Reducing the Subscription Latency in Reliable Causal Publish-Subscribe Systems

Filipa Pedrosa and Luís Rodrigues

INESC-ID, Instituto Superior Técnico, Universidade de Lisboa

Publish-Subscribe

Publish-Subscribe

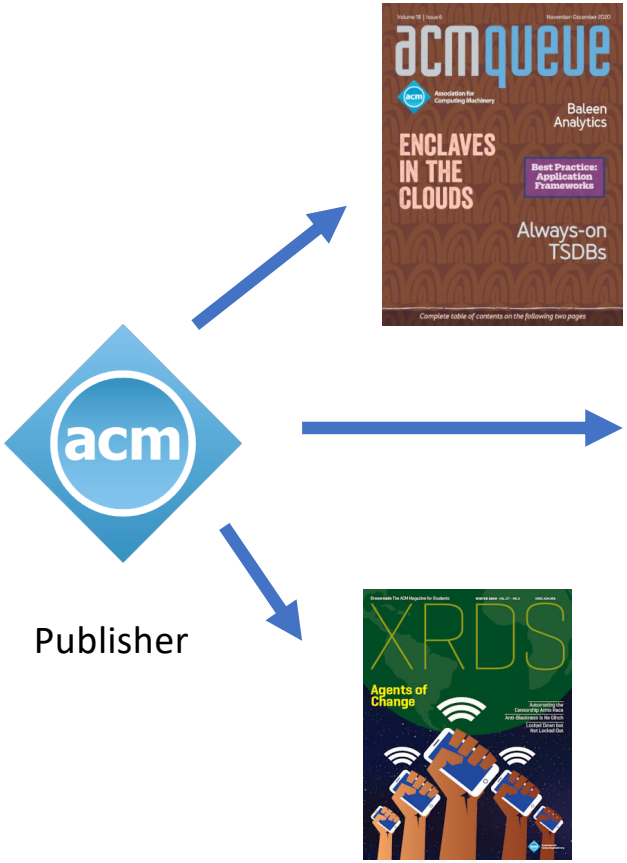

Publisher

Subscriber

Subscription Latency

Subscription Latency

this paper: interplay between reliability and δ

Reliability vs Subscription Latency

- **Best-Effort Delivery**
  - Subscriber can miss some events
- **Gapless FIFO Delivery (GFD)**
  - After receiving some event **e** from publisher **p**, the subscriber receives all future events from that publisher **p**
- **Gapless Causal Delivery (GCD)**
  - After receiving some event **e** from publisher **p**, the subscriber receives all future events, that causally depend on **e**, from any publisher (and not only from **p**).

Gapless FIFO vs Gapless Causal Delivery

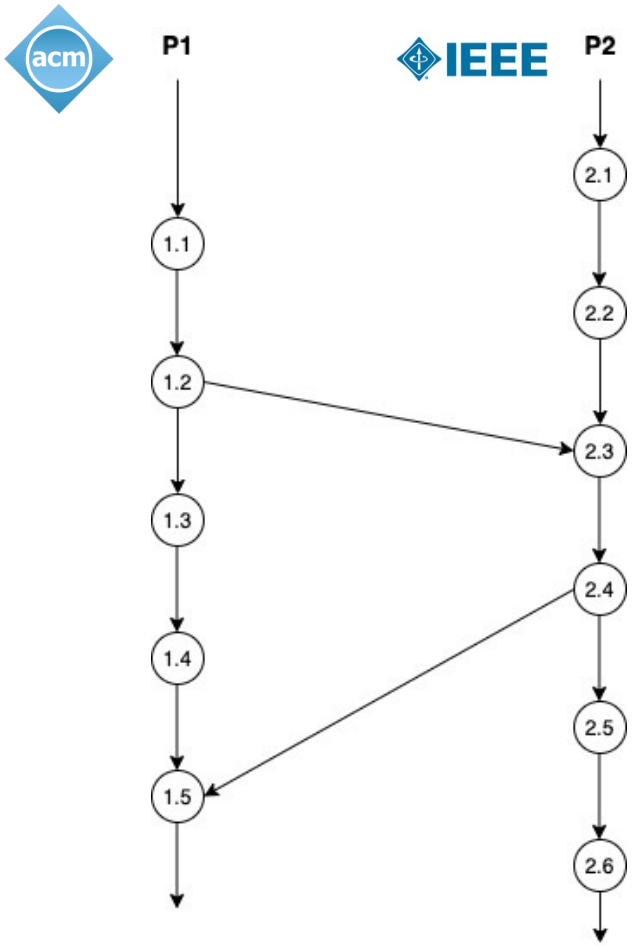

Gapless FIFO vs Gapless Causal Delivery

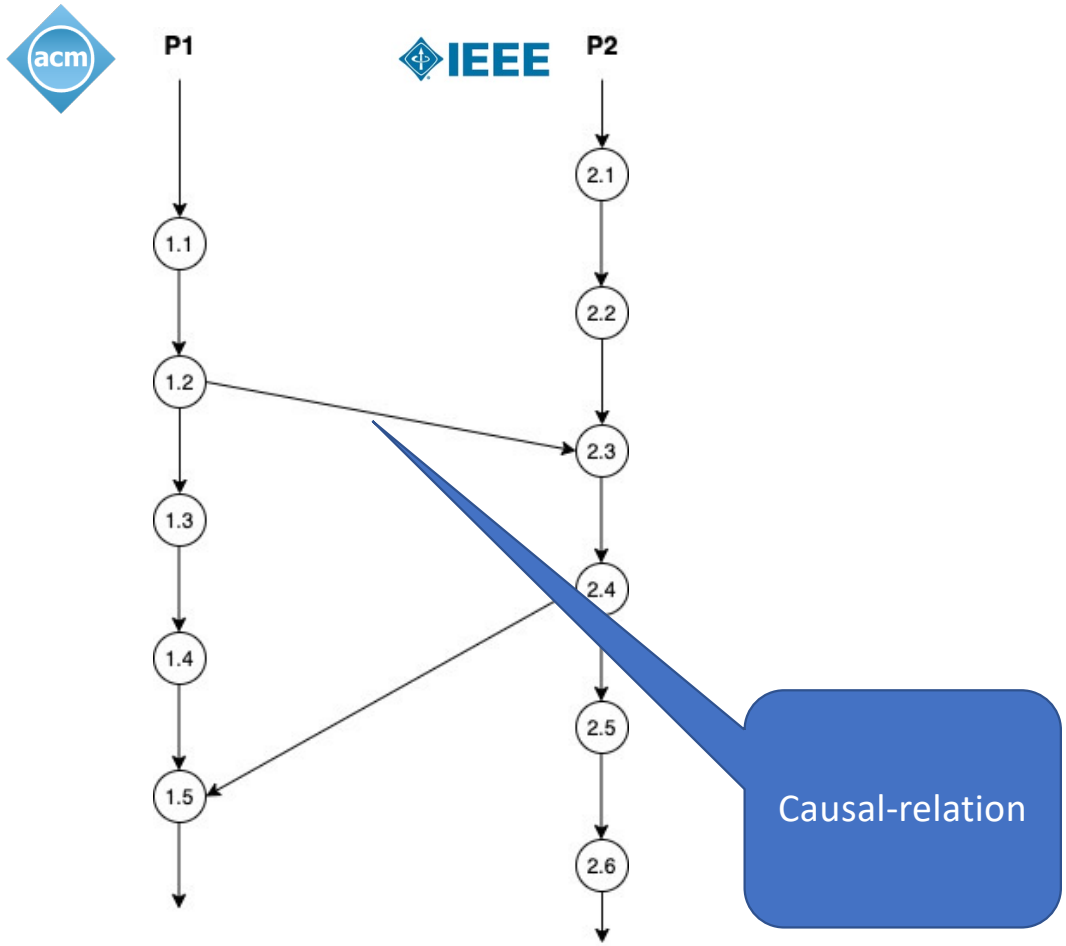

Gapless FIFO vs Gapless Causal Delivery

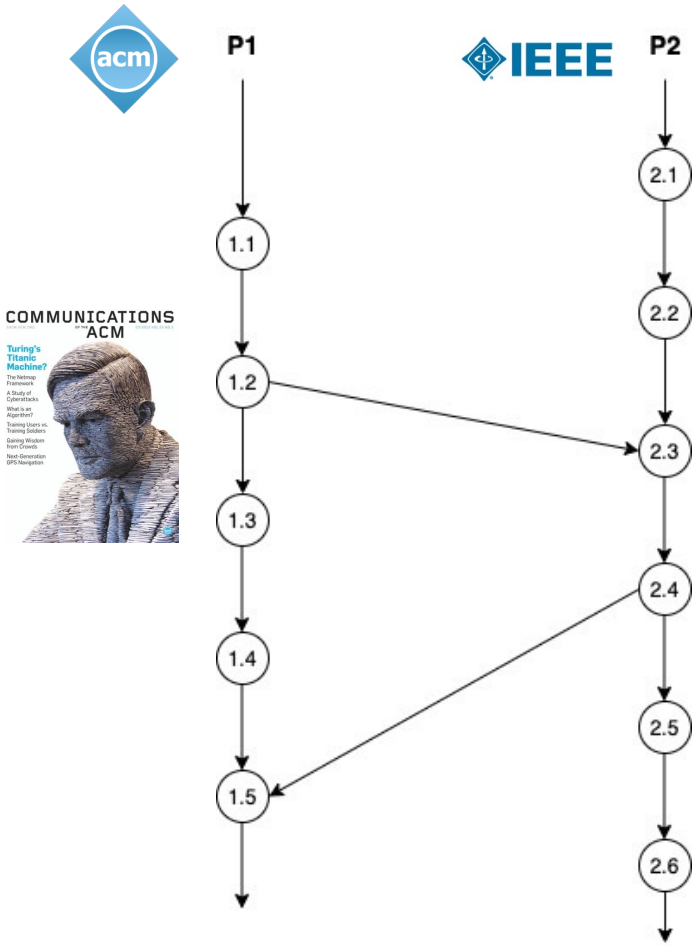

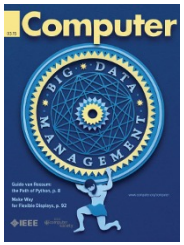

Paper in Computer cites a paper in Comm ACM

Gapless FIFO vs Gapless Causal Delivery

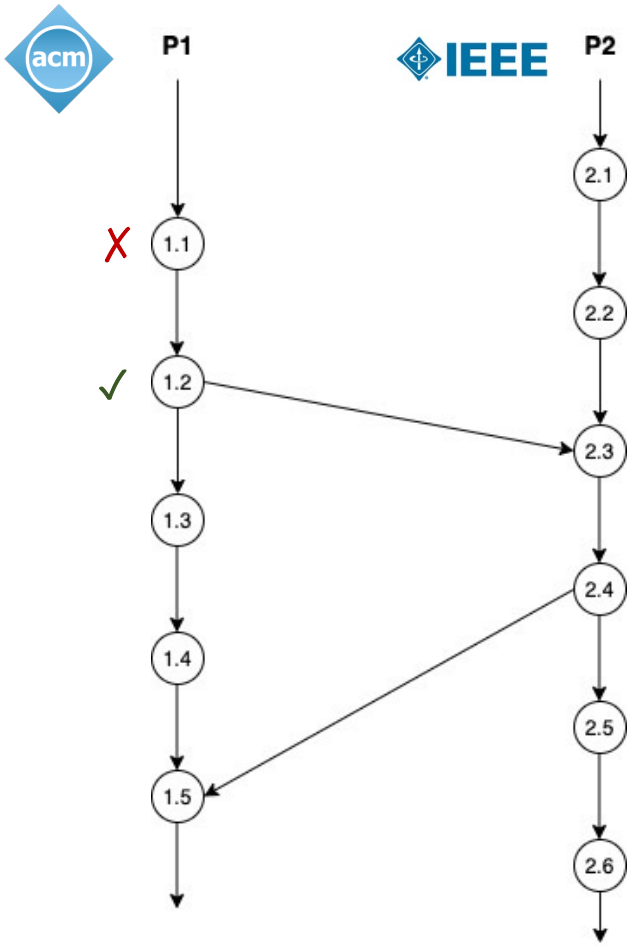

Gapless FIFO vs Gapless Causal Delivery

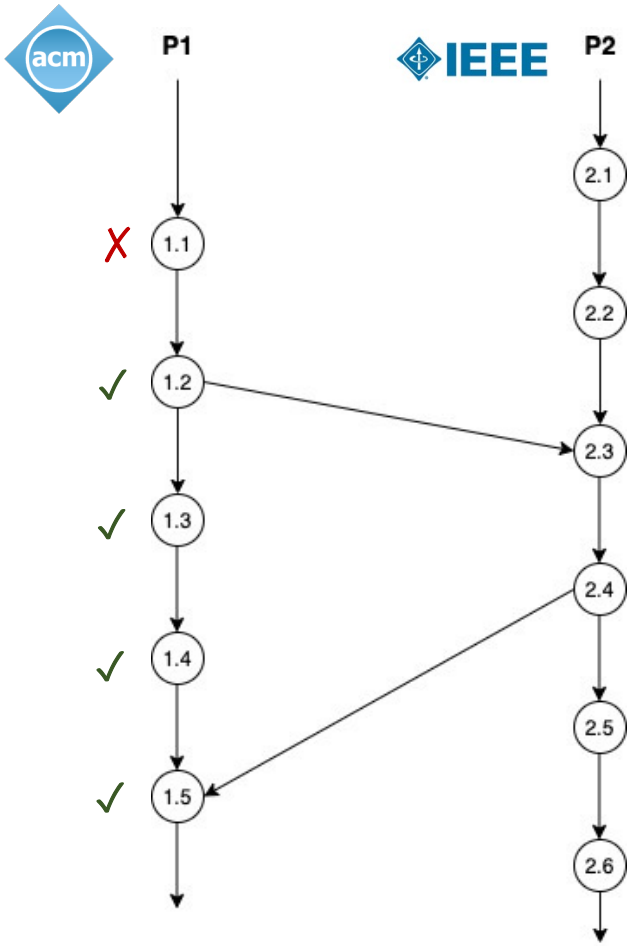

Gapless FIFO vs Gapless Causal Delivery

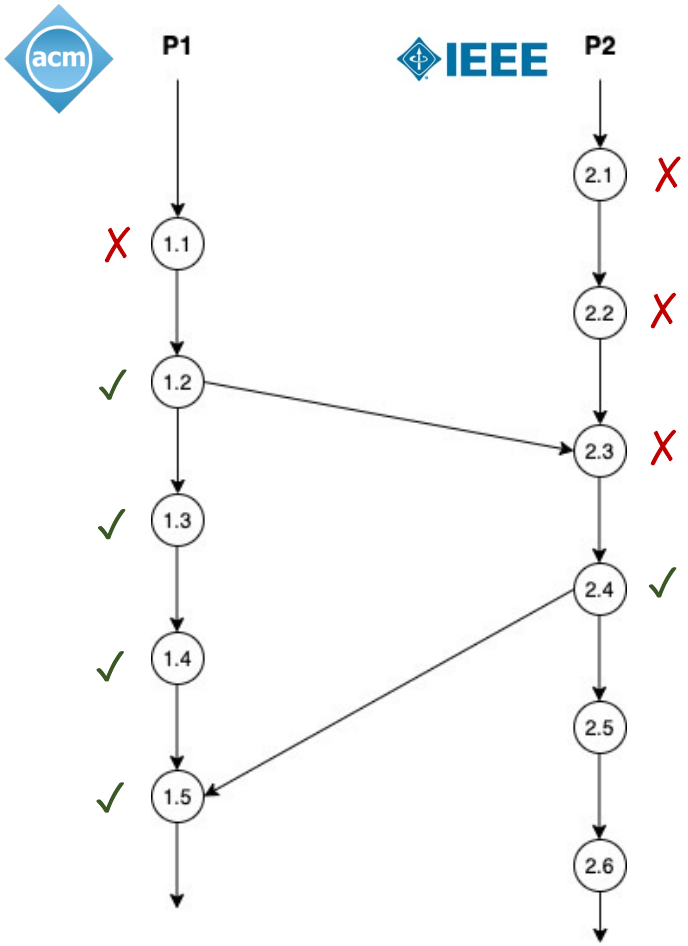

Gapless FIFO vs Gapless Causal Delivery

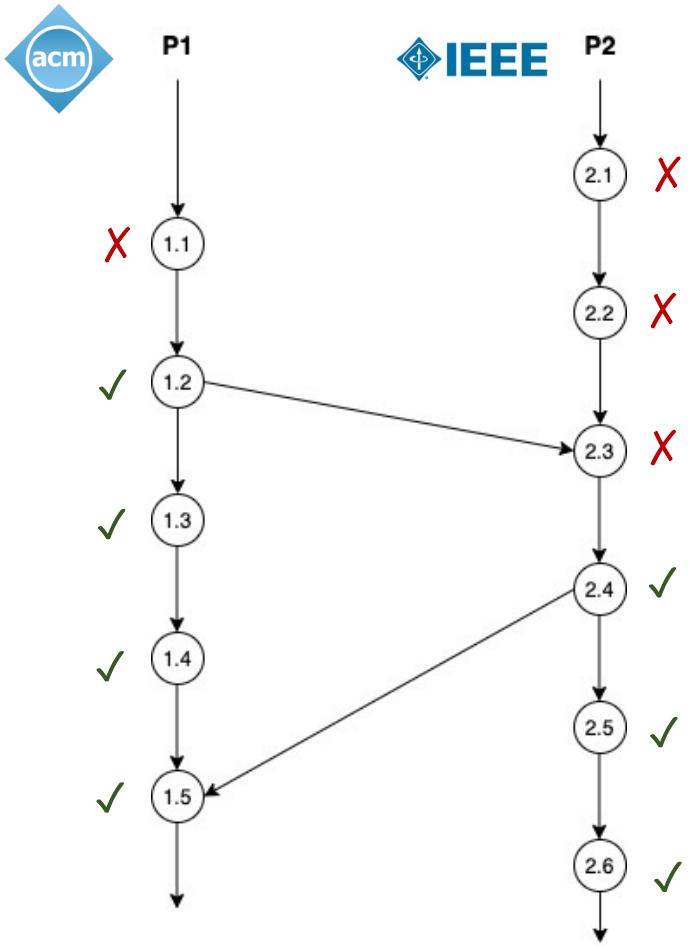

Gapless FIFO vs Gapless Causal Delivery

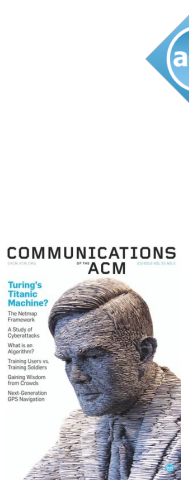

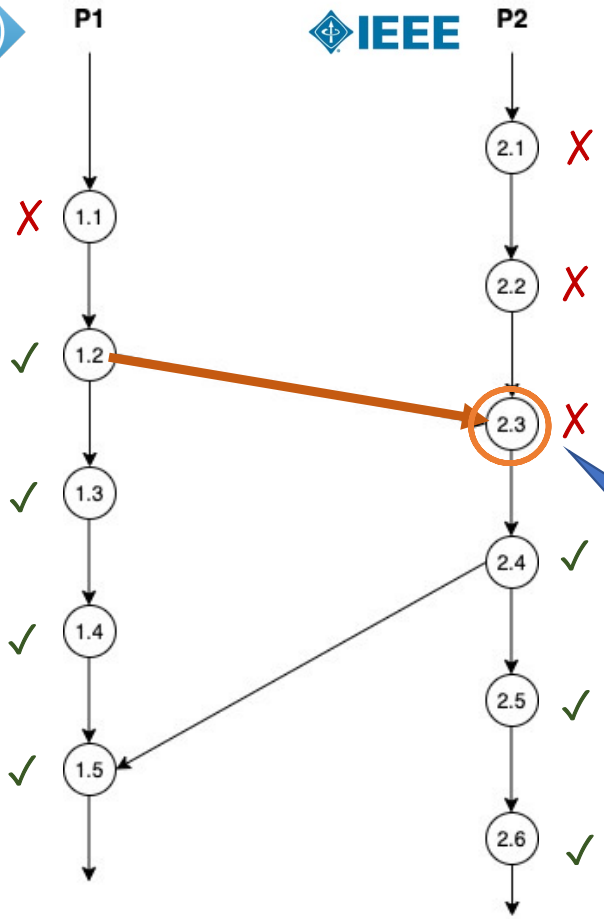

GAP in the causal history of the subscriber!

Reliability vs Subscription Latency

- FIFO is **stronger** than best-effort.
- FIFO brings **higher** subscription **latency** than best-effort.

Reliability vs Subscription Latency

- CAUSAL is **stronger** than FIFO.
- Does CAUSAL bring even higher latency subscription than FIFO?

Reliability vs Subscription Latency

- CAUSAL is **stronger** than FIFO.
- Does CAUSAL bring even higher latency subscription than FIFO?

- To answer this question, we need to understand how we can enforce GFD and GCD in practice.

Implementing GFD on a broker network

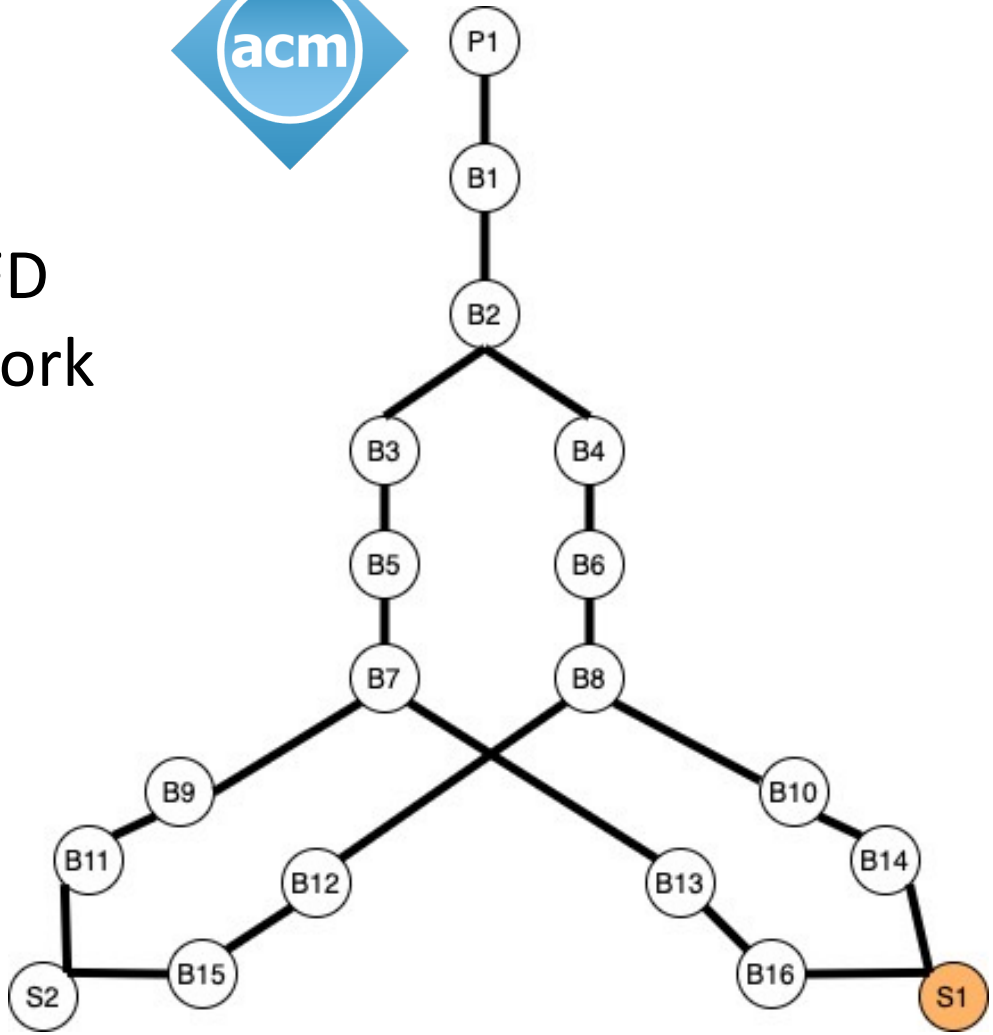

Alan

Implementing GFD on a broker network

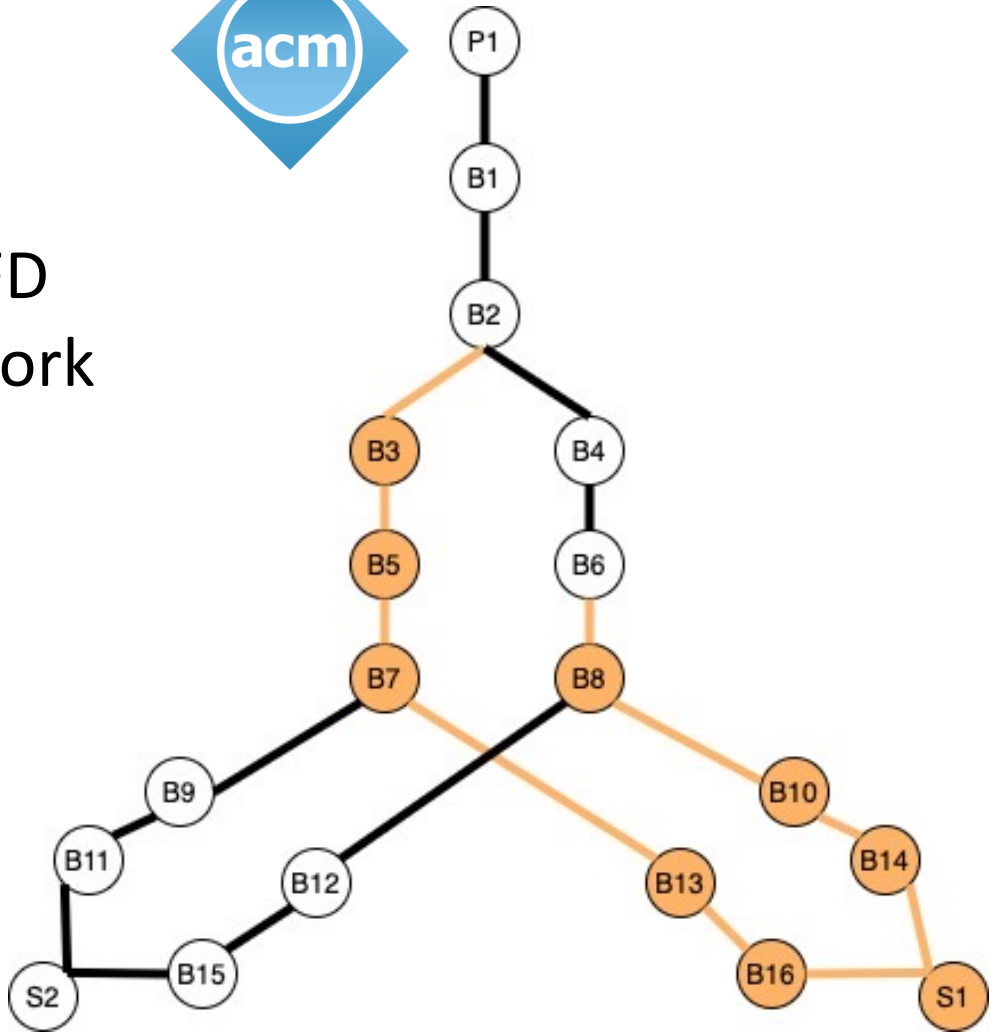

Subscrition

Alan

Implementing GFD on a broker network

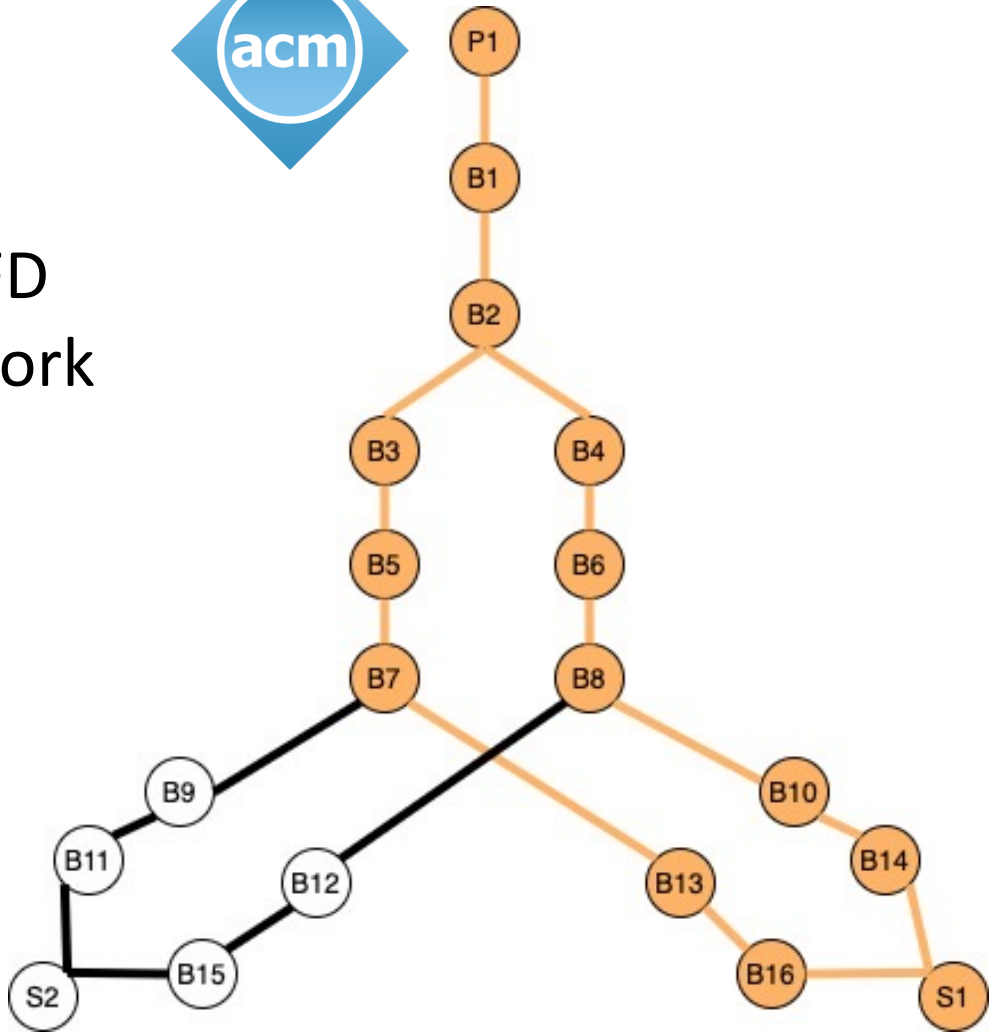

Subscrition

Alan

Implementing GFD on a broker network

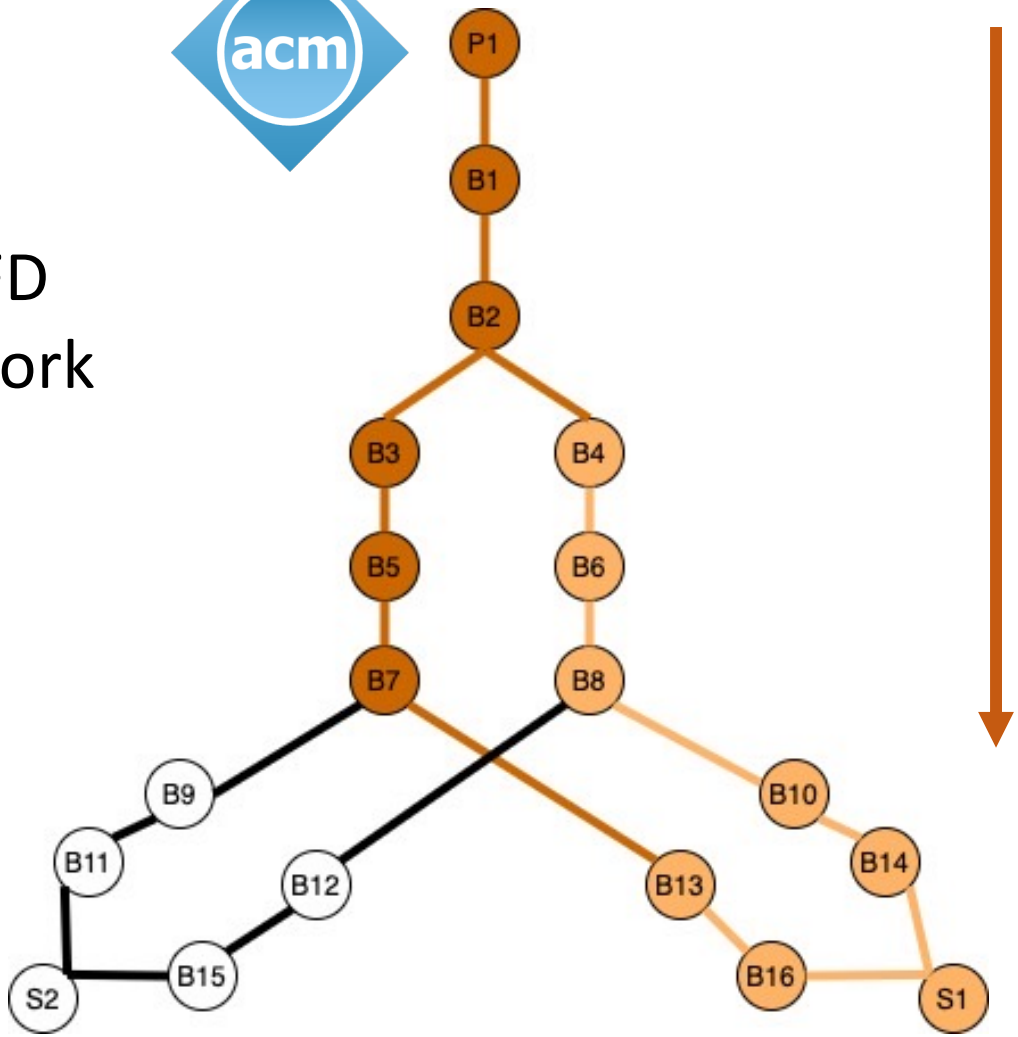

Subscription acknowledgement

Alan

Implementing GFD on a broker network

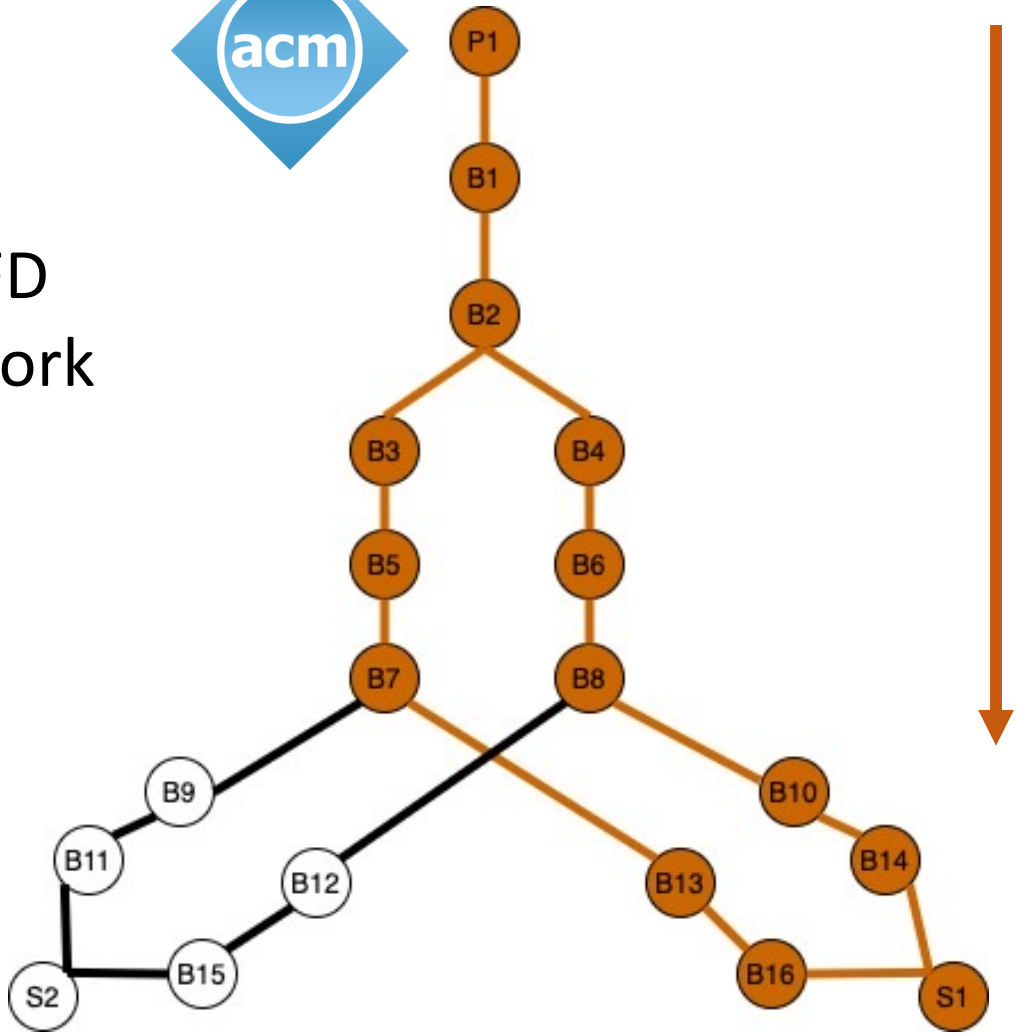

Subscription acknowledgement

Alan

Implementing GFD on a broker network

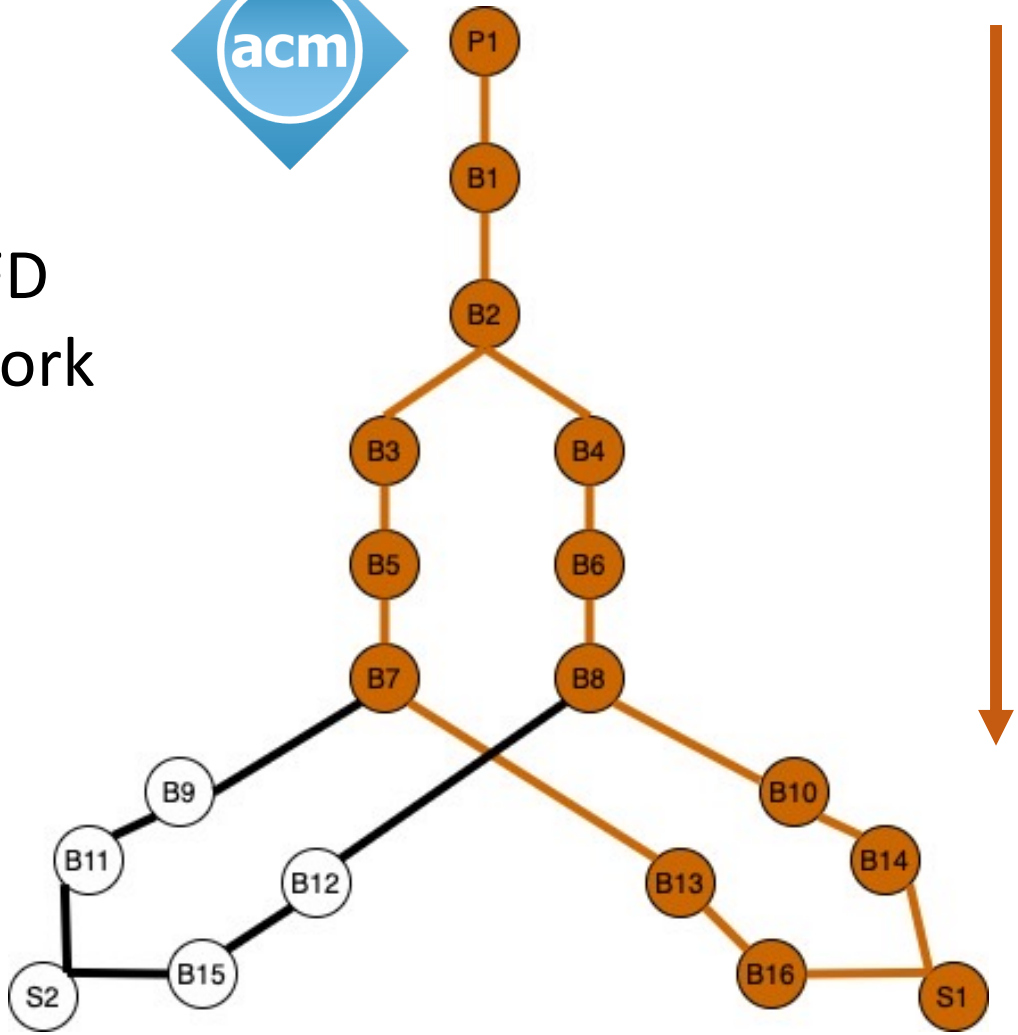

Subscription acknowledgement

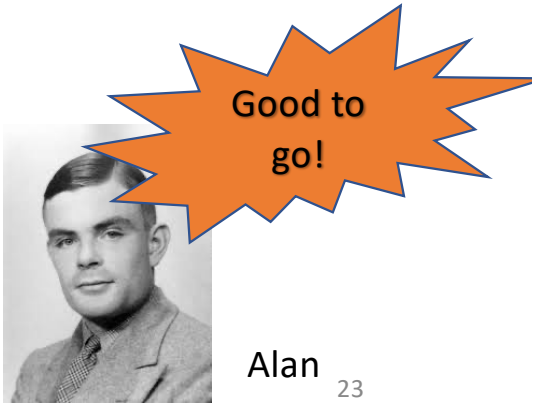

Alan

Enforcing GFD and GCD

- Necessary and sufficient conditions to enforce GFD and GCD.
  - A subscription **is stable on path** if it known by all nodes on that path.
  - **1st Result: For GFD, it is sufficient that ONE path is be stable.**

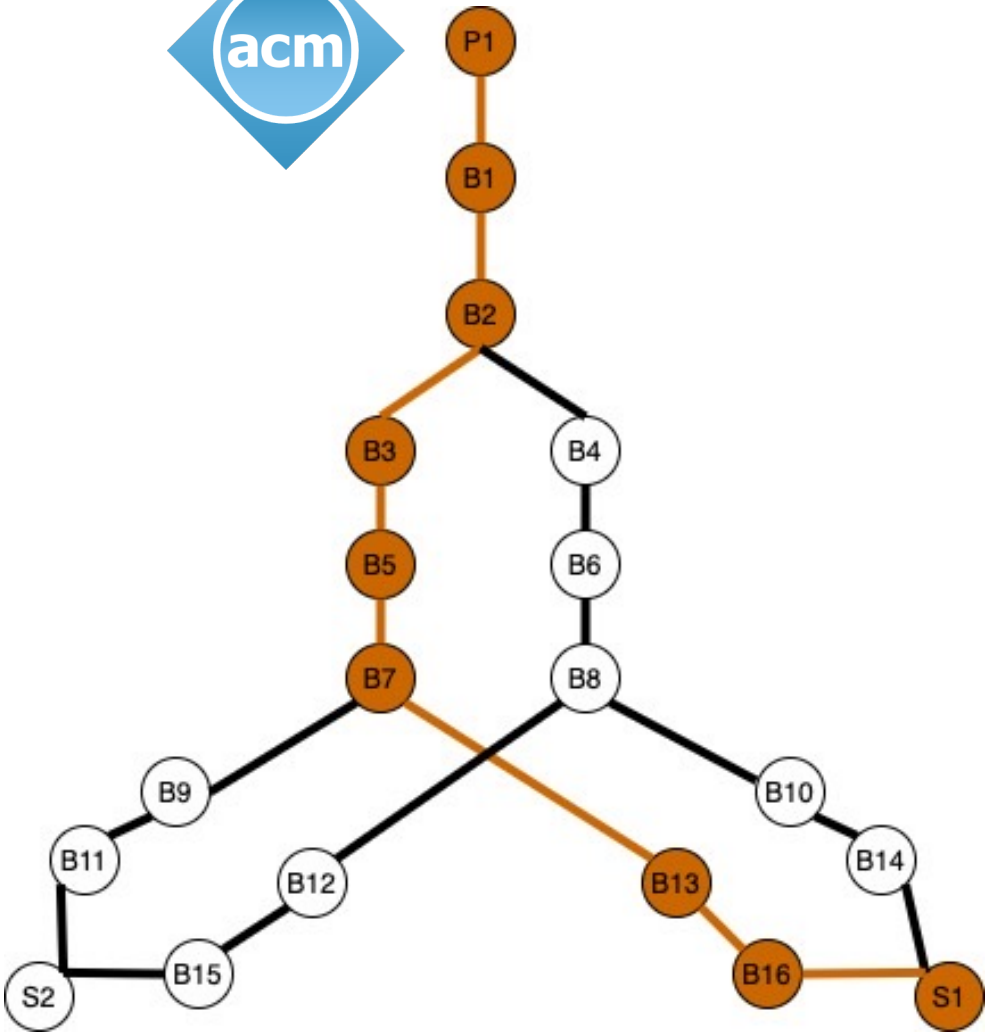

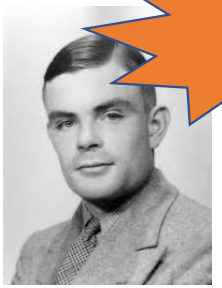

Good to go!

Alan

Enforcing GFD and GCD

- Necessary and sufficient conditions to enforce GFD and GCD
  - A subscription is stable on path if it known by all nodes on that path.
  - **1st Result: For GFD, it is sufficient that one path is be stable.**
  - **2nd Result: For GCD, it is also sufficient that one path is be stable!**

Enforcing GFD and GCD

- Necessary and sufficient conditions to enforce GFD and GCD.
  - A subscription is stable on path if it known by all nodes on that path.
  - **1st Result: For GFD, it is suffic… ath is …e stable.**
  - **2nd Result: Fo… …h is be stable!**

GCD is **NOT** inherently more expensive than GFD!

Leveraging Coverage

- If another subscription is already in place, can we used it to reduce the subscription latency?

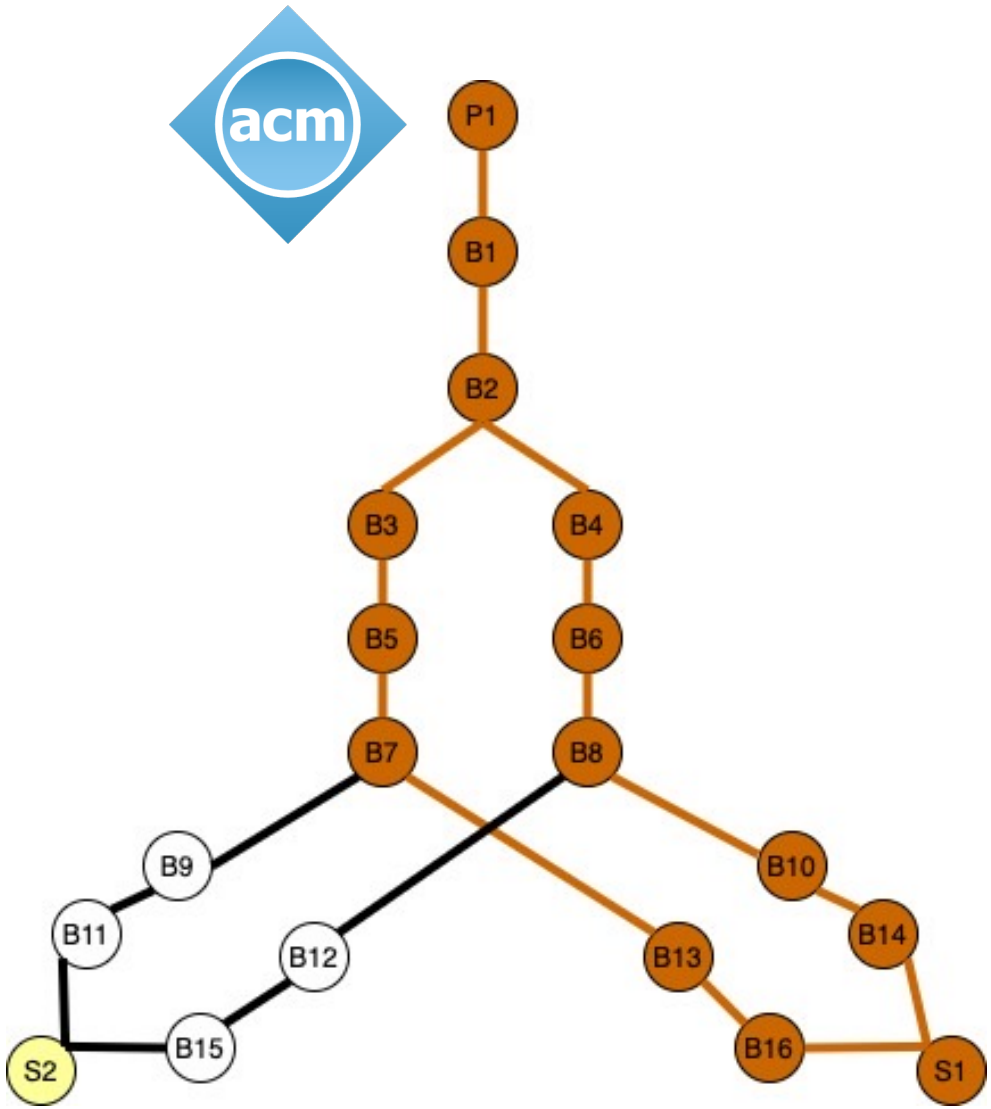

Luís

Alan

roundtrip to publisher required?

Luís

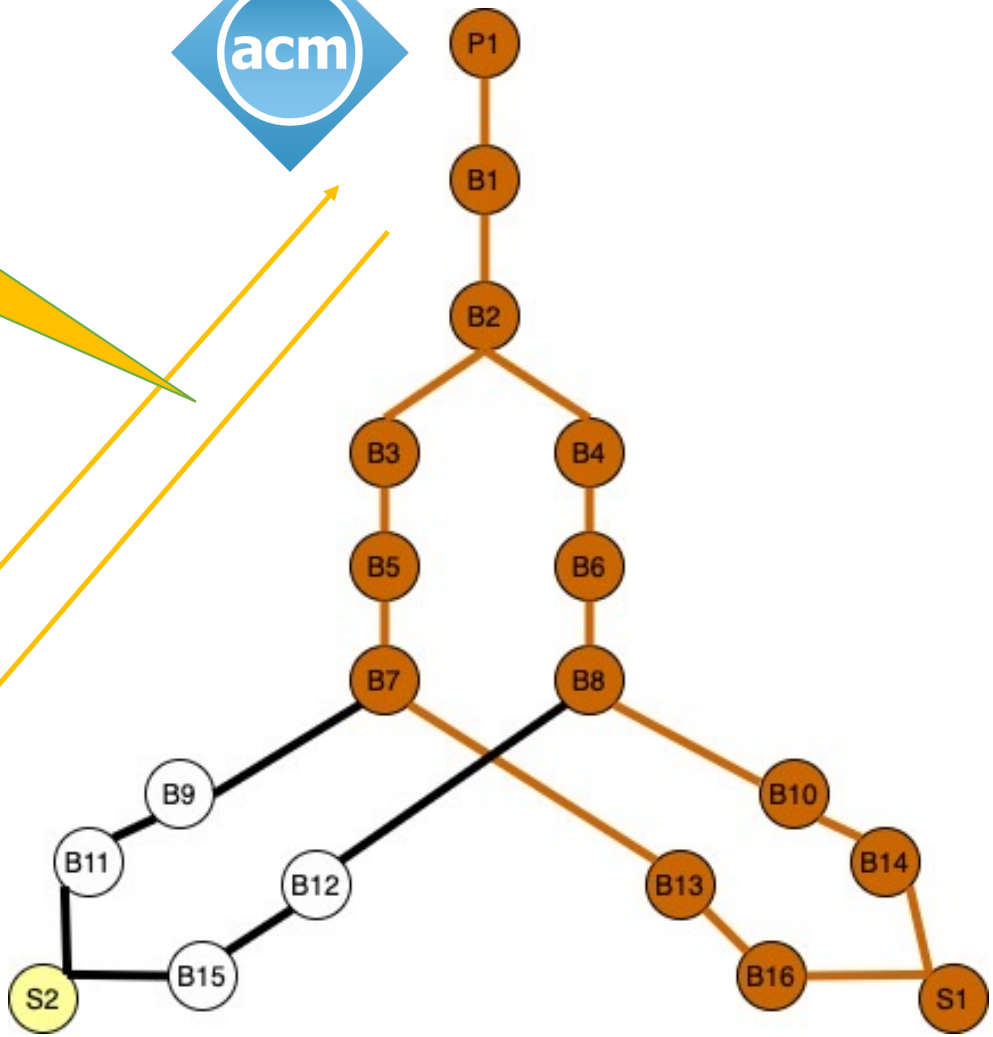

Alan

Luís

Alan

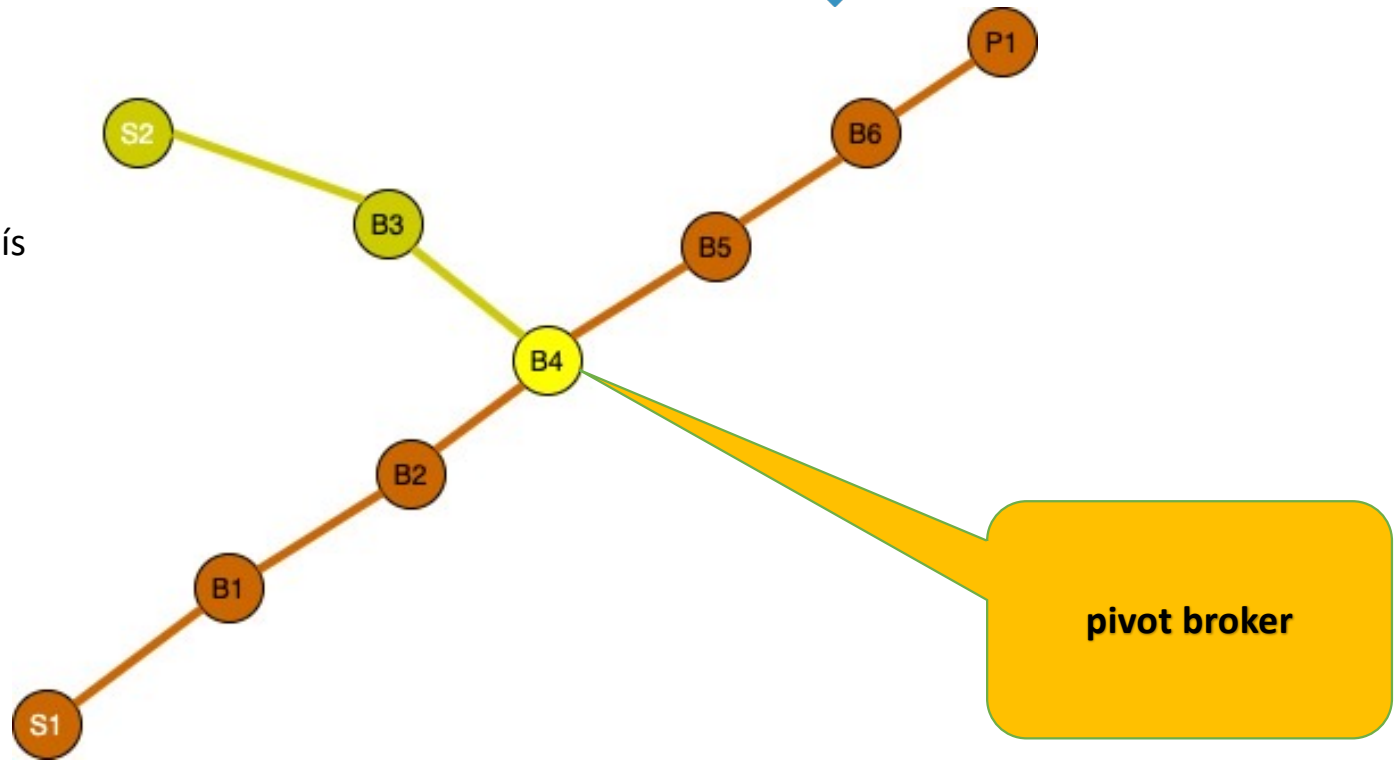

Luís

Alan

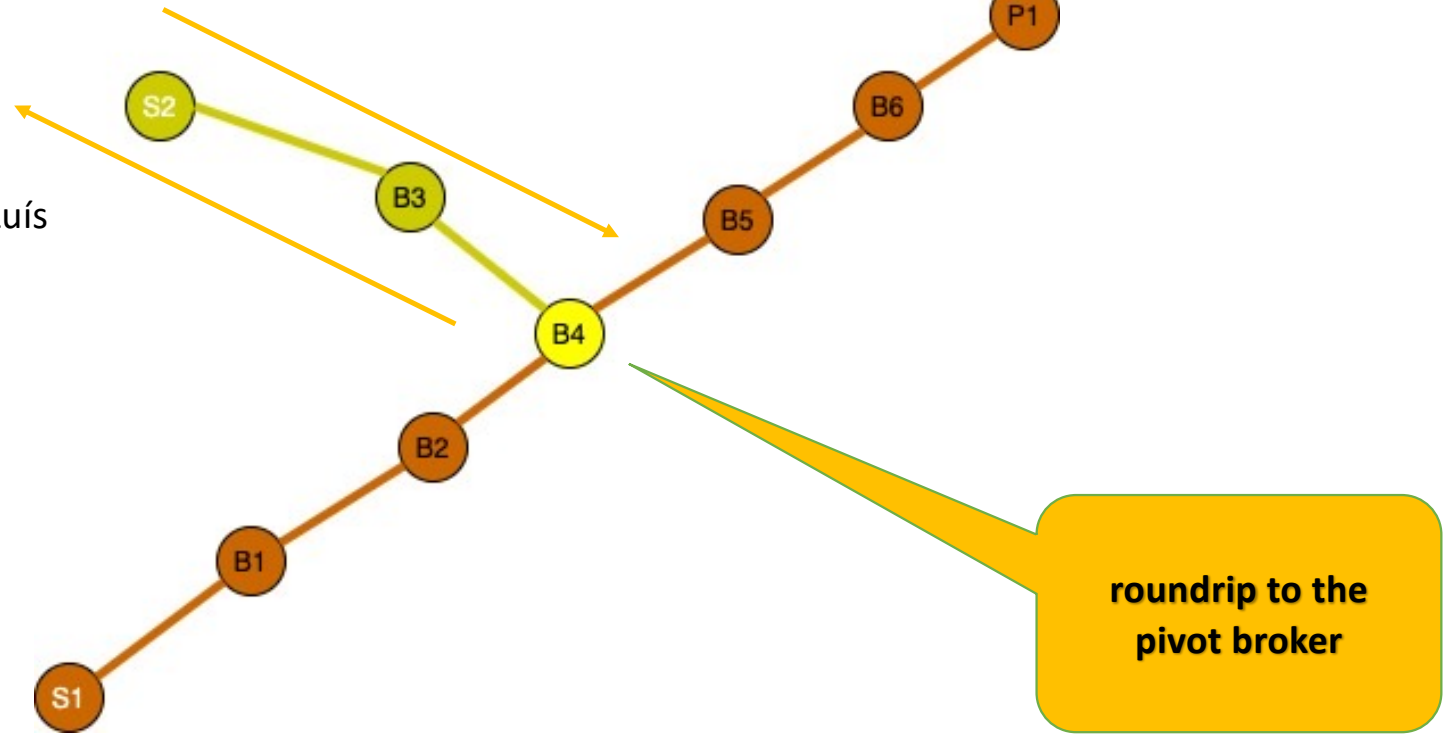

Luís

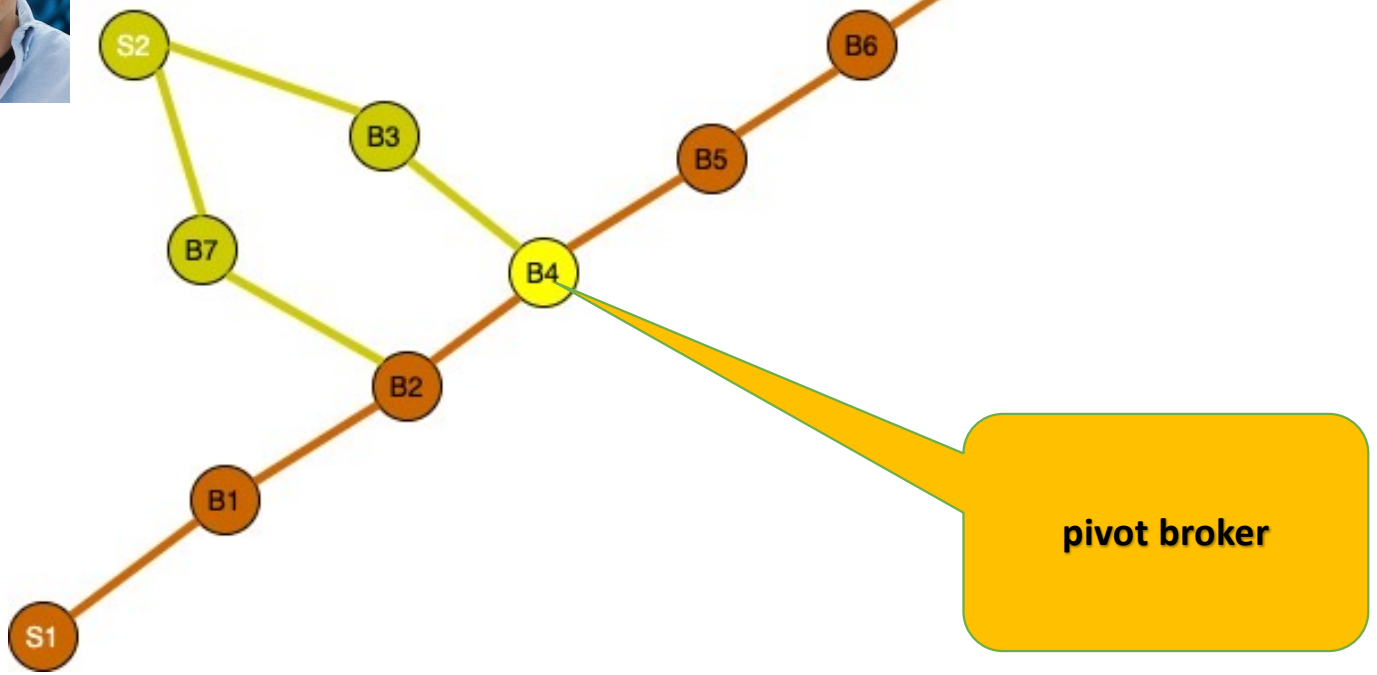

Alan

Luís

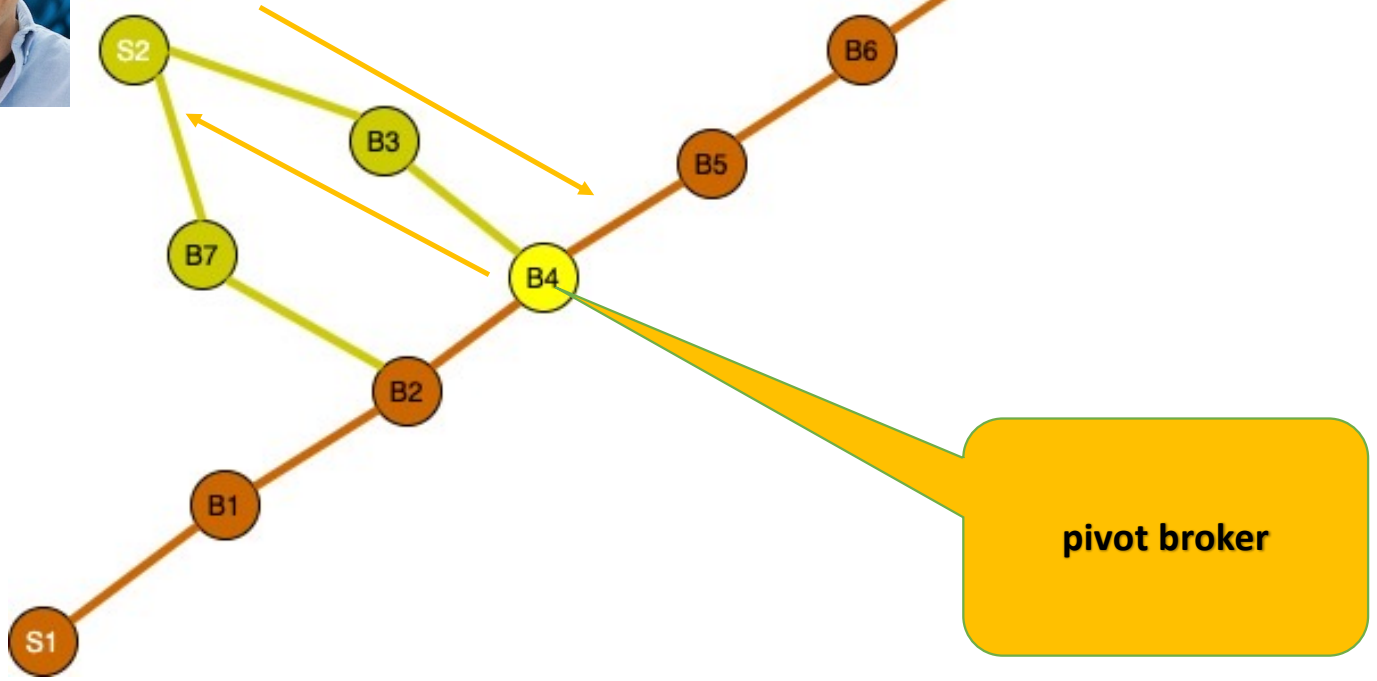

Alan

roundtrip to publisher required?

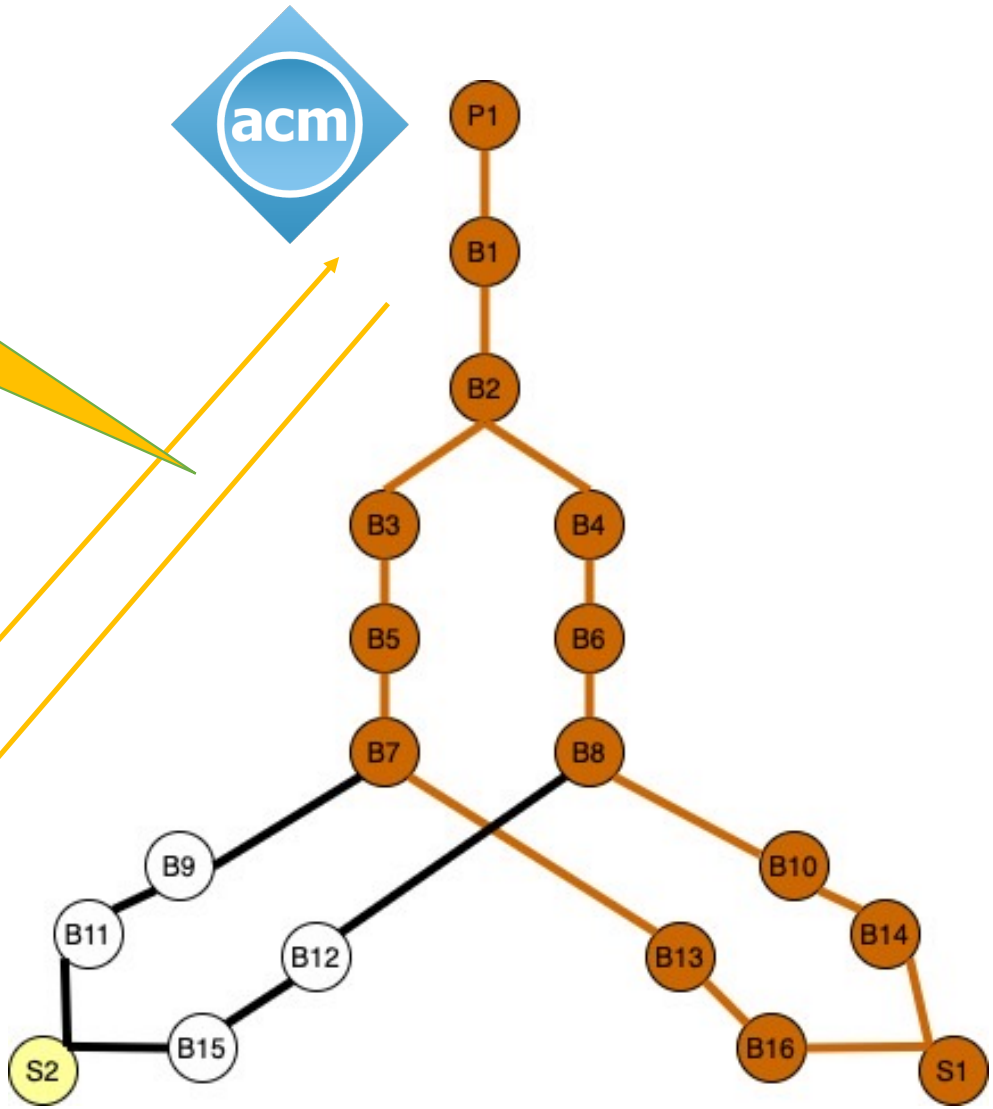

Luís

Alan

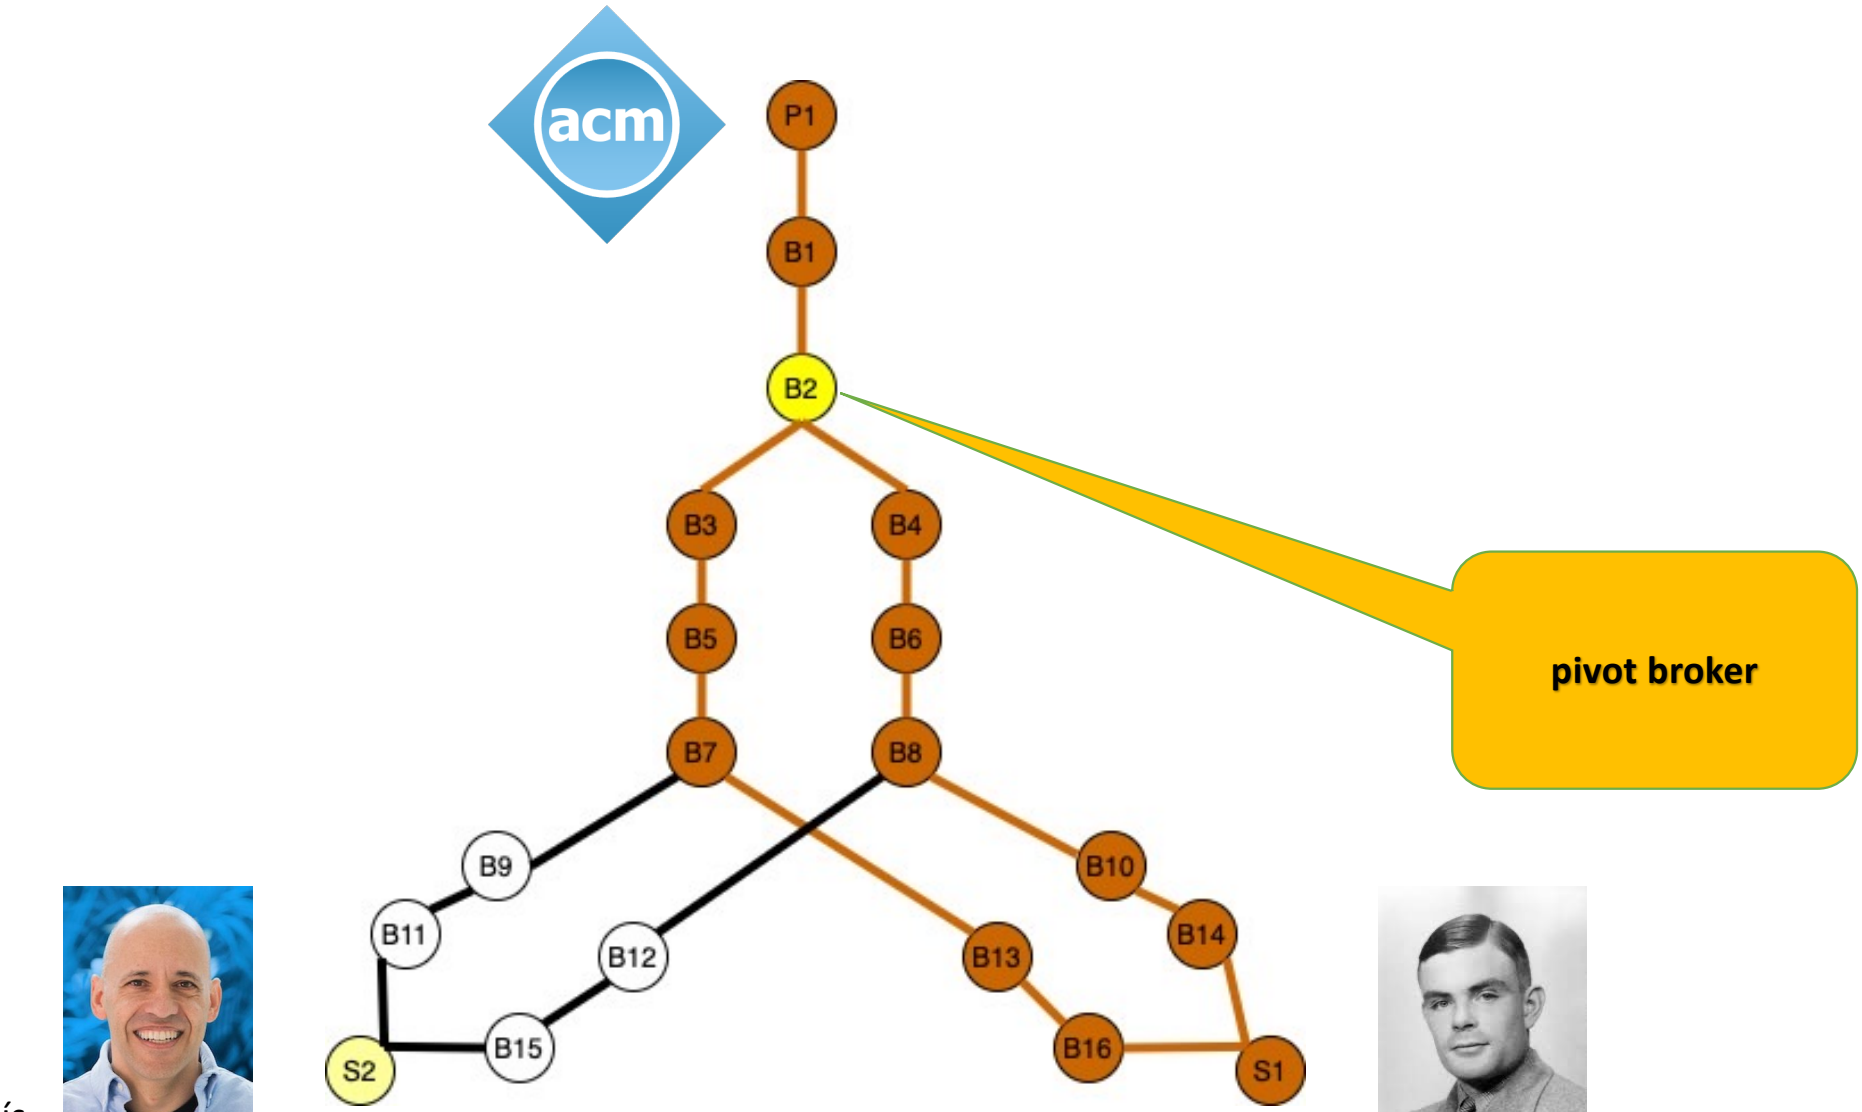
P1
B1
B2
B3
B4
B5
B6
B7
B8
B9
B10
B11
B12
B13
B14
B15
B16
S2
S1
pivot broker

Luís

Alan

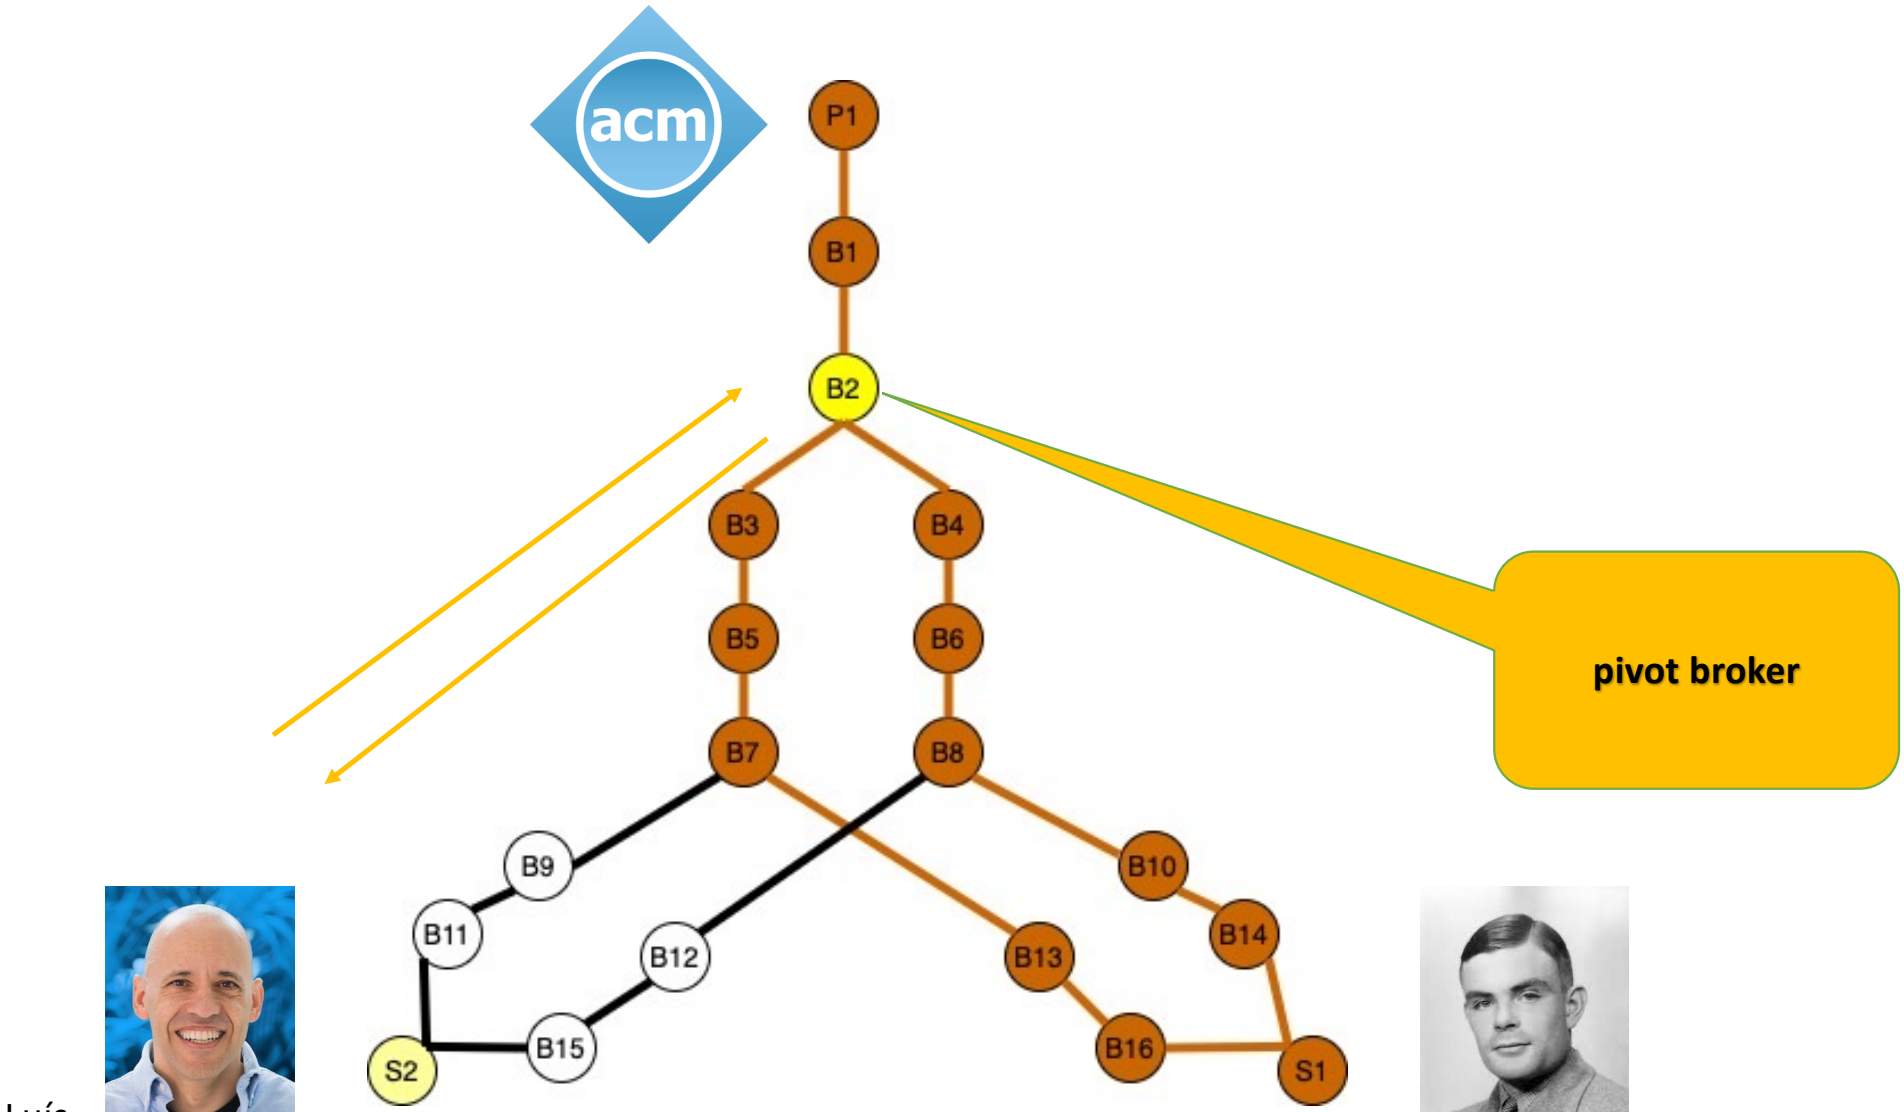
P1
B1
B2
B3
B4
B5
B6
B7
B8
B9
B10
B11
B12
B13
B14
B15
B16
S2
S1
pivot broker

Luís

Alan

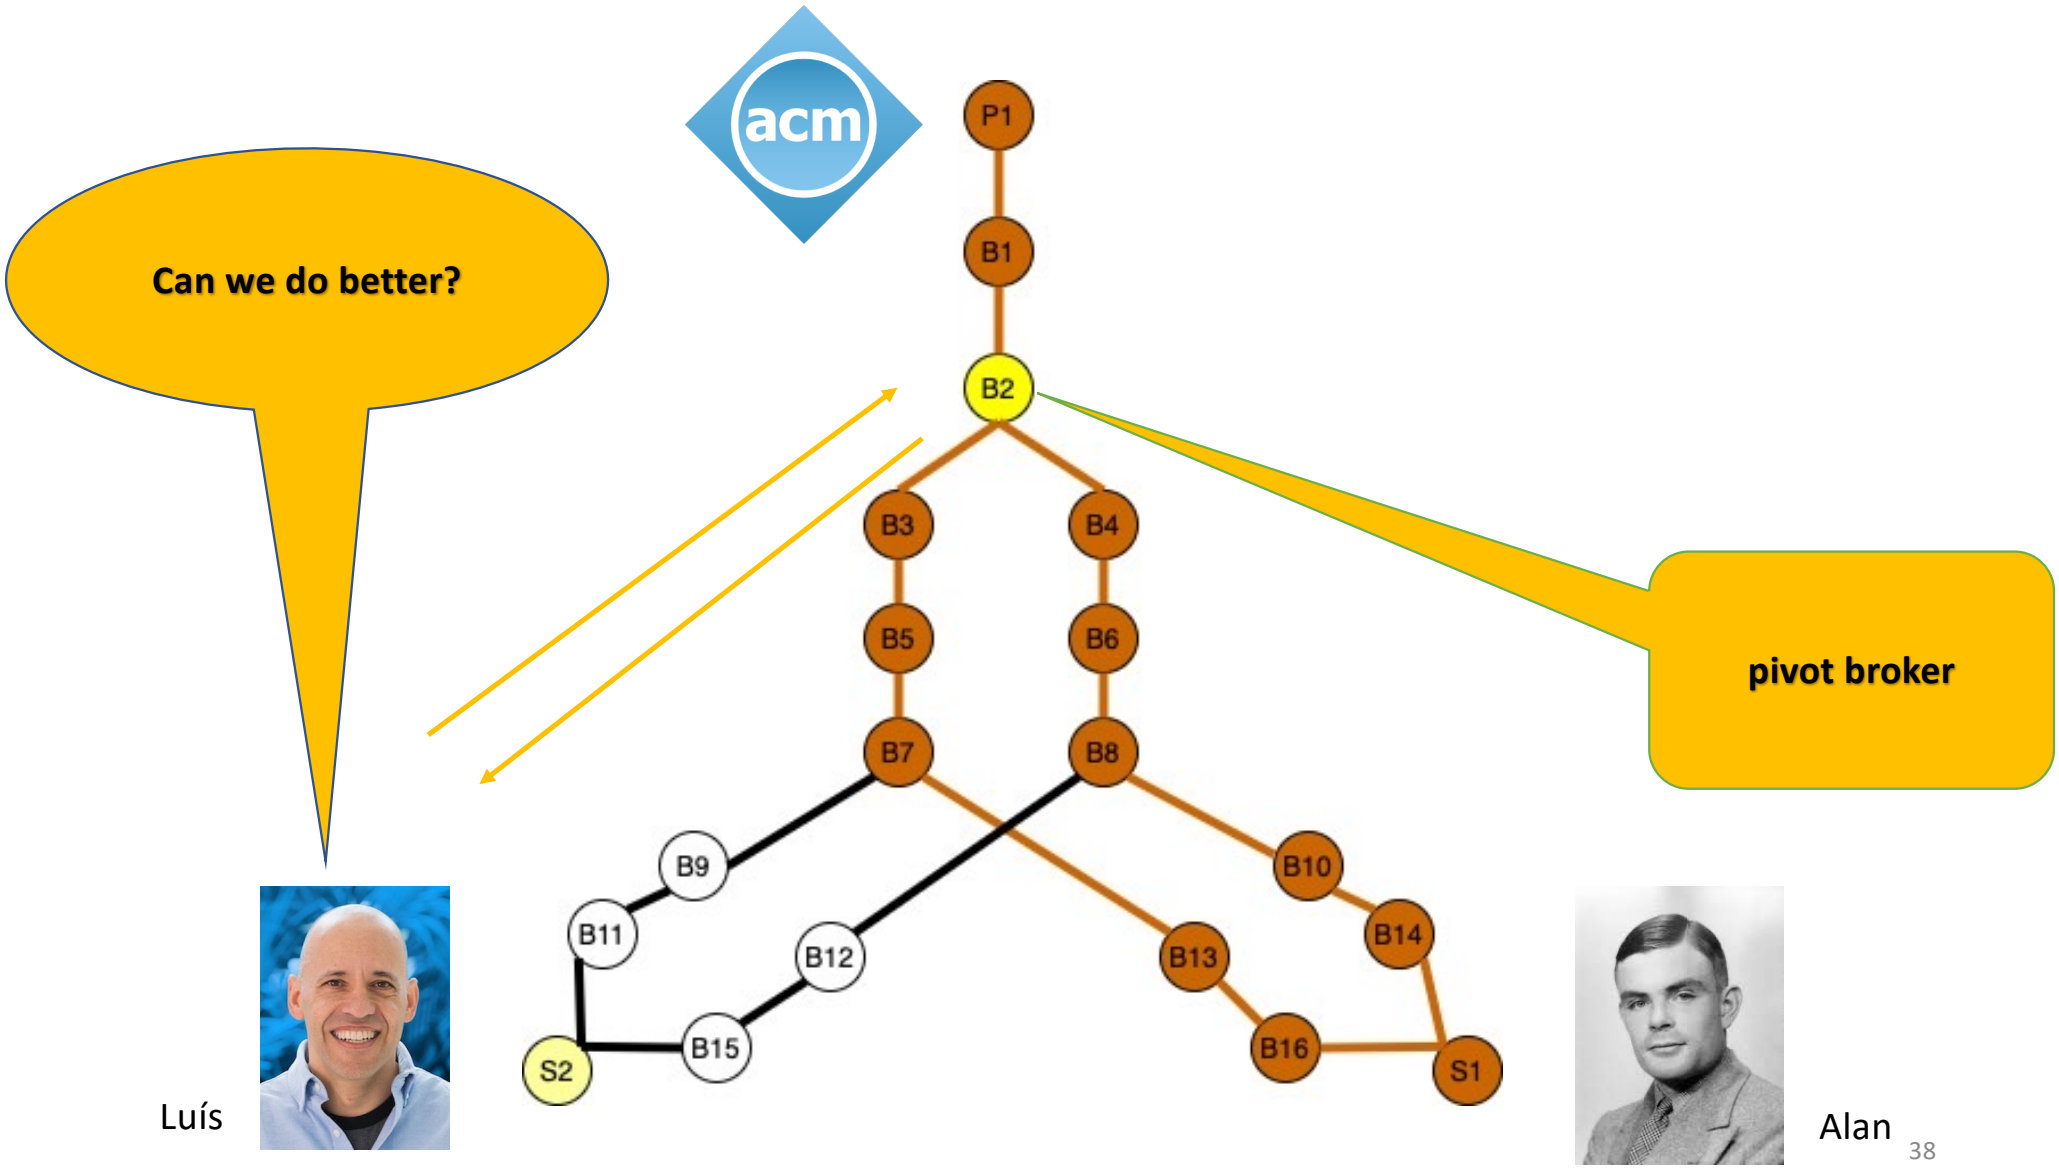
Can we do better?
P1
B1
B2
B3
B4
B5
B6
B7
B8
B9
B10
B11
B12
B13
B14
B15
B16
S2
S1
pivot broker
Luís
Alan

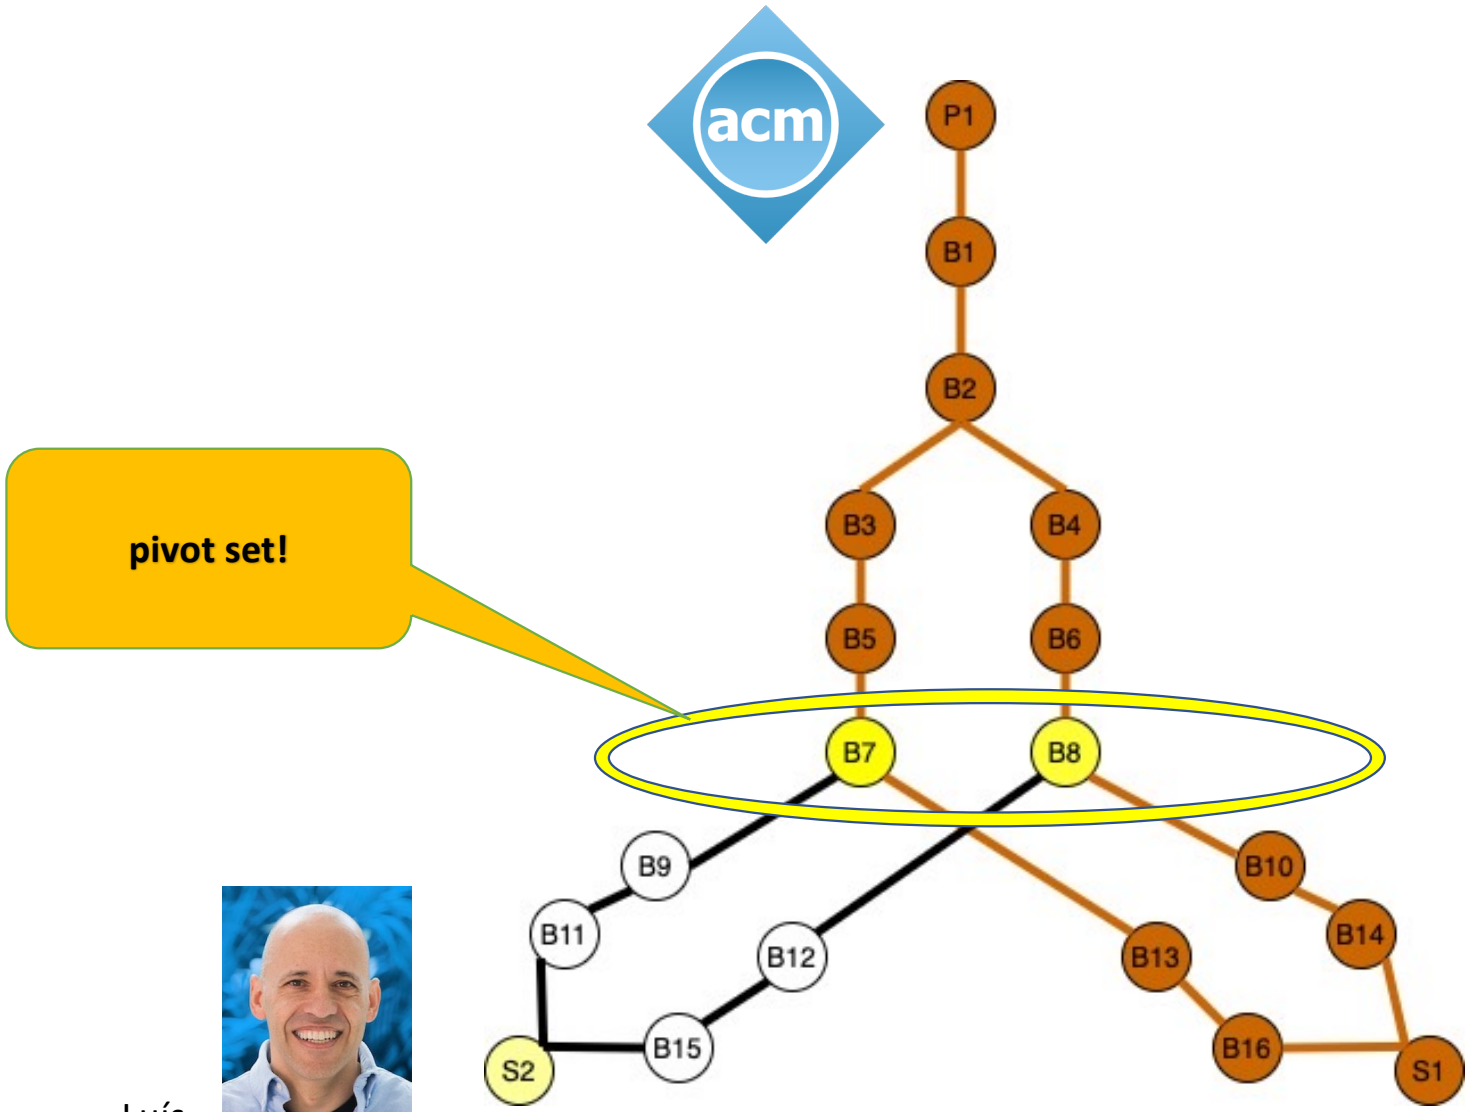
pivot set!
P1
B1
B2
B3
B4
B5
B6
B7
B8
B9
B10
B11
B12
B13
B14
B15
B16
S2
S1

Luís

Alan

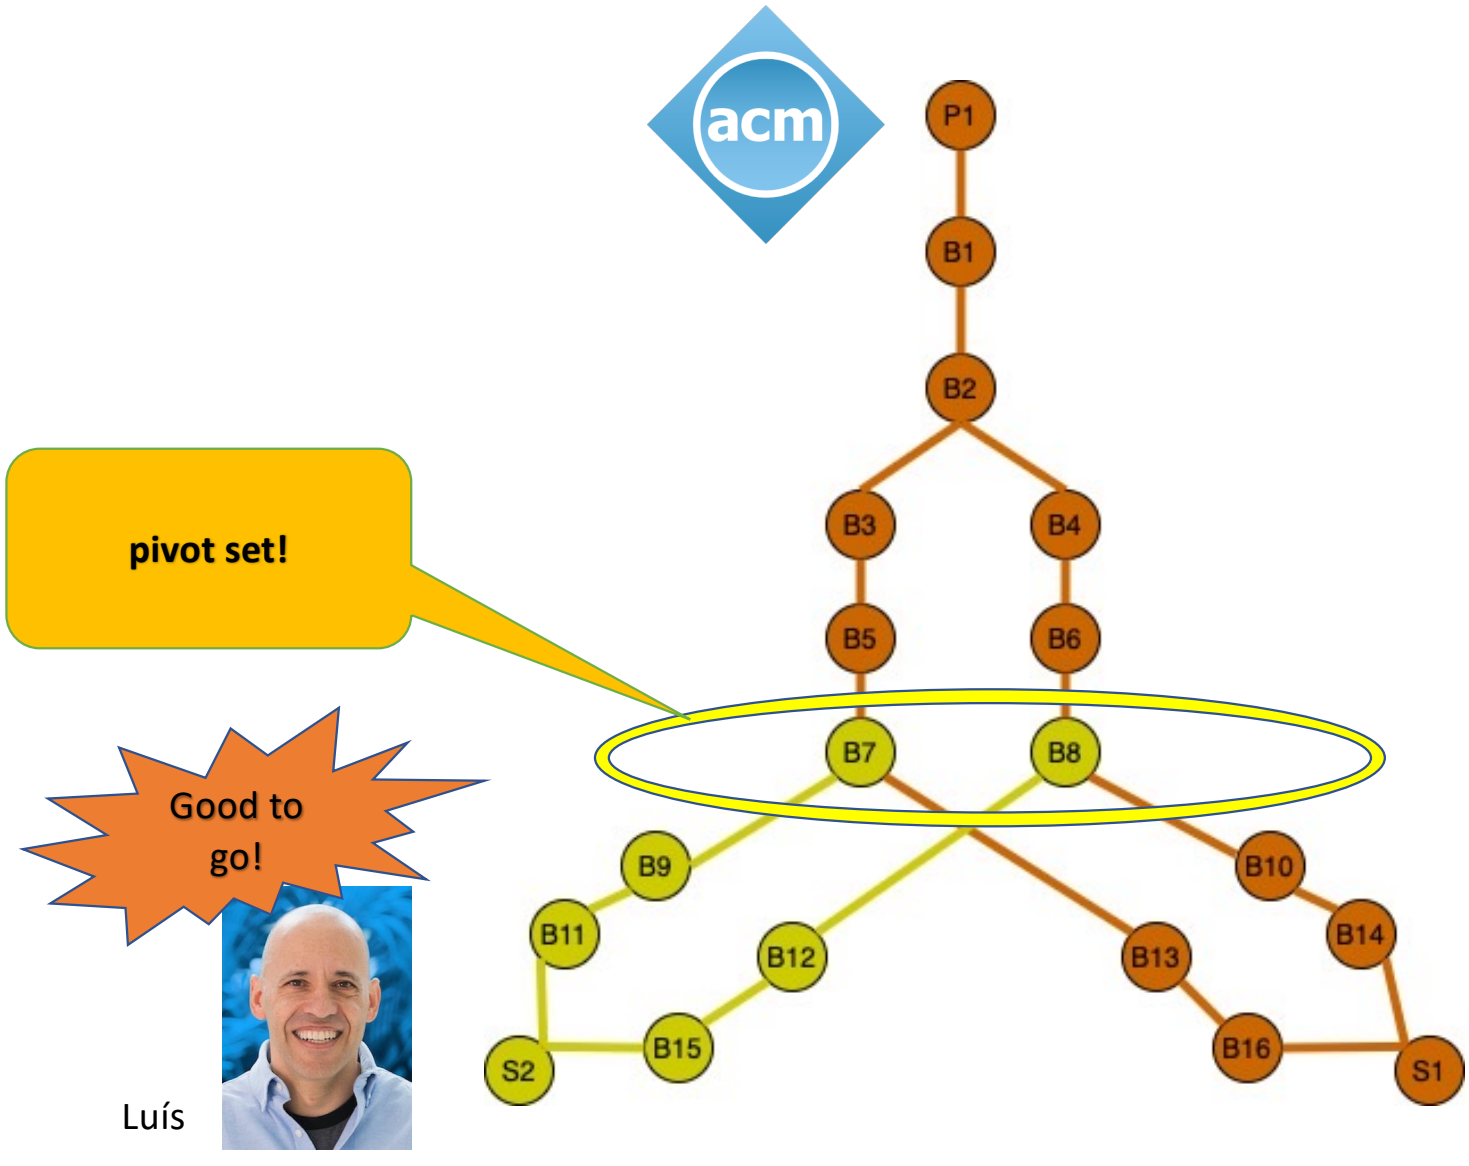
acm
P1
B1
B2
B3
B4
B5
B6
B7
B8
B9
B10
B11
B12
B13
B14
B15
B16
S1
S2
pivot set!
Good to go!

Luís

Alan

Publisher, pivot broker, and pivot set

- **Previous works**:
  - Subscription stable on one or more paths to the publisher
- **Our work**: subscription stable on
  - One path to the pivot broker
  - **OR**
  - One path to each member of the pivot set

Relevance of the pivot set

- For fault-tolerance you want to have multiple disjoint paths to the publisher.
- No pivot-broker

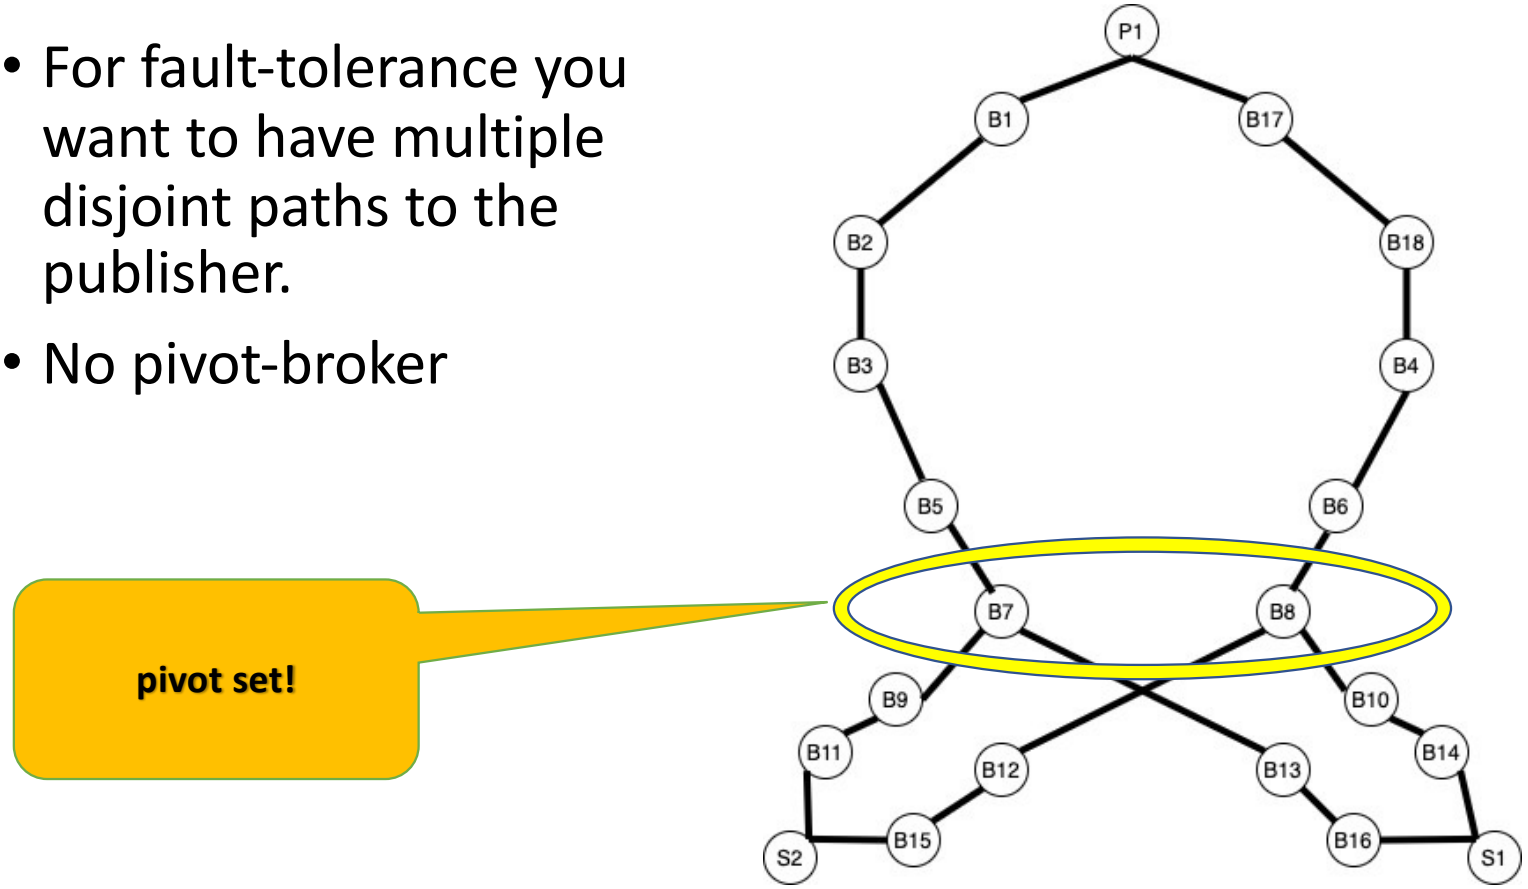

Finding the pivot set

- In general graphs it may be hard to find the pivot set.

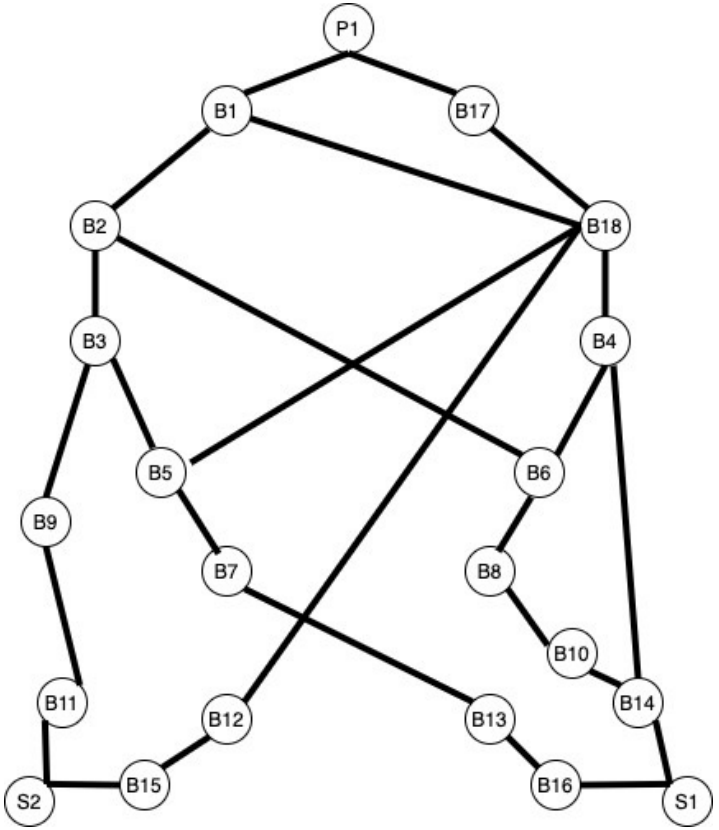

LoCaPS

- LoCaPS: publish-subscribe implementation that leverages our findings.

- It is possible to build the broker overlay that is fault-tolerant and where it is easy to find the pivot-set.

LoCaPS

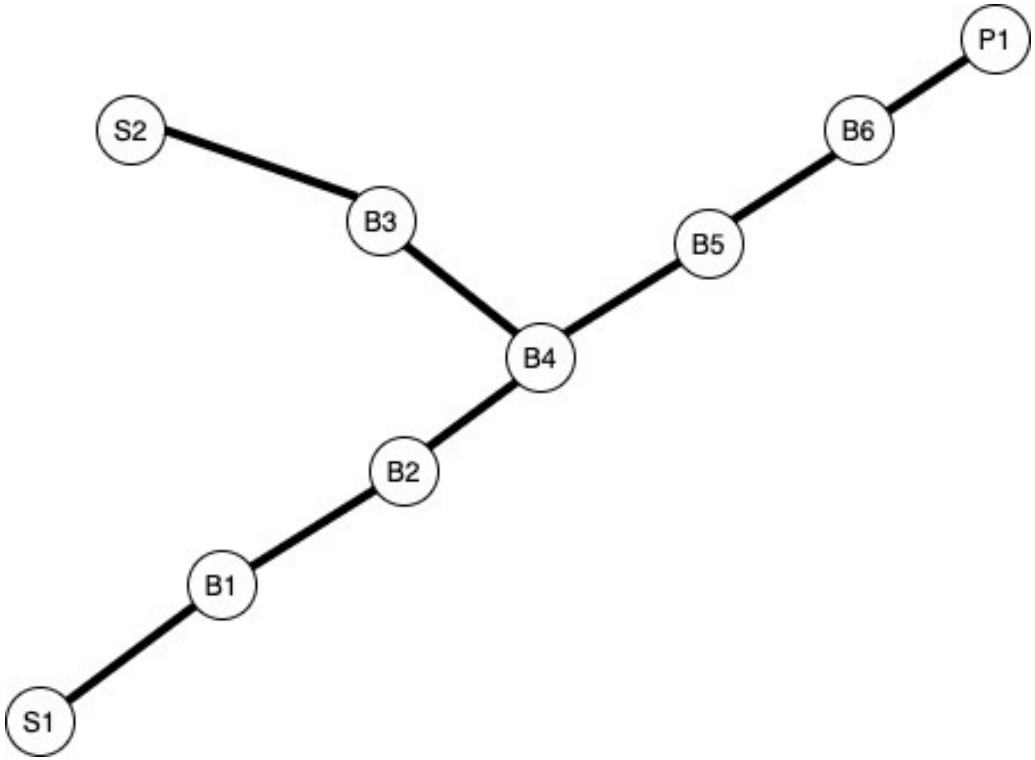

LoCaPS

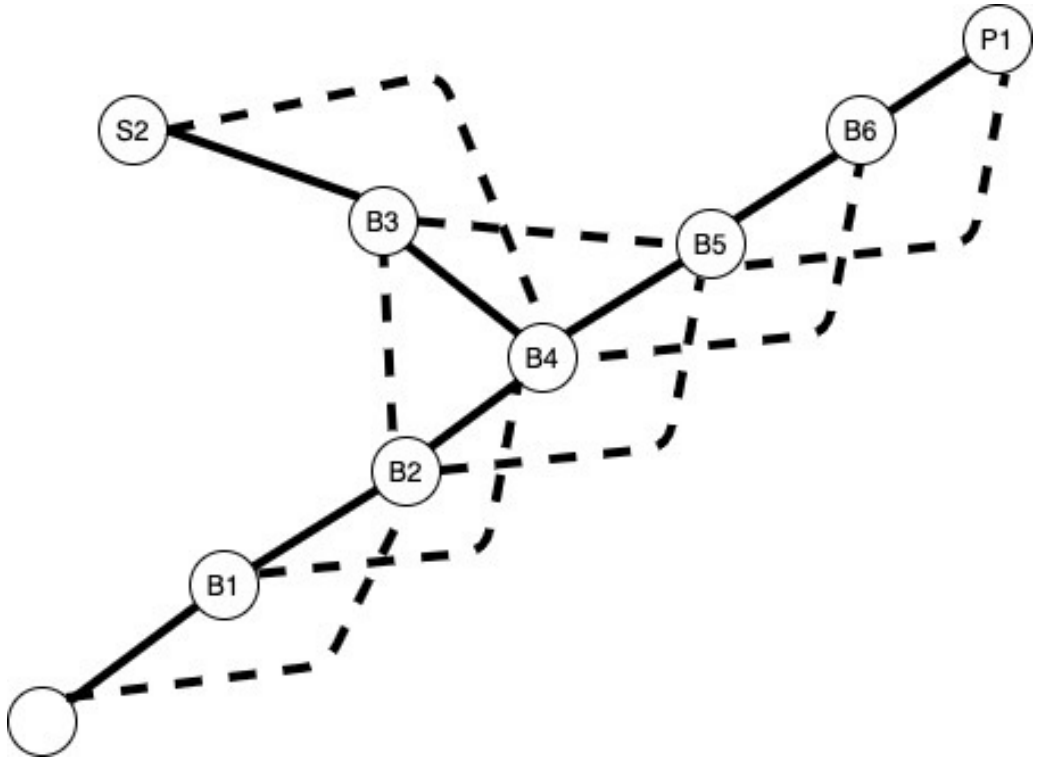

LoCaPS

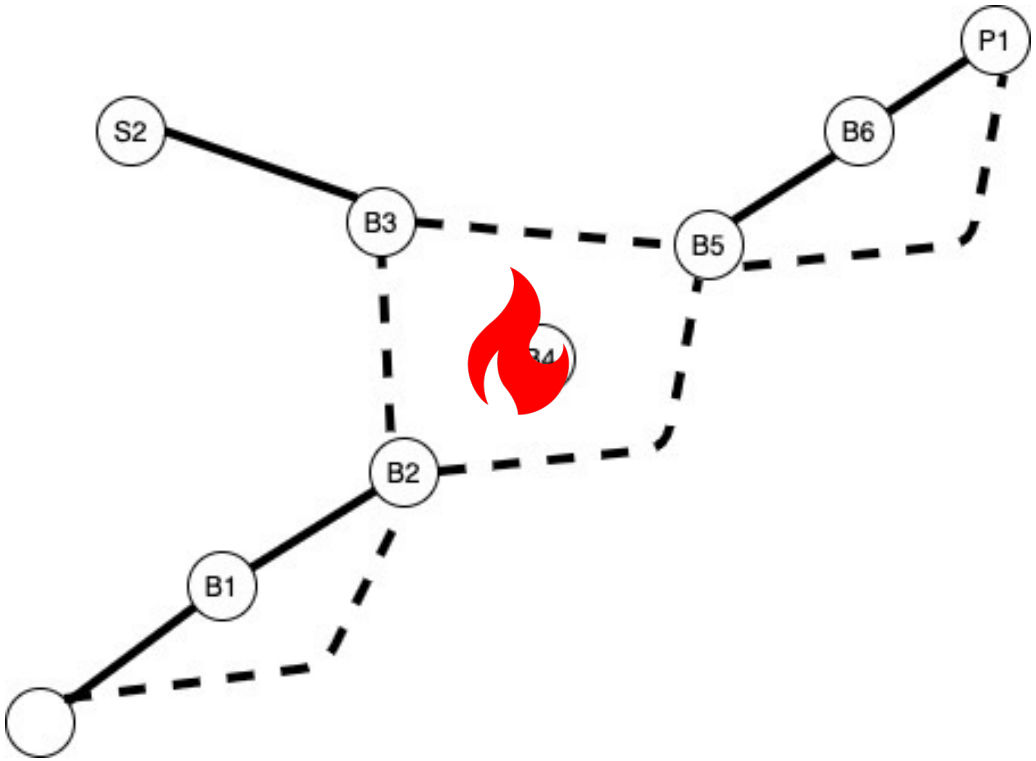

LoCaPS

Luís

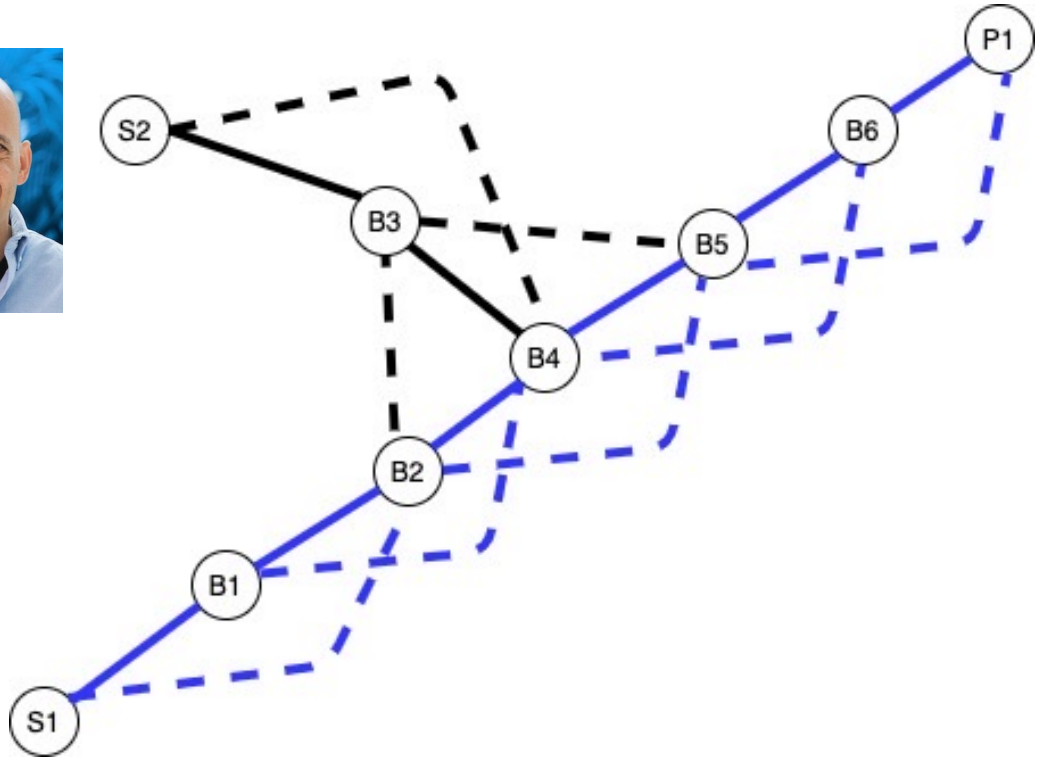

Alan

LoCaPS

Luís

Alan

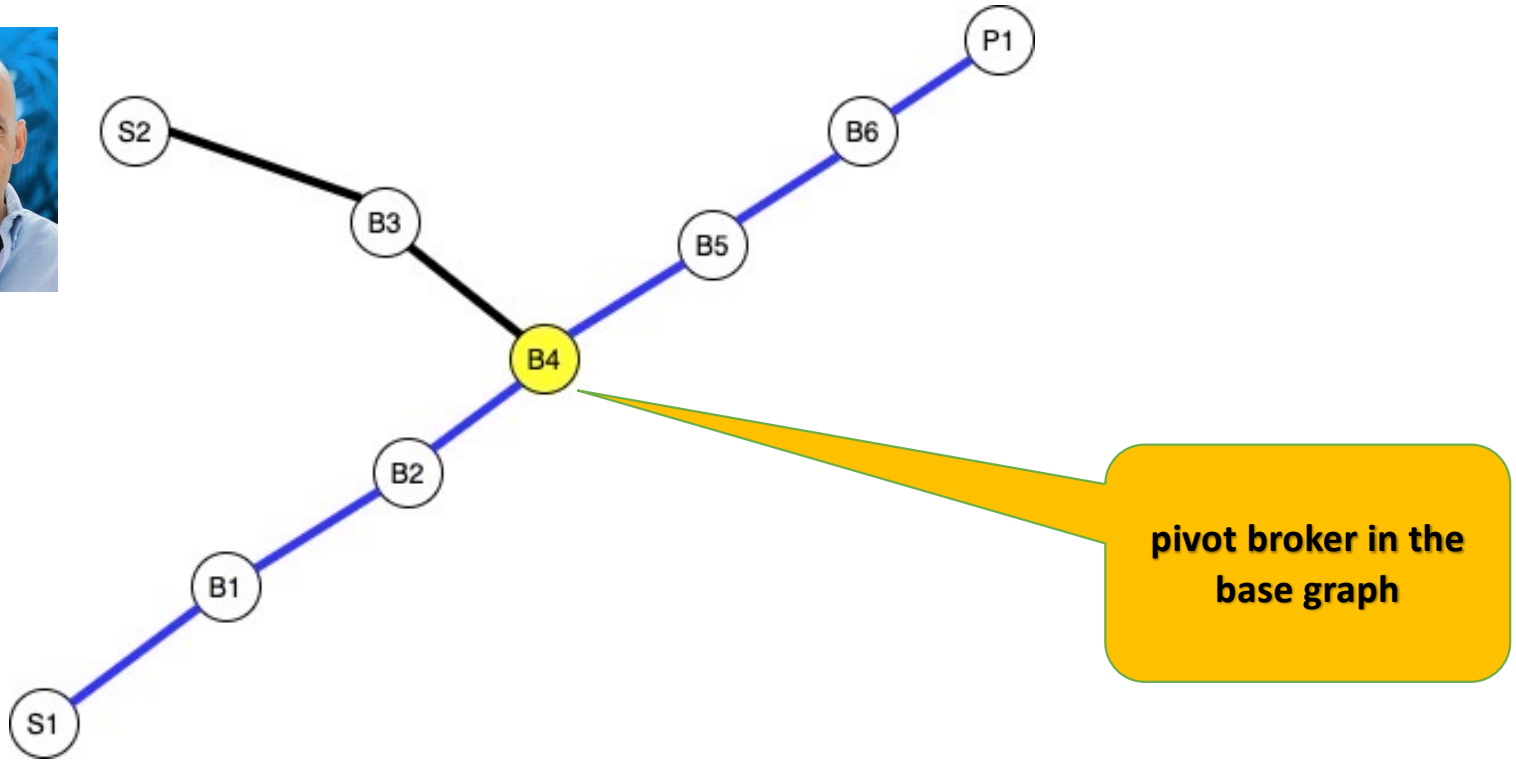

LoCaPS

Luís

Alan

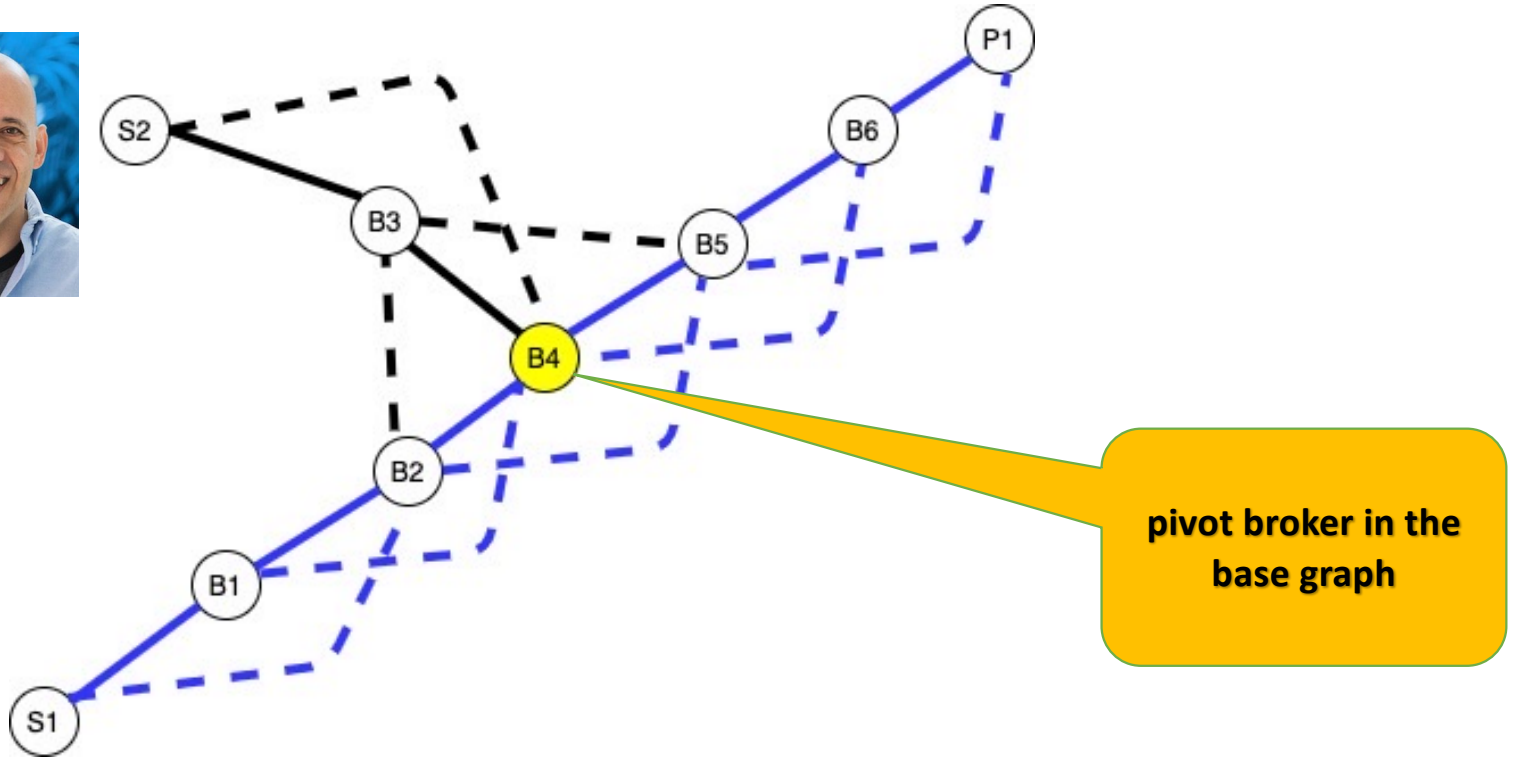

LoCaPS

Luís

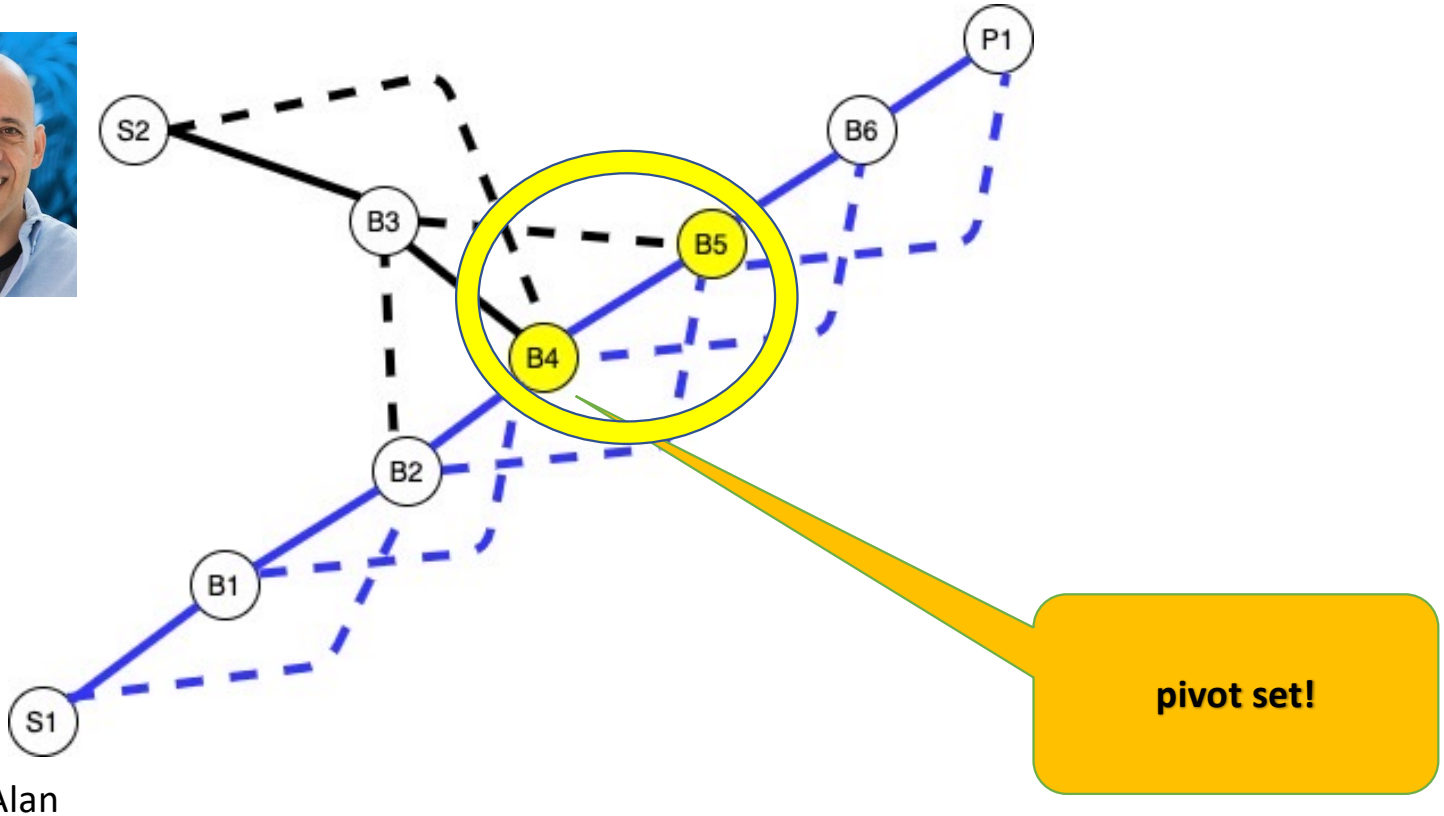

Alan

LoCaPS

Luís

Alan

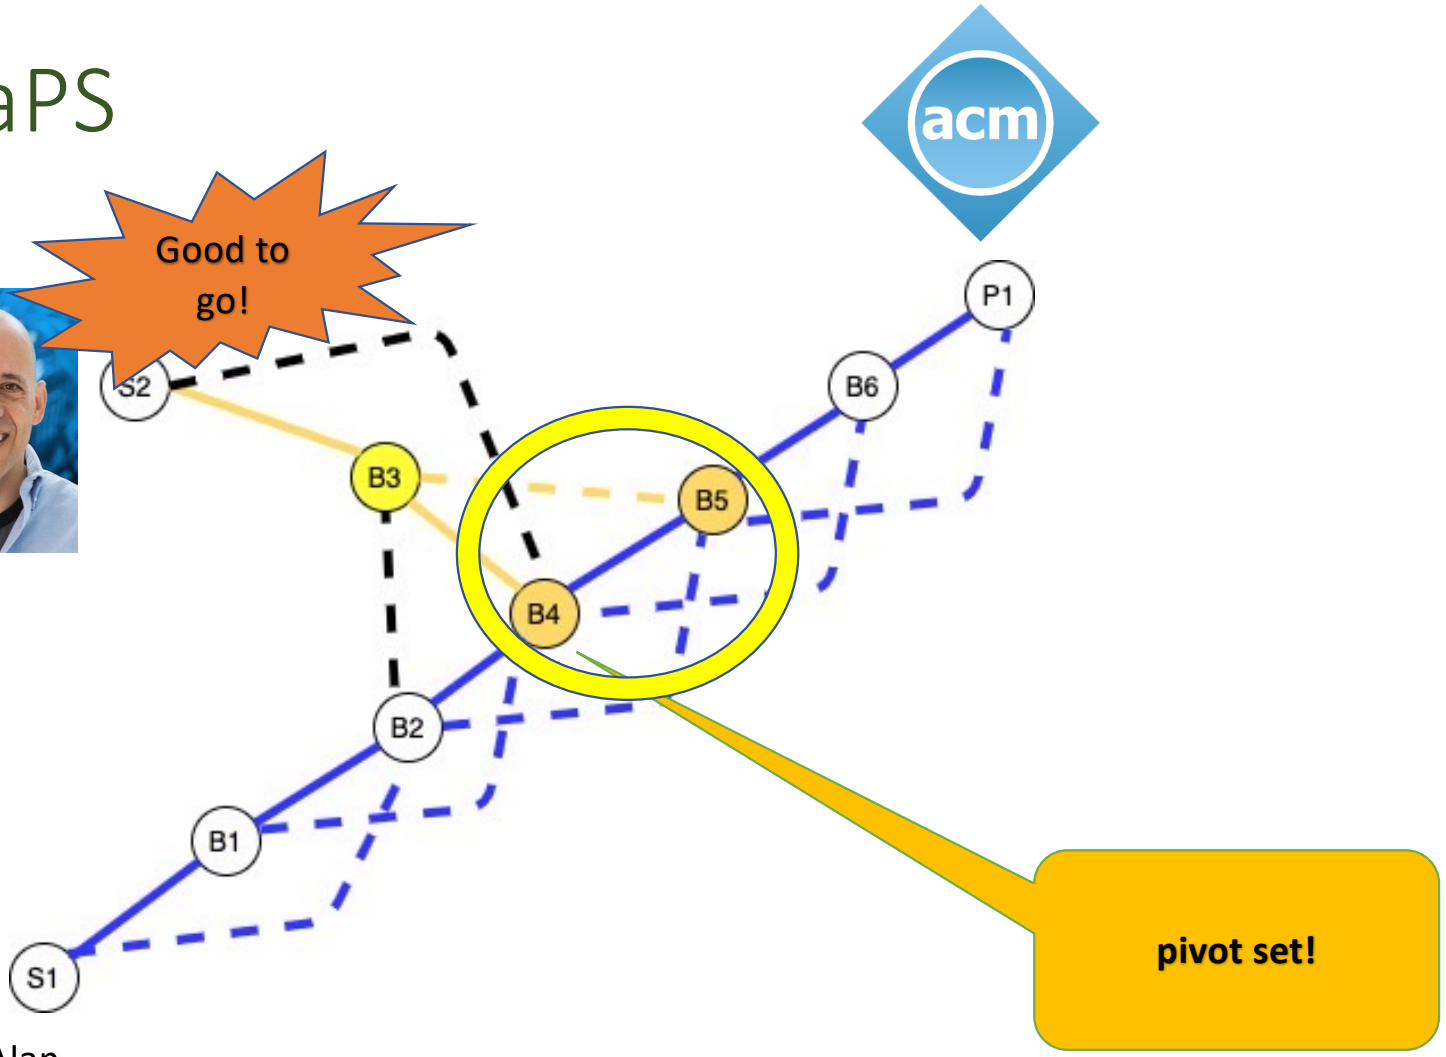

LoCaPS

- Evaluation against:
  - Delta: R. Kazemzadeh and H. Jacobsen. 2011. Partition-tolerant distributed publish/subscribe systems. In SRDS. IEEE, Madrid, Spain.
  - Gryphon: Y. Zhao, D. Sturman, and S. Bhola. 2004. Subscription Propagation in Highly-Available Publish/Subscribe Middleware. In Middleware. ACM, Toronto, Canada.

LoCaPS

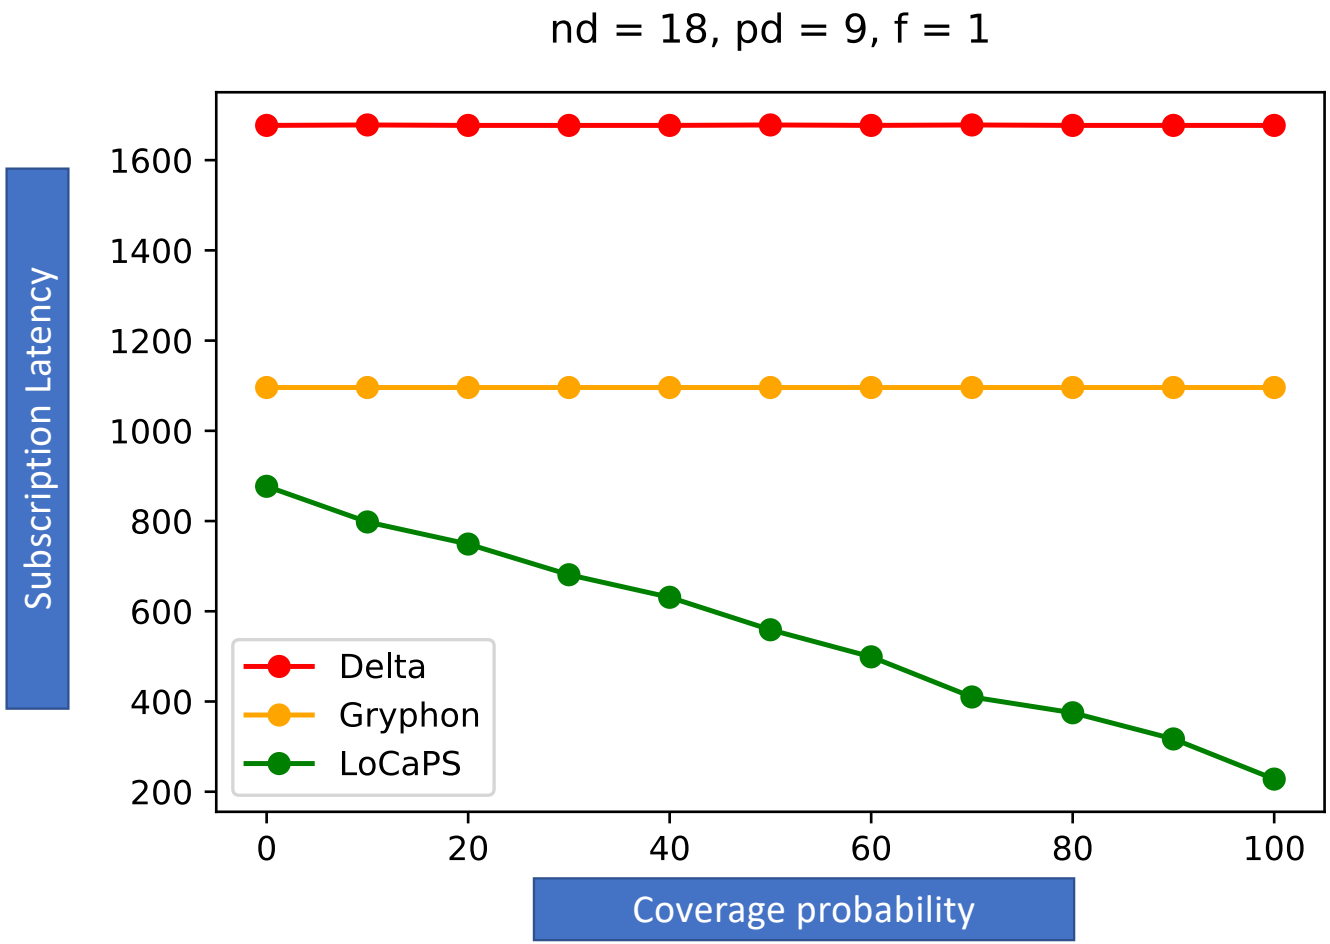
nd = 18, pd = 9, f = 1
Subscription Latency
1600
1400
1200
1000
800
600
400
200
0
20
40
60
80
100
Coverage probability
Delta
Gryphon
LoCaPS

LoCaPS

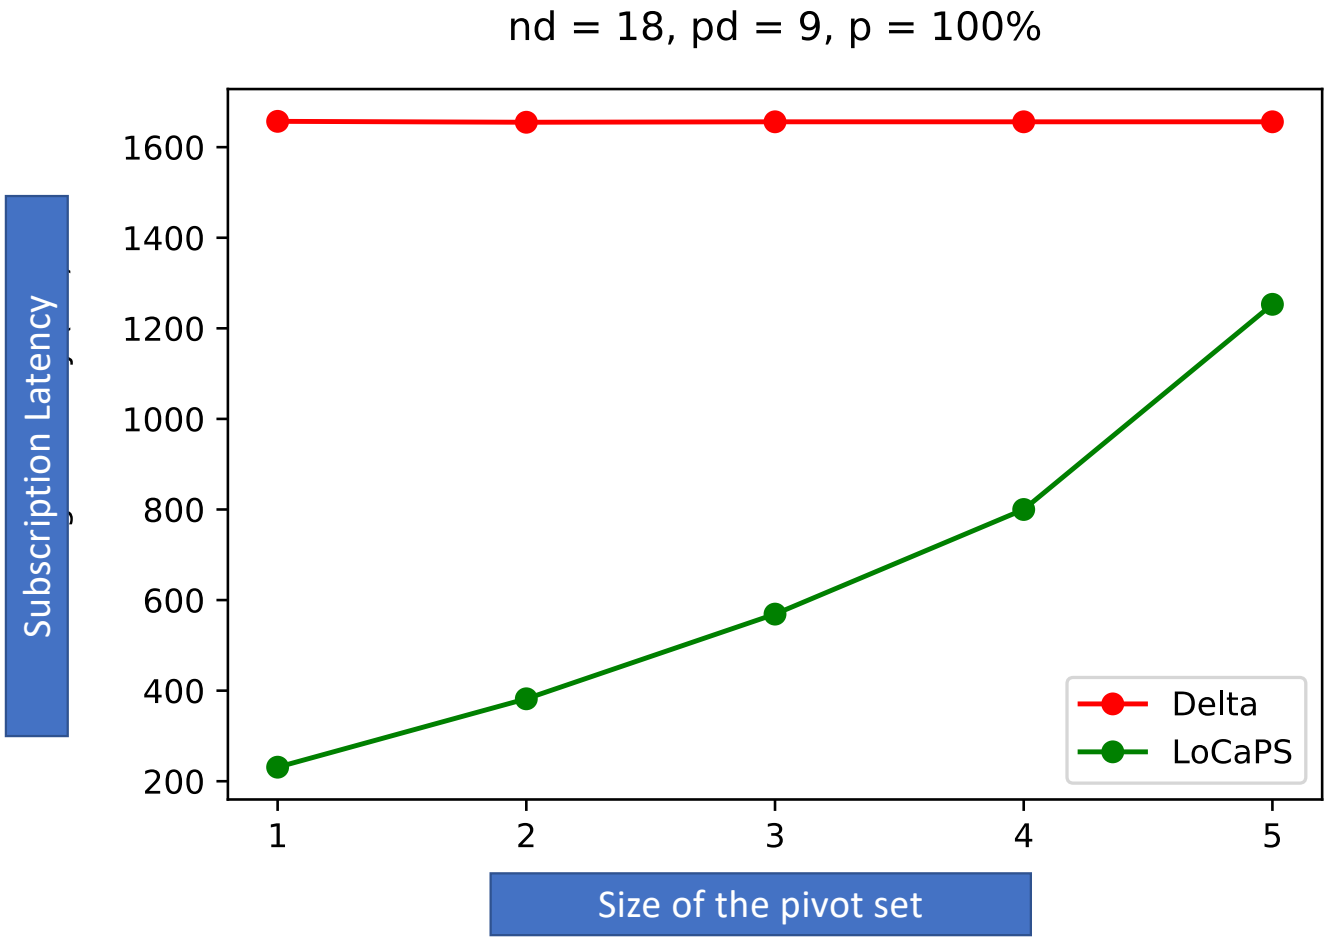
nd = 18, pd = 9, p = 100%
Subscription Latency
1600
1400
1200
1000
800
600
400
200
1
2
3
4
5
Delta
LoCaPS
Size of the pivot set

Conclusions

- We have studied the necessary and sufficient conditions that need to be met to offer different reliability semantics to subscribers, namely Gapless FIFO delivery and Gapless Causal delivery.
- We shown that Gapless Causal delivery can be implemented as efficiently as Gapless FIFO delivery.
- Unlike previous systems, LoCaPS can leverage existing subscriptions to reduce the latency of a new subscription.

- **More details and more evaluation results on the paper!**