

MultiPub: Latency and Cost-Aware Global-Scale Cloud Publish/Subscribe

ICDCS 2017

Julien Gascon-Samson, Bettina Kemme, Jörg Kienzle

Post-Doctoral Fellow @ University of British Columbia
Current Advisor: Dr Karthik Pattabiraman
Department of Electrical and Computer Engineering
Vancouver, Canada



Electrical and
Computer
Engineering



Work completed under the advisement of:
Dr Bettina Kemme and Dr Jörg Kienzle
School of Computer Science, McGill University
Montreal, Canada



Tuesday June 6th, 2017

Topic-Based Publish/Subscribe



2


- Provides an elegant way to decouple content producers (publishers) from content consumers (subscribers)
- Very common flavour of pub/sub
- Subscription language: a key (topic name)
- Publications tagged with a topic T , sent to all subscribers of T

Topic-Based Publish/Subscribe - Twitter Example

Twitter Publishers



Twitter Followers



Subscribe

 $\{T=CBCNews\}$

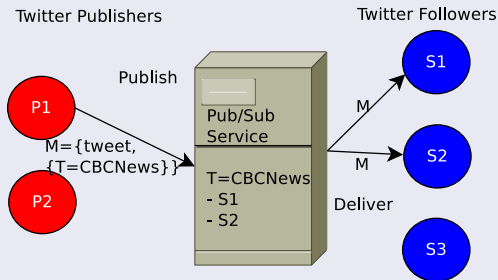
Topic-Based Publish/Subscribe



2

- Provides an elegant way to decouple content producers (publishers) from content consumers (subscribers)
- Very common flavour of pub/sub
- Subscription language: a key (topic name)
- Publications tagged with a topic T , sent to all subscribers of T

Topic-Based Publish/Subscribe - Twitter Example



Applications of Topic-Based Pub/Sub




3

Traffic alert systems



Weather alert systems



Mobile notif. frameworks



Social networks



Chat/IM systems



Multiplayer Games



Applications of Topic-Based Pub/Sub



3

Traffic alert systems



Weather alert systems



Cloud:

Mobile notif. frameworks



Social networks



Chat/IM systems



Multiplayer Games



Applications of Topic-Based Pub/Sub



3

Traffic alert systems



Weather alert systems



Mobile notif. frameworks



Social networks



Chat/IM systems



Multiplayer Games



Cloud:

- Scalability & Elasticity

Applications of Topic-Based Pub/Sub



3

Traffic alert systems



Weather alert systems



Mobile notif. frameworks



Social networks



Chat/IM systems



Multiplayer Games



Cloud:

- Scalability & Elasticity
- Availability & Fault Tolerance

Applications of Topic-Based Pub/Sub



3

Traffic alert systems



Weather alert systems



Mobile notif. frameworks



Social networks



Chat/IM systems



Multiplayer Games



Cloud:

- Scalability & Elasticity
- Availability & Fault Tolerance

Challenge:

- 1 Guarantee suitable delivery times
- 2 Reducing cloud-induced costs

EC2 Cloud Regions - Cost vs Latency



4

Region	Location	\$/GB	Region	Location	\$/GB
us-east-1	N. Virginia	0.09	ap-northeast-1	Tokyo	0.14
us-west-1	N. California	0.09	ap-northeast-2	Seoul	0.126
us-west-2	Oregon	0.09	ap-southeast-1	Singapore	0.12
eu-west-1	Ireland	0.09	ap-southeast-2	Sydney	0.14
eu-central-1	Frankfurt	0.09	sa-east-1	Sao Paulo	0.25

EC2 Cloud Regions - Cost vs Latency



4

Region	Location	\$/GB	Region	Location	\$/GB
us-east-1	N. Virginia	0.09	ap-northeast-1	Tokyo	0.14
us-west-1	N. California	0.09	ap-northeast-2	Seoul	0.126
us-west-2	Oregon	0.09	ap-southeast-1	Singapore	0.12
eu-west-1	Ireland	0.09	ap-southeast-2	Sydney	0.14
eu-central-1	Frankfurt	0.09	sa-east-1	Sao Paulo	0.25



Example Scenario



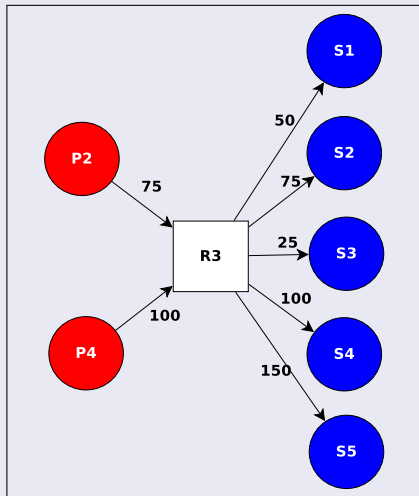
- Large-Scale MMOG: World of Warcraft
 - Assuming ~ 5 M active subscribers
 - Bandwidth usage - 10-20 mb/hour per player
 - Assuming a playing time of 2 hours/day per player
 - Total yearly outgoing bandwidth: 50 TB
 - Assuming 90% of all delivery times should be ≤ 150 ms
 - Topic-based pub/sub model

Latency & Cost Analysis



6

Solution 1: Minimal Cost




Latency & Cost Analysis



6

Solution 1: Minimal Cost



Latencies


 $75+25$ $75+50$ $100+25$ $75+75$ $100+50$ $100+75$ $75+100$ $100+100$ $75+150$ $100+150$

Latency & Cost Analysis



6

Solution 1: Minimal Cost




Latencies	Latencies
75+25	100
75+50	125
100+25	125
75+75	150
100+50	150
100+75	175
75+100	175
100+100	200
75+150	225
100+150	250

Latency & Cost Analysis



Solution 1: Minimal Cost



Cost Computation

R	\$/gb	% Publications
R_1	0.25	
R_2	0.126	
R_3	0.09	100%
R_4	0.14	
R_5	0.14	

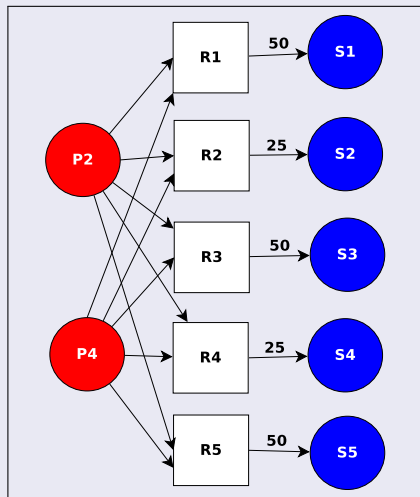
- Assuming: total bandwidth 50TB
- Yearly cost: 4.5M\$

Latency & Cost Analysis



6

Solution 2: Min. Delivery Times

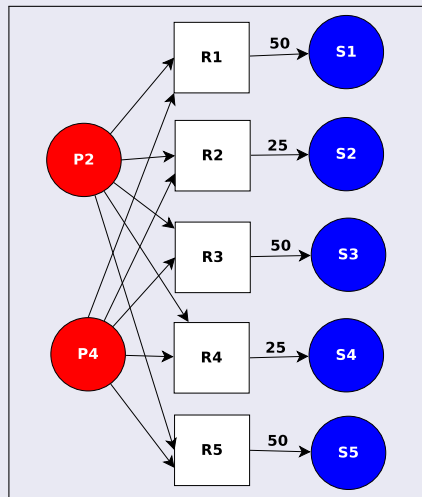


Latency & Cost Analysis



6

Solution 2: Min. Delivery Times



Latencies

25+25

25+25

50+25

50+25

50+50

50+50

75+50

75+50

75+50


100+50

Latency & Cost Analysis



6

Solution 2: Min. Delivery Times




Latencies	Latencies
25+25	50
25+25	50
50+25	75
50+25	75
50+50	100
50+50	100
75+50	125
75+50	125
75+50	125
100+50	150

Latency & Cost Analysis



6

Solution 2: Min. Delivery Times



Cost Computation

R	\$/gb	% Publications
R_1	0.25	20%
R_2	0.12	20%
R_3	0.09	20%
R_4	0.14	20%
R_5	0.14	20%


- Assuming: total bandwidth 50TB
- Yearly cost: ~~4.5M\$~~ 7.4M\$

Latency & Cost Analysis



6

Solution 3: Best Tradeoff




Latency & Cost Analysis



6

Solution 3: Best Tradeoff



Latencies

50+25

50+50

75+25

75+25

75+50

75+50

100+25

75+75

75+75


100+50

Latency & Cost Analysis



6

Solution 3: Best Tradeoff




Latencies	Latencies
50+25	75
50+50	100
75+25	100
75+25	100
75+50	125
75+50	125
100+25	125
75+75	150
75+75	150
100+50	150

Latency & Cost Analysis



6

Solution 3: Best Tradeoff



Cost Computation

R	\$/gb	% Publications
R_1	0.25	
R_2	0.12	
R_3	0.09	60%
R_4	0.14	
R_5	0.14	40%


- Assuming: total bandwidth 50TB
- Yearly cost: 4.5M\$ 7.4M\$ 5.5M\$

Optimization Problem



7

	90% delivery time	Yearly Costs
1) Cost-Minimizing	225 ms	4.5M\$
2) Latency-Minimizing	125 ms	7.4M\$
3) Best Tradeoff	150 ms	5.5M\$



- *Minimizing cloud costs, while meeting the needs of latency-constrained applications*

Conclusion & Future Work



Please have a look at our paper for evaluation results.

Summary of Contributions

- Goal: minimizing the cost of topic-based pub/sub-based applications in the cloud while respecting latency constraints
- Optimization problem: assign topics to cloud regions
- Two different delivery schemes
- Full model implementation & several simulation experiments

Conclusion & Future Work



Please have a look at our paper for evaluation results.

Summary of Contributions

- Goal: minimizing the cost of topic-based pub/sub-based applications in the cloud while respecting latency constraints
- Optimization problem: assign topics to cloud regions
- Two different delivery schemes
- Full model implementation & several simulation experiments

Future Work


- Clustering and heuristic approaches
- Supporting content-based pub/sub
- Taking VM rental costs into consideration
- Multiple clouds

Results




9

Delivery Time



Costs




Content-Based Publish/Subscribe



10

- Also referred to as *attribute-based* pub/sub
 - Lots of research attention over the past years
 - Subscription language: predicate over metadata or over the data itself
 - Matching process: non-trivial (can be CPU-intensive)

Content-Based Publish/Subscribe - Stock Auction Example




Content-Based Publish/Subscribe



10

- Also referred to as *attribute-based* pub/sub
 - Lots of research attention over the past years
 - Subscription language: predicate over metadata or over the data itself
 - Matching process: non-trivial (can be CPU-intensive)

Content-Based Publish/Subscribe - Stock Auction Example



Publication Delivery Time



	R_1	R_2	R_3	R_4	R_5	R_6	R_7	R_8	R_9	R_{10}
R_1	-	33	42	41	45	74	92	109	115	60
R_2	34	-	11	74	84	52	68	88	79	93
R_3	42	11	-	70	79	45	60	81	81	82
R_4	41	74	70	-	10	107	122	105		
R_5	45	84	79	11	-	117	135	126		
R_6	73	52	45	107	117	-	17	38		
R_7	89	68	60	122	133	17	-	34	67	144
R_8	109	88	81	97	126	38	34	-	88	164
R_9	115	79	81	155	162	52	67	88	-	171
R_{10}	60	93	91	96	100	129	144	165	171	-




Typical Delivery Approach



12

One server/region assigned for a given topic T



MultiPub



13

