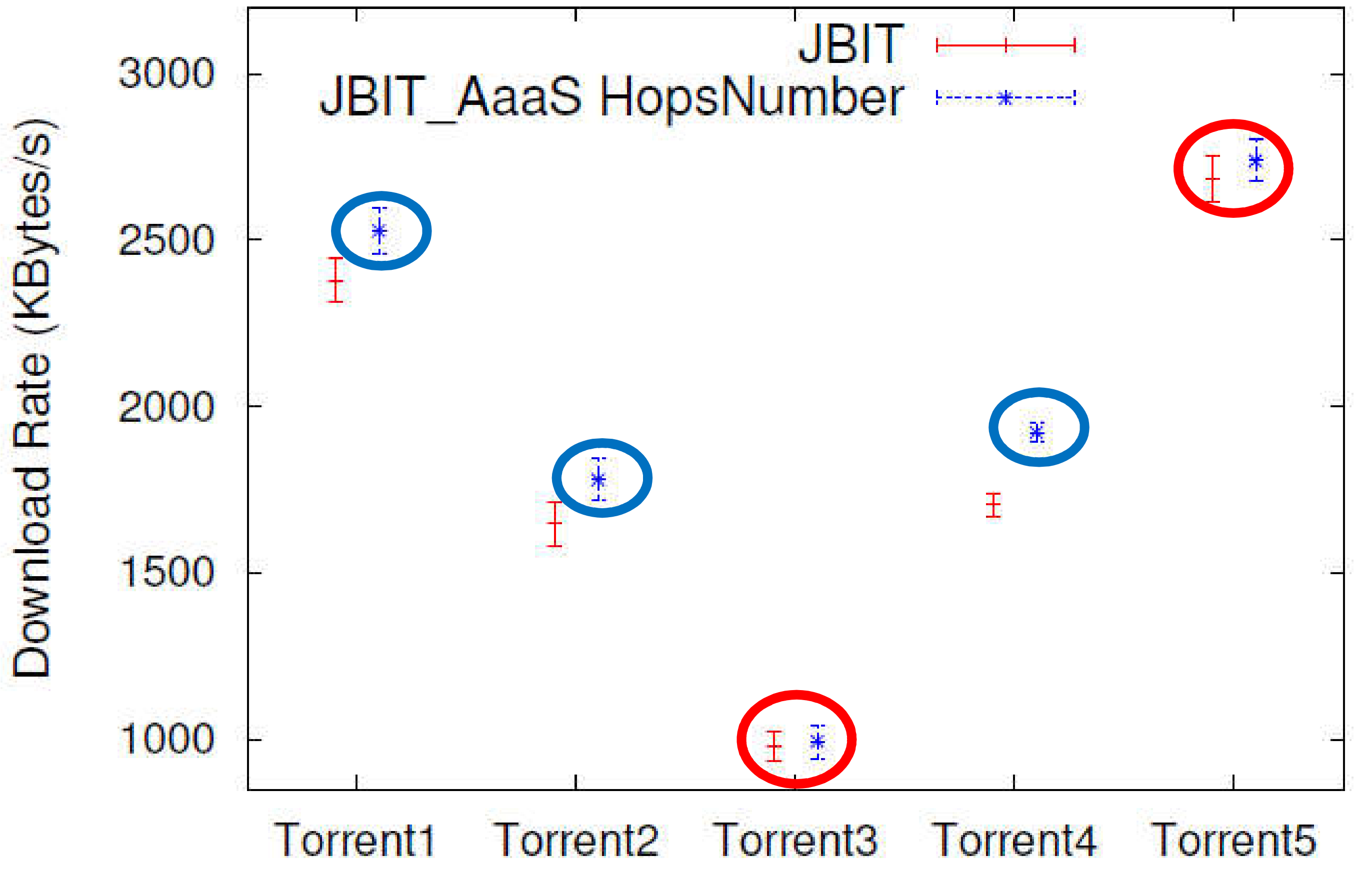

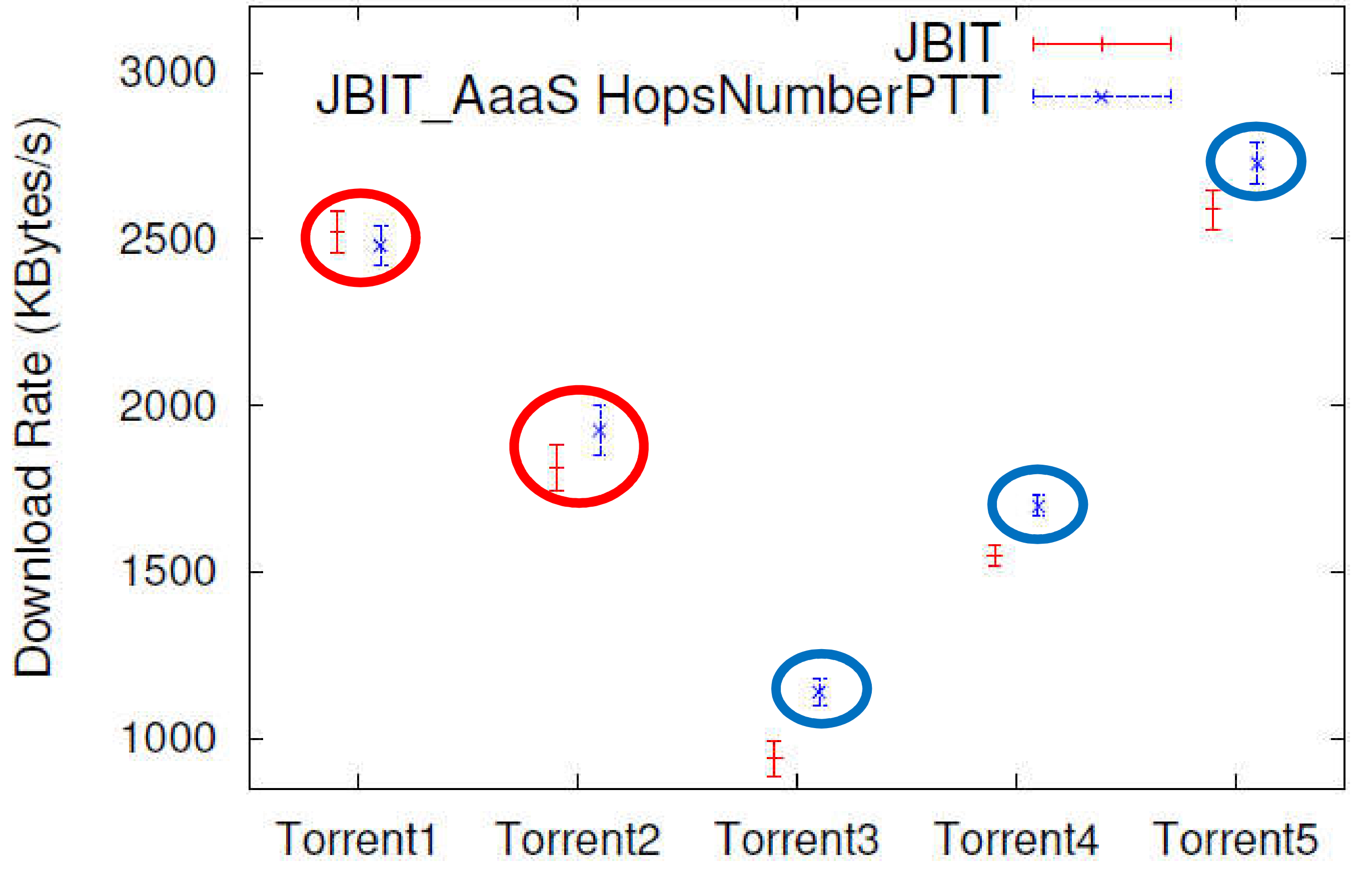

- ✓ torrent1, torrent2 and torrent4: We can state at 95% of confidence that JBIT_AaaS have a higher download rate than JBIT at random.
- ✓ torrent3, torrent4 and torrent5: our approach (JBIT_AaaS) shows statistically significantly better download rate than a standard Bit-Torrent protocol.
- ✓ Others torrents: it is not possible to state that JBIT_AaaS has a better performance than JBIT or vice versa.

Experimental Evaluation :

Destination Peer Statistics (1/2)

Torrent File	JBIT					JBIT_AaaS (HopsNumber)				
	Peers			ASes		Peers			ASes	
	Nº	Download* Rate (KB/s)	R^{+}	Nº	Depth*	Nº	Download* Rate (KB/s)	R^{+}	Nº	Depth*
Torrent1	221	98.9 (±7.4)	0.66	125	3.7 [3.2]	166	96.4 (±7.6)	0.63	75	3.6 [2.9]
Torrent2	74	99.6 (±13.9)	0.64	50	3.7 [3.1]	48	102.8 (±16.9)	0.74	28	3.5 [2.8]
Torrent3	116	72.6 (±14.1)	0.90	90	3.9 [3.3]	77	88.2 (±19.1)	0.68	54	3.5 [2.8]
Torrent4	54	108.9 (±18.6)	0.72	41	3.7 [3.2]	30	116.2 (±22.8)	0.86	23	3.5 [2.7]
Torrent5	281	104.9 (±6.7)	0.66	161	3.8 [3.3]	230	112.9 (±16.8)	0.86	105	3.5 [2.8]

- ✓ JBIT splits each download rate into a large and random number of peers. JBIT_AaaS connects to the same set of peers continuously.
- ✓ JBIT: each connection obtains a relatively low individual transfer rate.
- ✓ JBIT_AaaS obtains a CC value greater than peers randomly selected by JBIT.
- ✓ # ASes reached by JBIT_AaaS proves to be lower than the number of ASes selected by a JBIT.
- ✓ Depth varies between 3.7 and 3.9 (JBIT) and Depth values lie between 3.5 and 3.6 (JBIT_AaaS).

Experimental Evaluation :

Destination Peer Statistics (2/2)

Torrent File	JBIT					JBIT_AaaS (HopsNumberPTT)				
	Peers			ASes		Peers			ASes	
	Nº	Download* Rate (KB/s)	R^+	Nº	Depth*	Nº	Download* Rate (KB/s)	R^+	Nº	Depth*
Torrent1	199 ↑	104.7 (±8.0) ↑	0.64	113 ↑	3.6 [3.2]	157 ↓	96.8 (±8.0)	0.67 ↑	68 ↓	3.6 [3.0]
Torrent2	75 ↑	103.4 (±13.1) ↑	0.60	48 ↑	3.6 [3.1]	50 ↓	94.1 (±18.0)	0.82 ↑	33 ↓	3.6 [2.9]
Torrent3	112 ↑	79.9 (±14.1)	0.73 ↑	80 ↑	3.8 [3.3] ↑	78 ↓	86.6 (±10.9) ↑	0.65	55 ↓	3.6 [2.9] ↓
Torrent4	59 ↑	101.8 (±18.0)	0.70	45 ↑	3.6 [3.2] ↑	37 ↓	108.5 (±20.3) ↑	0.80 ↑	23 ↓	3.4 [2.8] ↓
Torrent5	268 ↑	99.9 (±7.3)	0.56	159 ↑	3.7 [3.1] ↑	246 ↓	101.4 (±7.7) ↑	0.69 ↑	106 ↓	3.5 [3.0] ↓

- ✓ JBIT splits each download rate into a large and random number of peers. JBIT_AaaS connects to the same set of peers continuously.
- ✓ JBIT: each connection obtains a relatively low individual transfer rate.
- ✓ JBIT_AaaS obtains a CC value greater than peers randomly selected by JBIT.
- ✓ # ASes reached by JBIT_AaaS proves to be lower than the number of ASes selected by a JBIT.
- ✓ Depth varies between 3.6 and 3.8 (JBIT) and Depth values lie between 3.4 and 3.6 (JBIT_AaaS).

Experimental Evaluation :
Total Download Time

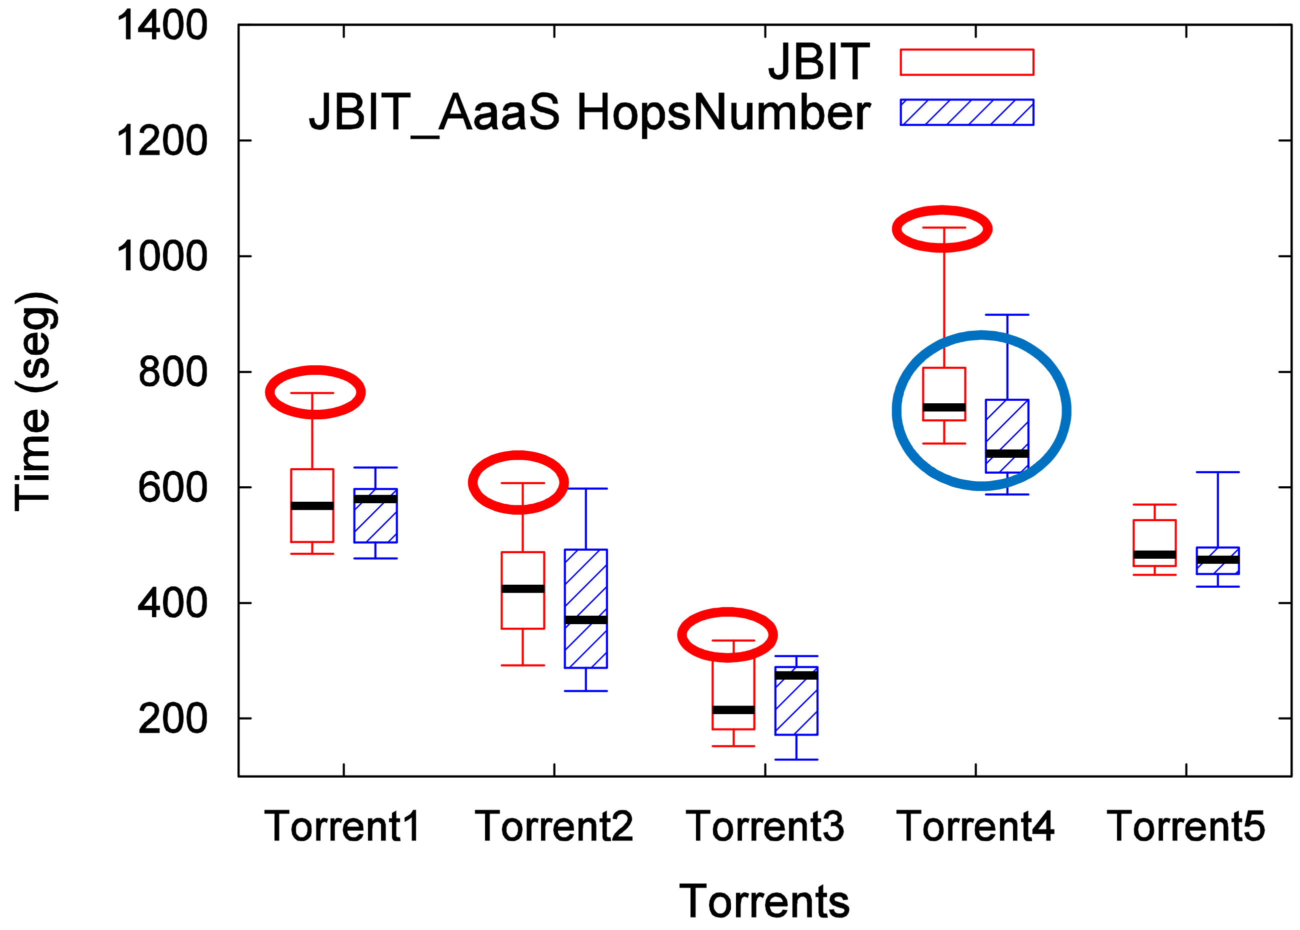

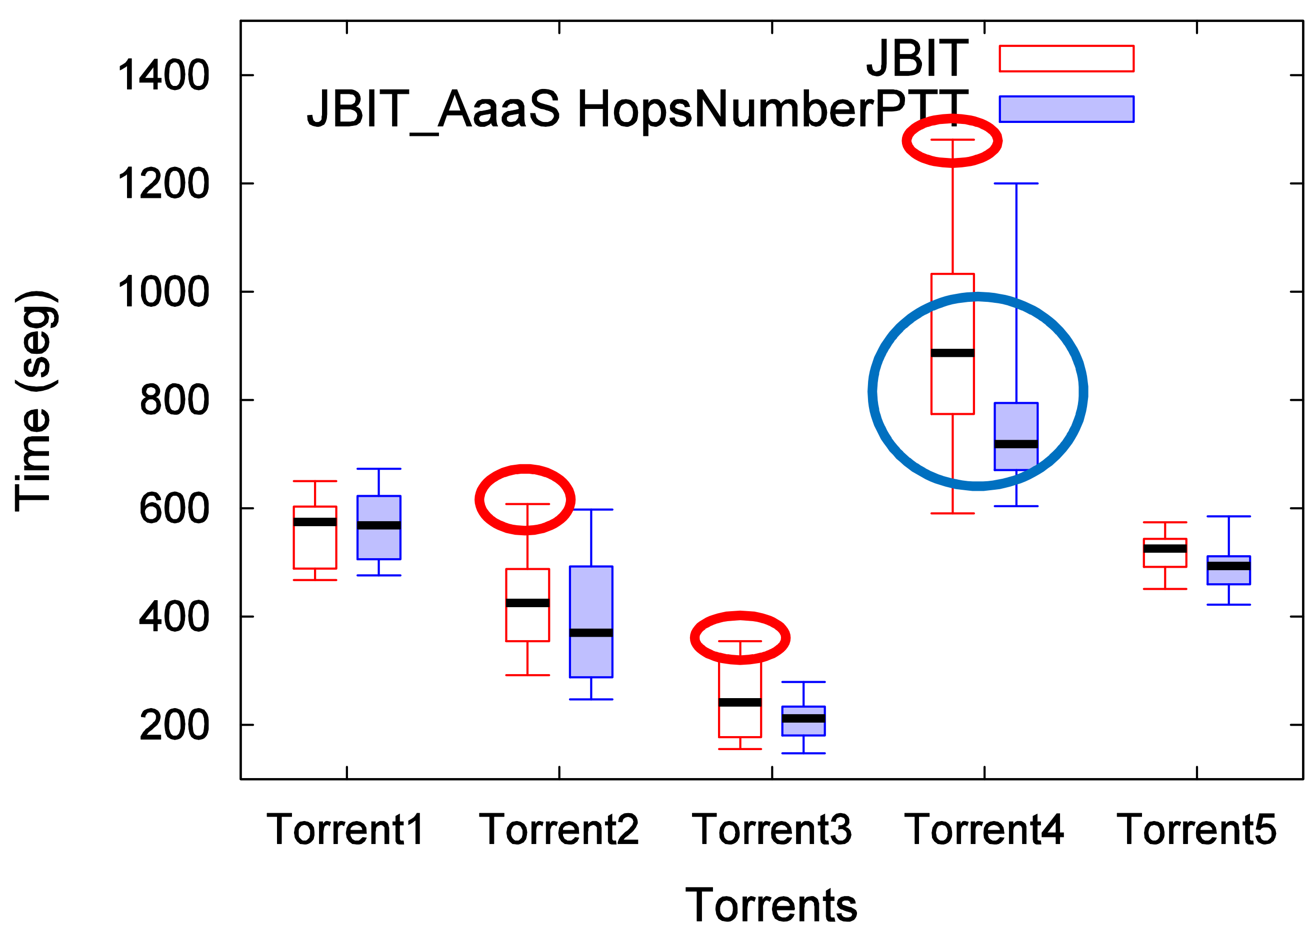

Results Analysis

✓ The maximum values were obtained by JBIT.

✓ The difference of median values are generally in favour of JBIT_AaaS.

- HopsNumber: The greatest median difference is calculated in torrent4 (-80 seg.)
- HopsNumberPTT: Torrent4 achieves the best difference (-169 seg.).

Agenda

1. Introduction
2. Background
3. Related work
4. Objectives
5. Design of ALTO-as-a-Service (AaaS)
6. Prototype Implementation
7. Experimental Evaluation
8. **Future Work / Conclusions**

Future work

- ✓ The Cost Map rankings are based on **relatively static** AS-Path distance and do not consider more dynamic information such as actual bandwidth, latency, packet loss rate, etc.
  - Thus, **dynamic updates** of cost maps based on public **Internet quality measurements** (e.g., SIMET, RIPE TTM) are in our research agenda.
- ✓ The AaaS workflow **does not handle periodic updates** (e.g. when a new dataset is retrieved) and a single snapshot was used in order to built the ALTO information.
  - Therefore, the creation of **a procedural flow** where the AaaS can **support scenarios with multiple datasets** is at the top in our future research activities.
- ✓ Implement the remaining ALTO services (e.g., full Map-Service, EPS and ECS).
- ✓ Prototype performance optimization (e.g., Neo4j query tuning techniques, Linux fs configuration, security).

Ongoing work

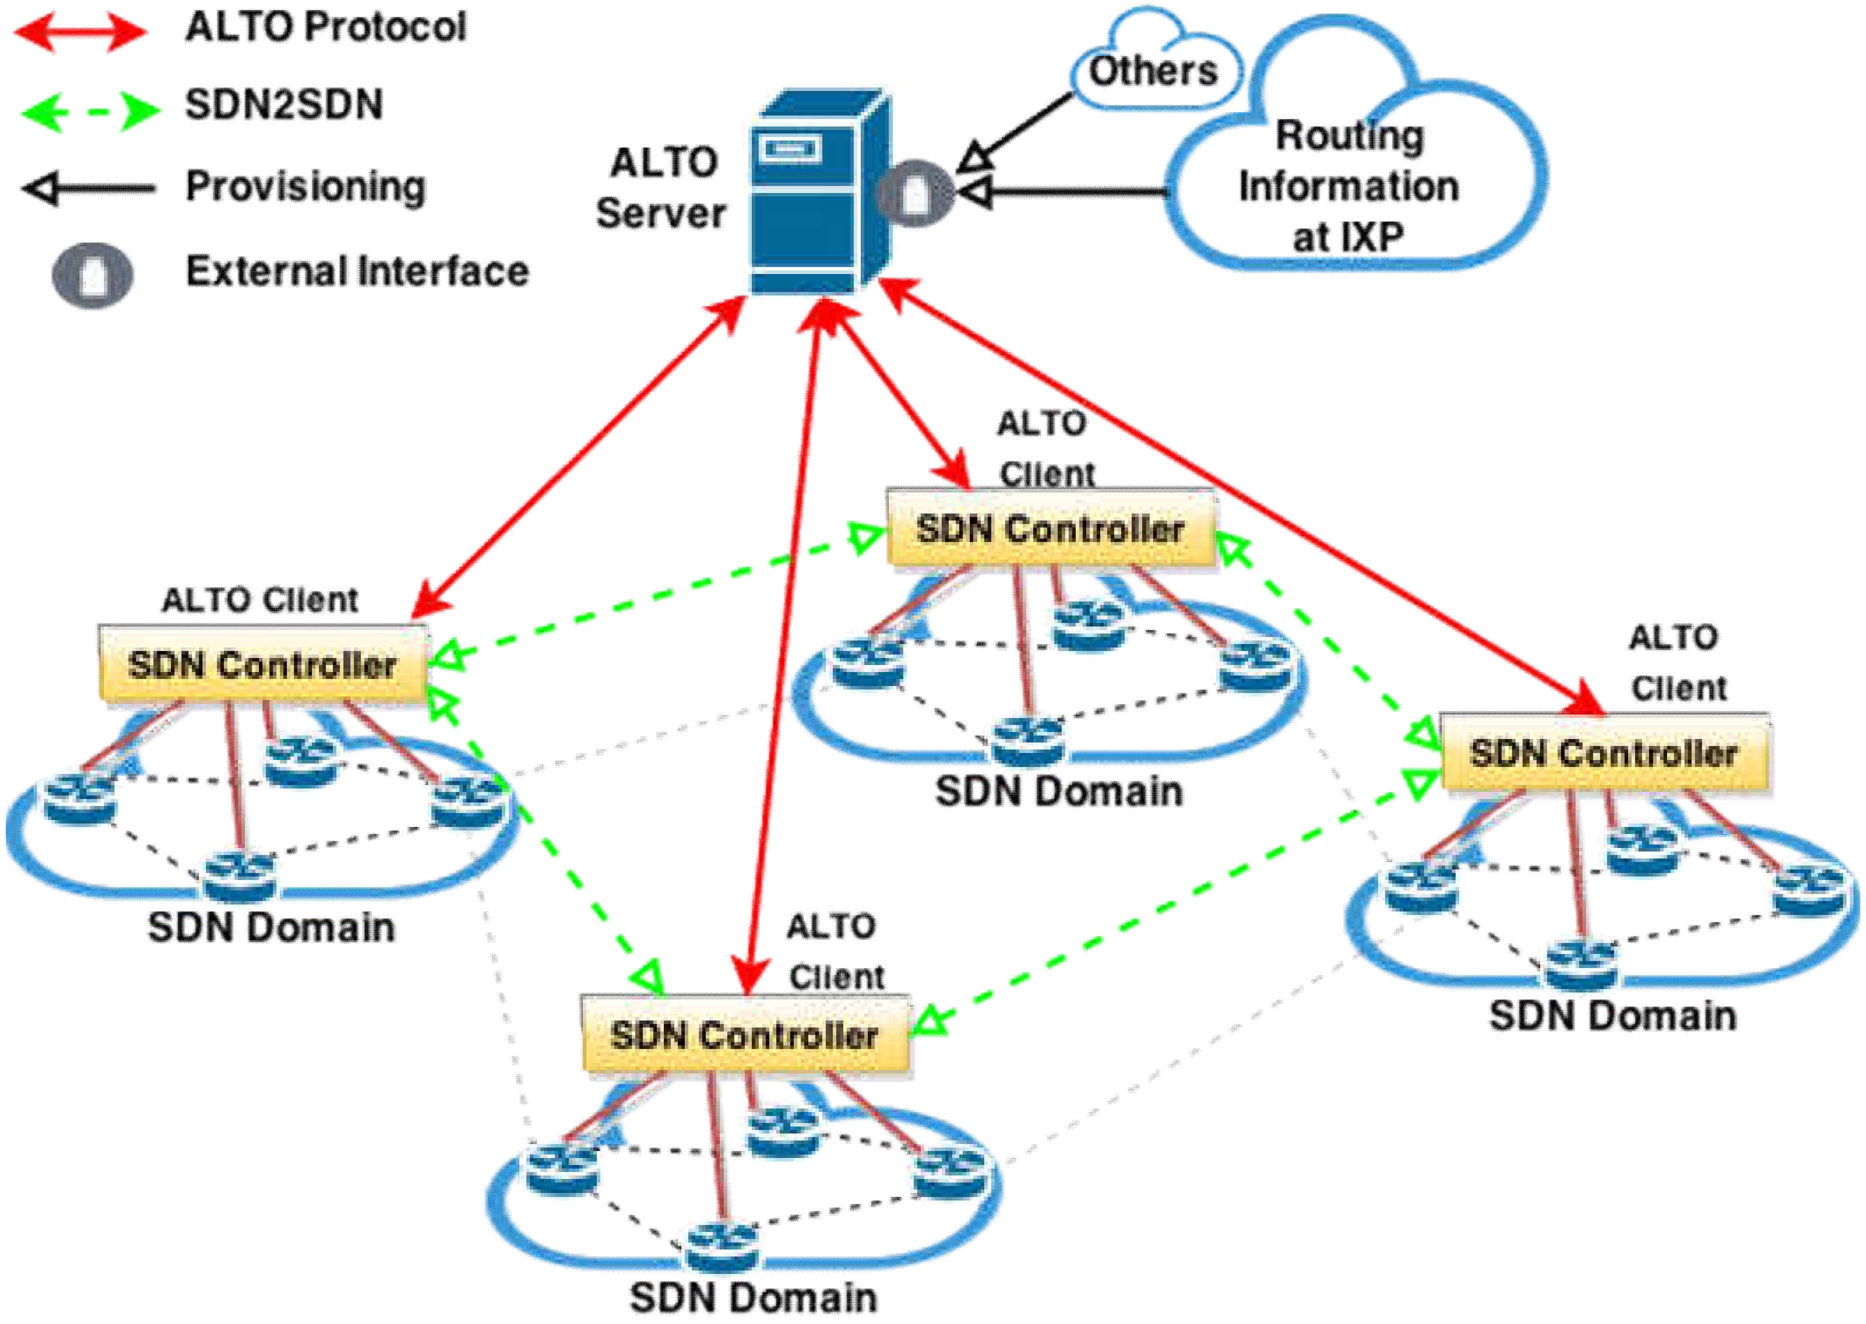

ALTO – SDN Use Cases

Conclusions

- ✓ This is the first work that explores the use of Inter-domain routing data publicly available at IXPs to create abstract topology and cost maps following the recently standardized IETF ALTO protocol.
- ✓ Our proposed architecture (AaaS) encompasses the whole process of ALTO service delivery.
- ✓ Our AaaS prototype is based on the popular Neo4j graph database and the OpenDaylight controller and validated the potential of applications to leverage the network awareness provided by ALTO servers.
- ✓ Our approach can reduce the amount of traffic that crosses network boundaries, resulting in ISPs' cost reductions and ALTO service providers (in our case IX.br operators) can benefit from increased IXP traffic exchange.

Publications

- ✓ **LACHOS, D. A.** ; BRITO, S. H. B. ; FONTES, R. R. ; ROTHENBERG, C. E. . Delivering Application-Layer Traffic Optimization Services based on Public Routing Data at Internet eXchange Points. In: *XXXIV Simpósio Brasileiros de Redes de Computadores SBRC 2016*, Salvador, BA, Brazil, Jun 2016.

- ✓ BRITO, S. H. B.; SANTOS, M. A. S. ; FONTES, R. R. ; **LACHOS, D. A.** ; ROTHENBERG, C. E. . Dissecting the Largest National Ecosystem of Public Internet eXchange Points in Brazil. In: *Passive and Active Measurements PAM 2016*, Hereklion, Crete, Greece, 2016.

- ✓ BRITO, S. H. B.; SANTOS, M. A. S. ; FONTES, R. R. ; **LACHOS, D. A.** ; ROTHENBERG, C. E. . Anatomia do Ecossistema de Pontos de Troca de Tráfego Públicos na Internet do Brasil. In: *XXXIII Simpósio Brasileiros de Redes de Computadores SBRC 2015*, Vitória, ES, Brazil, May 2015.

- ✓ **LACHOS, D.A.**; ROTHENBERG, C. E. . Delivering Application-Layer Traffic Optimization Services based on Public Routing Information at Internet eXchange Points. In: *VIII Encontro dos Alunos e Docentes do Departamento de Engenharia de Computação e Automação Industrial 2015*, Campinas, SP, Brazil, Set 2015.

Thanks! Obrigado! (More) Questions?

dlachosp@dca.fee.unicamp.br

http://lattes.cnpq.br/5466177320244302

INTRIG INFORMATION & NETWORKING TECHNOLOGIES RESEARCH & INNOVATION GROUP

FEEC Faculdade de Engenharia Elétrica e de Computação

UNICAMP

BACKUP

References

[1] S. H. B. Brito, M. A. S. Santos, R. dos Reis Fontes, D. A. L. Perez, and C. E. Rothenberg, *“Anatomia do ecossistema de pontos de troca de trafego publicos na internet do brasil,”* In XXXIII Simposio Brasileiro de Redes de Computadores (SBRC). May 18-22, 2015. Vitoria, ES, Brazil.

[2] A. Khan, T. Kwon, H.-c. Kim, and Y. Choi, *“As-level topology collection through looking glass servers,”* in Proceedings of the 2013 Conference on Internet Measurement Conference, IMC '13, (New York, NY, USA), pp. 235-242, ACM, 2013.

[3] UCLA IRL. http://irl.cs.ucla.edu/topology.

[4] CAIDA ARK. http://www.caida.org/projects/ark.

[5] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, “iplane: An information plane for distributed services,” in Proceedings of the 7th symposium on Operating systems design and implementation, pp. 367-380,USENIX Association, 2006.

[6] Internet Routing Registry. http://www.irr.net.

References

[7] CHOFFNES, D. R.; BUSTAMANTE, F. E. Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems. In: ACM. ACM SIGCOMM Computer Communication Review. [S.l.], 2008. v. 38, n. 4, p. 363–374. Citado 2 vezes nas páginas 29 and 31.

[8] BONAVENTURE, O.; SAUCEZ, D.; DONNET, B. The case for an informed path selection service. [S.l.], 2008. Citado 2 vezes nas páginas 29 and 31.

[9] XIE, H.; YANG, Y. R.; KRISHNAMURTHY, A.; LIU, Y. G.; SILBERSCHATZ, A. P4p: Provider portal for applications. In: ACM. ACM SIGCOMM Computer Communication Review. [S.l.], 2008. v. 38, n. 4, p. 351–362. Citado 2 vezes nas páginas 30 and 31.

[10] V. K. Gurbani, D. Goergen, R. State, and T. Engel, "Making historical connections: Building application layer trac optimization (ALTO) network and cost maps from public broadband data," in 10th International Conference on Network and Service Management, CNSM 2014 and Workshop, Rio de Janeiro, Brazil, November 17-21, 2014 (D. Raz, M. Nogueira, E. R. M. Madeira, B. Jennings, L. Z. Granville, and L. P. Gaspary, eds.), pp. 193-198, IEEE, 2014.

References

[11] PINTHONG, N.; LILAKIATSAKUN, W. Performance of bittorrent-like p2p file sharing systems inspired by alto. In: TENCON IEEE Region 10 Conference (31194). [S.l.: s.n.], 2013. Citado 2 vezes nas páginas 28 and 31.

[12] GUANXIU, L.; SUQI, Y.; XINLI, H. A Novel ALTO Scheme for BitTorrent-Like P2P File Sharing Systems. In: Proceedings of the 2013 Third International Conference on Intelligent System Design and Engineering Applications. [S.l.: s.n.], 2013. Citado 2 vezes nas páginas 28 and 31.

[13] SHIBUYA, M.; HEI, Y.; OGISHI, T. ISP-friendly peer selection mechanism with ALTO-like server. In: Network Operations and Management Symposium (APNOMS). [S.l.: s.n.], 2011. Citado 2 vezes nas páginas 29 and 31.

Design of ALTO-as-a-Service (AaaS) :

Creating and Delivering ALTO Information

Create the ALTO information and populate the graph DB

1. Grouping of prefixes into PID (Network Map) by ASes, IXPs, BGP communities, points of presence, just to cite a few examples.

2. Defining the preferences / costs between the groups PID (Cost Map) expressed on a path cost such as physical distance between IXPs, topological distances between ASes, etc.

Delivering ALTO Services

The last step is deploying an ALTO web server implementing the client-server protocol delivering the REST/JSON APIs to ALTO clients as defined by RFC7285.

Experimental Evaluation :

System Performance Profiling (2/2)

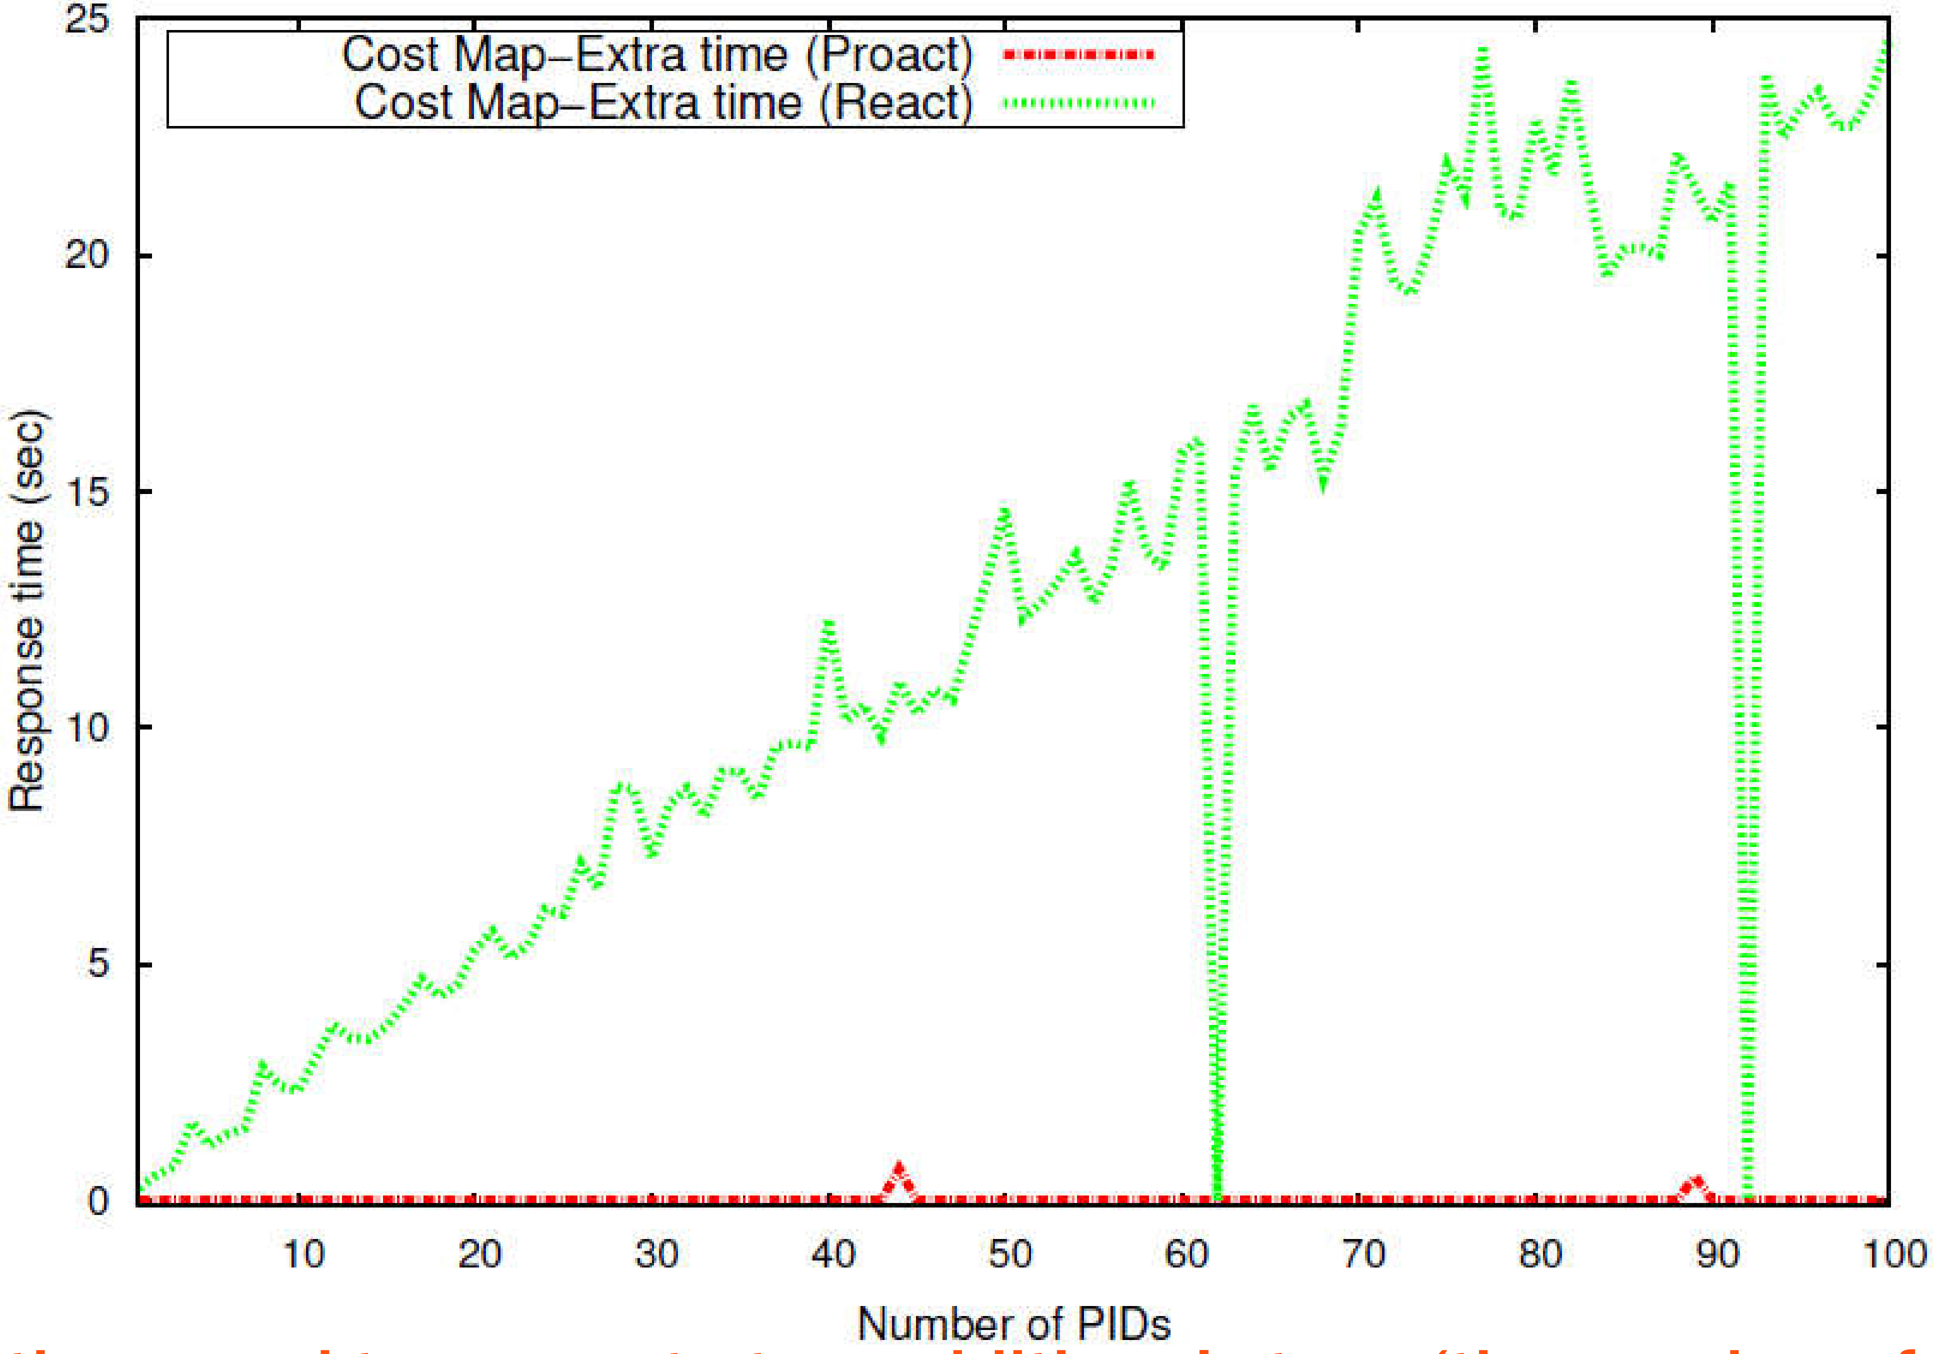

Processing time used to compute two additional steps (the number of hops and insert it into database).

Experimental Evaluation :

Use Case Scenario

	h1	h2	h3	h4	h5	h6	h7	h8	h9	h10
h1	—	120.5	120.6	160.6	160.6	180.5	240.8	300.7	360.9	401.0
	—	(±0.4)	(±0.3)	(±0.5)	(±0.5)	(±0.7)	(±0.5)	(±0.8)	(±0.8)	(±1.0)
h2	120.5	—	120.4	160.6	160.6	200.7	240.9	320.9	360.9	401.1
	(±0.3)	—	(±0.2)	(±0.4)	(±0.4)	(±0.3)	(±0.6)	(±0.7)	(±0.8)	(±0.6)
h3	120.5	120.5	—	160.5	140.5	200.6	220.6	320.8	340.8	380.8
	(±0.4)	(±0.5)	—	(±0.6)	(±0.4)	(±0.5)	(±0.4)	(±0.5)	(±0.4)	(±0.8)

Processing Network Latency (ms) in a scenario with no traffic expressed as RTT AVG and RTT MDEV.

Future work

- ✓ The Cost Map rankings are based on relatively static AS-Path distance and do not consider more dynamic information such as actual bandwidth, latency, packet loss rate, etc.
  - Thus, dynamic updates of cost maps based on public Internet quality measurements (e.g., SIMET, RIPE TTM) are in our research agenda.

- ✓ Implement the remaining ALTO services (e.g., full Map-Service, EPS and ECS).

- ✓ Prototype performance optimization (e.g., Neo4j query tuning techniques, Linux fs configuration, security).

Conclusions

- ✓ This is the first work that explores the use of Inter-domain routing data publicly available at IXPs to create abstract topology and cost maps following the recently standardized IETF ALTO protocol.
- ✓ Our proposed architecture (AaaS) encompasses the whole process of ALTO service delivery.
- ✓ Our proof-of-concept implementation is based on the popular Neo4j graph database and the OpenDaylight controller and validated the potential of applications to leverage the network awareness provided by ALTO servers.
- ✓ At the same time, ISPs and ALTO service providers (in our case IX.br operators) can benefit from increased and localized IXP traffic exchange.