École Doctorale Mathématiques et Informatique – Spécialité Informatique

HIGH-QOE PRIVACY-PRESERVING VIDEO STREAMING

Diffusion Vidéo avec une Meilleure Qualité d'Expérience et Respectant la Vie Privée

Simon **Da Silva**

université de BORDEAUX LaBRI

2020 - 10 - 07

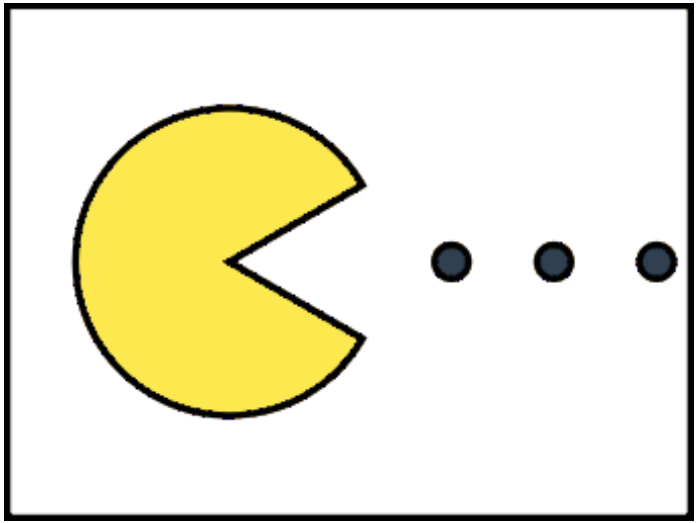

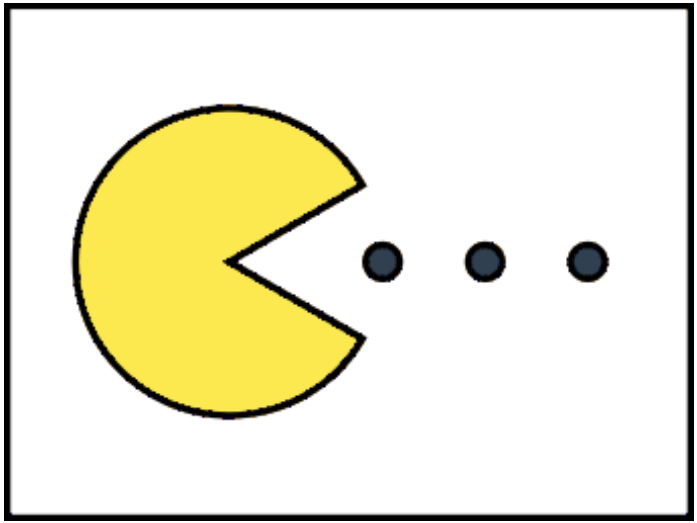

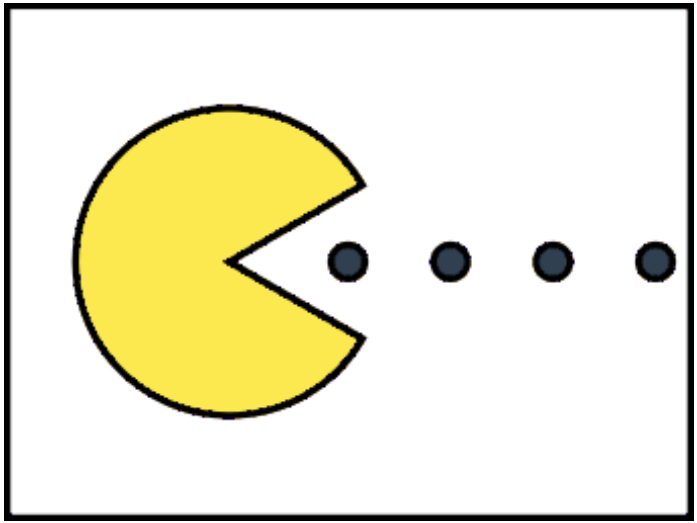

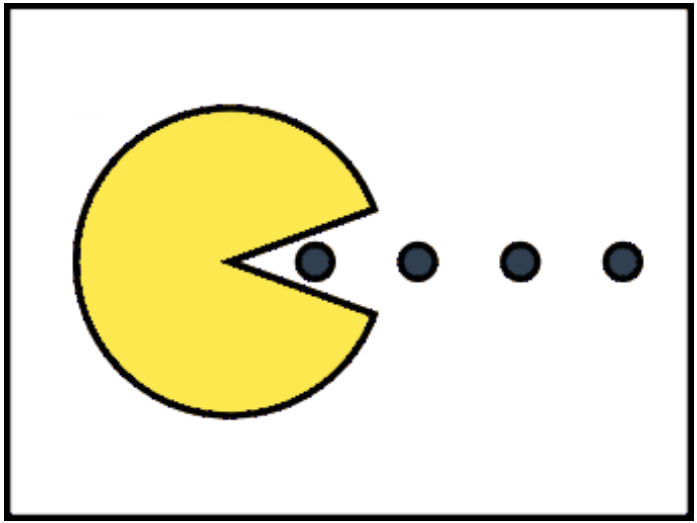

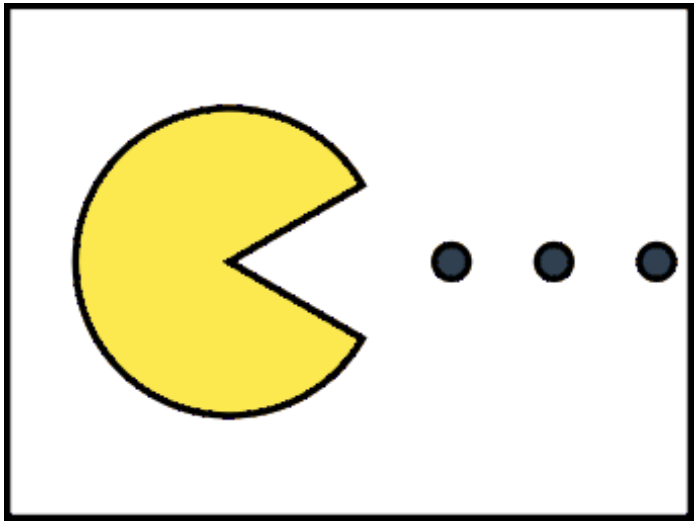

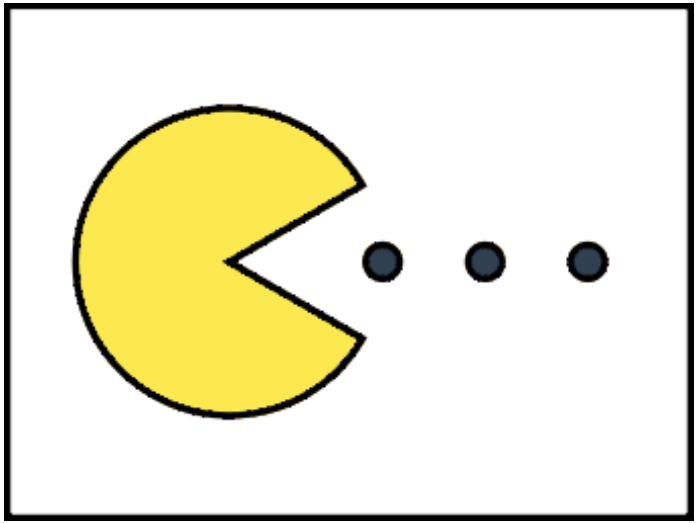

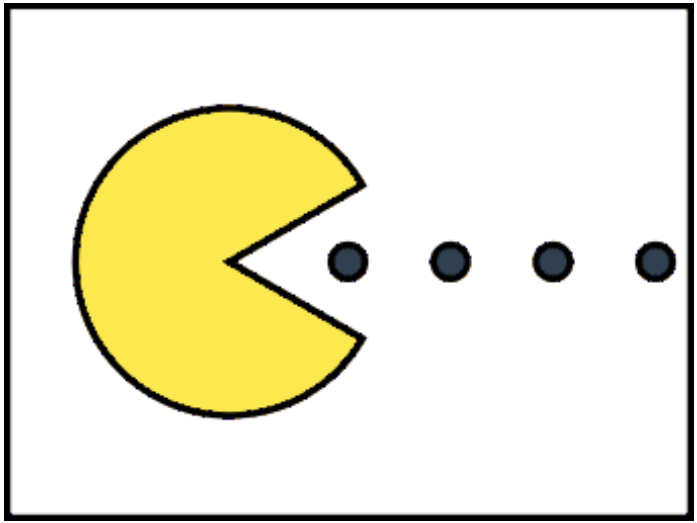

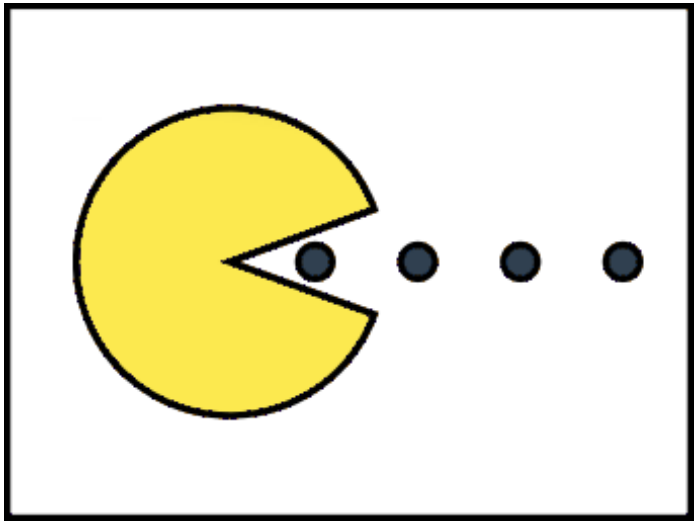

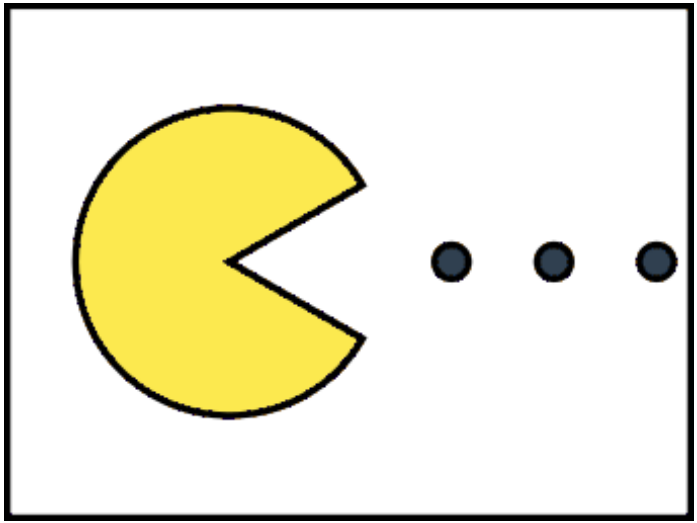

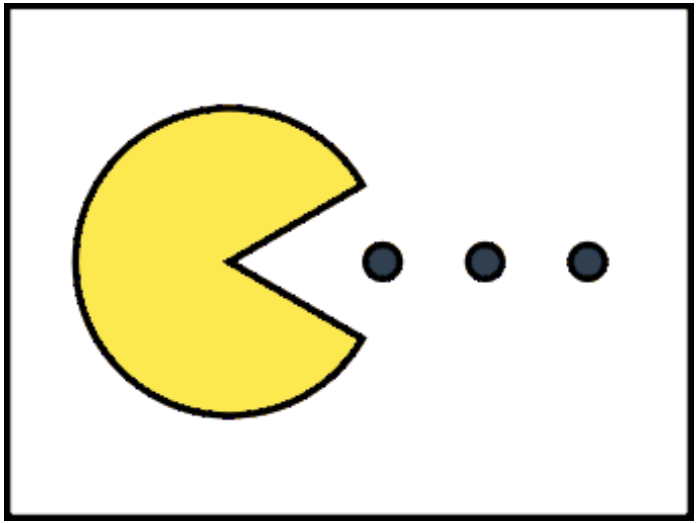

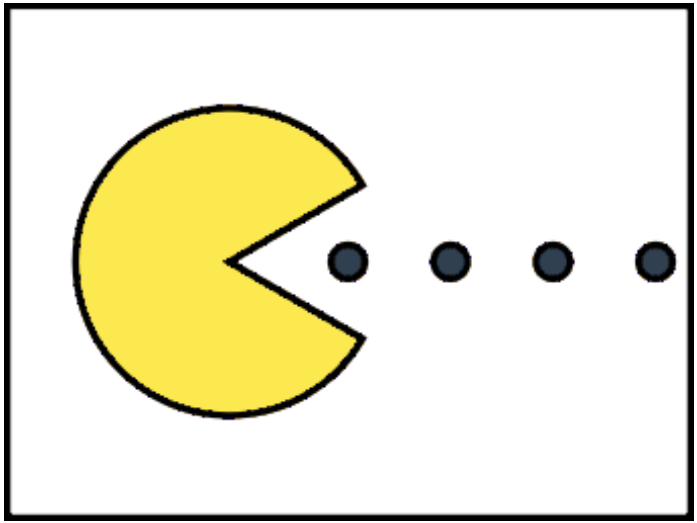

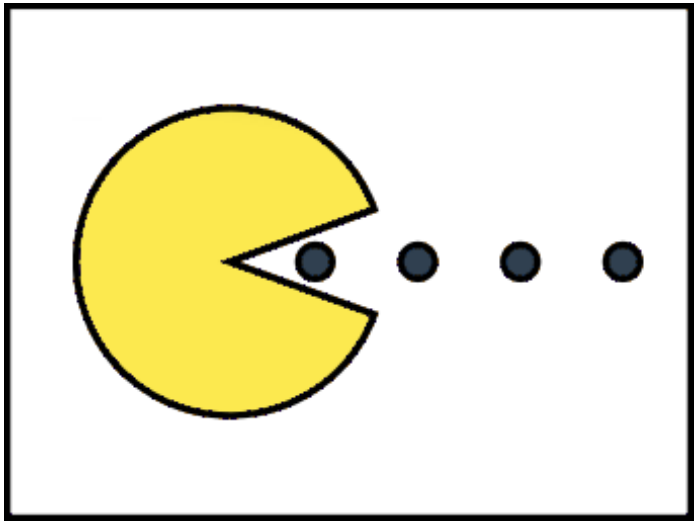

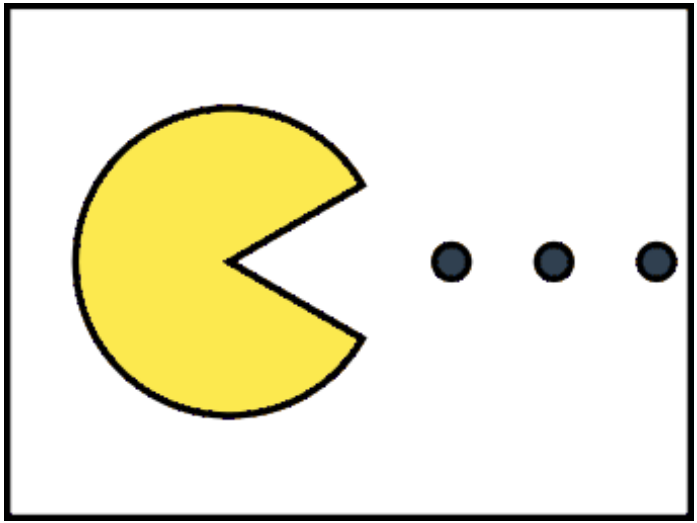

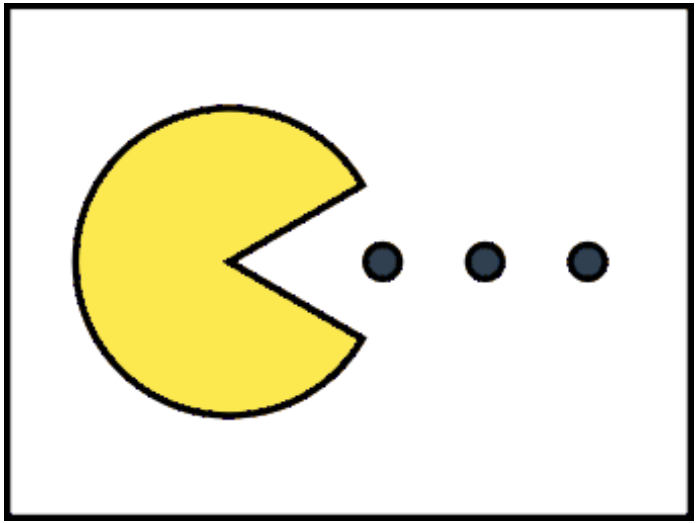

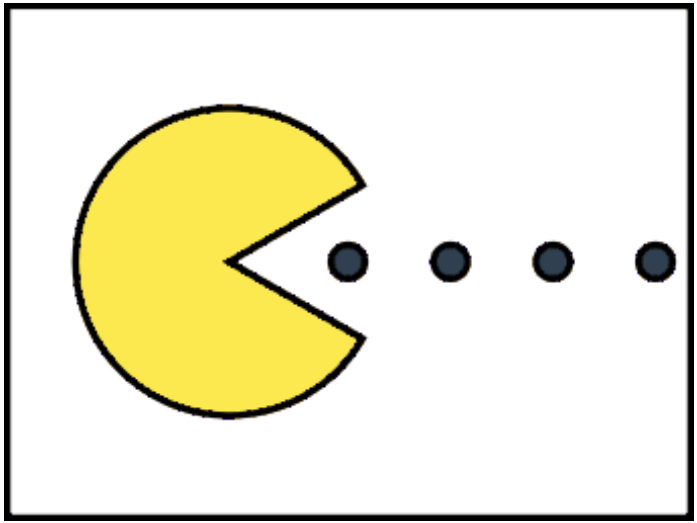

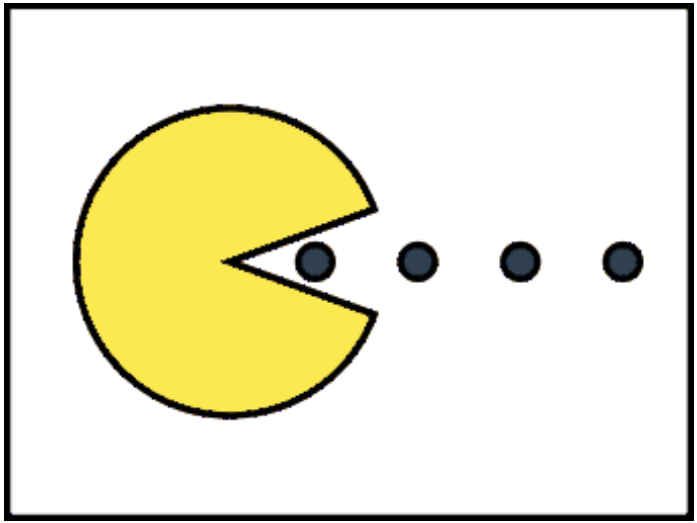

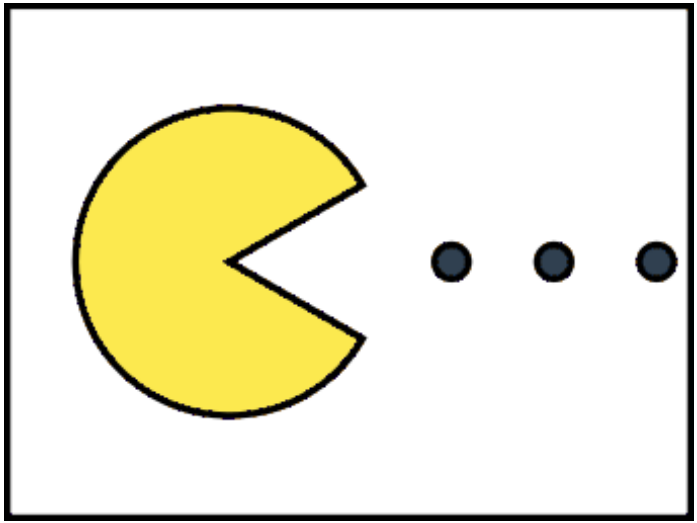

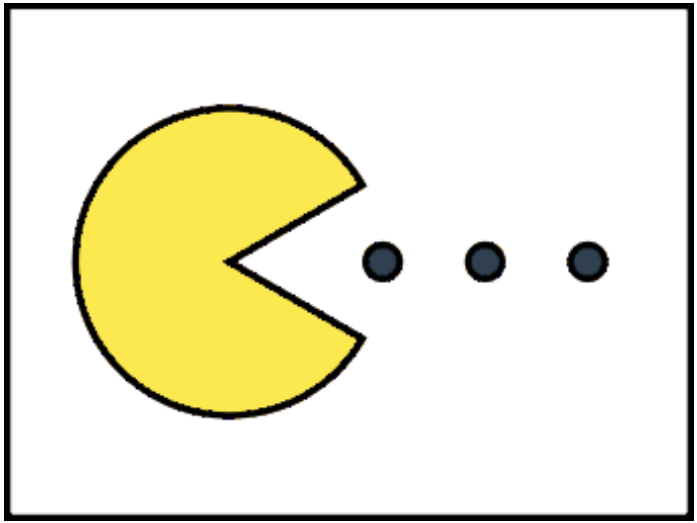

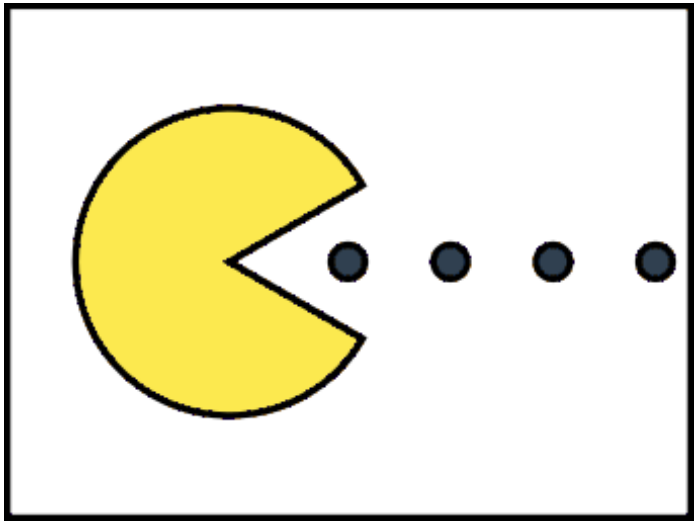

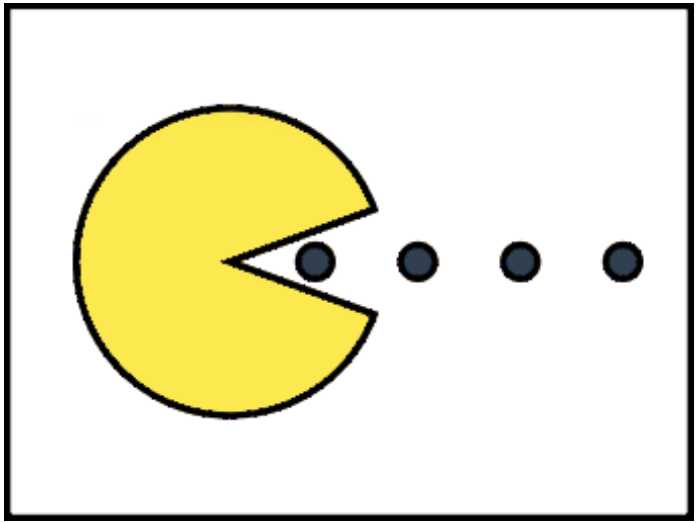

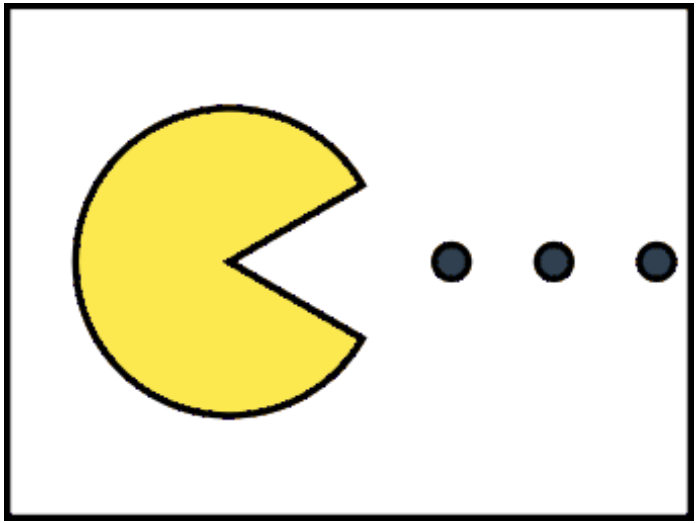

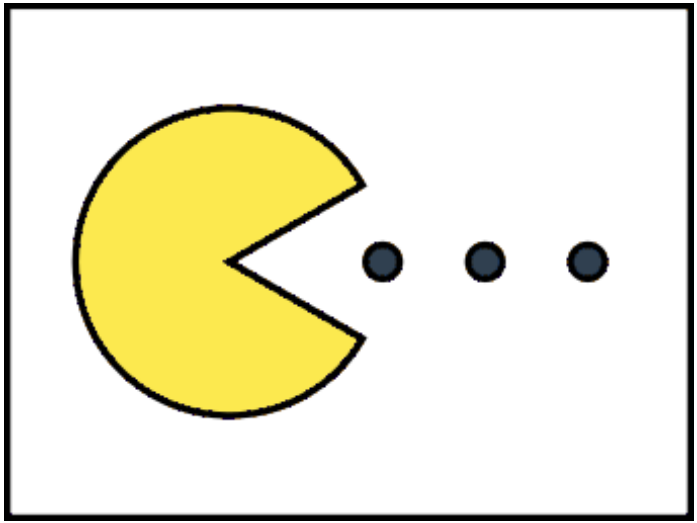

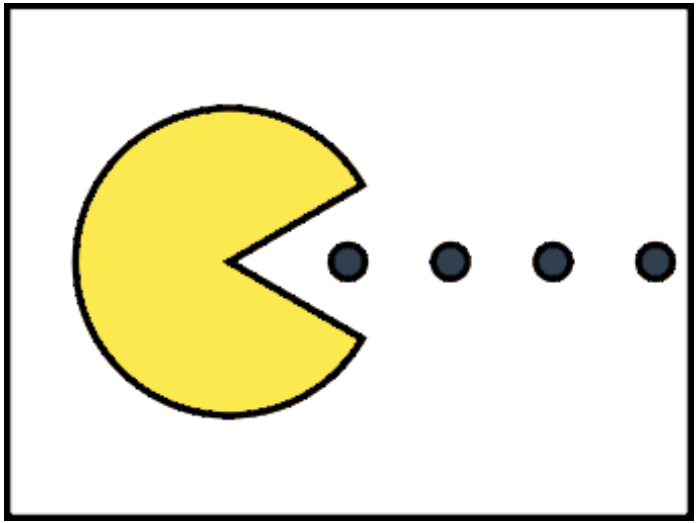

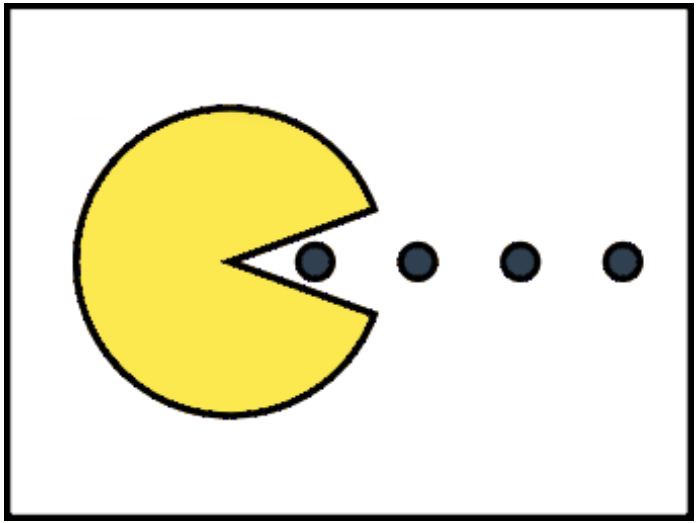

Video encoding

Displayed content:

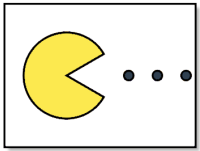

Frame #1

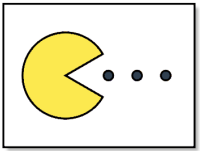

Frame #2

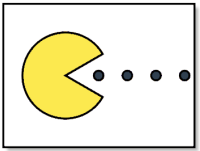

Frame #3

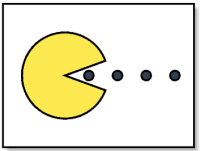

Frame #4

Displayed content:

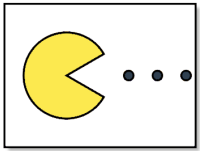
Frame #1

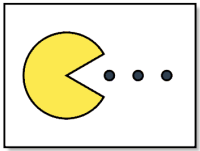
Frame #2

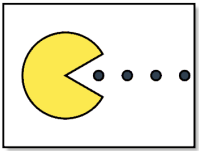
Frame #3

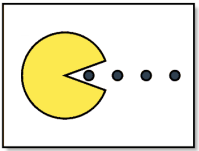
Frame #4

Encoded content:

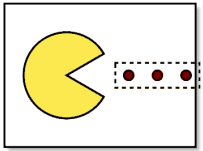
I-frame

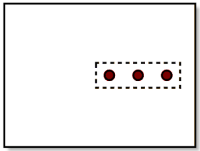
P-frame

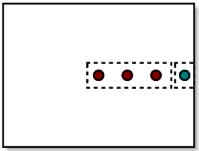
B-frame

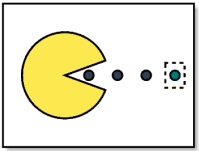
I-frame

Video encoding

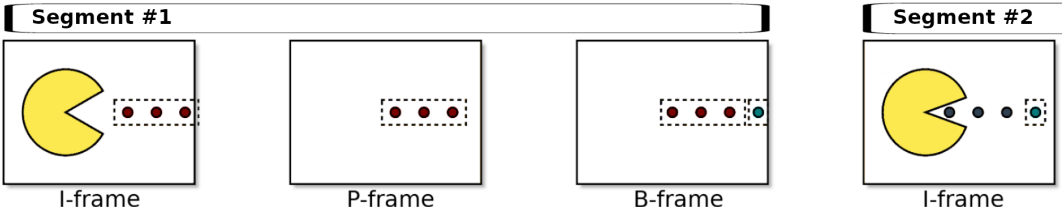

Video content consumption evolves...

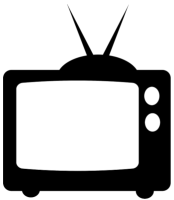

Video content consumption evolves...

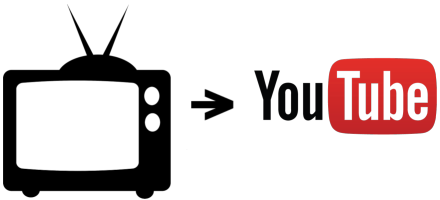

Video content consumption evolves...

Video content consumption evolves...

Video content consumption evolves...

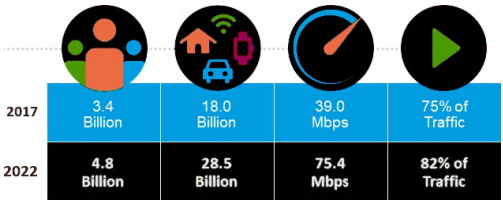

Cisco VNI 2017-2022

Video content consumption evolves...

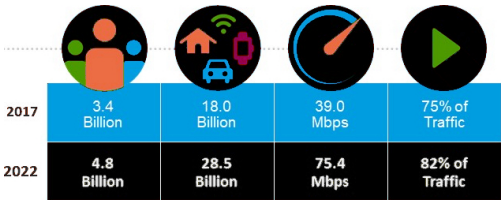

- High QoE

Video content consumption evolves...

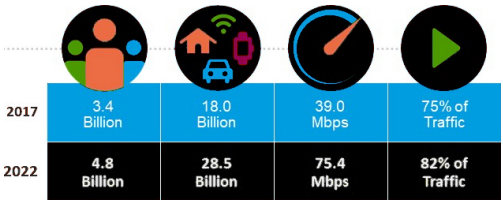

- High QoE
- Reliability

Video content consumption evolves...

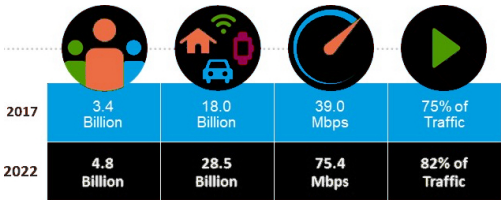

- High QoE
- Reliability
- Low cost

Video content consumption evolves...

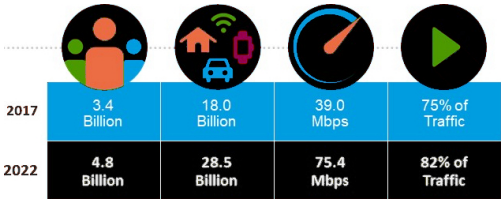

- High QoE
- Reliability
- Low **scalability** cost

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

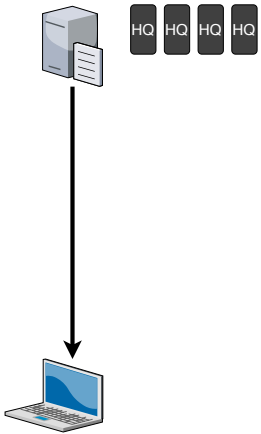

université de BORDEAUX

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

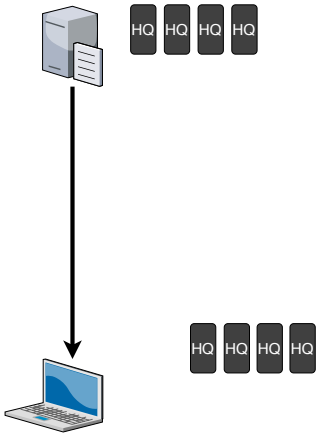

université de BORDEAUX

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

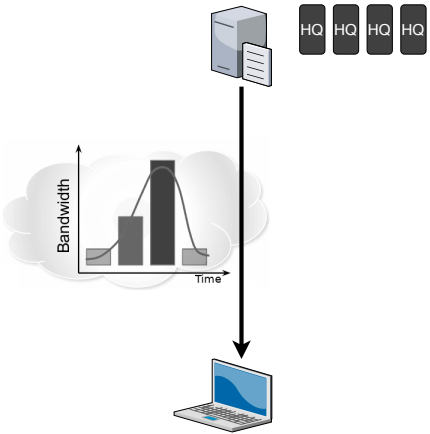

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

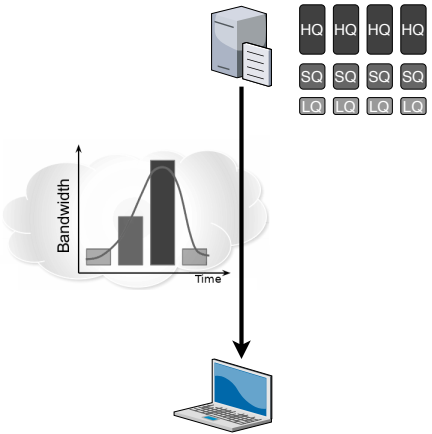

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

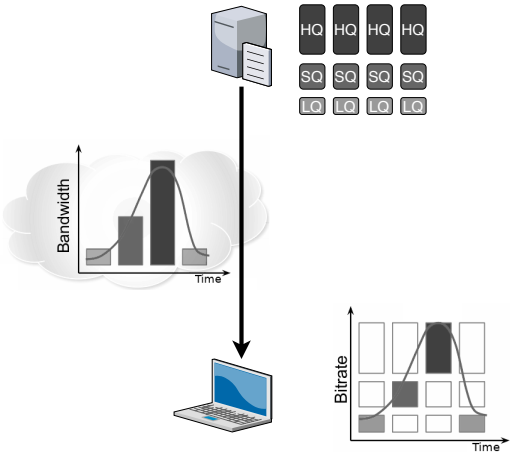

université de BORDEAUX

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

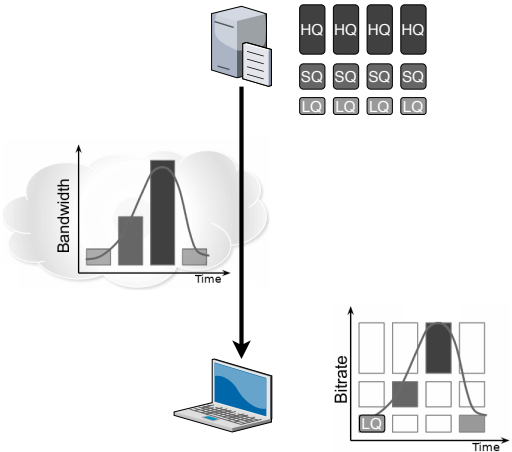

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

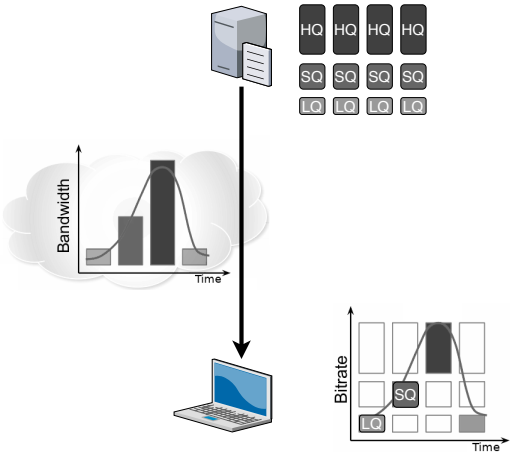

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

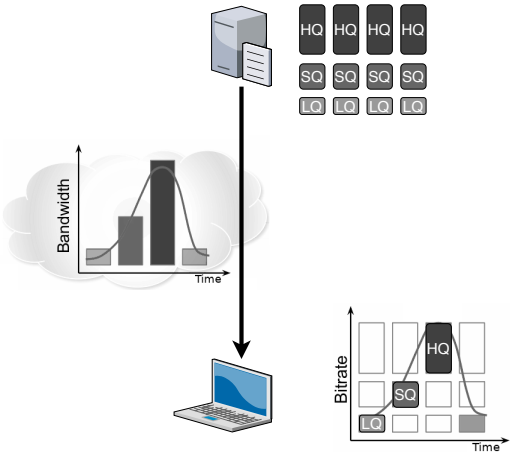

High QoE: HTTP Adaptive Streaming - MPEG-DASH

- High QoE
- Reliability
- Low scalability cost

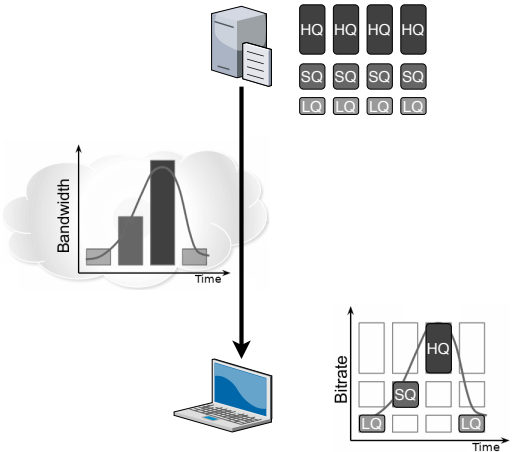

Reliability: Content Delivery Networks (CDN)

- High QoE
- Reliability
- Low scalability cost

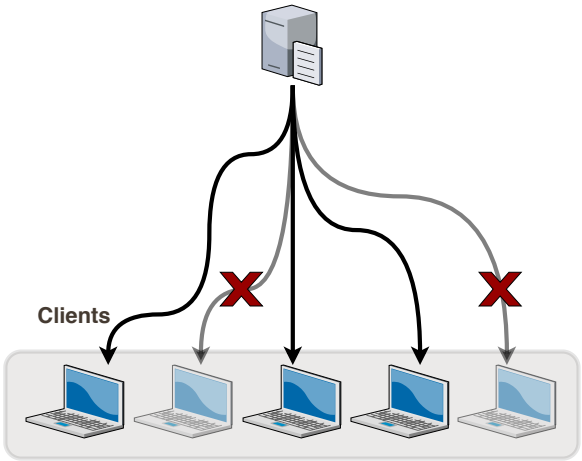

université de BORDEAUX

Reliability: Content Delivery Networks (CDN)

- High QoE
- Reliability
- Low scalability cost

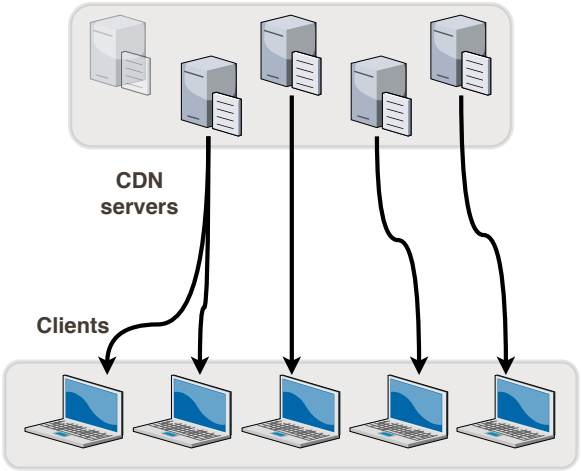

Reliability: Content Delivery Networks (CDN)

- High QoE
- Reliability
- Low scalability cost

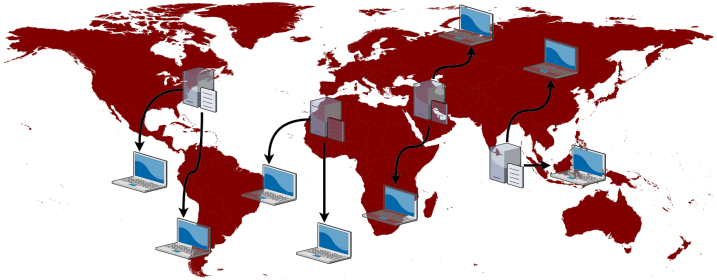

université de BORDEAUX

Low scalability cost?

- High QoE
- Reliability
- Low scalability cost

Low scalability cost?

Live content caching requires fast computation and decision

- High QoE
- Reliability
- Low scalability cost

Low scalability cost?

- High QoE
- Reliability
- Low scalability cost

Live content caching requires fast computation and decision

Content replication algorithms are NP-complete [KRR02]

Low scalability cost?

- High QoE
- Reliability
- Low scalability cost

Live content caching requires fast computation and decision

Content replication algorithms are NP-complete [KRR02]

CDN operators statically provision servers in advance [Pas12]

Low scalability cost?

- High QoE
- Reliability
- Low scalability cost

Live content caching requires fast computation and decision

Content replication algorithms are NP-complete [KRR02]

CDN operators statically provision servers in advance [Pas12]

CDN users are bound to one server [AGH+12]

université de BORDEAUX

Plan

université de BORDEAUX

MUSLIN: Multi-Source Live Streaming

- High QoE
- Reliability
- Low scalability cost

MUSLIN: Multi-Source Live Streaming

- High QoE
- Reliability
- Low scalability cost

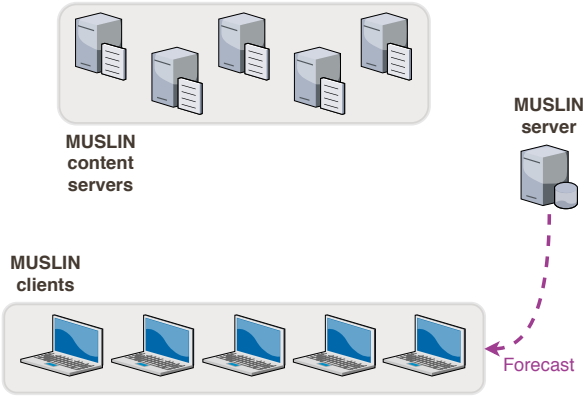

MUSLIN: Multi-Source Live Streaming

- High QoE
- Reliability
- Low scalability cost

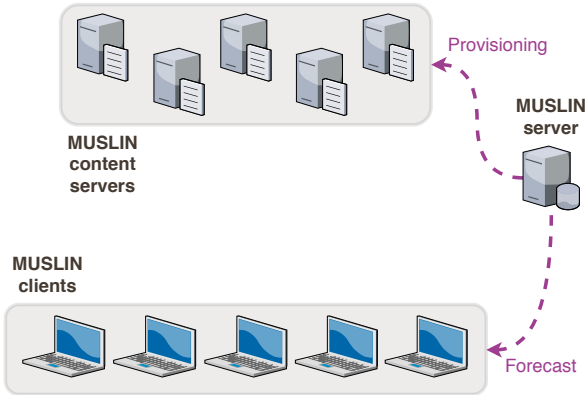

MUSLIN: Multi-Source Live Streaming

- High QoE
- Reliability
- Low scalability cost

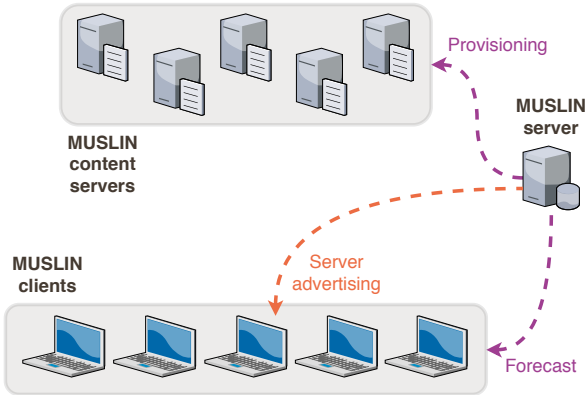

MUSLIN: Multi-Source Live Streaming

- High QoE
- Reliability
- Low scalability cost

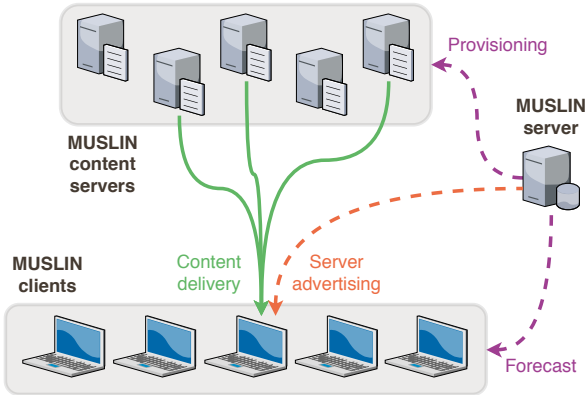

MUSLIN: Multi-Source Live Streaming

- High QoE
- Reliability
- Low scalability cost

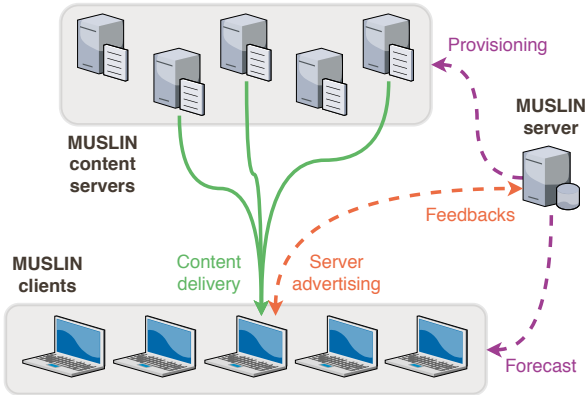

MS-Stream: Multi-Source Streaming over HTTP

- High QoE
- Reliability
- Low scalability cost

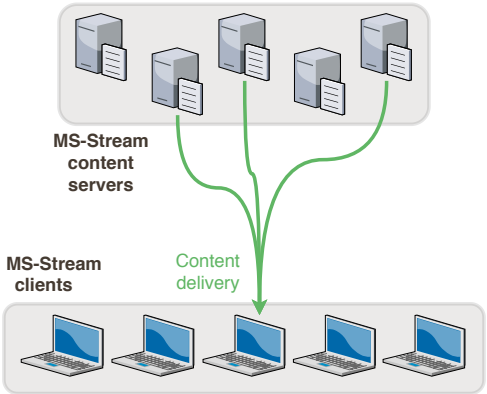

MS-Stream: Multi-Source Streaming over HTTP

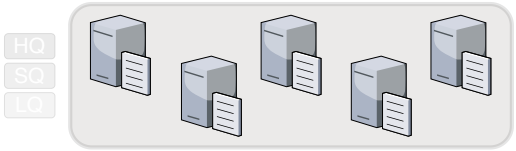

MS-Stream: Multi-Source Streaming over HTTP

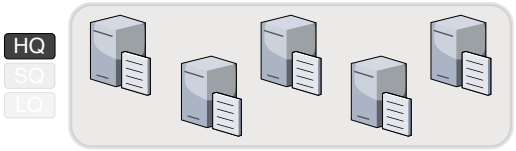

MS-Stream: Multi-Source Streaming over HTTP

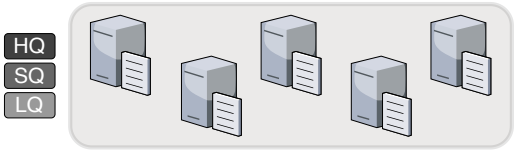

MS-Stream: Multi-Source Streaming over HTTP

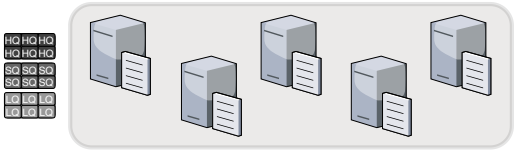

MS-Stream content servers

MS-Stream clients

MS-Stream: Multi-Source Streaming over HTTP

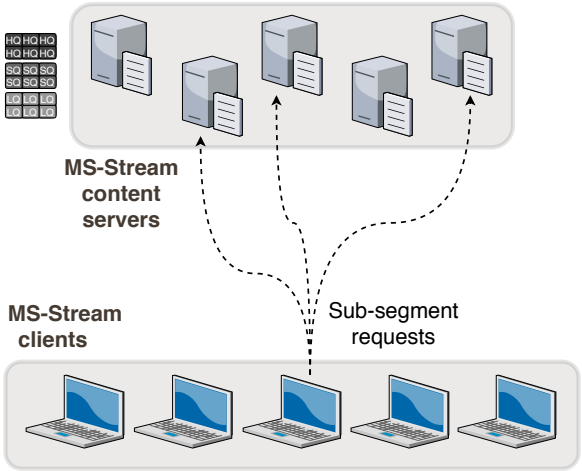

MS-Stream: Multi-Source Streaming over HTTP

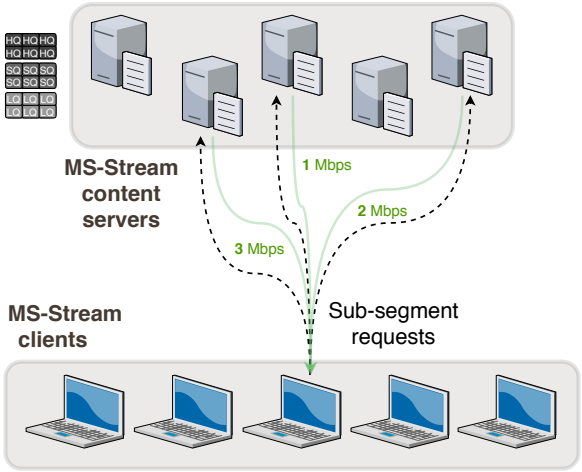

MS-Stream: Multi-Source Streaming over HTTP

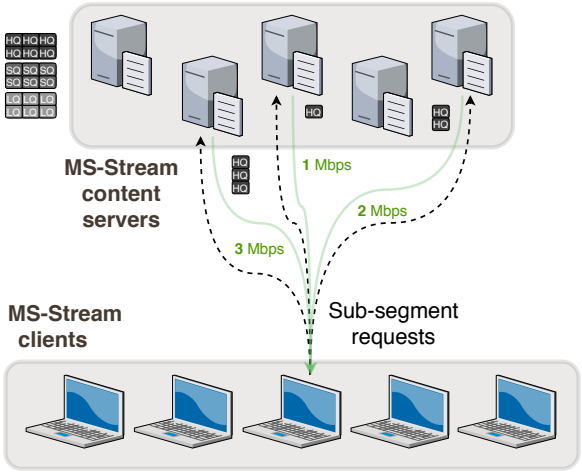

MS-Stream: Multi-Source Streaming over HTTP

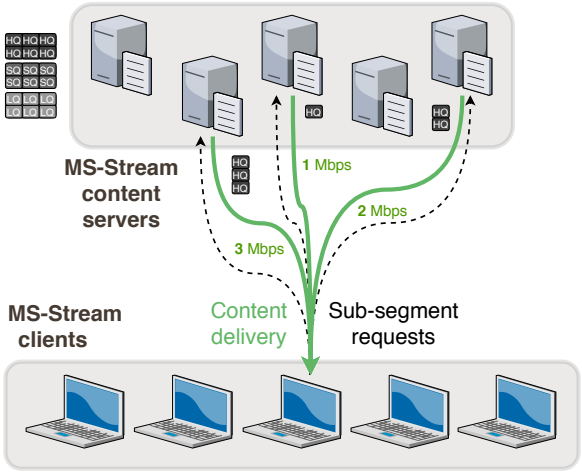

MS-Stream: Multi-Source Streaming over HTTP

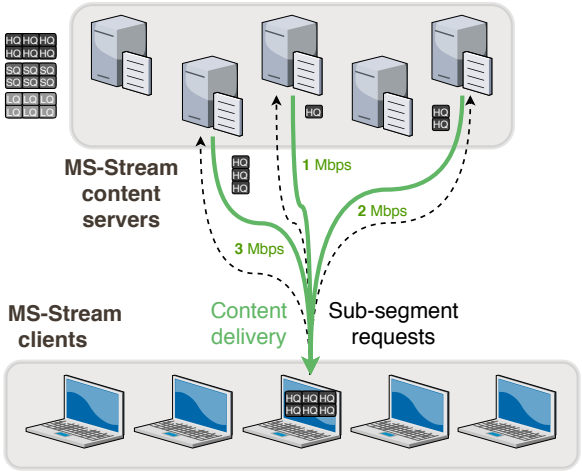

MS-Stream: Multi-Source Streaming over HTTP

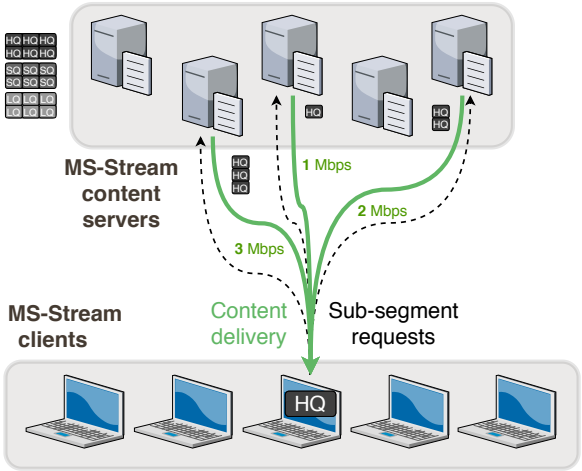

MS-Stream: Multi-Source Streaming over HTTP

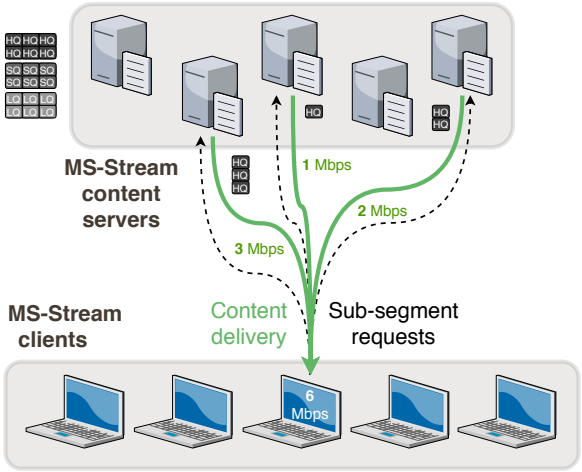

MUSLIN: Multi-Source Live Streaming

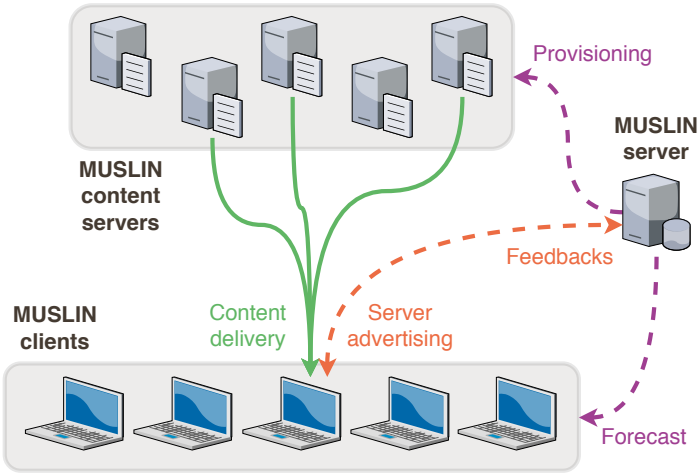

Server provisioning

Server provisioning

1. Audience forecast
   - Current audience
   - Past trend

Server provisioning

1. Audience forecast
   - Current audience
   - Past trend
2. Throughput estimation
   - Target quality
   - Network bandwidth overhead
   - Average video bitrate
   - Failure rate

Server provisioning

1. Audience forecast
   - Current audience
   - Past trend
2. Throughput estimation
   - Target quality
   - Network bandwidth overhead
   - Average video bitrate
   - Failure rate
3. Provisioning decision (Server Ranking Score RS_s)
   - Clients location
   - Server failure rate
   - Observed bandwidth

Server provisioning

1. Audience forecast

$$\widehat{v_{t+T}} = v_t + \Delta v$$

Server provisioning

1. Audience forecast

$$\widehat{v_{t+T}} = v_t + \Delta v$$

2. Throughput estimation

$$D = C * \widehat{v_{t+T}} * (Q + O)$$

$$\left(C = \frac{Q}{B} * (1 + FR) \right)$$

Q: target quality - O: network bandwidth overhead - B: average video bitrate

Server provisioning

1. Audience forecast

$$\widehat{v_{t+T}} = v_t + \Delta v$$

2. Throughput estimation

$$D = C * \widehat{v_{t+T}} * (Q + O)$$

$$\left(C = \frac{Q}{B} * (1 + FR) \right)$$

Q: target quality - O: network bandwidth overhead - B: average video bitrate

3. Provisioning decision

Server Ranking Score RS_s for each server s:

$$RS_s = (N_s * (1 - FR_s) * OBW_s)^{\frac{1}{3}}$$

Server advertising

Server advertising

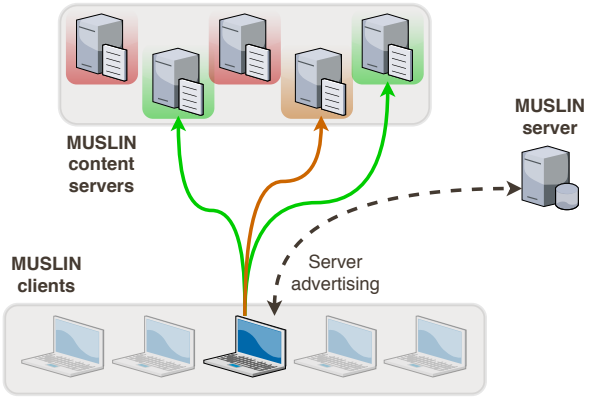

Server advertising

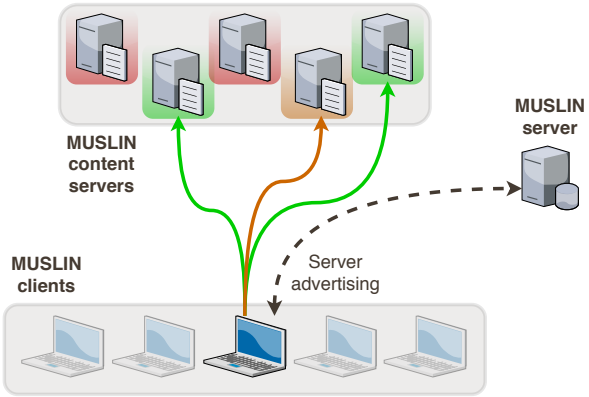

MUSLIN Ranking Score RS_{sc}:

- Server - client distance
- Server failure rate
- Observed bandwidth

Implementation and scalability

Implementation and scalability

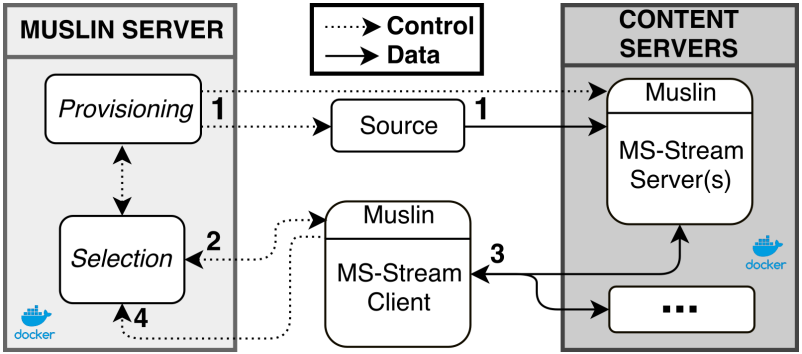

1 Content replication - **2** Server advertising - **3** Content delivery - **4** Clients feedbacks

Implementation and scalability

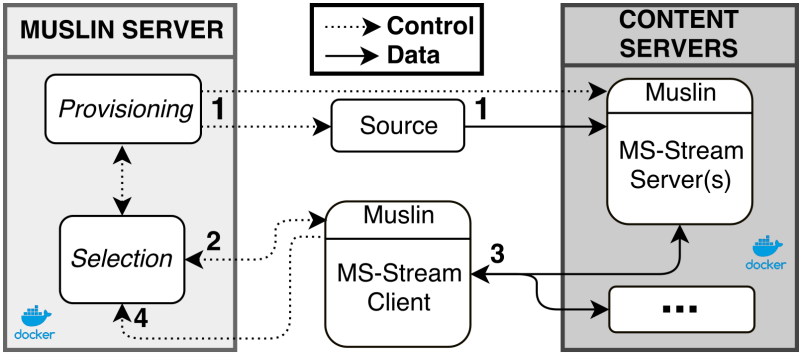

1 Content replication - **2** Server advertising - **3** Content delivery - **4** Clients feedbacks

Feedback request probability: $\Pr = \min(1, N/v_t)$

Evaluation setup

Evaluation setup

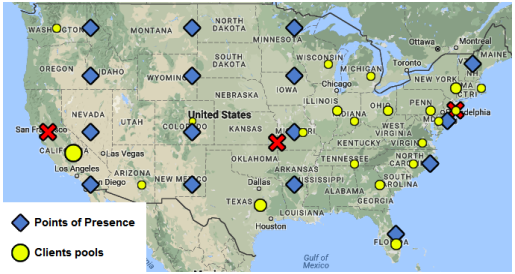

- 16 servers (Points of Presence)
- 21 client pools locations

Evaluation setup

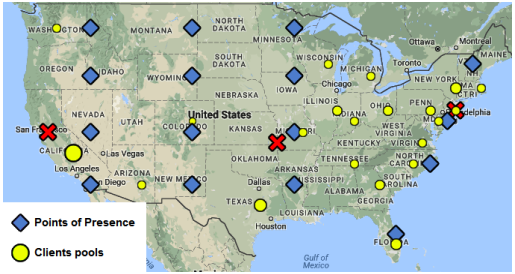

- 16 servers (Points of Presence)
- 21 client pools locations

Table: Evaluated policies

Provisioning and Forecast	Selection	Delivery
MUSLIN	MUSLIN	MS-Stream

Evaluation setup

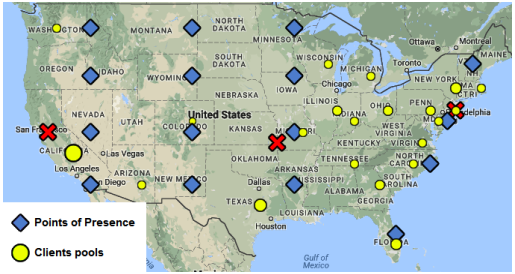

- 16 servers (Points of Presence)
- 21 client pools locations

Table: Evaluated policies

Provisioning and Forecast	Selection	Delivery
MUSLIN	MUSLIN	MS-Stream
Geographical oracle		
Geographical oracle		
Geographical oracle		

Evaluation setup

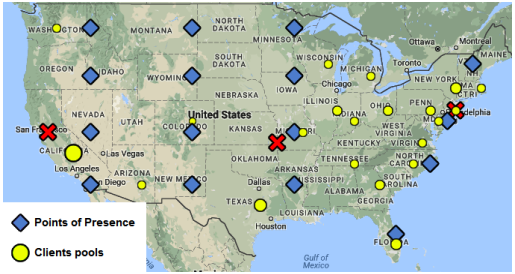

- 16 servers (Points of Presence)
- 21 client pools locations

Table: Evaluated policies

Provisioning and Forecast	Selection	Delivery
MUSLIN	MUSLIN	MS-Stream
Geographical oracle	CDN	
Geographical oracle	Random	
Geographical oracle	Round Robin	

Evaluation setup

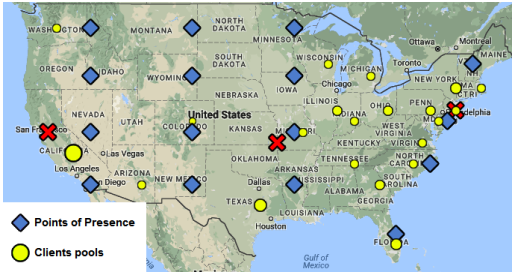

- 16 servers (Points of Presence)
- 21 client pools locations

Table: Evaluated policies

Provisioning and Forecast	Selection	Delivery
MUSLIN	MUSLIN	MS-Stream
Geographical oracle	CDN	MS-Stream
Geographical oracle	Random	MS-Stream
Geographical oracle	Round Robin	MS-Stream

0.3 - 6.4 Mbps video

Evaluation setup

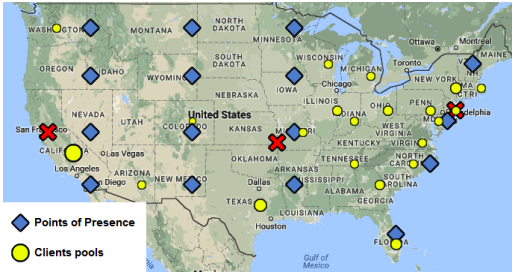

- 16 servers (Points of Presence)
- 21 client pools locations

Table: Evaluated policies

Provisioning and Forecast	Selection	Delivery
MUSLIN	MUSLIN	MS-Stream
Geographical oracle	CDN	MS-Stream
Geographical oracle	Random	MS-Stream
Geographical oracle	Round Robin	MS-Stream

0.3 - 6.4 Mbps video
Actual live audience trace

Evaluation results

Compared to a best-case CDN setup (geographical oracle):

Evaluation results

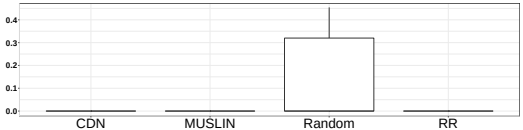

Figure: Rebufferings (per minute)

Compared to a best-case CDN setup (geographical oracle):

- 0 rebufferings

Evaluation results

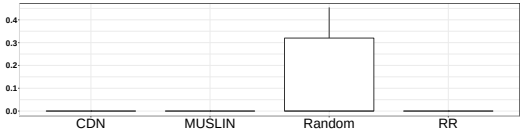

Figure: Rebufferings (per minute)

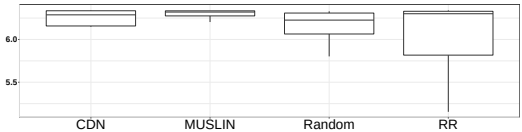

Figure: Displayed bitrate (Mbps)

Compared to a best-case CDN setup (geographical oracle):

- 0 rebufferings
- + 1.6% displayed bitrate

Evaluation results

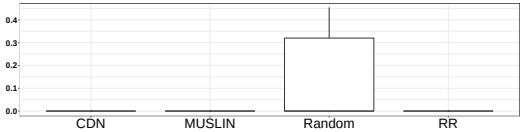

Figure: Rebufferings (per minute)

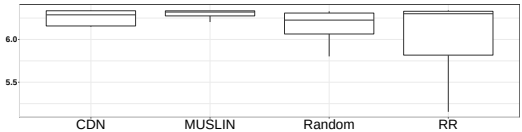

Figure: Displayed bitrate (Mbps)

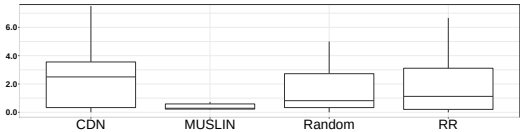

Figure: Quality changes (per minute)

Compared to a best-case CDN setup (geographical oracle):

- 0 rebufferings
- + 1.6% displayed bitrate
- - 625% quality changes

Evaluation results

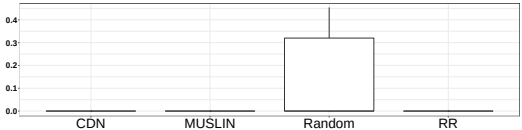

Figure: Rebufferings (per minute)

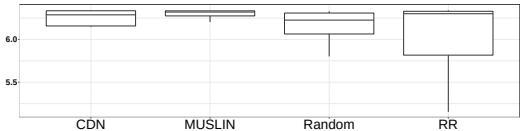

Figure: Displayed bitrate (Mbps)

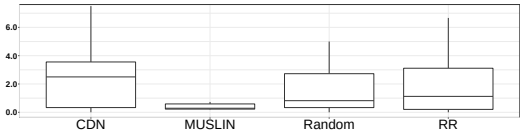

Figure: Quality changes (per minute)

Compared to a best-case CDN setup (geographical oracle):

- 0 rebufferings
- + 1.6% displayed bitrate
- - 625% quality changes
- - 18% provisionned server time

Evaluation results

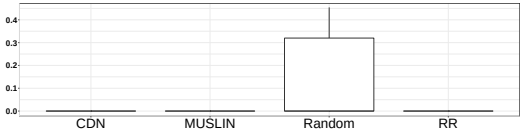

Figure: Rebufferings (per minute)

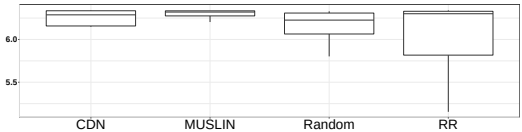

Figure: Displayed bitrate (Mbps)

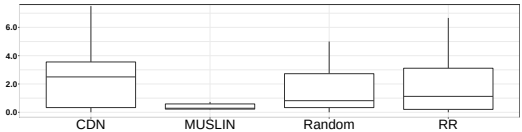

Figure: Quality changes (per minute)

Compared to a best-case CDN setup (geographical oracle):

- 0 rebufferings
- + 1.6% displayed bitrate
- - 625% quality changes
- - 18% provisionned server time

QoE fairness (F index):

T. Hoßfeld et al. Definition of QoE Fairness in Shared Systems. IEEE Communications Letters (2017)

Evaluation results

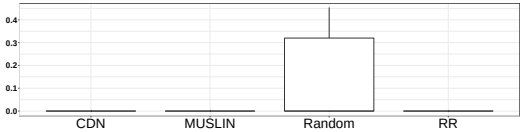

Figure: Rebufferings (per minute)

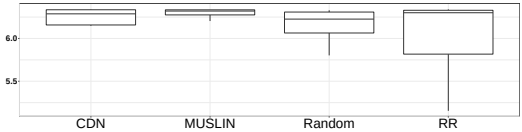

Figure: Displayed bitrate (Mbps)

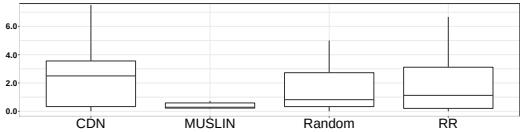

Figure: Quality changes (per minute)

Compared to a best-case CDN setup (geographical oracle):

- 0 rebufferings
- + 1.6% displayed bitrate
- - 625% quality changes
- - 18% provisionned server time

QoE fairness (F index):

- + 23.6% rebuffering fairness
- + 19.6% bitrate fairness
- + 52% quality changes fairness

T. Hoßfeld et al. Definition of QoE Fairness in Shared Systems. IEEE Communications Letters (2017)

université de BORDEAUX

MUSLIN - Conclusion

MUSLIN - Conclusion

- High QoE
- Reliability

Multiple-Source Streaming

MUSLIN - Conclusion

- High QoE
- Reliability

RS_{sc}-based server advertising

Multiple-Source Streaming

MUSLIN - Conclusion

- High QoE
- Reliability

RS_{sc}-based server advertising | Multiple-Source Streaming | Real-time clients feedbacks

MUSLIN - Conclusion

- High QoE
- Reliability

RS_{sc}-based server advertising | Multiple-Source Streaming | Real-time clients feedbacks

- Scalability

Dynamic server provisioning and content replication

MUSLIN - Conclusion

- High QoE
- Reliability

RS_{sc}-based server advertising | Multiple-Source Streaming | Real-time clients feedbacks

- Scalability

Dynamic server provisioning and content replication

Fairness: All users get a similar QoE

MUSLIN - Conclusion

- High QoE
- Reliability

RS_{sc}-based server advertising | Multiple-Source Streaming | Real-time clients feedbacks

- Scalability

Dynamic server provisioning and content replication

Fairness: All users get a similar QoE

S. Da Silva, J. Bruneau-Queyreix, M. Lacaud, D. Négru, L. Réveillère.

MUSLIN: A QoE-Aware CDN Resources Provisioning and Advertising System for Cost-Efficient Multi-Source Live Streaming.
International Journal of Network Management (IJNM '19).
MUSLIN: Achieving High, Fairly Shared QoE Through Multi-Source Live Streaming. *Packet Video Workshop (PV '18).*
MUSLIN demo: High QoE Fair Multi-Source Live Streaming. *ACM Multimedia Systems Conference (MMSys '18).*
DASH-IF Excellence in DASH Award, 3rd place, *MMSys '18*

J. Bruneau-Queyreix, S. Da Silva, M. Lacaud, D. Négru.

Vers une meilleure diffusion vidéo sur Internet. *Interstices, INRIA, 2018.*

université de BORDEAUX

Low scalability cost: MUSLIN

- High QoE
- Reliability
- Low scalability cost

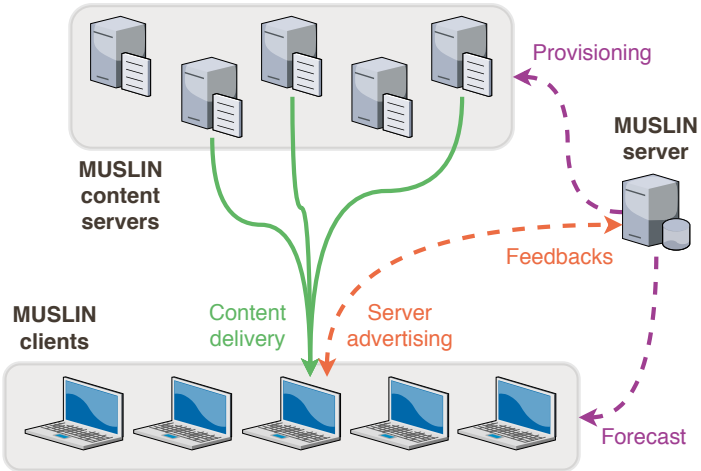

Low scalability cost: Edge-assisted CDN

- High QoE
- Reliability
- Low scalability cost

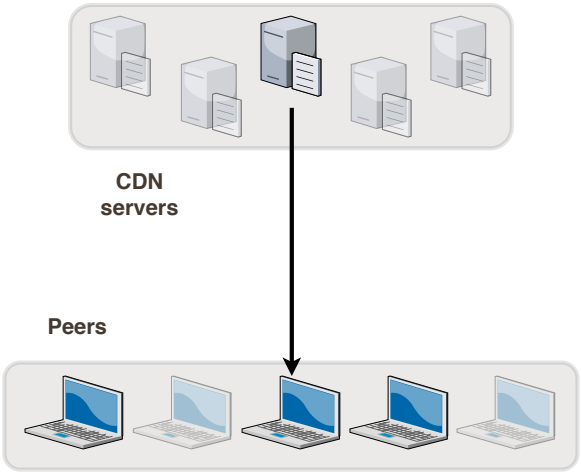

université de BORDEAUX

Low scalability cost: Edge-assisted CDN

- High QoE
- Reliability
- Low scalability cost

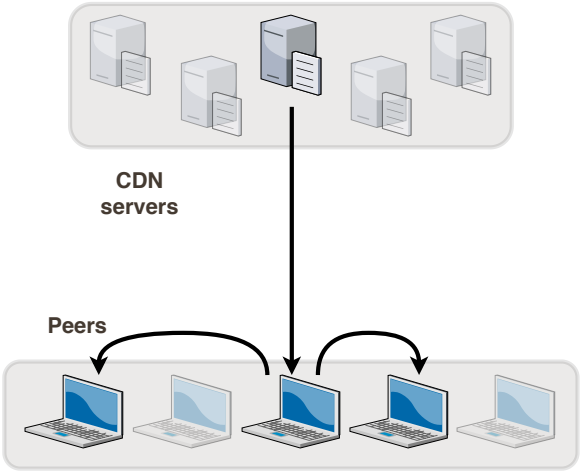

université de BORDEAUX

Multi-source adaptive streaming with MS-Stream

- High QoE
- Reliability
- Low scalability cost

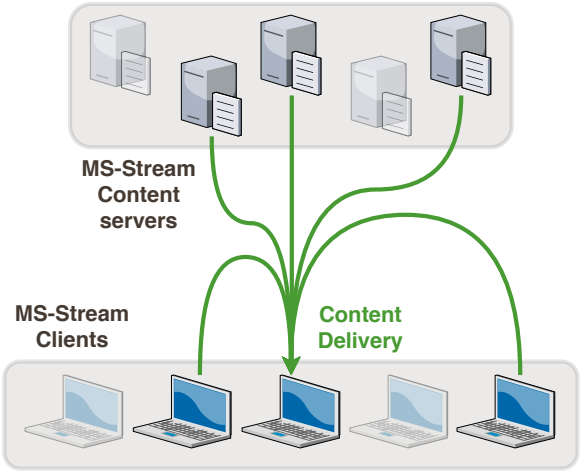

Multi-source adaptive streaming with MS-Stream

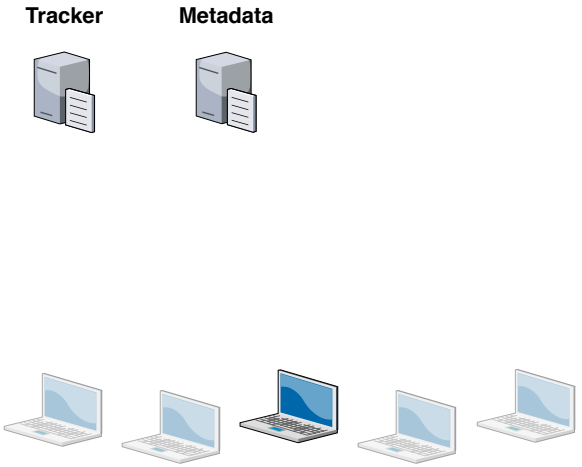

Multi-source adaptive streaming with MS-Stream

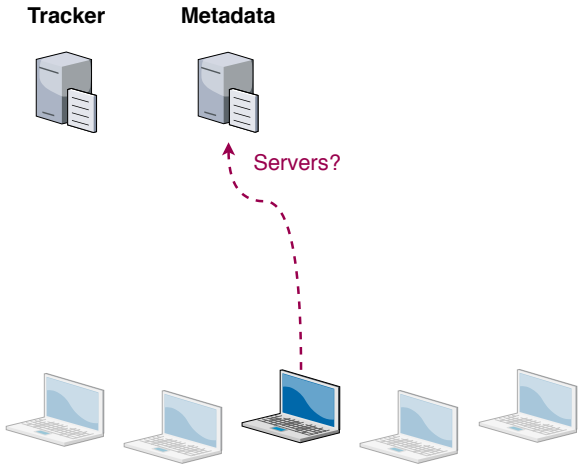

Multi-source adaptive streaming with MS-Stream

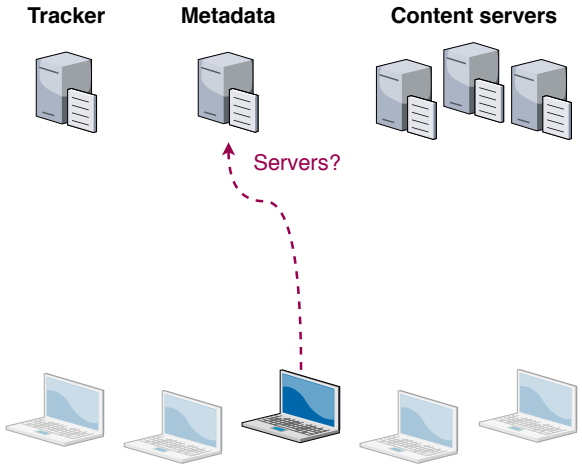

Multi-source adaptive streaming with MS-Stream

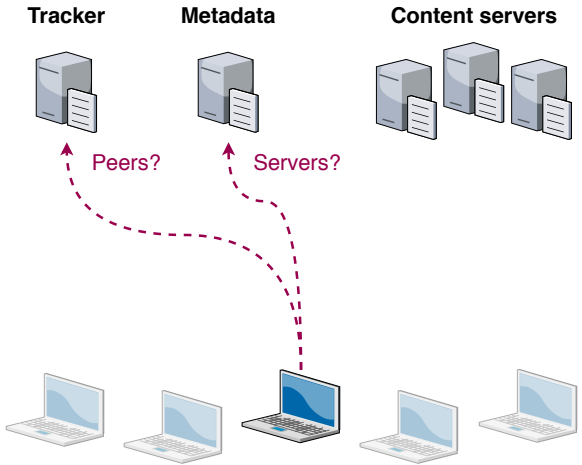

Multi-source adaptive streaming with MS-Stream

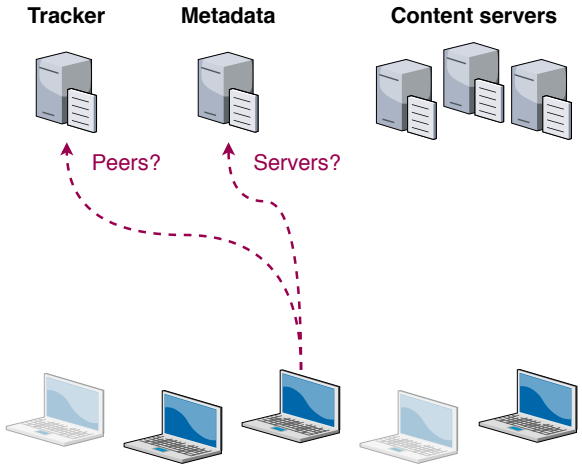

Multi-source adaptive streaming with MS-Stream

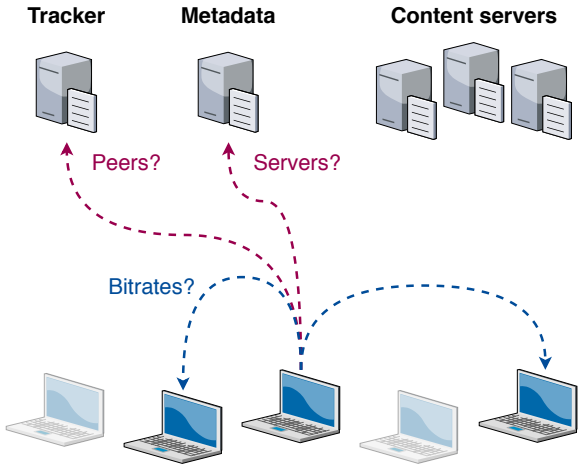

Multi-source adaptive streaming with MS-Stream

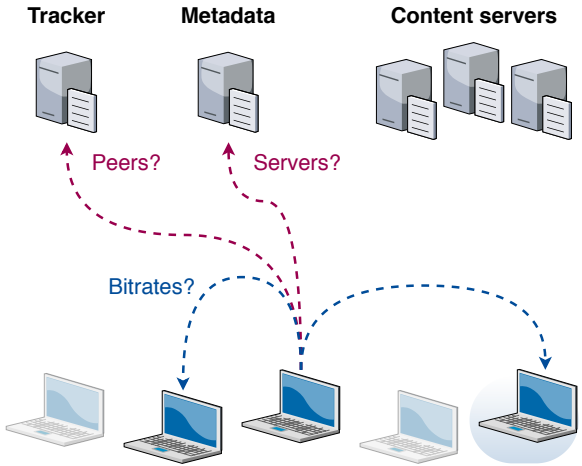

Multi-source adaptive streaming with MS-Stream

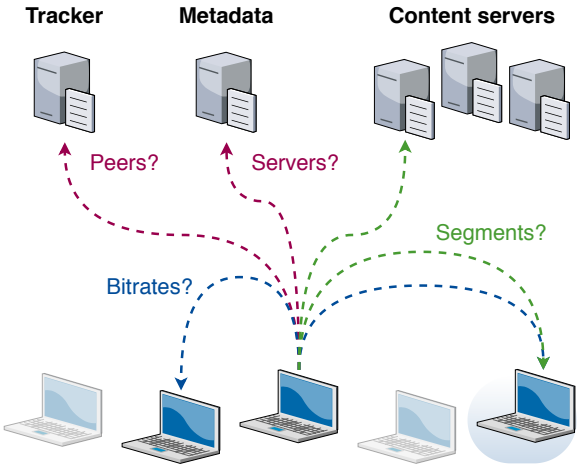

Multi-source adaptive streaming with MS-Stream

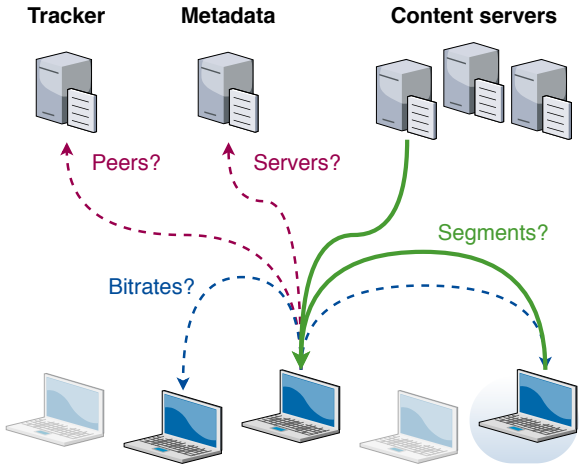

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Popcorn [GCM+16] and Tor [DMS04] require non-colluding servers

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Popcorn [GCM+16] and Tor [DMS04] require non-colluding servers
Costly cryptographic mechanisms

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Popcorn [GCM+16] and Tor [DMS04] require non-colluding servers
Costly cryptographic mechanisms => high latency (RTT $\geq$ 1s) + scalability issues

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Popcorn [GCM+16] and Tor [DMS04] require non-colluding servers
Costly cryptographic mechanisms => high latency (RTT $\geq$ 1s) + scalability issues

Differential privacy (P3LS [DBYEV19])

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Popcorn [GCM+16] and Tor [DMS04] require non-colluding servers
Costly cryptographic mechanisms => high latency (RTT $\geq$ 1s) + scalability issues

Differential privacy (P3LS [DBYEV19]) => bandwidth & storage overhead + privacy issues

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Popcorn [GCM+16] and Tor [DMS04] require non-colluding servers
Costly cryptographic mechanisms => high latency (RTT $\geq$ 1s) + scalability issues

Differential privacy (P3LS [DBYEV19]) => bandwidth & storage overhead + privacy issues

SGX proxies (X-Search [MBF+17a], Cyclosa [PGM+18])

Edge-assisted CDN and privacy

People are mostly reluctant to seed content because of **privacy issues**.

Popcorn [GCM+16] and Tor [DMS04] require non-colluding servers
Costly cryptographic mechanisms => high latency (RTT $\geq$ 1s) + scalability issues

Differential privacy (P3LS [DBYEV19]) => bandwidth & storage overhead + privacy issues

SGX proxies (X-Search [MBF+17a], Cyclosa [PGM+18]) => bandwidth & latency overhead ($\simeq$ 1s)

Privacy objective

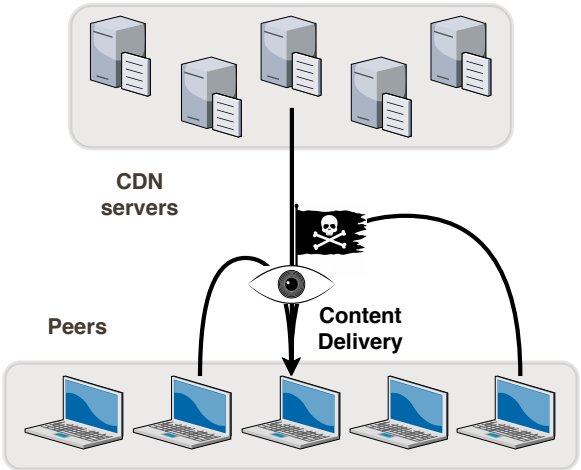

Adversary model

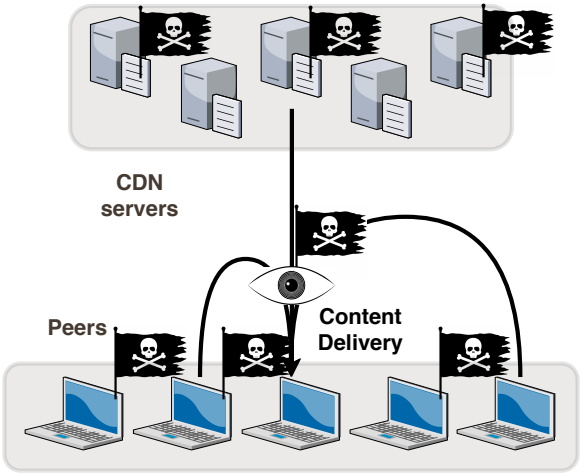

Plan

université de BORDEAUX

PRIVATUBE

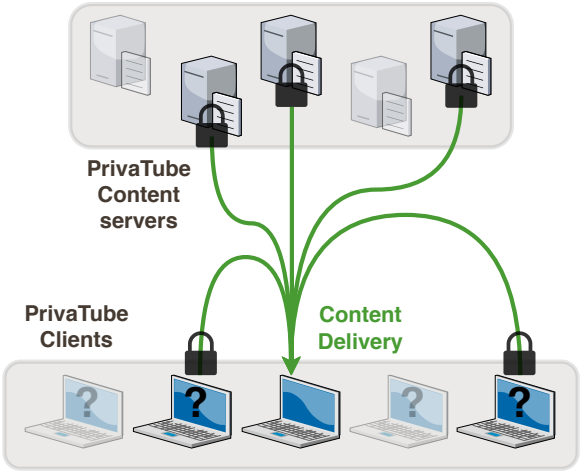
PrivaTube
Content
servers
PrivaTube
Clients
Content
Delivery

PrivaTube HTTP proxies

PrivaTube HTTP proxies

- Queries encryption
  e.g. Game of Thrones
  $\rightarrow$ *R2F8MgbYgGhyb2*

PrivaTube HTTP proxies

- Queries encryption
  e.g. Game of Thrones
  $\rightarrow$ *R2F8MgbYgGhyb2*
- Hostnames & IP addresses encryption
  e.g. Alice $\rightarrow$ *9jB4q*
  e.g. Server1 $\rightarrow$ *7Tl13nV*

PrivaTube HTTP proxies

- Queries encryption
  e.g. Game of Thrones
  $\rightarrow$ *R2F8MgbYgGhyb2*
- Hostnames & IP addresses encryption
  e.g. Alice $\rightarrow$ *9jB4q*
  e.g. Server1 $\rightarrow$ *7Tl13nV*
- Video stream encryption

SGX enclaves

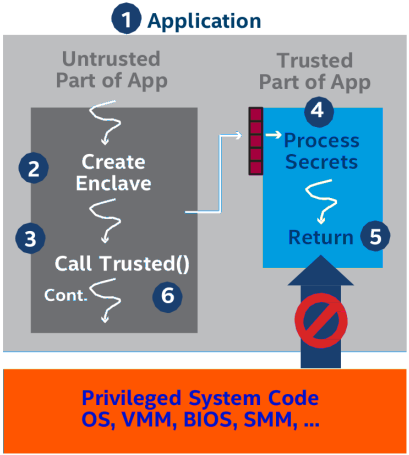

https://software.intel.com - Intel SGX Product brief

PrivaTube architecture

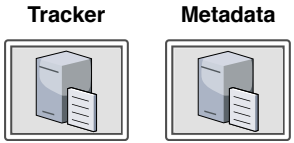

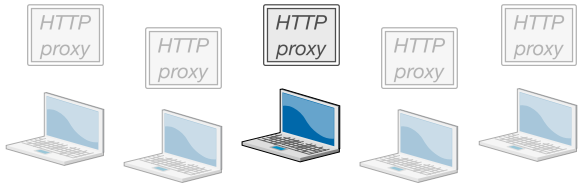

PrivaTube architecture

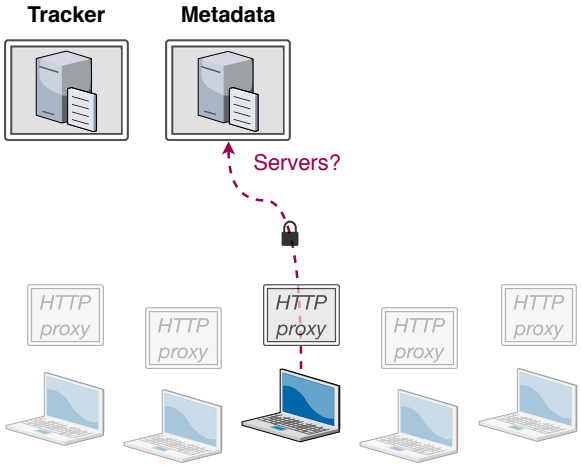

PrivaTube architecture

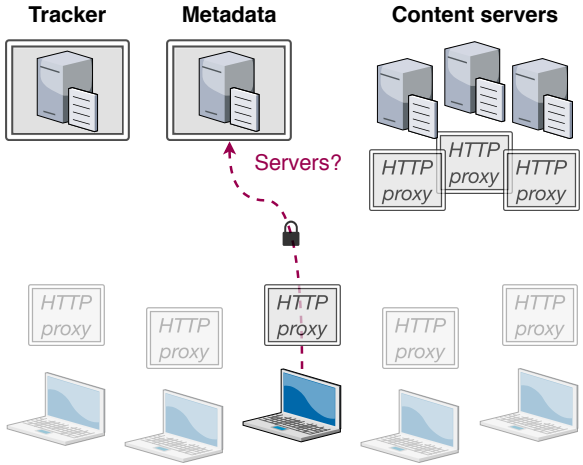

PrivaTube architecture

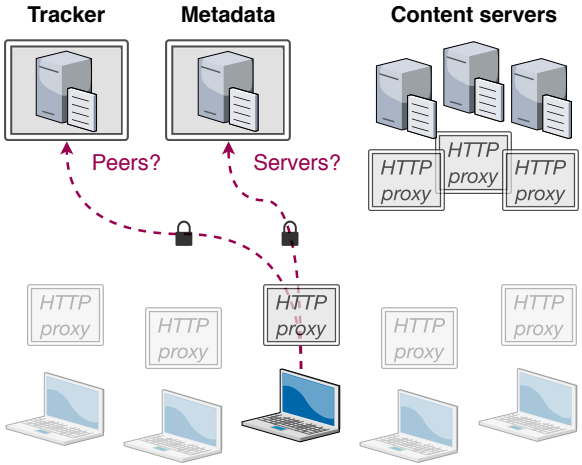

PrivaTube architecture

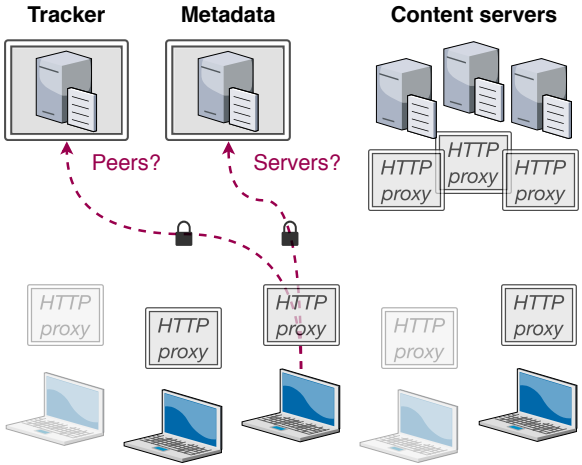

PrivaTube architecture

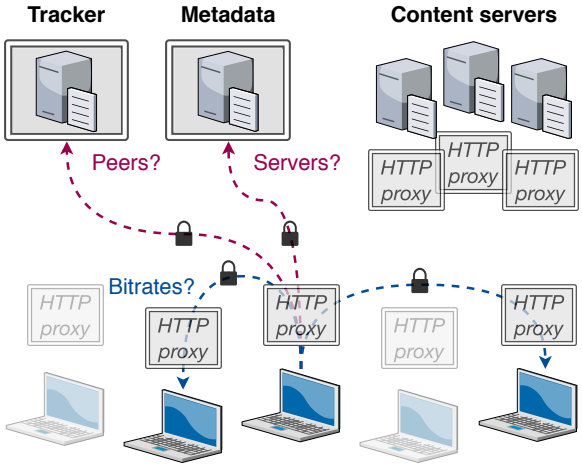

PrivaTube architecture

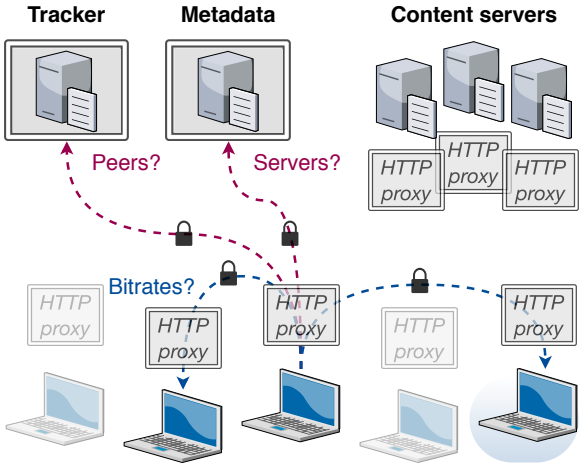

PrivaTube architecture

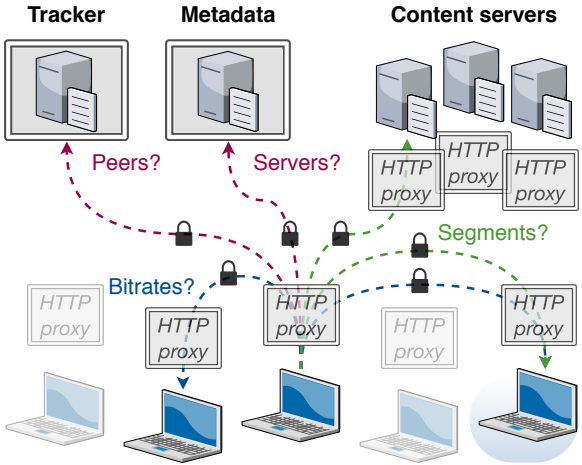

PrivaTube architecture

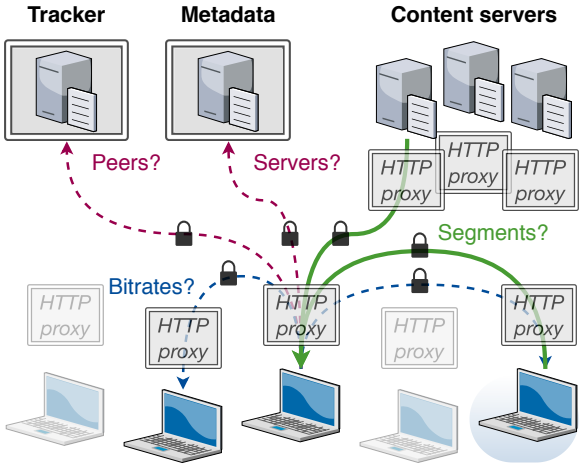

Adversary model

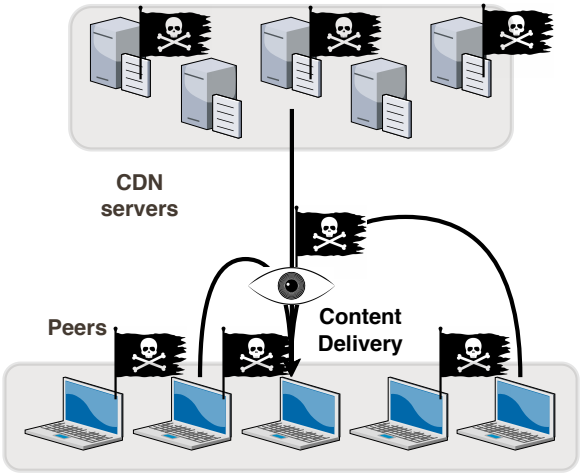

PrivaTube fake requests

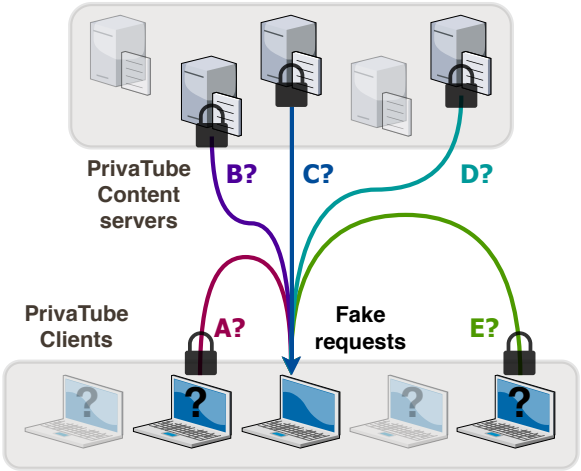

PrivaTube fake requests

Enforce privacy (probabilistic δ-unlinkability)

PrivaTube fake requests

Enforce privacy (probabilistic δ-unlinkability)

- **popularity**-aware fake requests policy

Fake requests according to contents popularity

PrivaTube fake requests

Enforce privacy (probabilistic δ-unlinkability)

- **popularity**-aware fake requests policy

Fake requests according to contents popularity

- Content pre-fetching

Evaluation: Streaming solutions

	Adaptive	Multi-source	Edge-assisted	Private
MPEG-DASH	✓	×	×	×
ClearTube	✓	✓	✓	×
PrivaTube	✓	✓	✓	✓

Micro-benchmarks setup

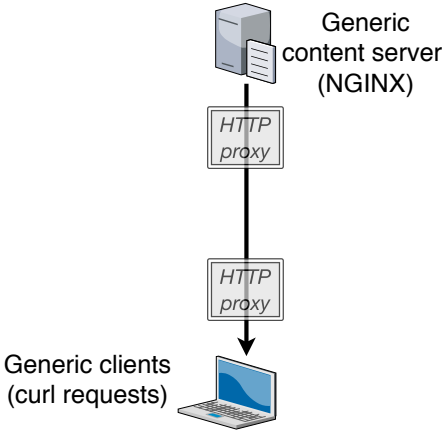

Download time

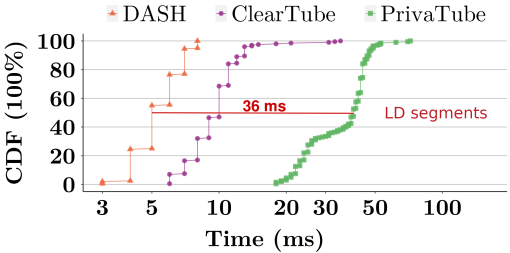

Download time

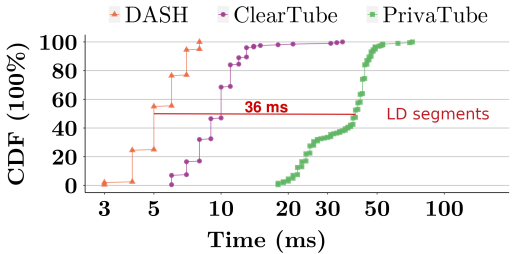

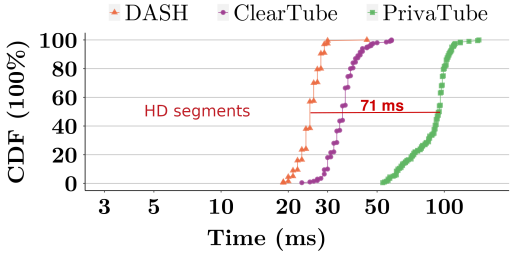

Throughput

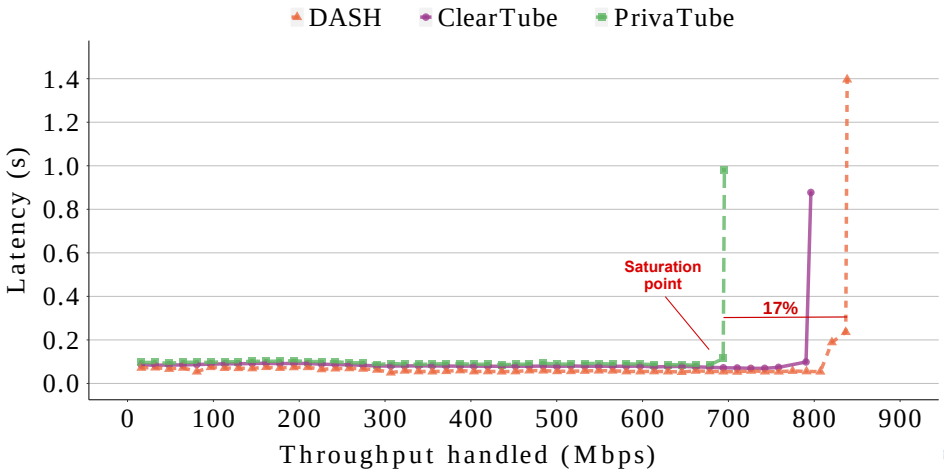
DASH
ClearTube
PrivaTube
Latency (s)
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
0
100
200
300
400
500
600
700
800
900
Throughput handled (Mbps)
Saturation point
17%

QoE evaluation setup

1. Low download time
2. High displayed bitrate

QoE evaluation setup

1. Low download time
2. High displayed bitrate

15 NUCs (computers):

- tracker + metadata + content
- 12 consuming clients

QoE evaluation

DASH

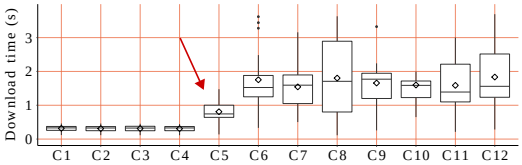

PrivaTube

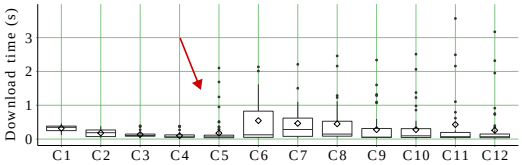

QoE evaluation

DASH

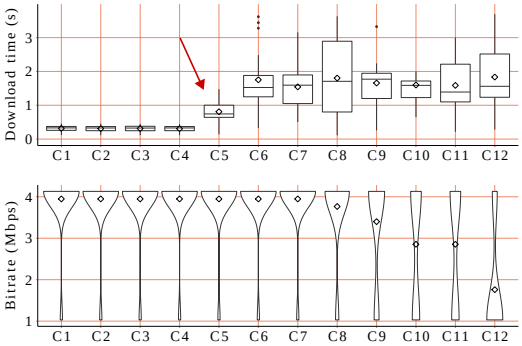

PRIVATUBE

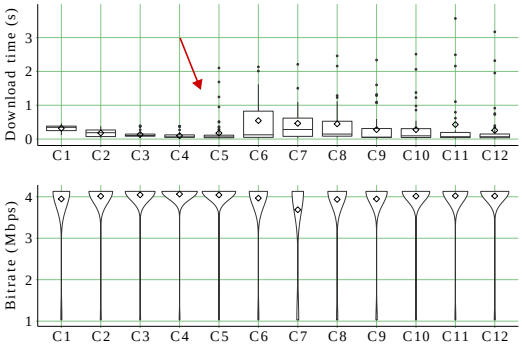

université de BORDEAUX

QoE evaluation

ClearTube

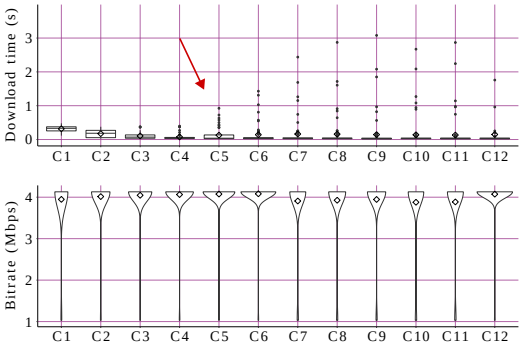

PrivaTube

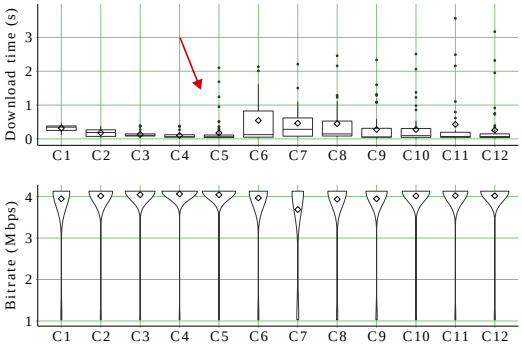

université de BORDEAUX

Privacy evaluation setup

1. Privacy (δ-unlinkability)
2. Content pre-fetching

Privacy evaluation setup

1. Privacy (δ-unlinkability)
2. Content pre-fetching

- **random**
- **popularity**-aware

MovieLens 20M dataset, year 2014
39,177 ratings from **7,763** users for **4,283** movies

Distribution of movies popularities in MovieLens

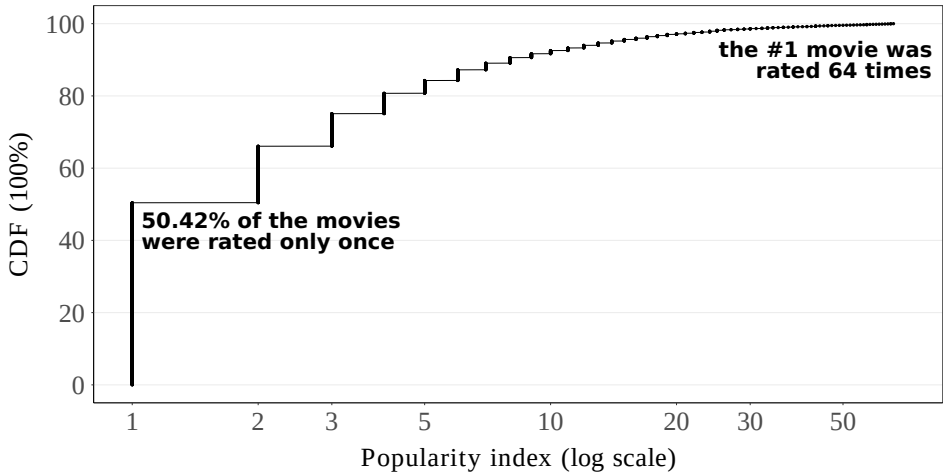

Movies copies, $\delta = 50\%$

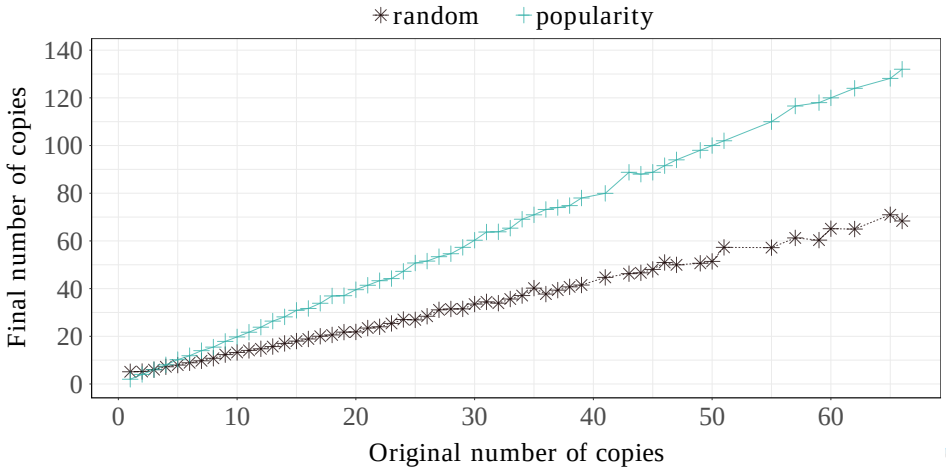

Movies copies increase factor, $\delta = 50\%$

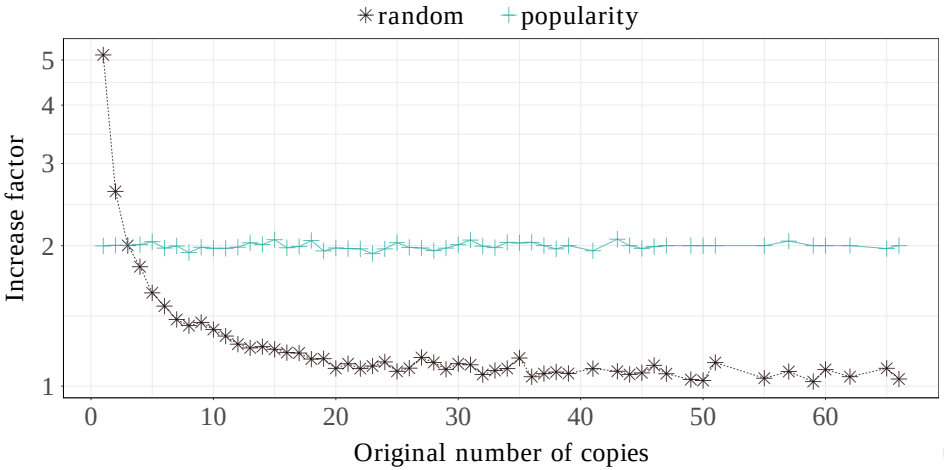

Replicas increase factor for various values of δ

# fake req.	δ = 66% × 0.5			δ = 50% × 1			δ = 25% × 3		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
random	2.3	1.7	1.4	3.5	2.5	2.6	8.6	6.0	6.9
popularity	1.5	1.5	0.5	2.0	2.0	0.6	4.0	4.0	1.1

PrivaTube - Conclusion

PrivaTube - Conclusion

- High QoE
- Reliability
- Scalability

Multiple-source, edge-assisted, adaptive video streaming

PrivaTube - Conclusion

- High QoE
- Reliability
- Scalability

Multiple-source, edge-assisted, adaptive video streaming

- Privacy

HTTP proxies inside SGX performing encryption and fake requests

PrivaTube - Conclusion

- High QoE
- Reliability
- Scalability

Multiple-source, edge-assisted, adaptive video streaming

- Privacy

HTTP proxies inside SGX performing encryption and fake requests

Pre-fetching: Fake requests according to contents popularity

PrivaTube - Conclusion

- High QoE
- Reliability
- Scalability

Multiple-source, edge-assisted, adaptive video streaming

- Privacy

HTTP proxies inside SGX performing encryption and fake requests

Pre-fetching: Fake requests according to contents popularity

S. Da Silva, S. Ben Mokhtar, S. Contiu, D. Négru, L. Réveillère, E. Rivière.

PrivaTube: Privacy-Preserving Edge-Assisted Video Streaming. *20th International Middleware Conference (Middleware '19).*
Priva-Stream: Private Collaborative Live Streaming. *19th International Middleware Conference Doctoral Symposium (Middleware '18).*
Priva-Stream: Private Collaborative Streaming. *19th International Middleware Conference Poster (Middleware '18).*

université de BORDEAUX

Video content consumption evolves...

- High QoE
- Reliability
- Scalability
- Privacy
-

Video content consumption evolves...

Video content consumption evolves...

- High QoE
- Reliability
- Scalability
- Privacy
- Recommendations

université de BORDEAUX

Recommendations

70+% of time spent on YouTube is from recommendations.

— YouTube Chief Product Officer Neal Mohan, *CES 2018*

Recommendations

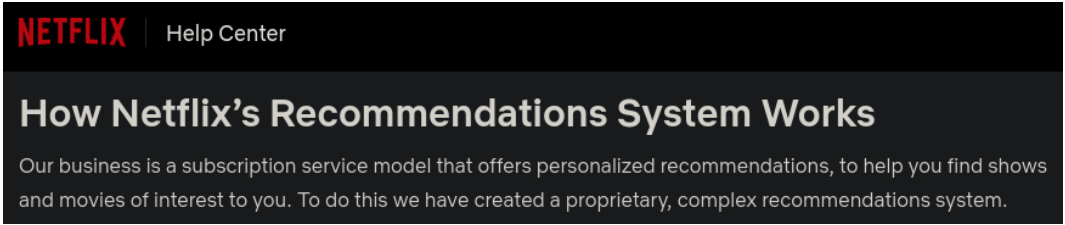

How Netflix's Recommendations System Works

Our business is a subscription service model that offers personalized recommendations, to help you find shows and movies of interest to you. To do this we have created a proprietary, complex recommendations system. This article provides a high level description of our recommendations system in plain language.

In addition to knowing **what you have watched** on Netflix, to best personalize the recommendations we also look at things like:

- the **time of day** you watch,
- the **devices** you are watching Netflix on, and
- **how long** you watch.

Privacy-preserving recommendations

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19])

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19])

Differentially private ([MM09, SJ14, SKSX18])

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19])

Differentially private ([MM09, SJ14, SKSX18])

Decentralized (PDMFRec [DTS+19])

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19]) => very slow ($\simeq$ 10s)

Differentially private ([MM09, SJ14, SKSX18])

Decentralized (PDMFRec [DTS+19])

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19]) => very slow ($\simeq$ 10s)

Differentially private ([MM09, SJ14, SKSX18]) => inaccurate

Decentralized (PDMFRec [DTS+19])

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19]) => very slow ($\simeq$ 10s)

Differentially private ([MM09, SJ14, SKSX18]) => inaccurate

Decentralized (PDMFRec [DTS+19]) => slow + less accurate + unreliable

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19]) => very slow ($\simeq$ 10s)

Differentially private ([MM09, SJ14, SKSX18]) => inaccurate

Decentralized (PDMFRec [DTS+19]) => slow + less accurate + unreliable

They all target specific recommendation algorithms, and require heavy client-side code

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19]) => very slow ($\simeq$ 10s)

Differentially private ([MM09, SJ14, SKSX18]) => inaccurate

Decentralized (PDMFRec [DTS+19]) => slow + less accurate + unreliable

They all target specific recommendation algorithms, and require heavy client-side code

Use an unmodified recommender system?

Privacy-preserving recommendations

Cryptography-based (X-Rec [GKP+17], CryptoRec [WTAR19]) => very slow ($\simeq$ 10s)

Differentially private ([MM09, SJ14, SKSX18]) => inaccurate

Decentralized (PDMFRec [DTS+19]) => slow + less accurate + unreliable

They all target specific recommendation algorithms, and require heavy client-side code

Recommendation-as-a-Service

(Darwin & Goliath [BGO19], Harness Universal Recommender [Actb], Mediego [Med19], Plista [Pli19b], Recombee [Pli19a], etc.)

université de BORDEAUX

Recommendation-as-a-Service - POST

Recommendation-as-a-Service - POST

Recommendation-as-a-Service - POST

Recommendation-as-a-Service - POST

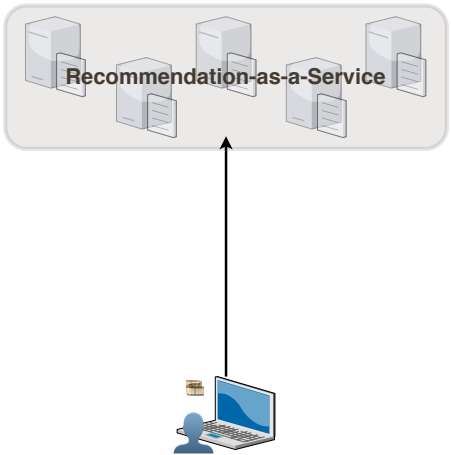

Recommendation-as-a-Service - POST

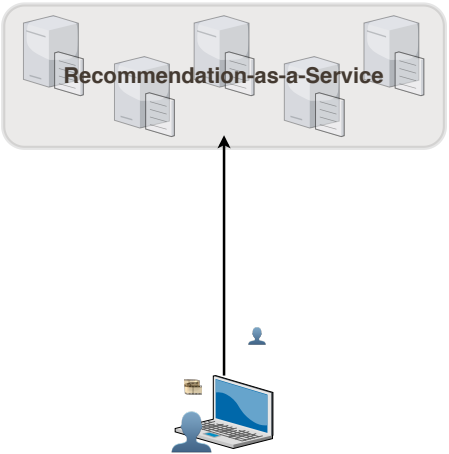

Recommendation-as-a-Service - POST

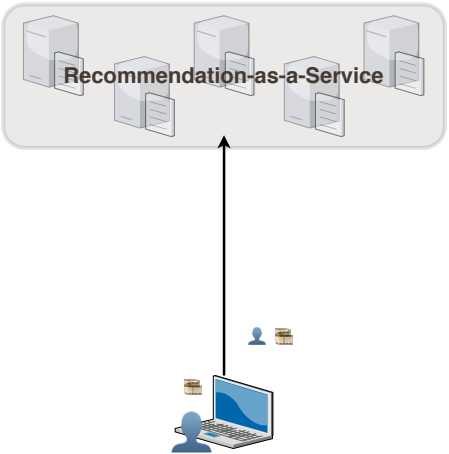

Recommendation-as-a-Service - GET

Recommendation-as-a-Service - GET

Recommendation-as-a-Service - GET

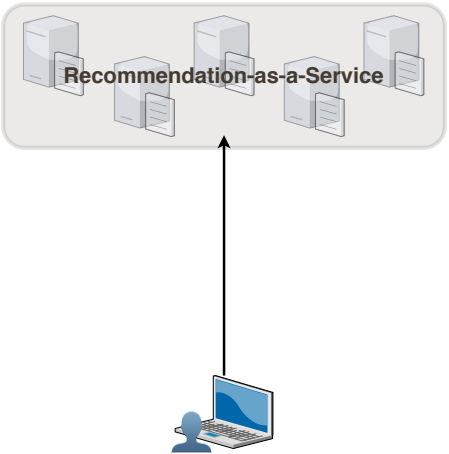

Recommendation-as-a-Service - GET

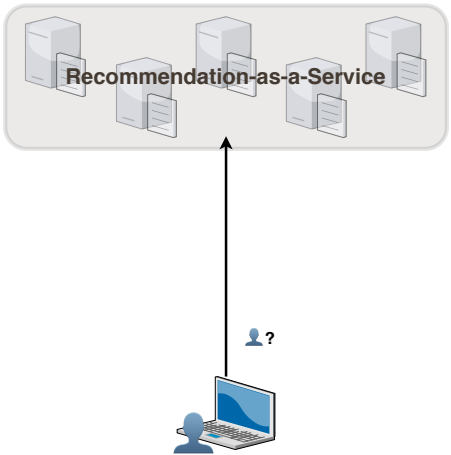

Recommendation-as-a-Service - GET

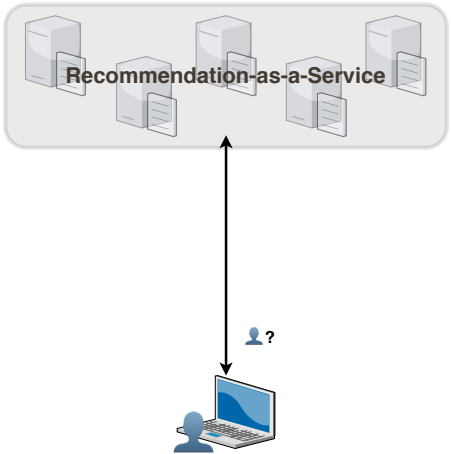

Recommendation-as-a-Service - GET

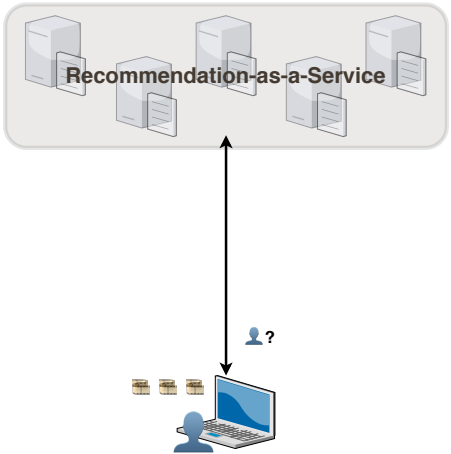

Plan

université de BORDEAUX

PProx

PProx

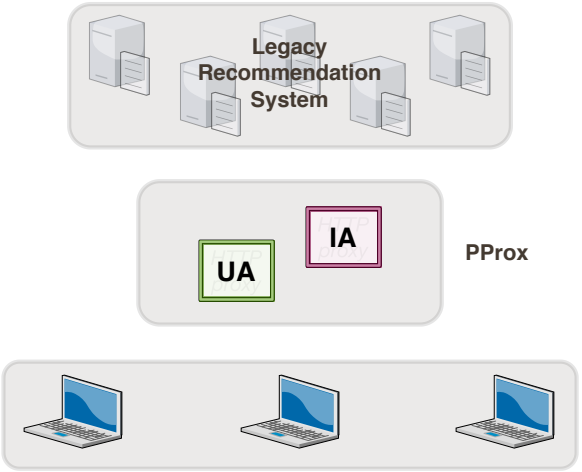

PProx

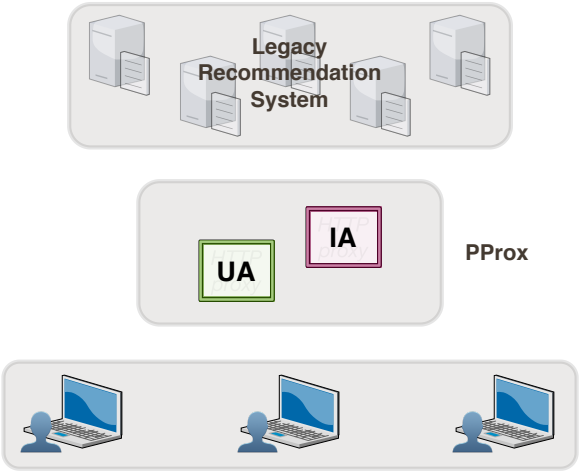

PProx

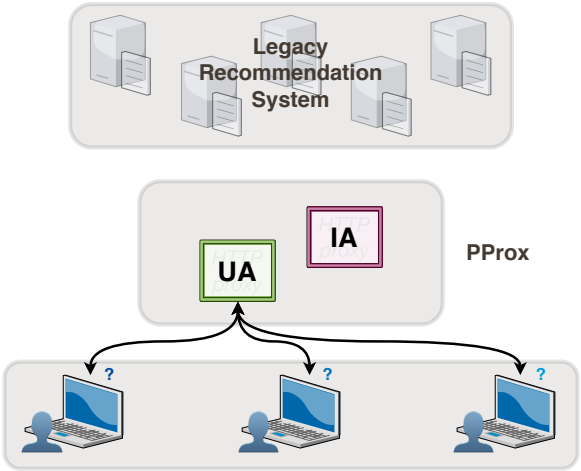

PProx

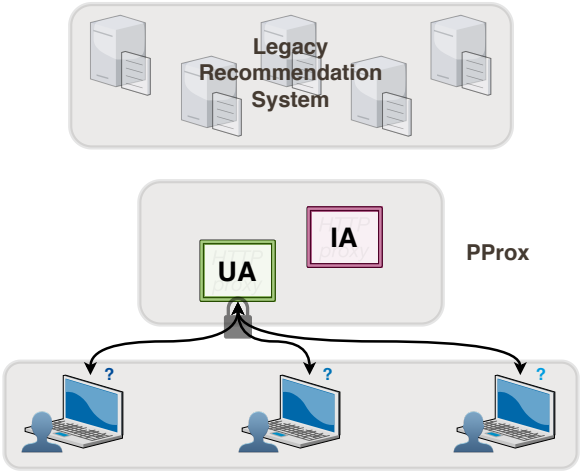

PProx

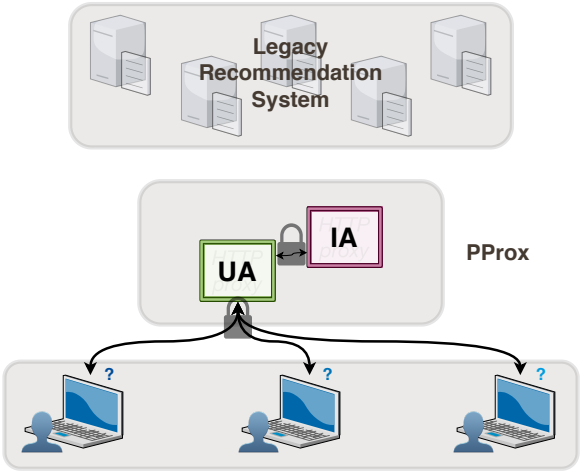

PProx

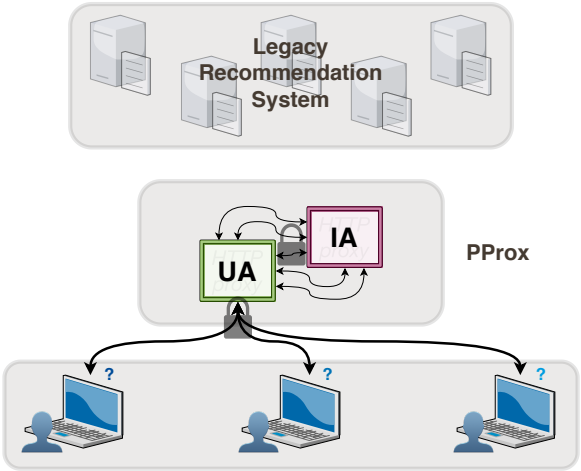

PProx

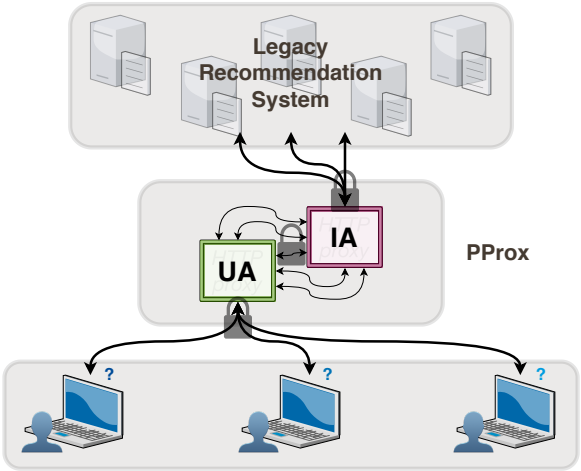

PProx

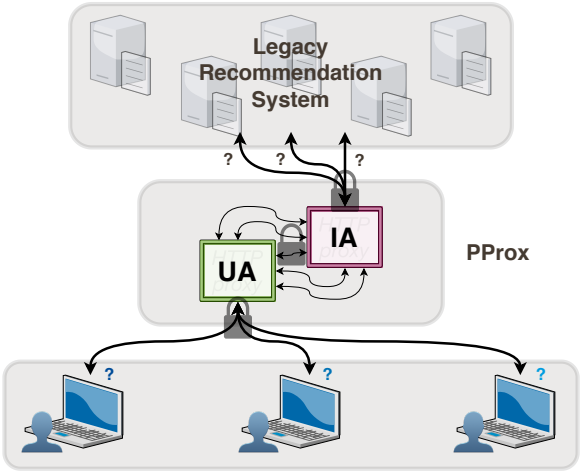

PProx security implementation - POST

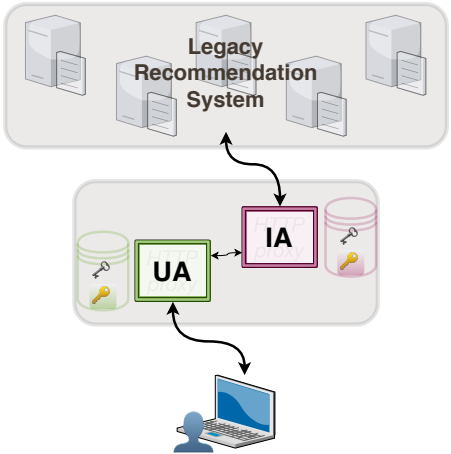

PProx security implementation - POST

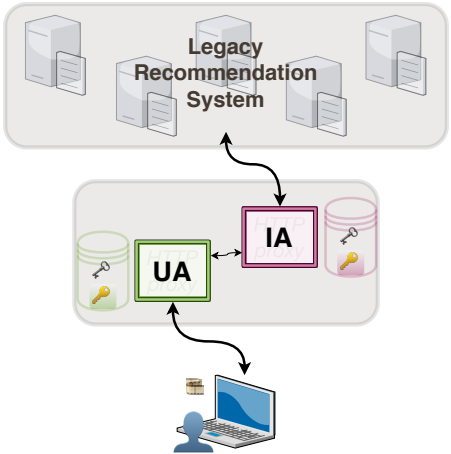

PProx security implementation - POST

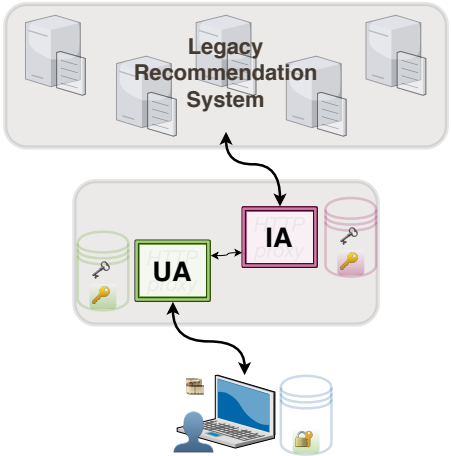

PProx security implementation - POST

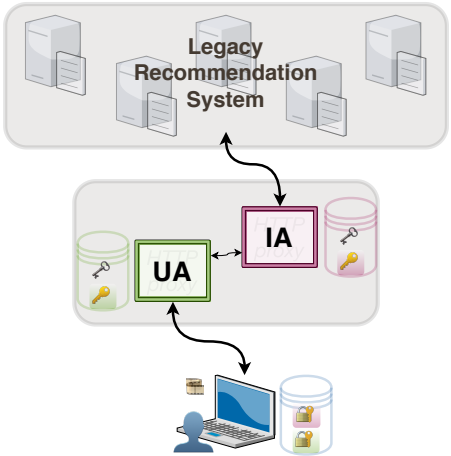

PProx security implementation - POST

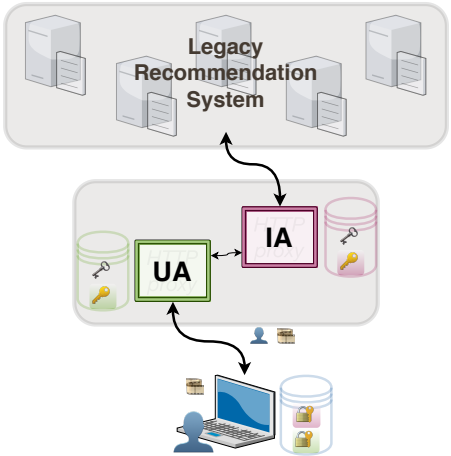

PProx security implementation - POST

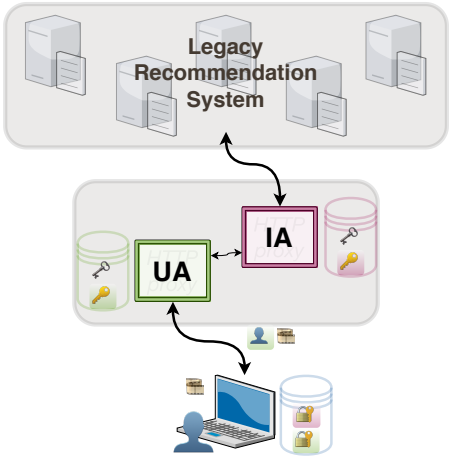

PProx security implementation - POST

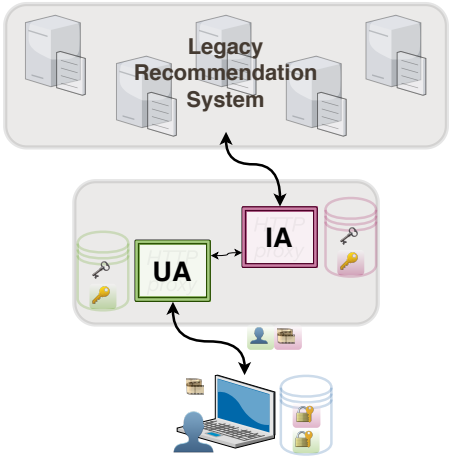

PProx security implementation - POST

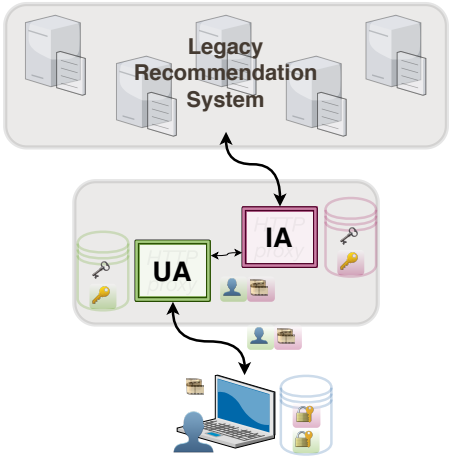

PProx security implementation - POST

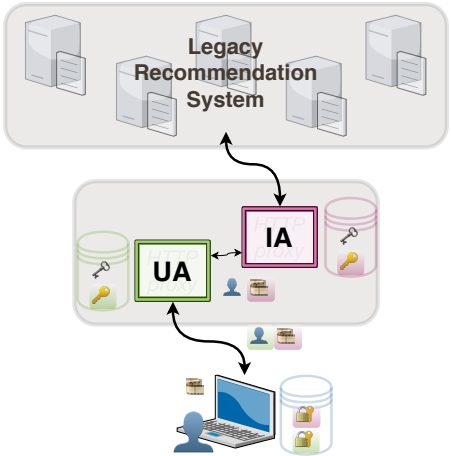

PProx security implementation - POST

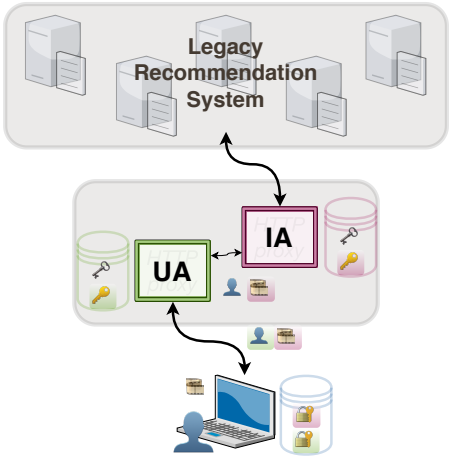

PProx security implementation - POST

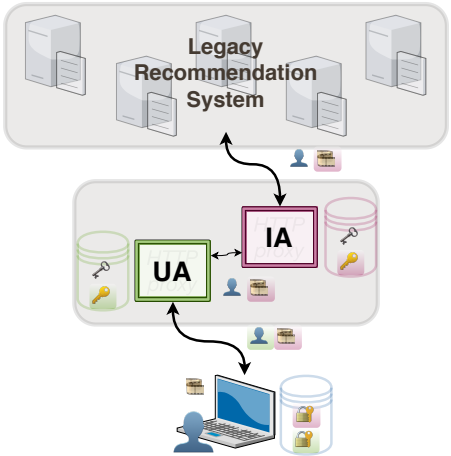

PProx security implementation - POST

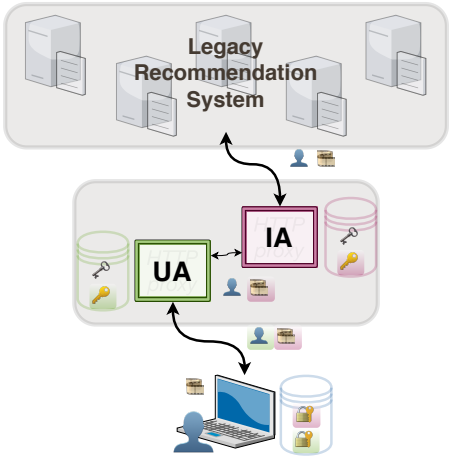

PProx security implementation - POST

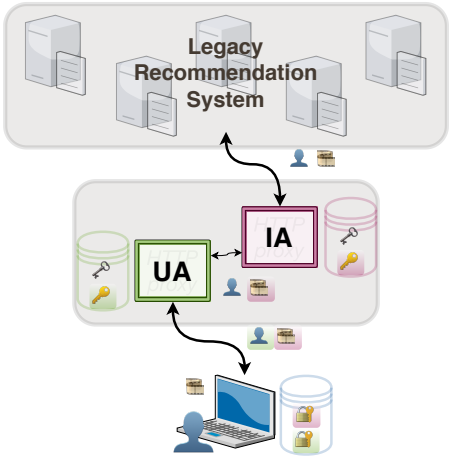

PProx security implementation - GET

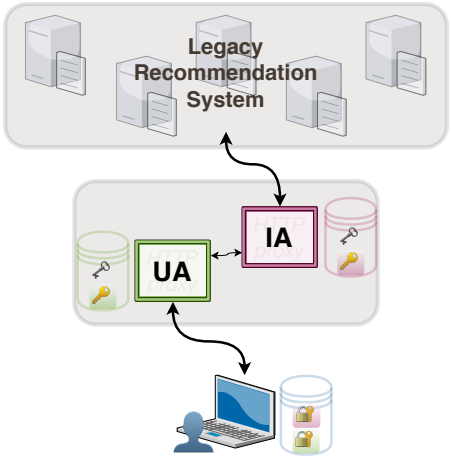

PProx security implementation - GET

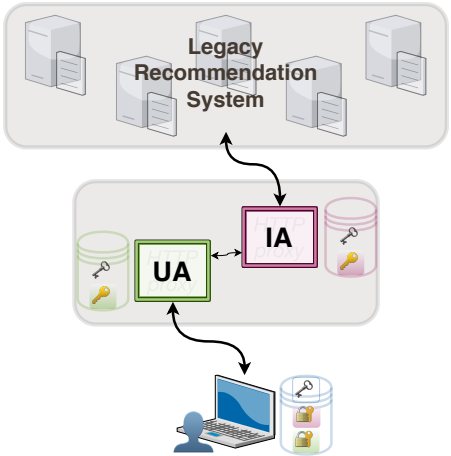

PProx security implementation - GET

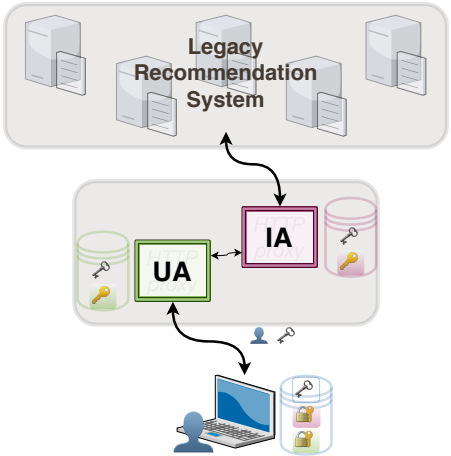

PProx security implementation - GET

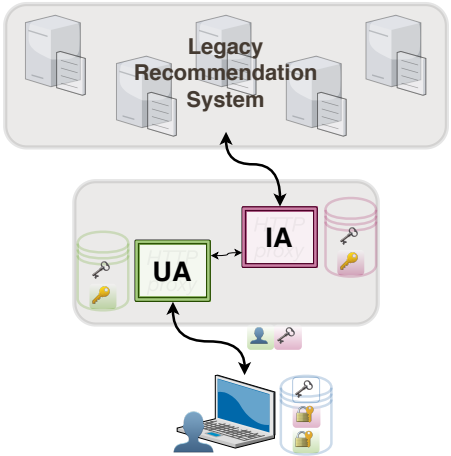

PProx security implementation - GET

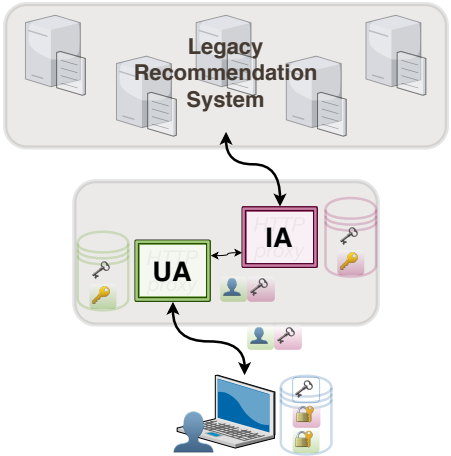

PProx security implementation - GET

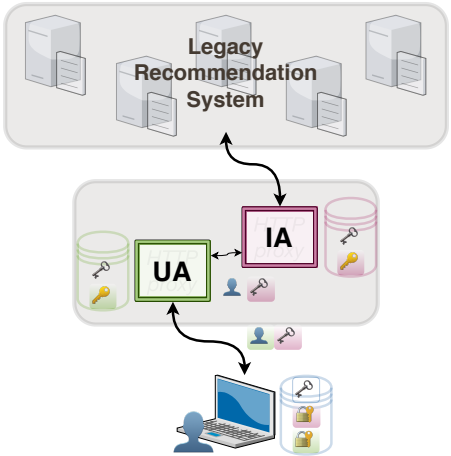

PProx security implementation - GET

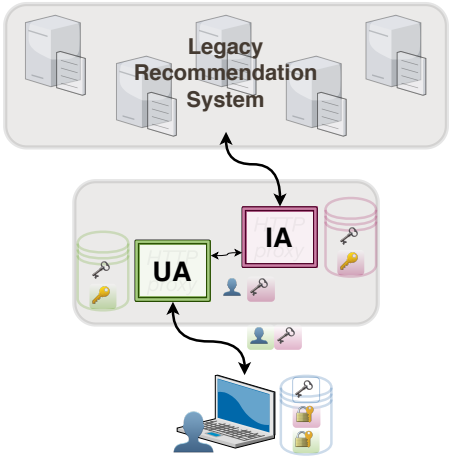

PProx security implementation - GET

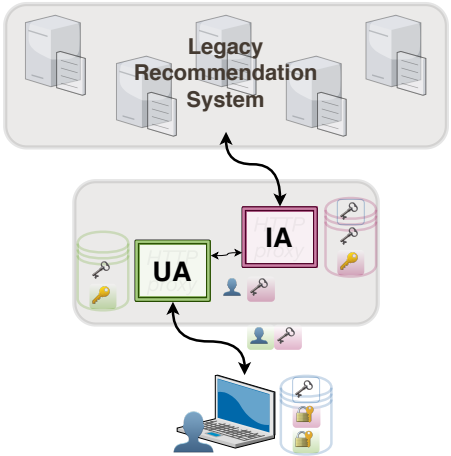

PProx security implementation - GET

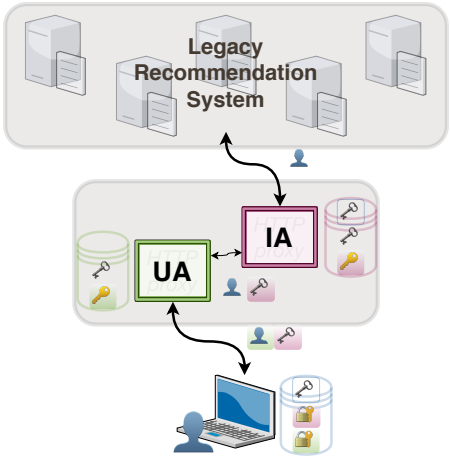

PProx security implementation - GET

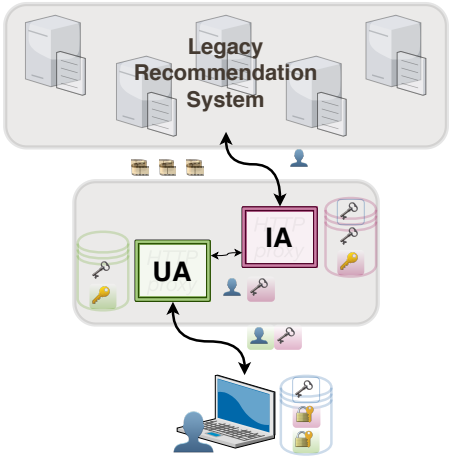

PProx security implementation - GET

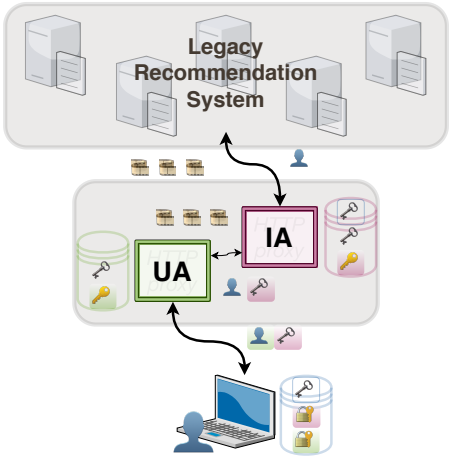

PProx security implementation - GET

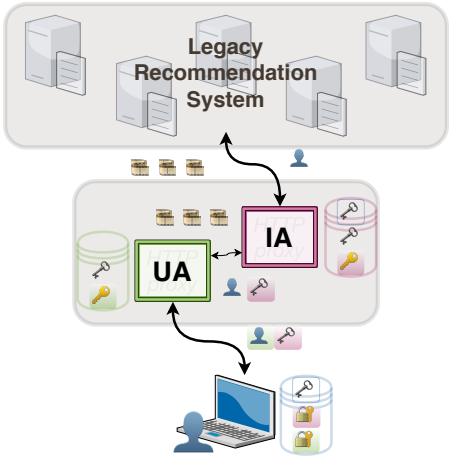

PProx security implementation - GET

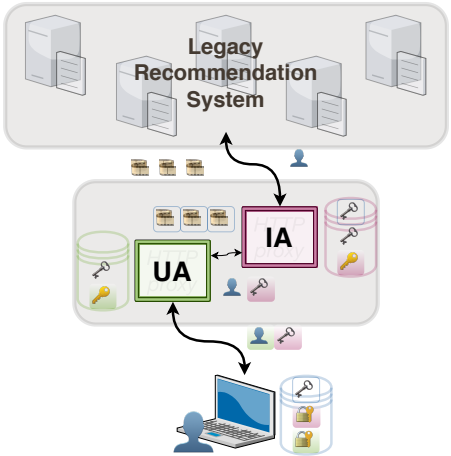

PProx security implementation - GET

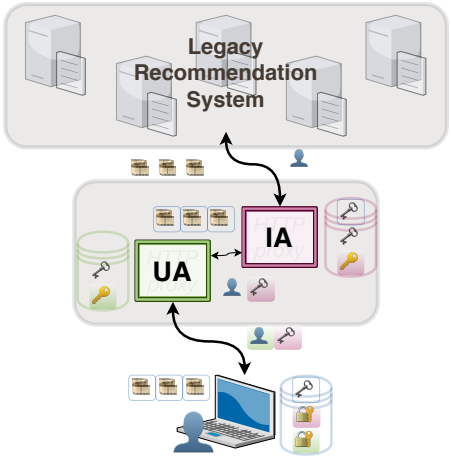

PProx security implementation - GET

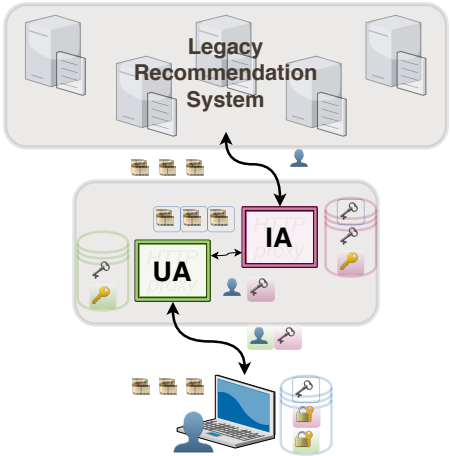

Micro-benchmark setup

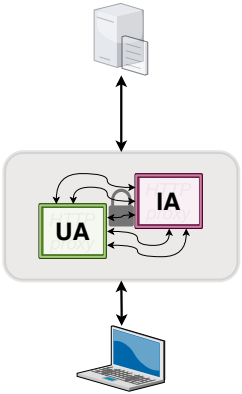

Table: Micro-benchmark configurations

	Encryption	SGX	Shuffling	RPS
m1	✗	✗	✗	250
m2	✓	✗	✗	250
m3	✓	✓	✗	250
m4	★	✓	✗	250
m5	✓	✓	5	250
m6	✓	✓	10	250

Impact of encryption and SGX

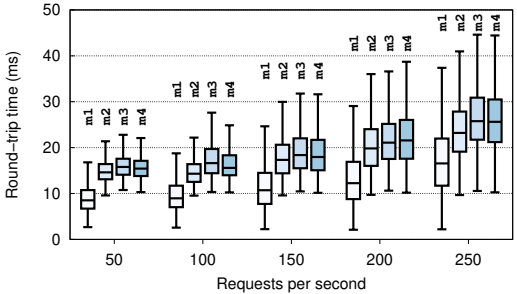

Table: Encryption and SGX configurations

	Encryption	SGX	Shuffling	RPS
m1	✗	✗	✗	250
m2	✓	✗	✗	250
m3	✓	✓	✗	250
m4	★	✓	✗	250

Impact of encryption and SGX

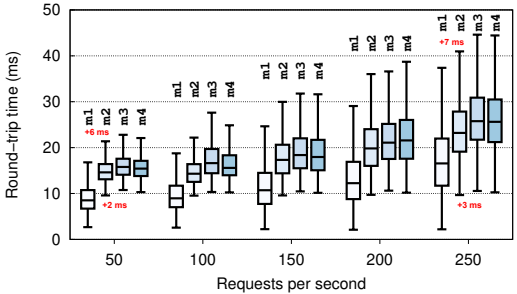

Table: Encryption and SGX configurations

	Encryption	SGX	Shuffling	RPS
m1	✗	✗	✗	250
m2	✓	✗	✗	250
m3	✓	✓	✗	250
m4	★	✓	✗	250

Impact of shuffling

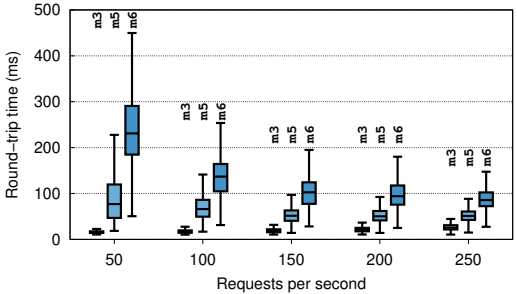

Table: Shuffling configurations

	Encryption	SGX	Shuffling	RPS
m3	✓	✓	✗	250
m5	✓	✓	5	250
m6	✓	✓	10	250

Impact of shuffling

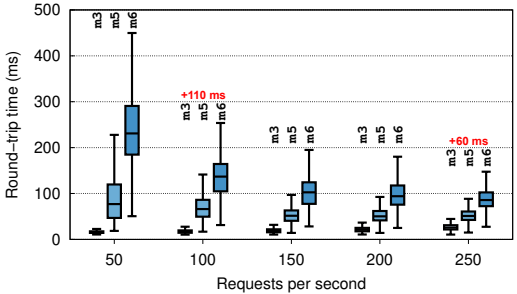

Table: Shuffling configurations

	Encryption	SGX	Shuffling	RPS
m3	✓	✓	✗	250
m5	✓	✓	5	250
m6	✓	✓	10	250

PProx scalability benchmark setup

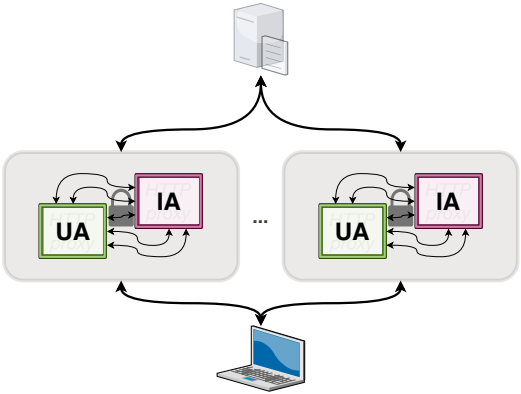

Table: PProx scalability configurations

	Enc.	SGX	S	UA	IA	RPS
m6	✓	✓	10	1	1	250
m7	✓	✓	10	2	2	500
m8	✓	✓	10	3	3	750
m9	✓	✓	10	4	4	1000

PProx scalability

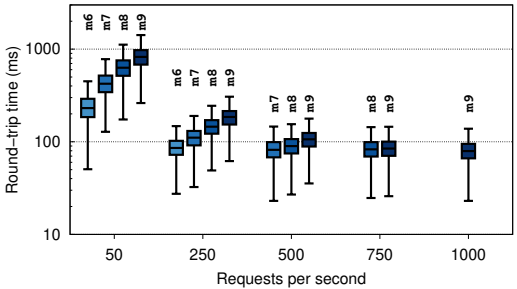

Table: PProx scalability configurations

	Enc.	SGX	S	UA	IA	RPS
m6	✓	✓	10	1	1	250
m7	✓	✓	10	2	2	500
m8	✓	✓	10	3	3	750
m9	✓	✓	10	4	4	1000

LRS scalability benchmark setup

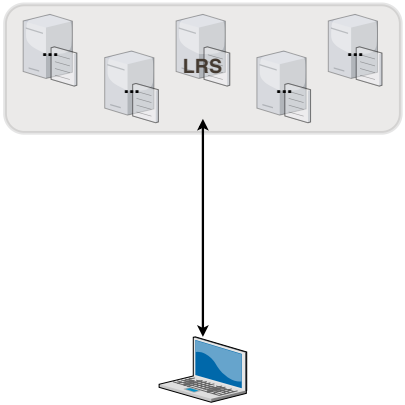

Legacy Recommendation System:
Harness *Universal Recommender* module

Table: Baseline LRS scalability configurations

	Enc.	SGX	S	UA	IA	LRS	RPS
b1	✗	✗	✗	✗	✗	3+4	250
b2	✗	✗	✗	✗	✗	6+4	500
b3	✗	✗	✗	✗	✗	9+4	750
b4	✗	✗	✗	✗	✗	12+4	1000

LRS scalability

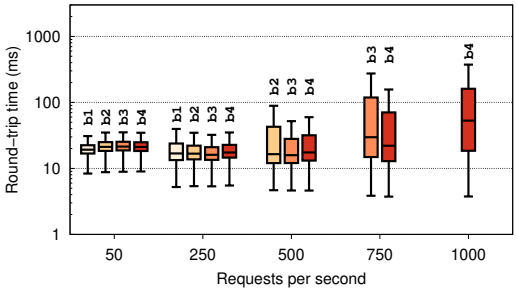

Legacy Recommendation System:
Harness *Universal Recommender* module

Table: Baseline LRS scalability configurations

	Enc.	SGX	S	UA	IA	LRS	RPS
b1	✗	✗	✗	✗	✗	3+4	250
b2	✗	✗	✗	✗	✗	6+4	500
b3	✗	✗	✗	✗	✗	9+4	750
b4	✗	✗	✗	✗	✗	12+4	1000

Macro-benchmark setup

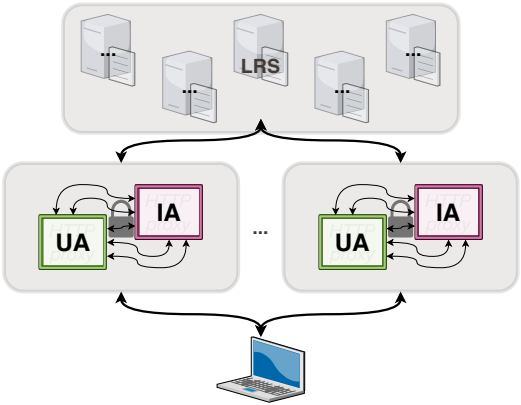

Table: Full macro-benchmark configurations

	Enc.	SGX	S	UA	IA	LRS	RPS
f1	✓	✓	10	1	1	3+4	250
f2	✓	✓	10	2	2	6+4	500
f3	✓	✓	10	3	3	9+4	750
f4	✓	✓	10	4	4	12+4	1000

Macro-benchmark

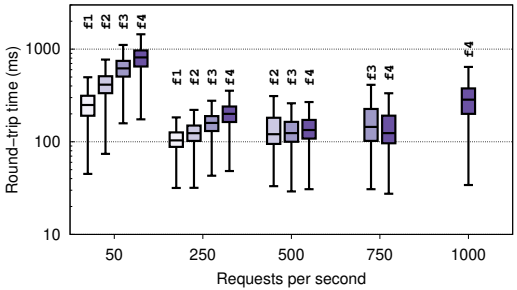

Table: Full macro-benchmark configurations

	Enc.	SGX	S	UA	IA	LRS	RPS
f1	✓	✓	10	1	1	3+4	250
f2	✓	✓	10	2	2	6+4	500
f3	✓	✓	10	3	3	9+4	750
f4	✓	✓	10	4	4	12+4	1000

PProx - Conclusion

PProx - Conclusion

- High QoE

Unaltered recommendations accuracy and low latency

PProx - Conclusion

- High QoE

Unaltered recommendations accuracy and low latency

- Reliability
- Scalability

Elastically scales, only requires 30% to 50% additional nodes

PProx - Conclusion

- High QoE

Unaltered recommendations accuracy and low latency

- Reliability
- Scalability

Elastically scales, only requires 30% to 50% additional nodes

- Privacy

HTTP proxies inside SGX performing encryption and shuffling

PProx - Conclusion

- High QoE

Unaltered recommendations accuracy and low latency

- Reliability
- Scalability

Elastically scales, only requires 30% to 50% additional nodes

- Privacy

HTTP proxies inside SGX performing encryption and shuffling

G. Rosinosky, S. Da Silva, S. Ben Mokhtar, D. Négru, L. Réveillère, E. Rivière.

PProx: Efficient Privacy for Recommendation-as-a-Service. *Pending*

université de BORDEAUX

Contributions summary

- High and fairly shared QoE
- Reduced infrastructure cost

Contributions summary

MUSLIN

- High and fairly shared QoE
- Reduced infrastructure cost

PRIVATUBE

- Unaltered QoE
- Reliability and scalability
- Strong privacy preservation

Contributions summary

MUSLIN

- High and fairly shared QoE
- Reduced infrastructure cost

PrivaTube

- Unaltered QoE
- Reliability and scalability
- Strong privacy preservation

PProx

- Unaltered recommendations accuracy
- Two-stage users and items pseudonymisation

Further research directions

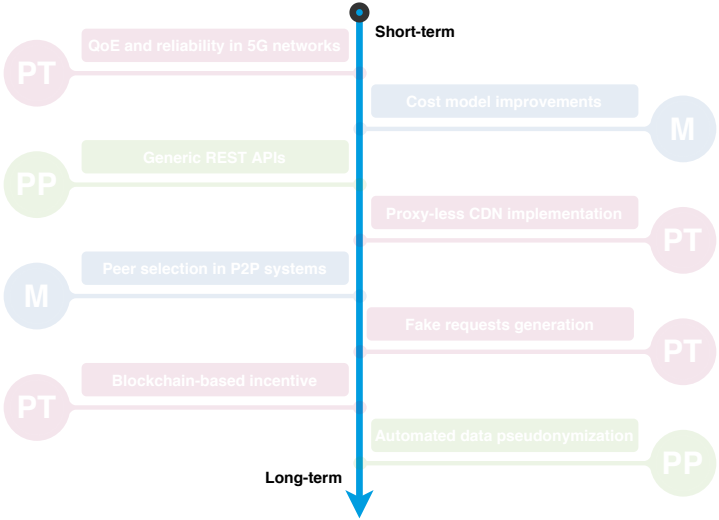

Further research directions

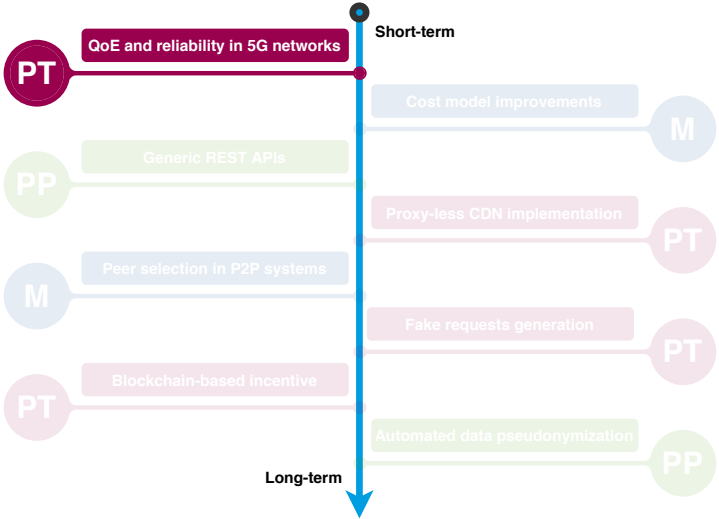

Further research directions

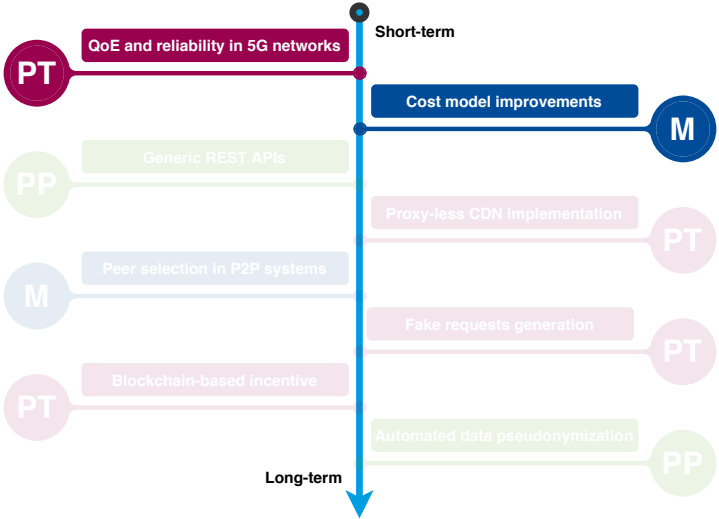

Further research directions

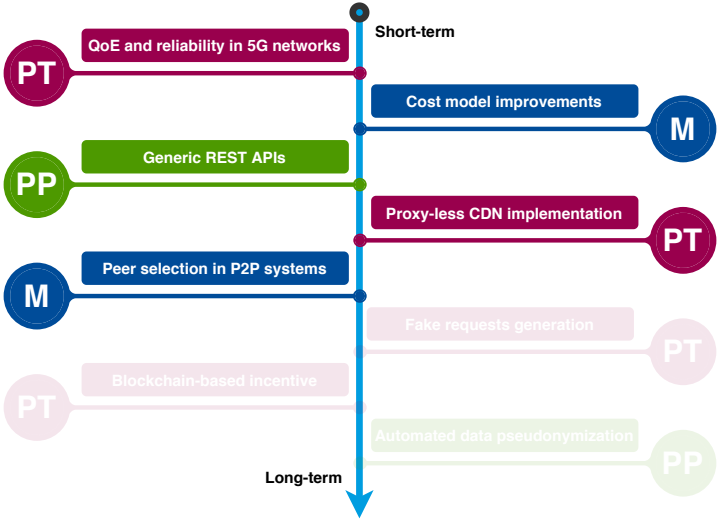

Further research directions

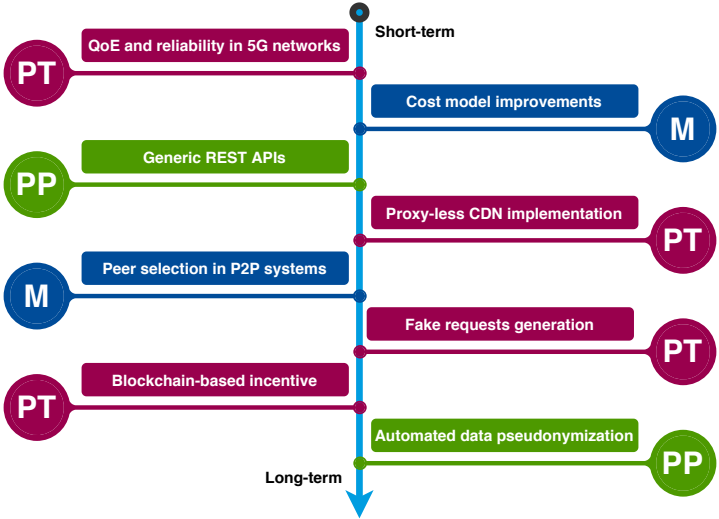

université de BORDEAUX

Conclusion

Strong privacy preservation should no longer come with low performance and low QoE.

Conclusion

Strong privacy preservation should no longer come with low performance and low QoE.

We hope generalizing MUSLIN, PrivaTube and PProx will enable a new generation of **High-QoE Privacy-Preserving Video Streaming** services with great performance.

MUSLIN
- Overview
- Implementation
- Evaluation

PrivaTube
- Overview
- Fake requests
- Evaluation

PProx
- Overview
- Security
- Evaluation

université de BORDEAUX