Practical, Real-time Centralized Control for CDN-based Live Video Delivery

Matt Mukerjee, David Naylor,
Junchen Jiang, Dongsu Han,
Srini Seshan, Hui Zhang

Carnegie Mellon University

CONVIVA

KAIST

Combating Latency in Wide Area Control Planes

- Centralization can provide major benefits
  - e.g., better performance, reliability, policy management, …
- Scalability is hard on WAN due to latency

Control Planes in the 4D* Model

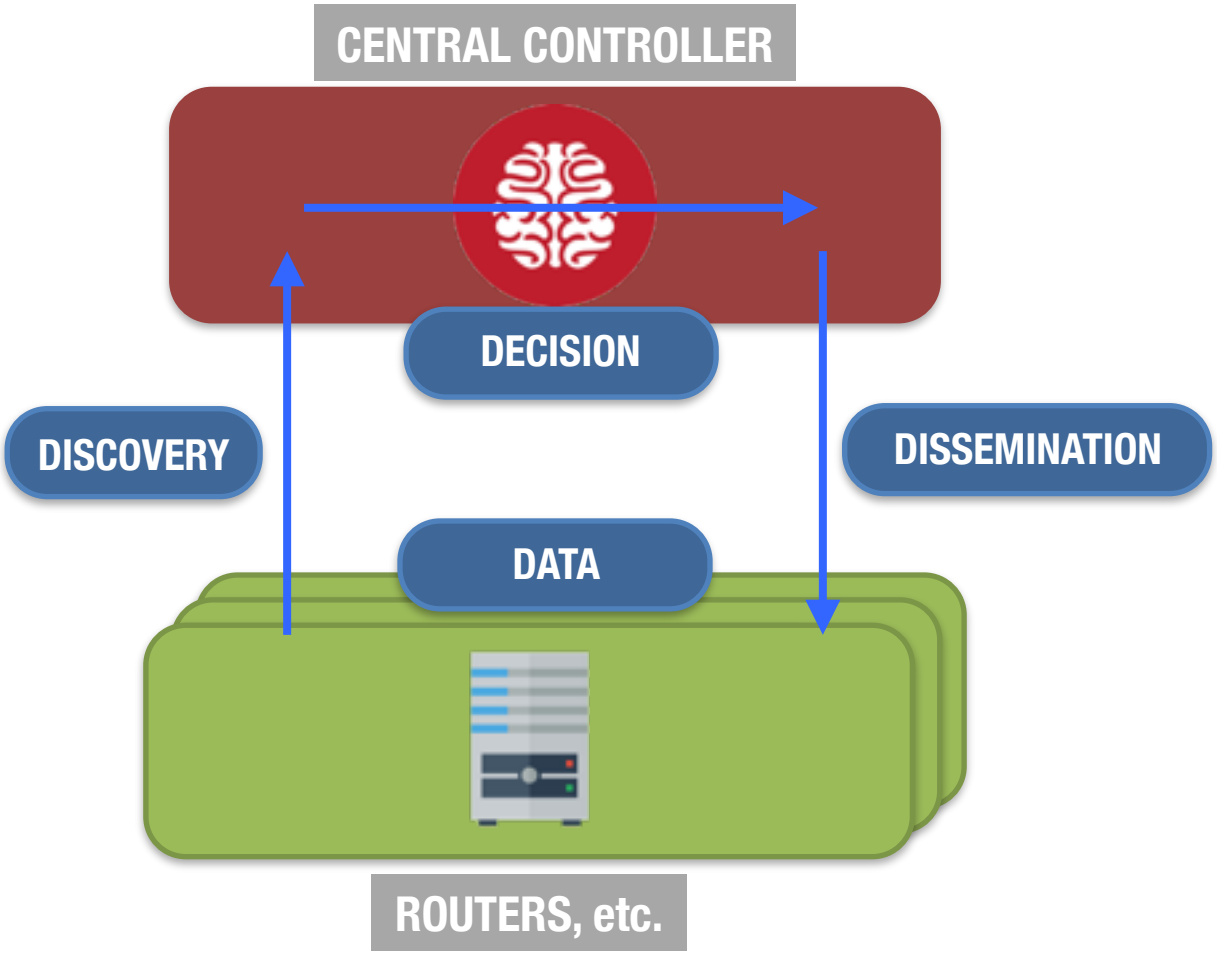

*Yan, Hong, et al. "Tesseract: A 4D Network Control Plane." NSDI. Vol. 7. 2007.

WAN Control Plane Latency

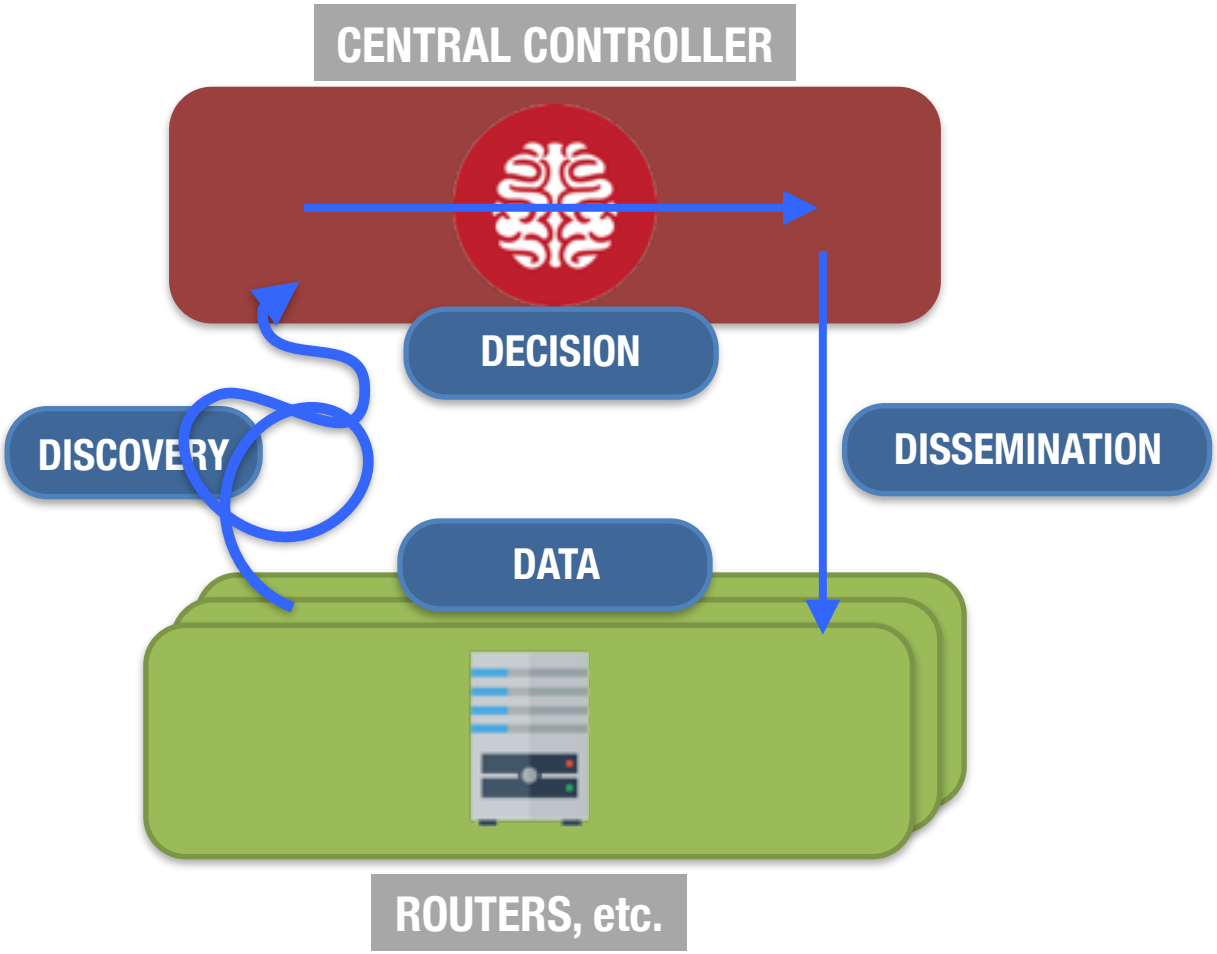

WAN Control Plane Latency

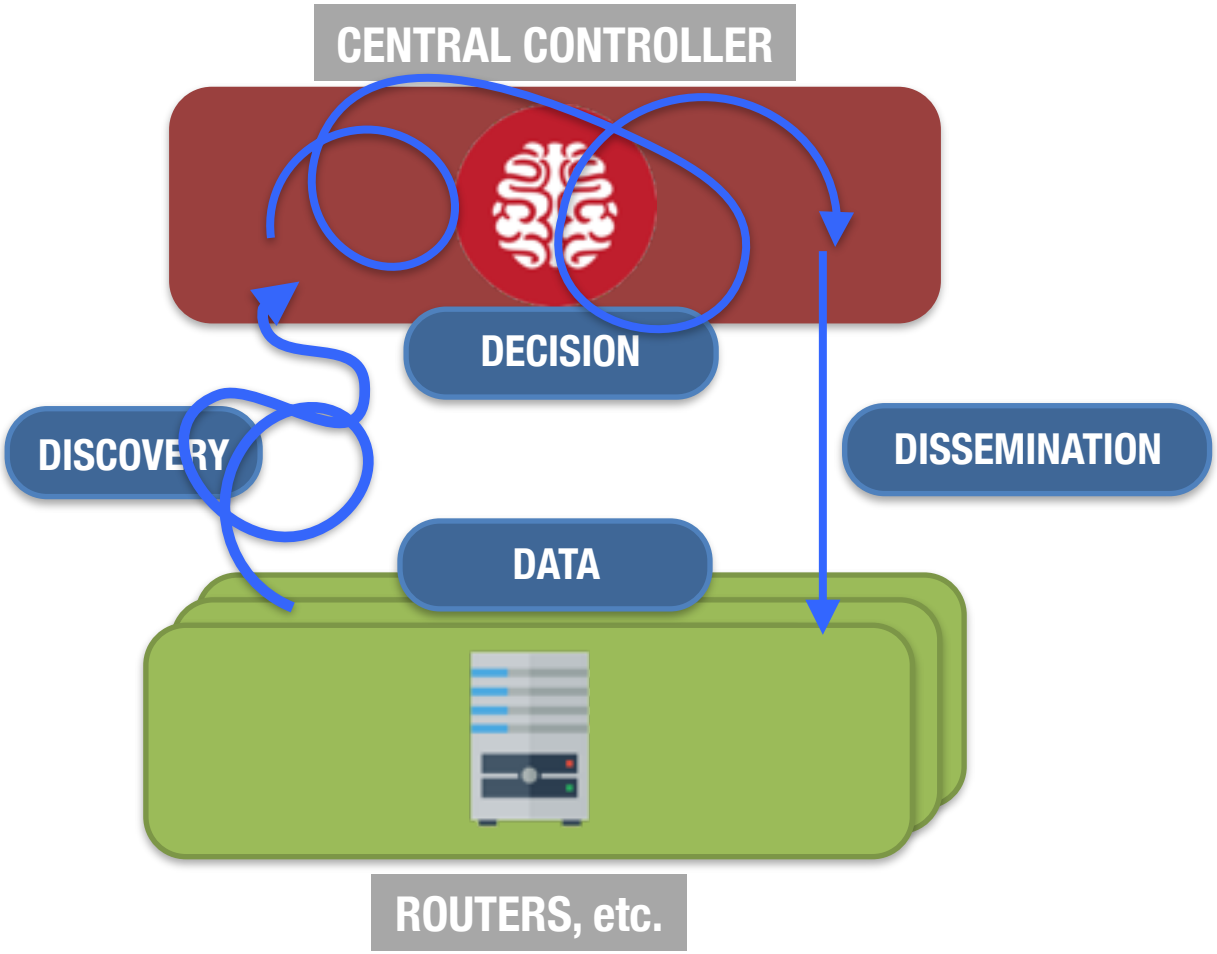

WAN Control Plane Latency

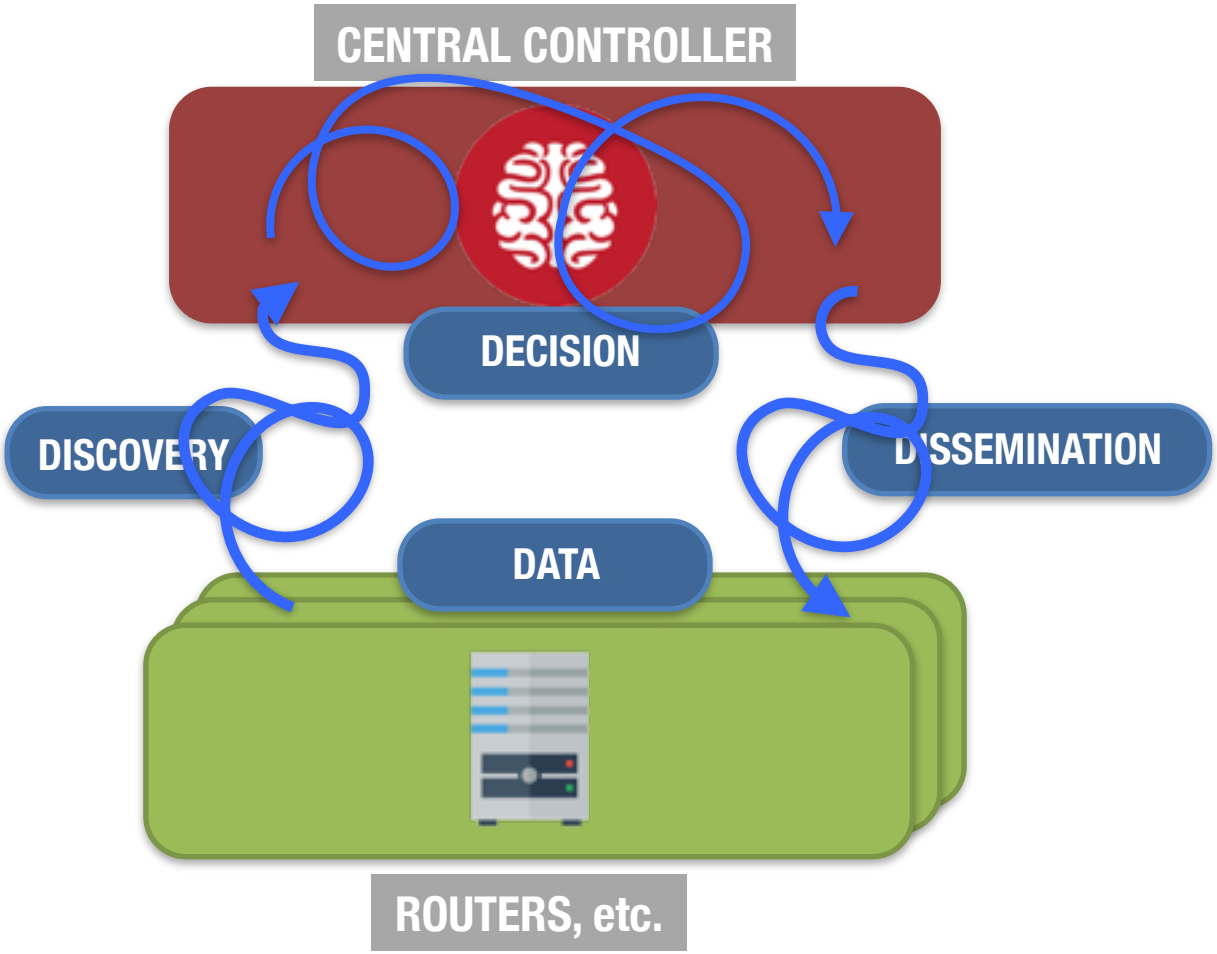

WAN Control Plane Latency

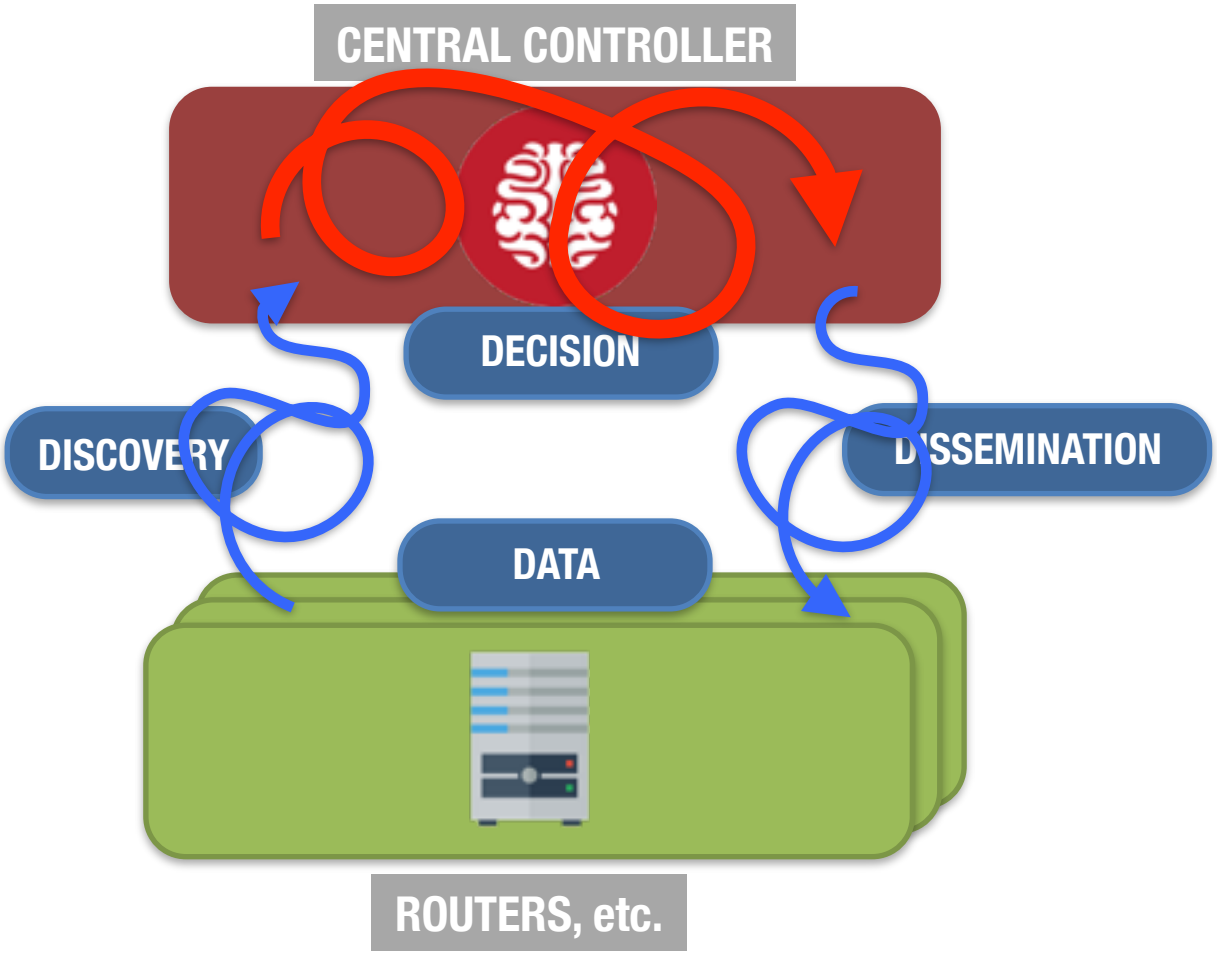

WAN Problems and Decision Planes

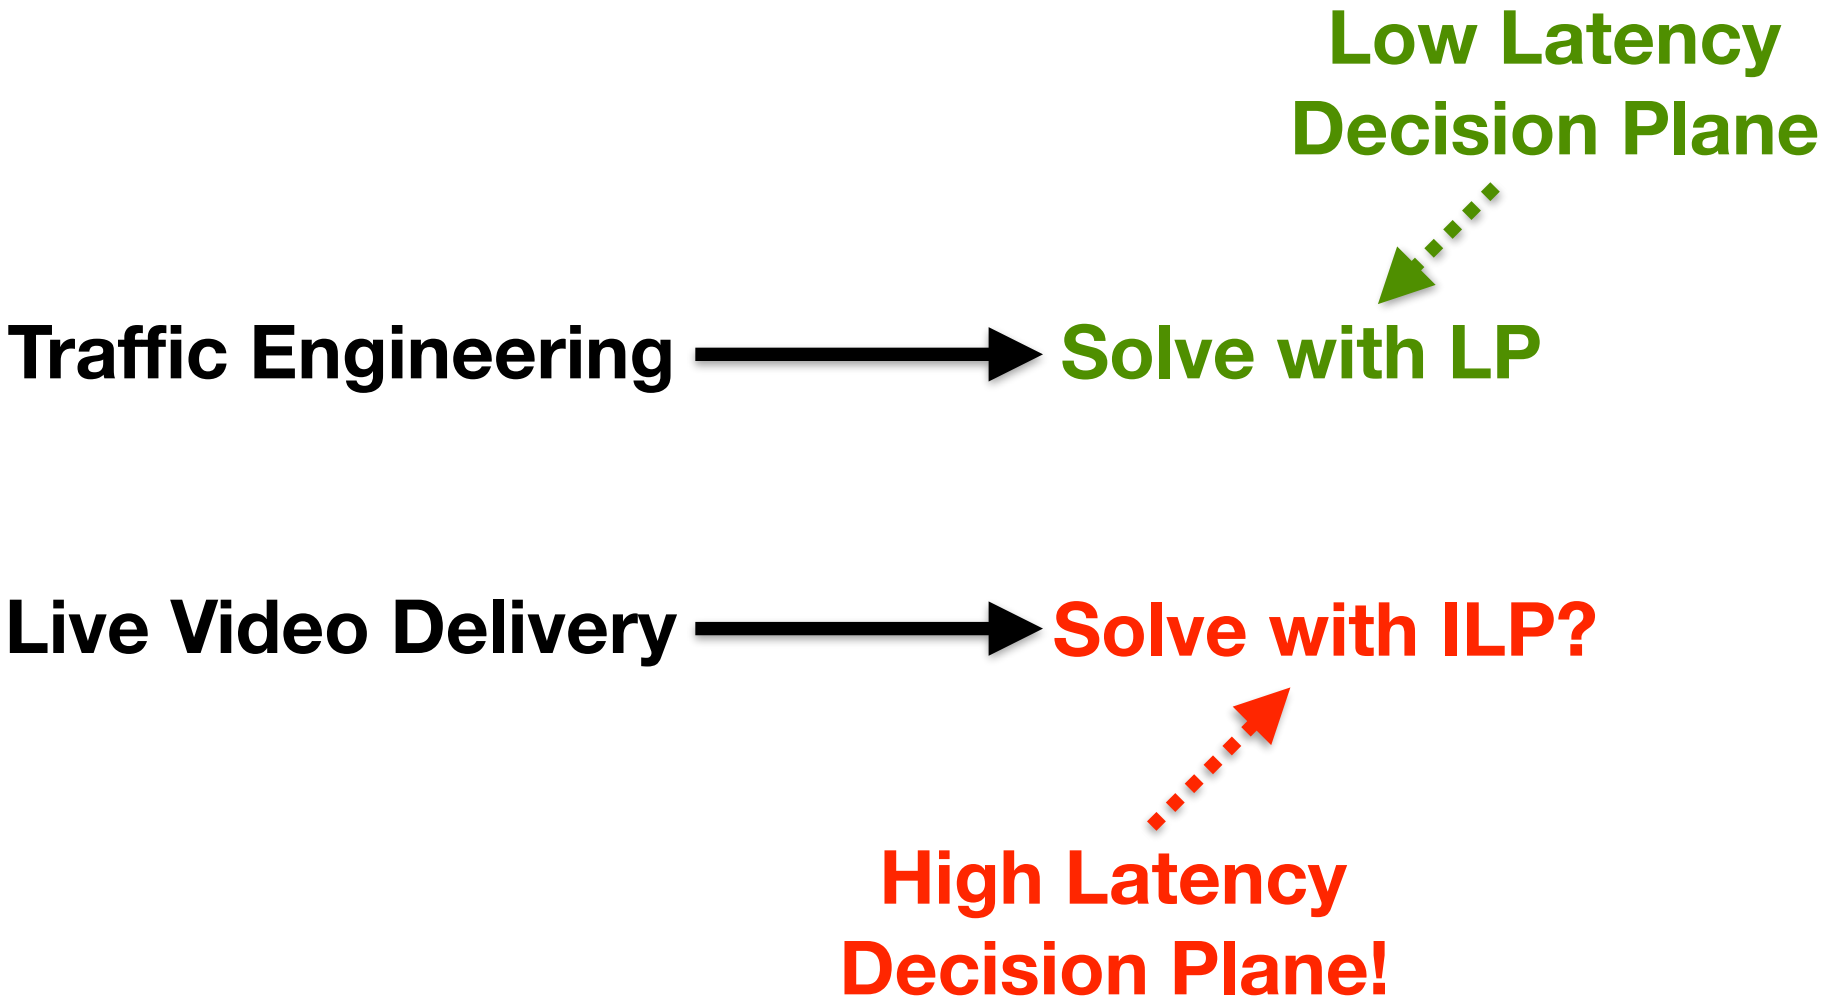

Attacking Decision Plane Latency

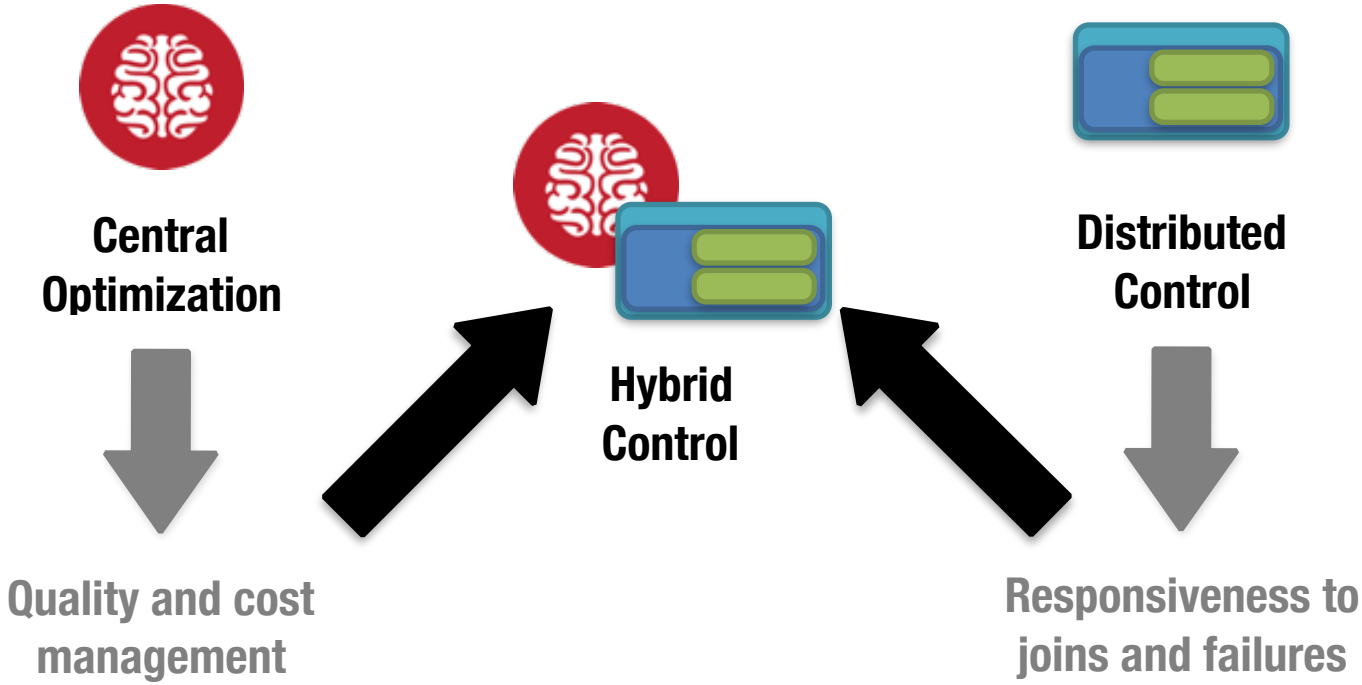

Outline

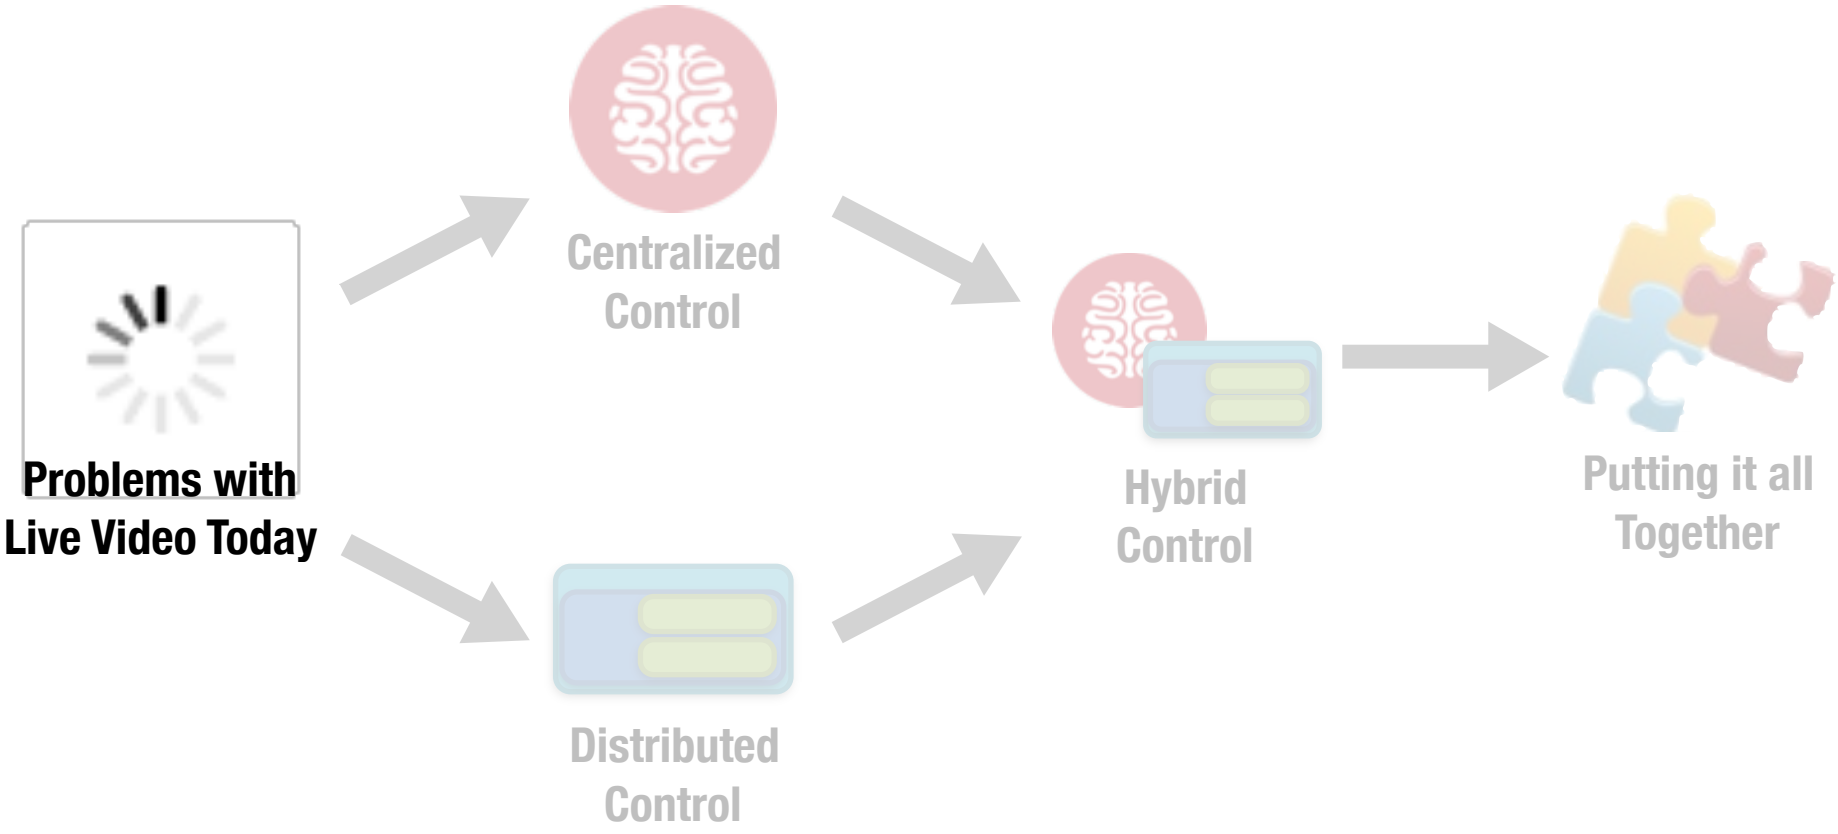
Problems with
Live Video Today
Centralized
Control
Distributed
Control
Hybrid
Control
Putting it all
Together

CDN Live Video Delivery Background

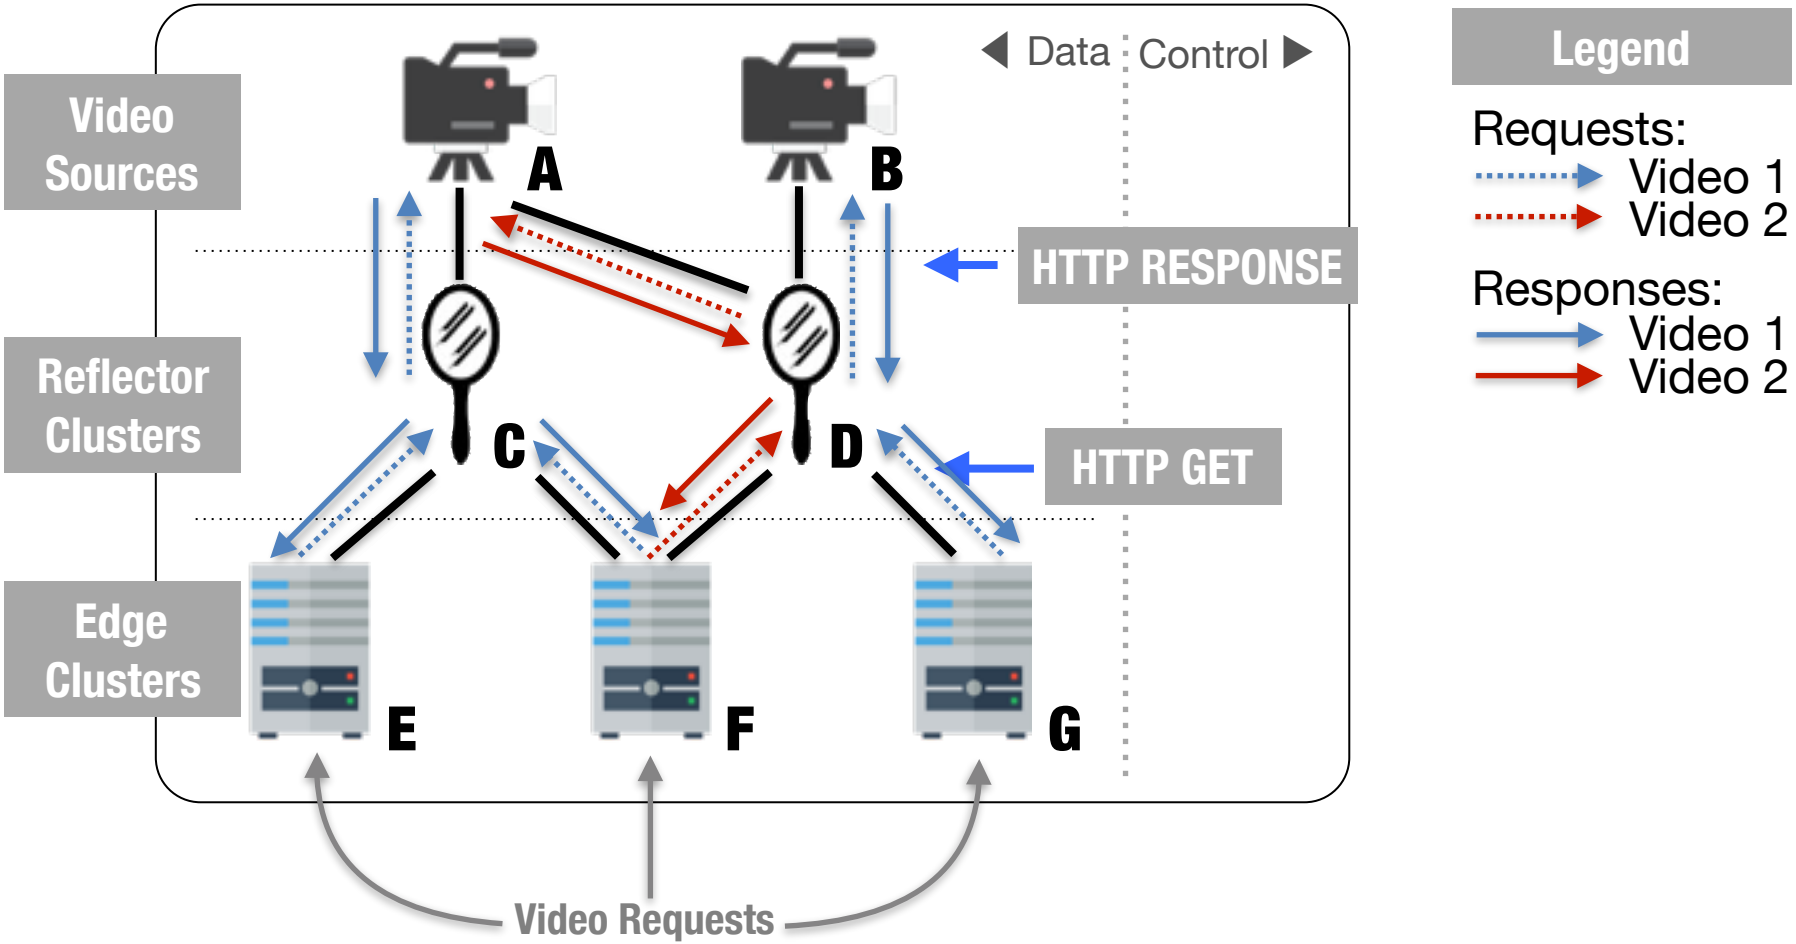

CDN Live Video Delivery Background

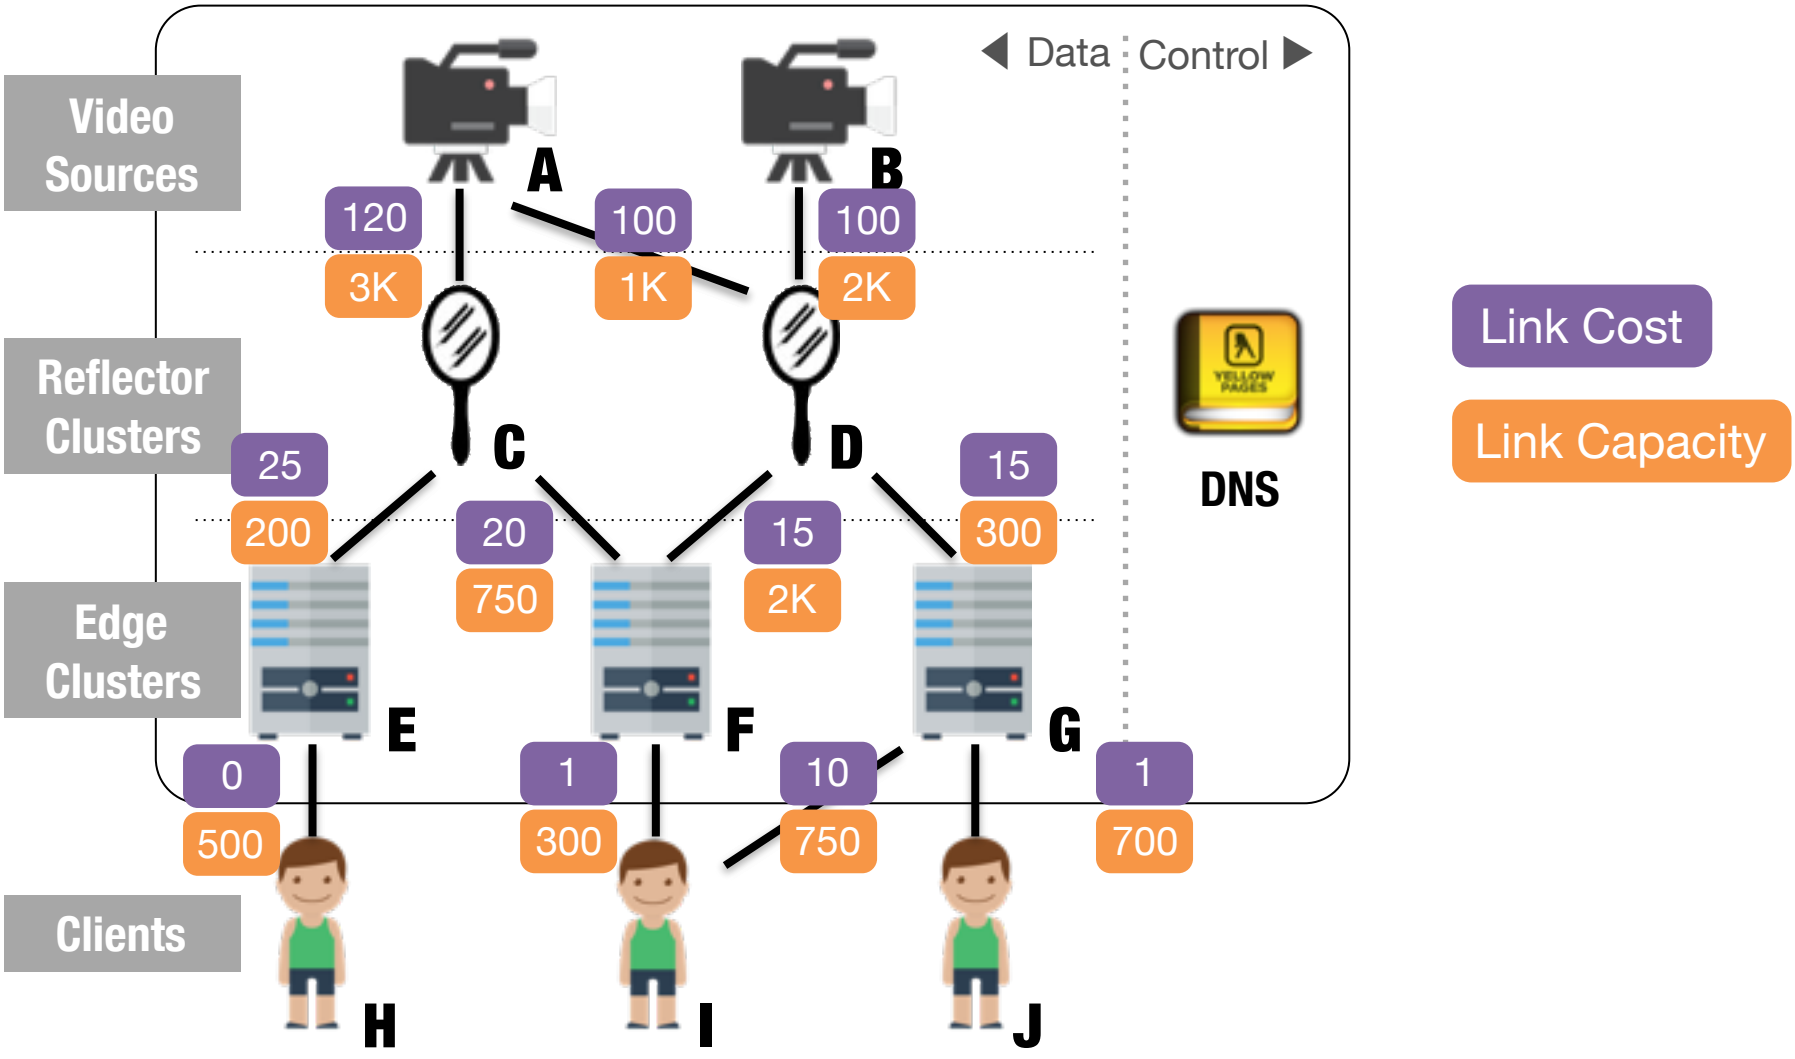

CDN Live Video Delivery Background

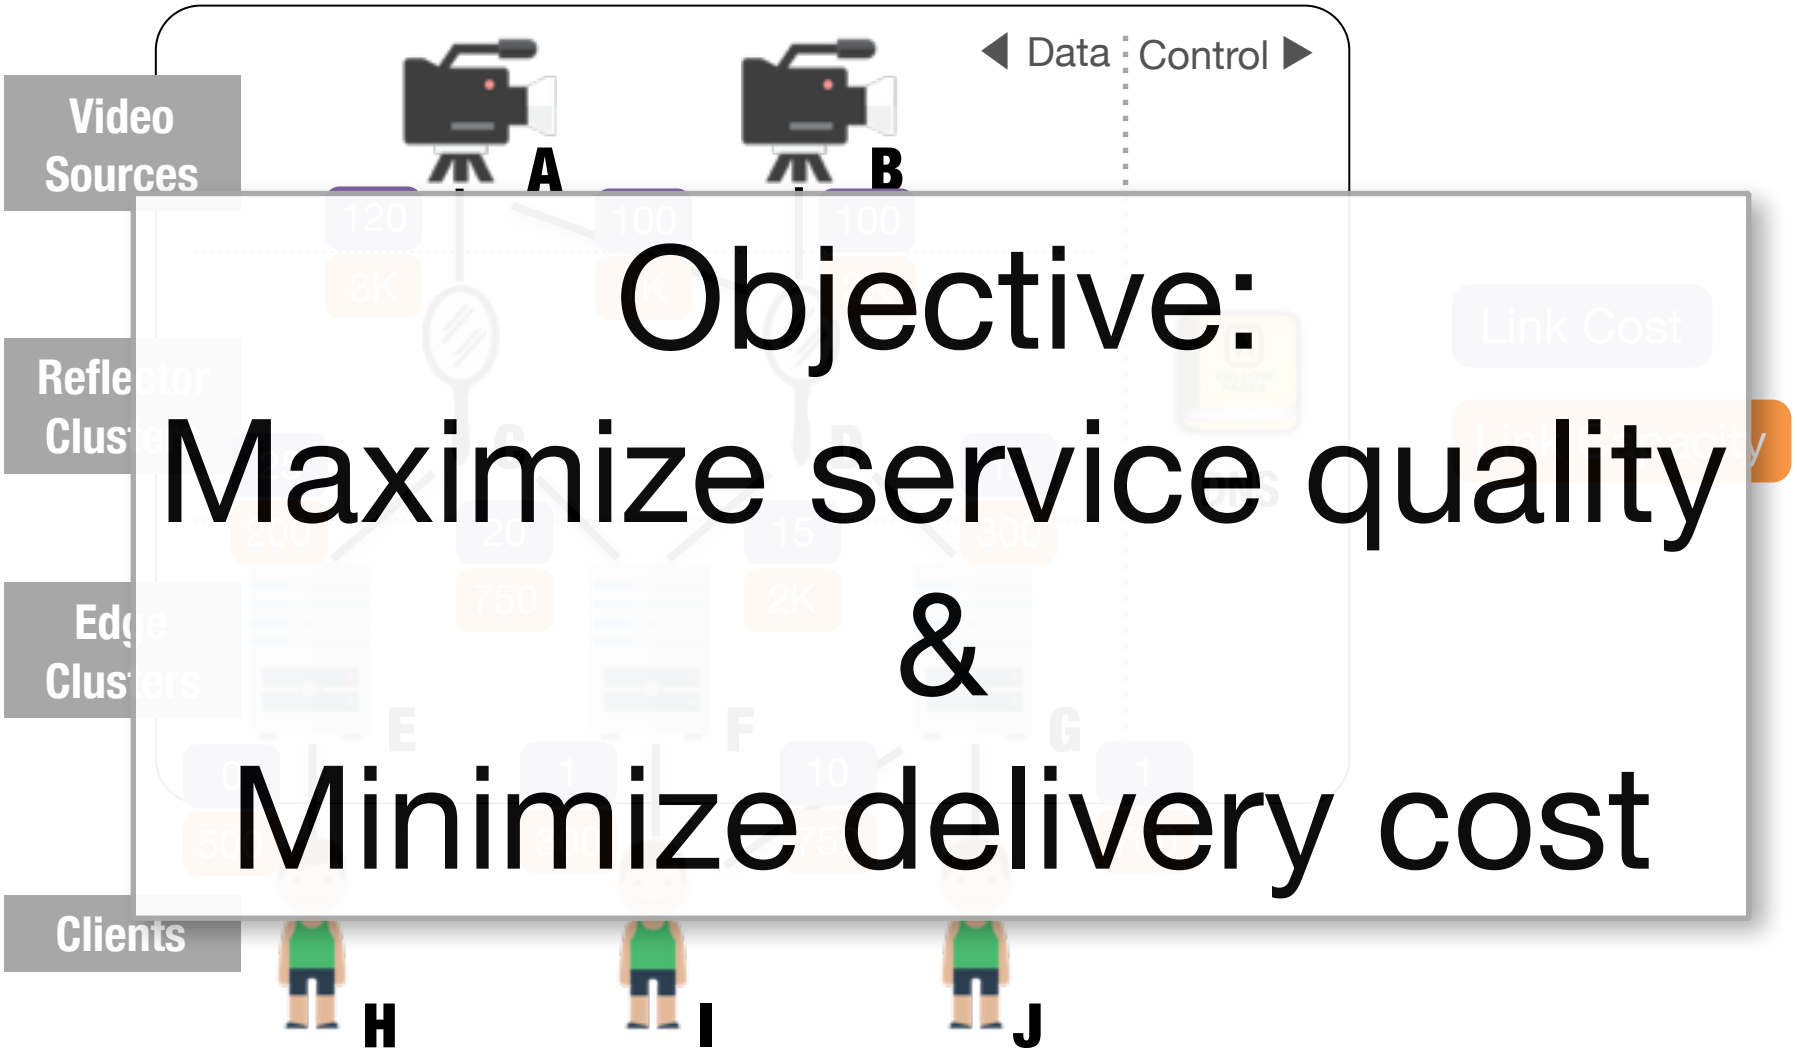

Objective:
Maximize service quality
&
Minimize delivery cost

Outline

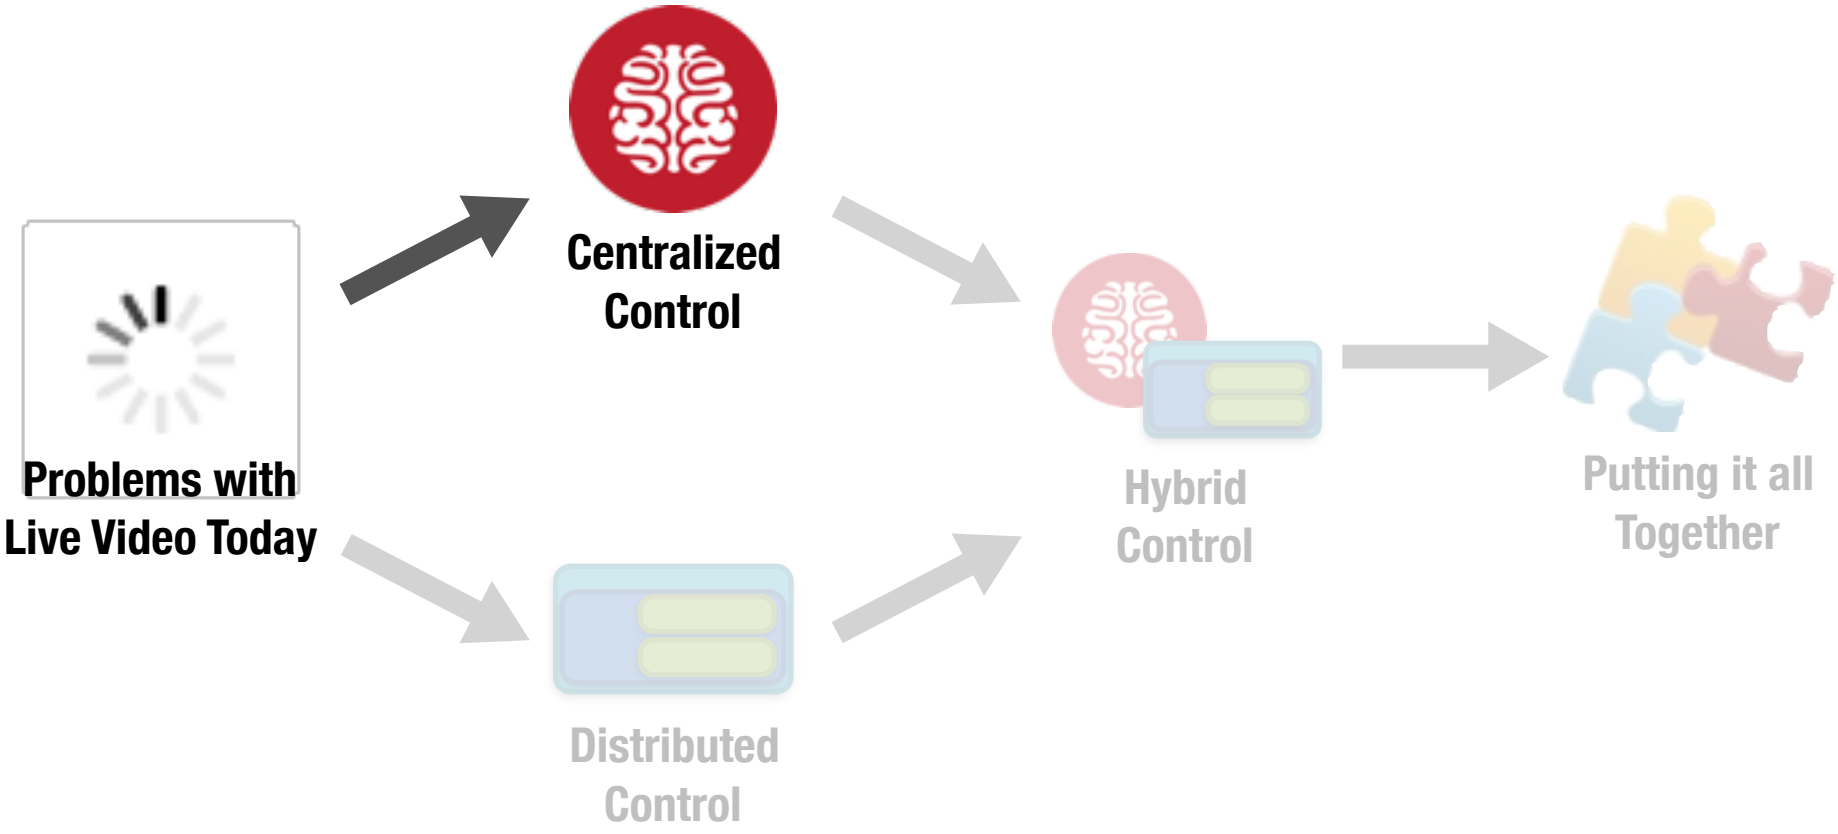
Problems with
Live Video Today
Centralized
Control
Distributed
Control
Hybrid
Control
Putting it all
Together

Motivating Centralized Optimization

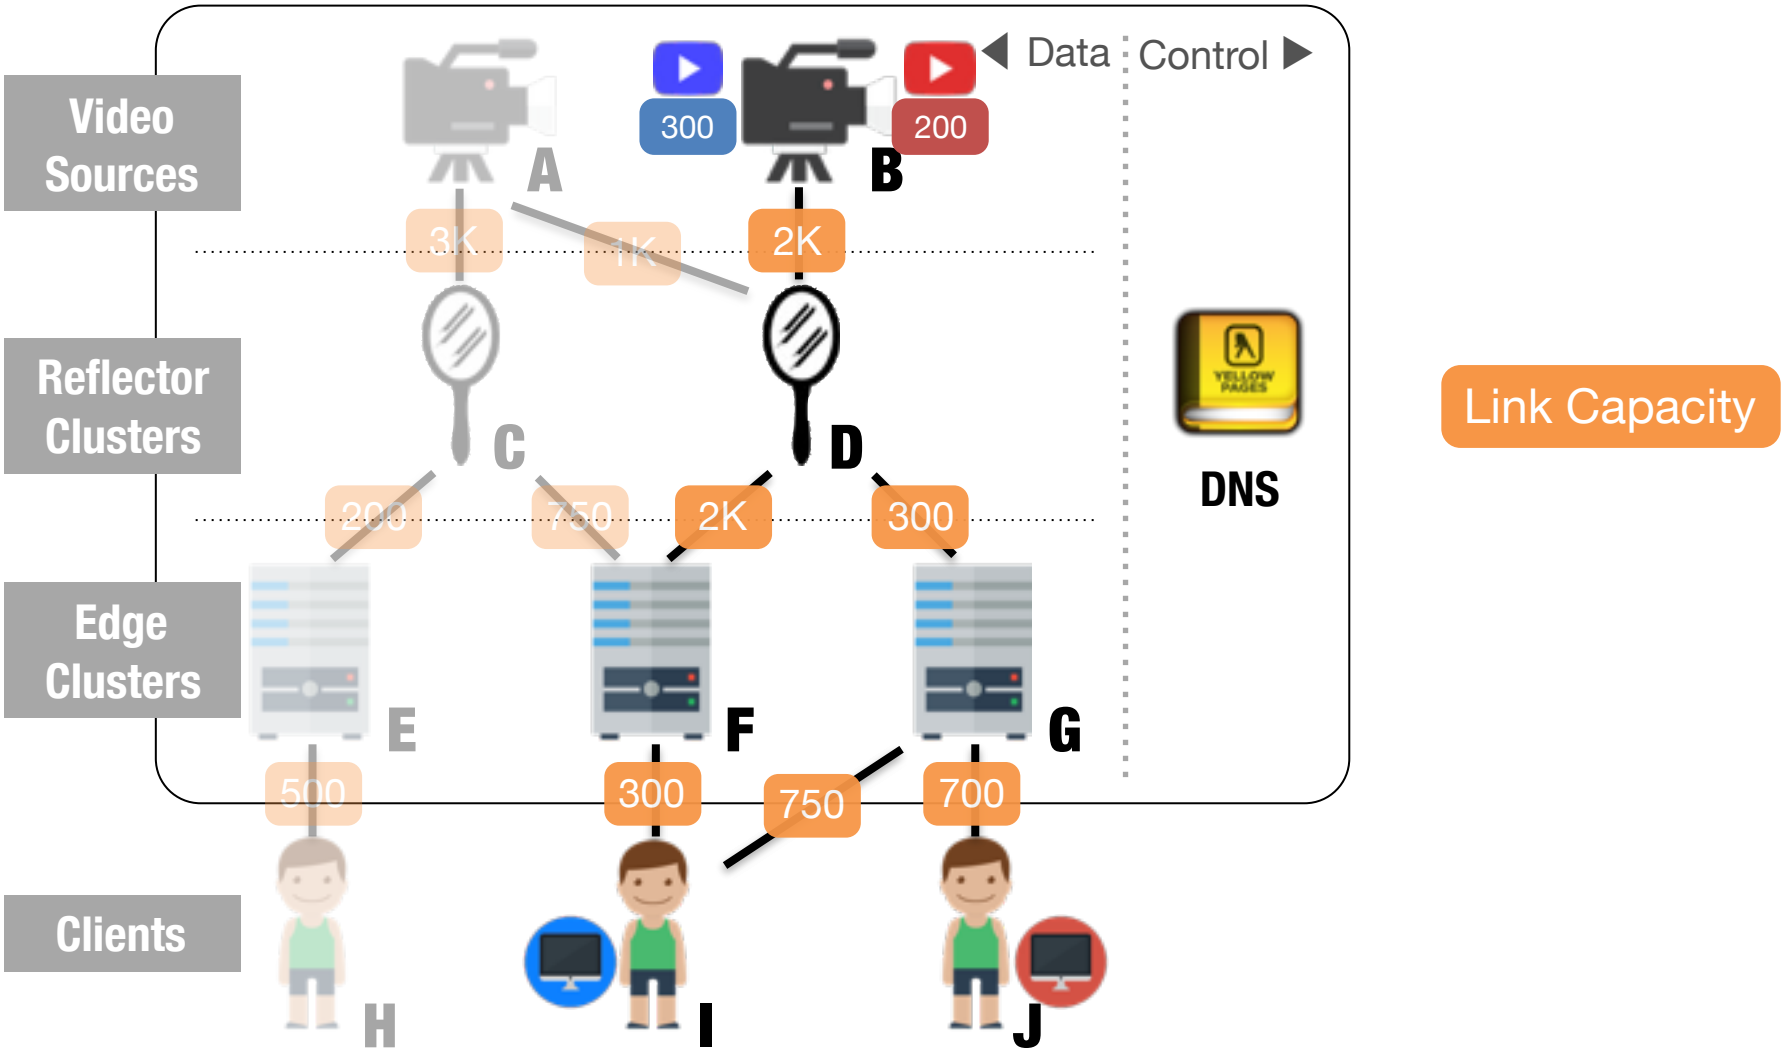

Link Capacity

Motivating Centralized Optimization

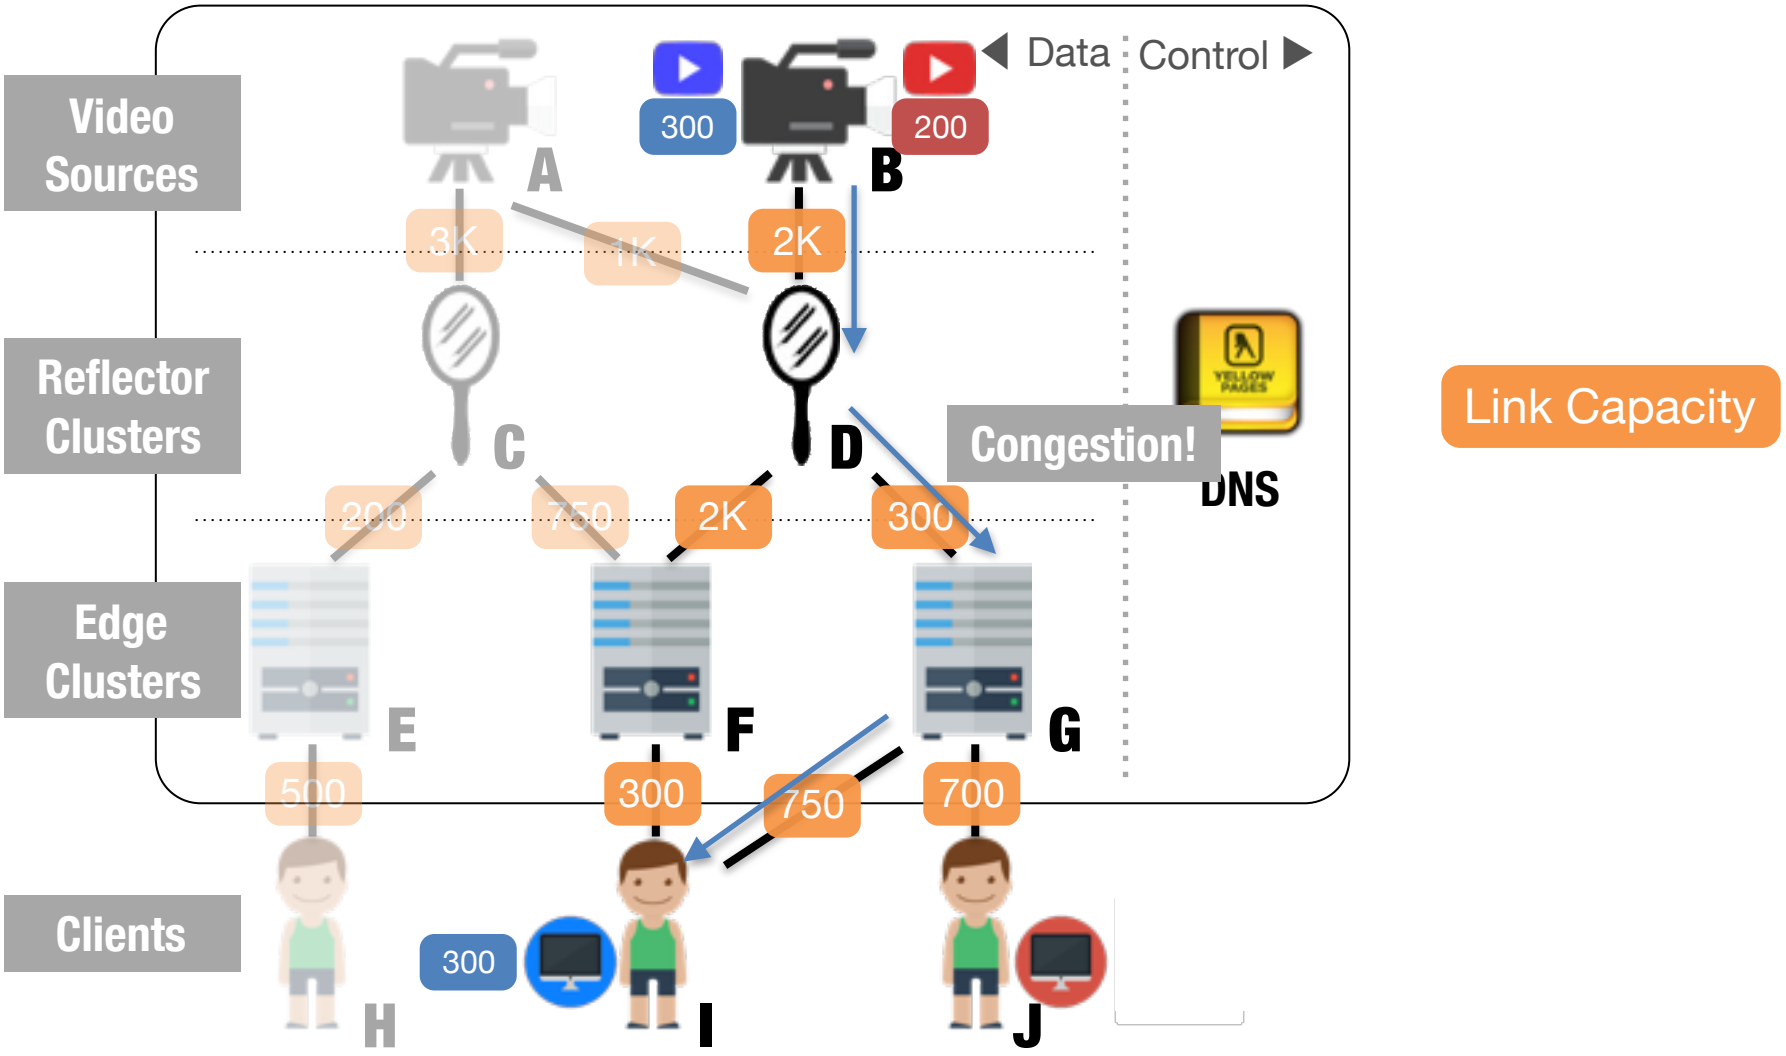

Motivating Centralized Optimization

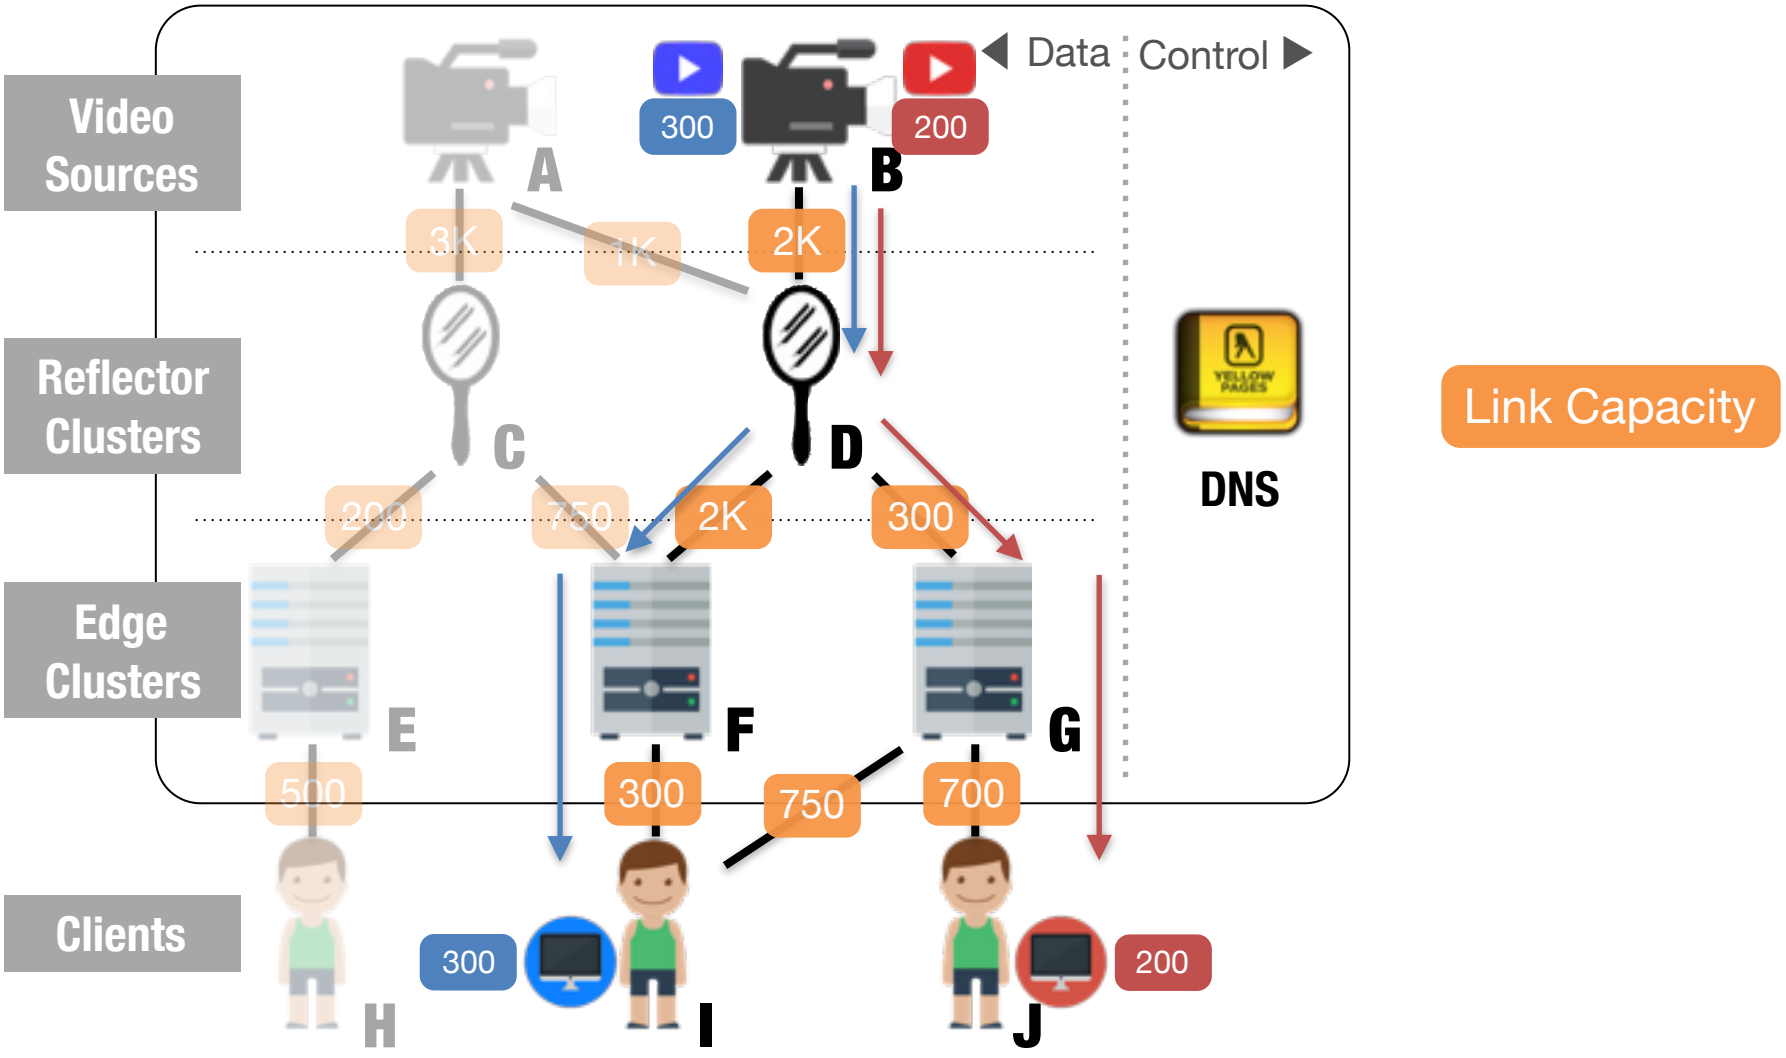

Motivating Centralized Optimization

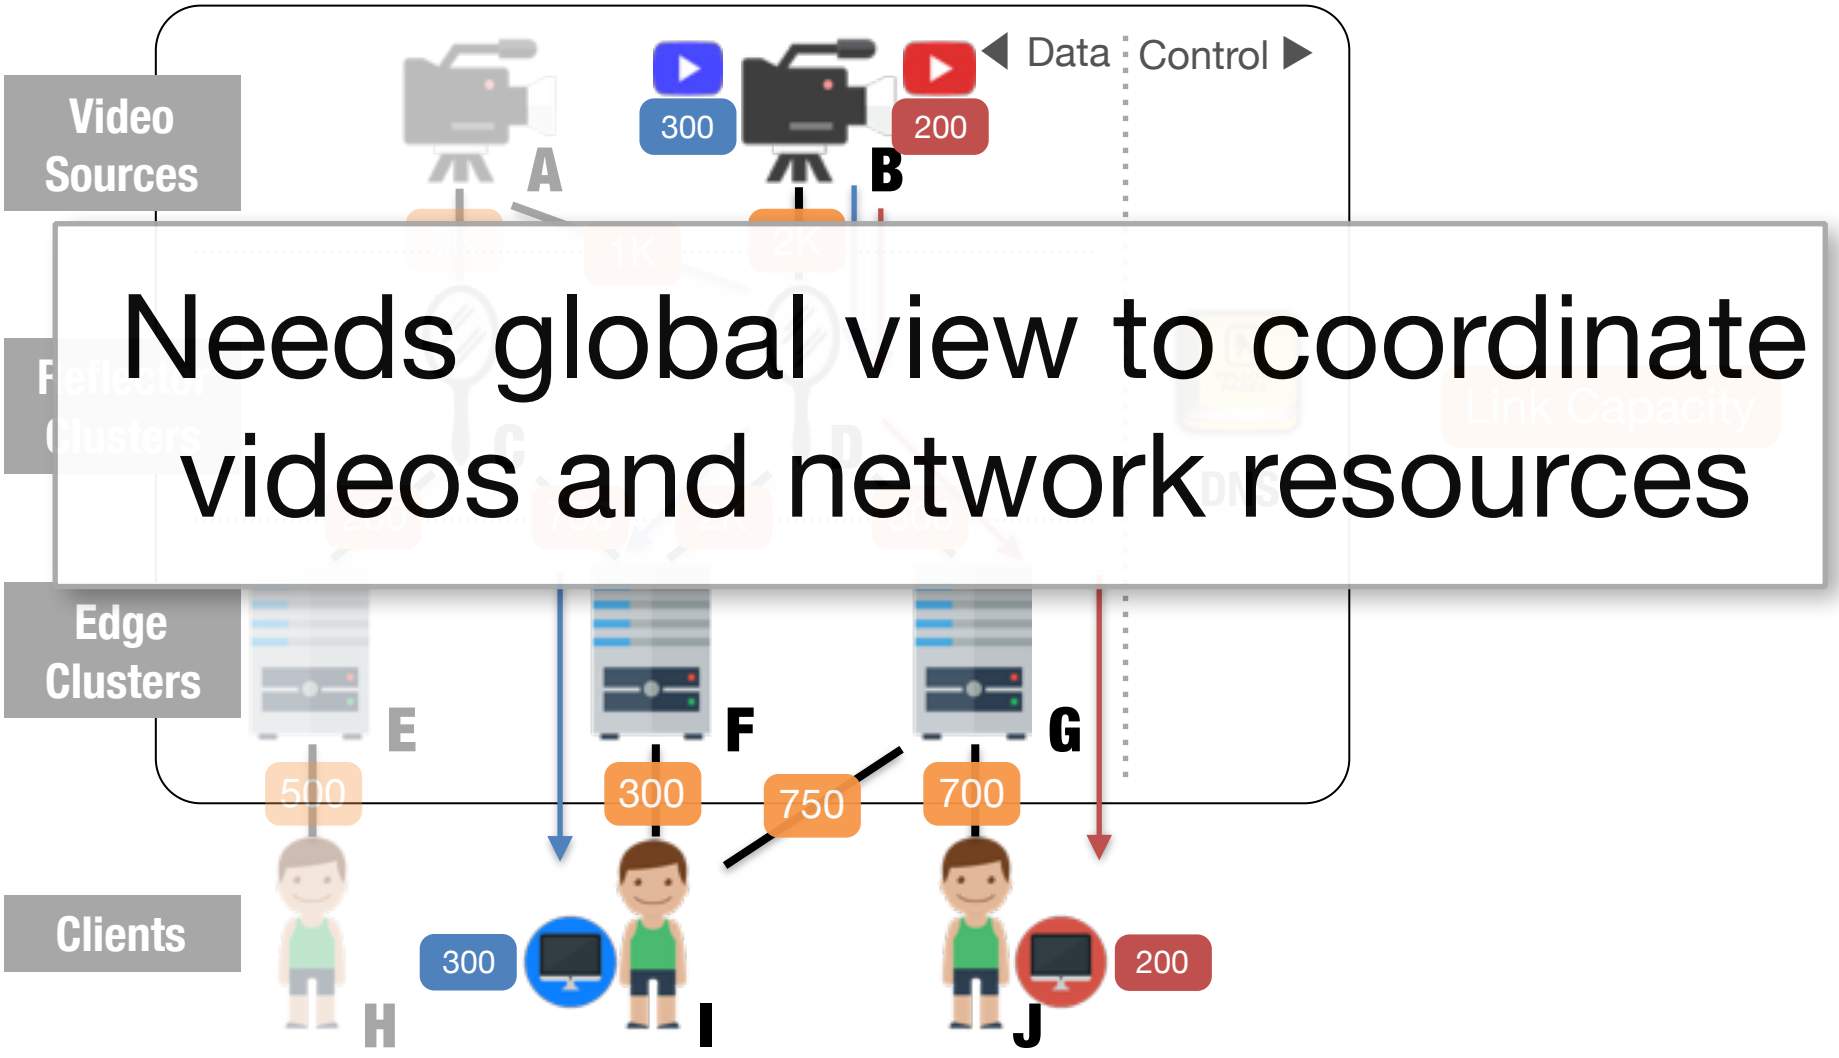

Needs global view to coordinate videos and network resources

Motivating Centralized Optimization

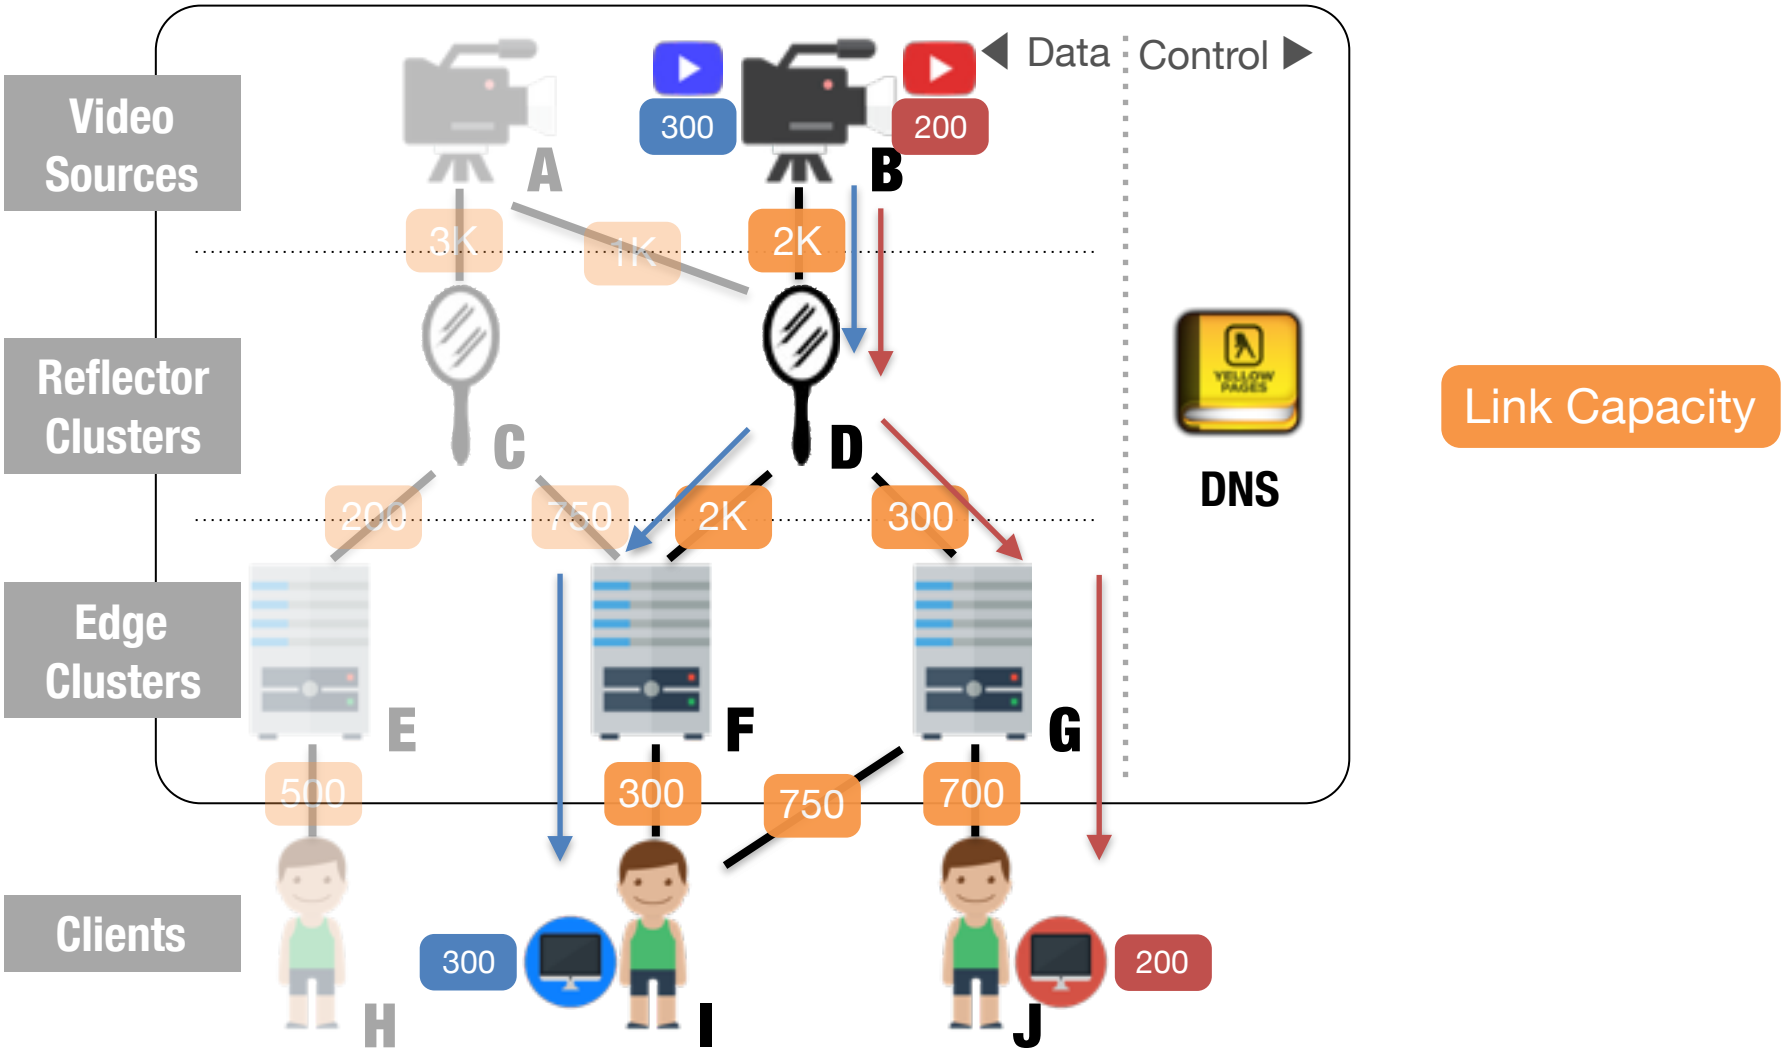

Motivating Centralized Optimization

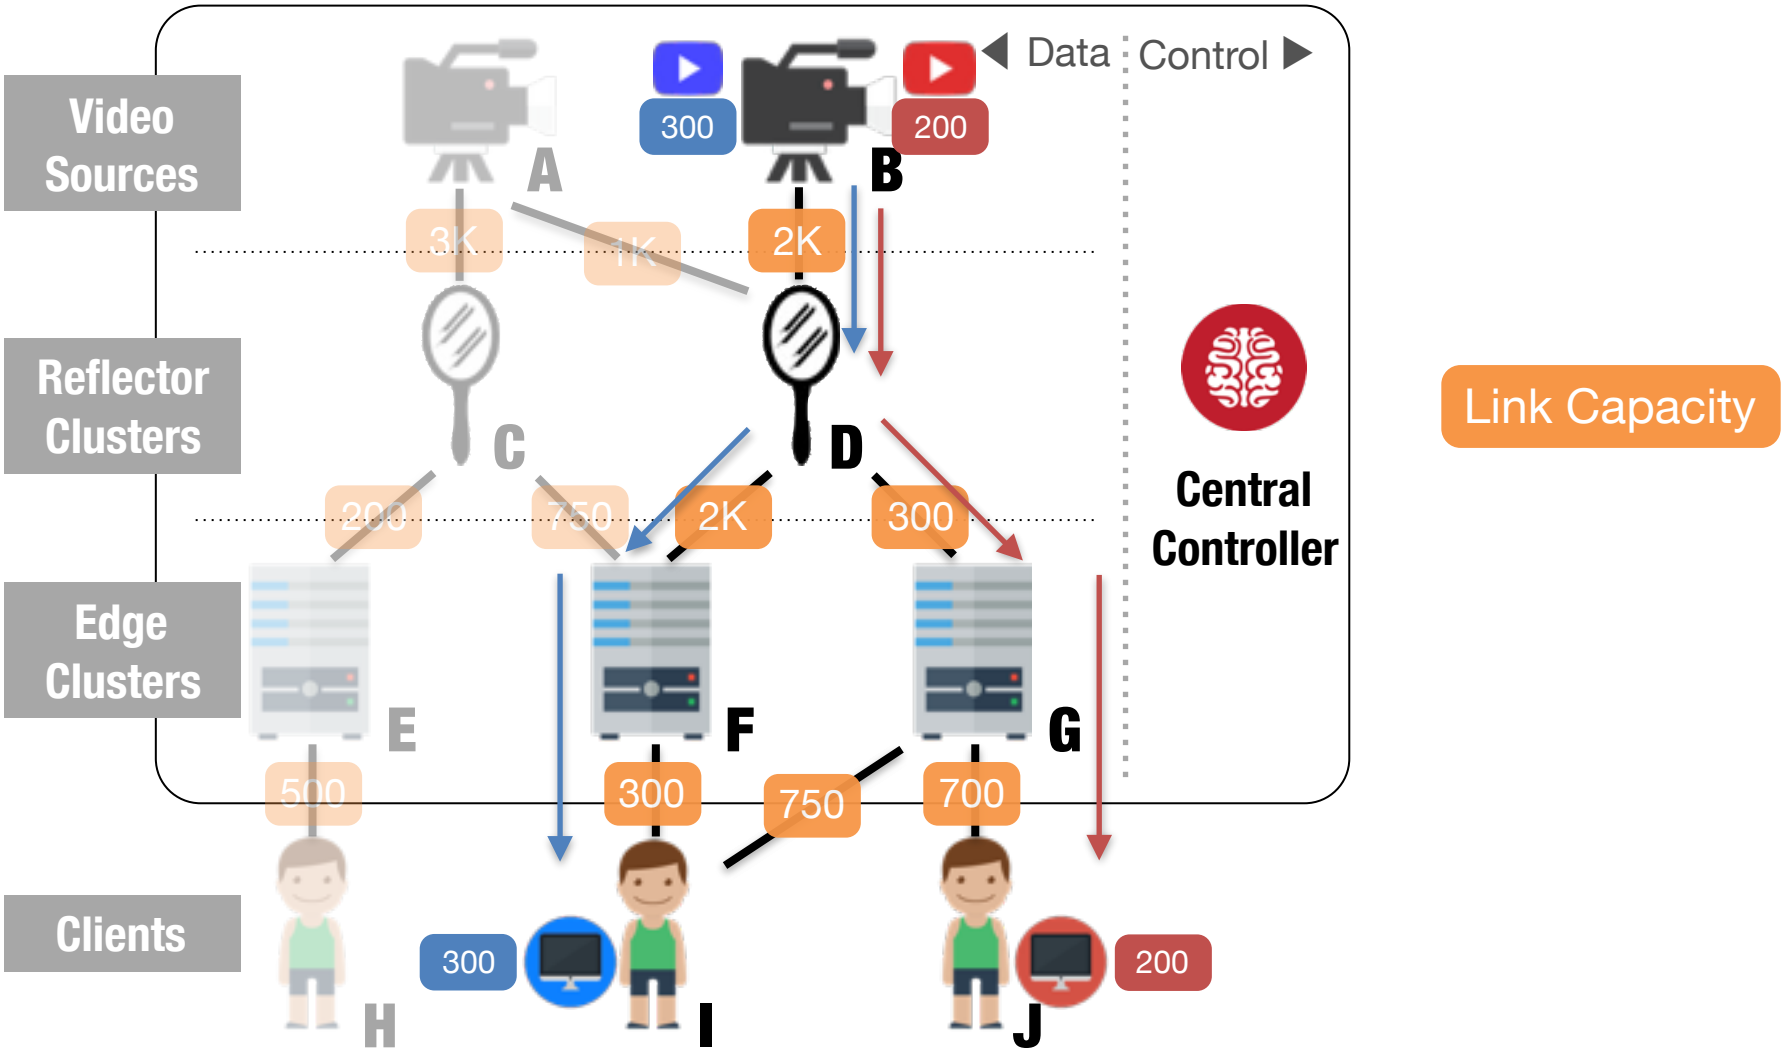

Solving Centralized Optimization

MAXIMIZE **SERVICE QUALITY**

MINIMIZE **DELIVERY COST**

SUBJECT TO **DON'T EXCEED LINK CAPACITY**

SENDER MUST HAVE RECEIVED VIDEO

Solving Centralized Optimization

SERVICE QUALITY

$$\max \; w_s \cdot \sum_{l \in L_{AS}, o \in O} \text{Priority}_o \cdot \text{Request}_{l,o} \cdot \text{Serves}_{l,o}$$

$$- \quad w_c \cdot \sum_{l \in L, o \in O} \text{Cost}(l) \cdot \text{Bitrate}(o) \cdot \text{Serves}_{l,o}$$

DELIVERY COST

subject to:

$$\forall l \in L, o \in O : \text{Serves}_{l,o} \in \{0, 1\}$$

DON'T EXCEED LINK CAPACITY

$$\forall l \in L : \sum_o \text{Bitrate}(o) \cdot \text{Serves}_{l,o} \leq \text{Capacity}(l)$$

$$\forall l \in L, o \in O : \sum_{l' \in \text{InLinks}(l)} \text{Serves}_{l',o} \geq \text{Serves}_{l,o}$$

SENDER MUST HAVE RECEIVED VIDEO

Centralized Optimization

Service Quality

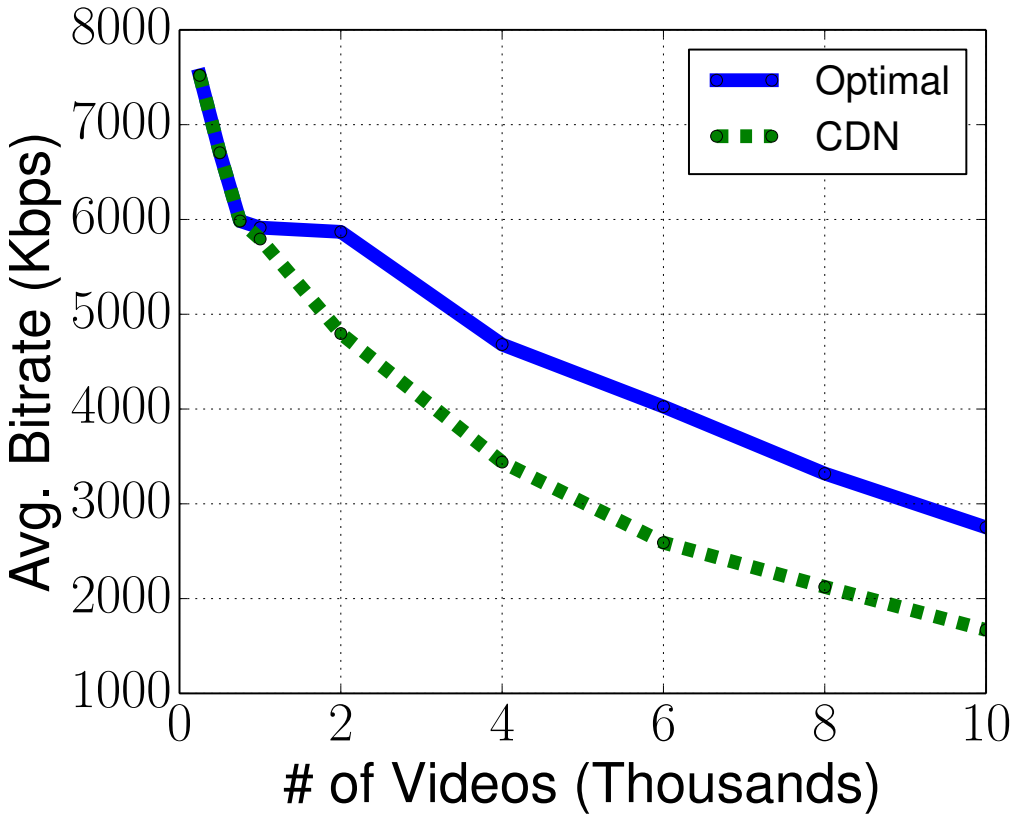

Simulation using Conviva traces, modeling user-generated content

Delivery Cost

(per request)

CDN
2.0x

OPTIMAL
1.0x

Simulation using Conviva traces, modeling large sports events

Effects of Latency in Decision Plane

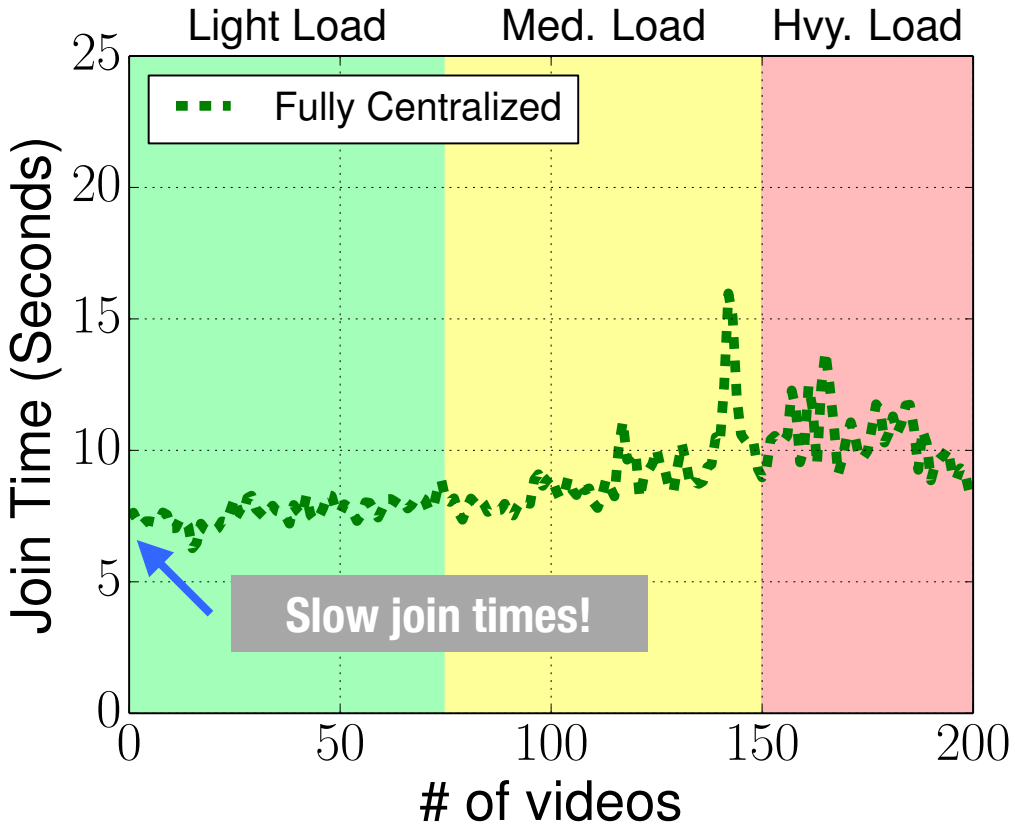

Experiments on EC2 nodes with a centralized controller at CMU across the Internet

Problems with Centralization

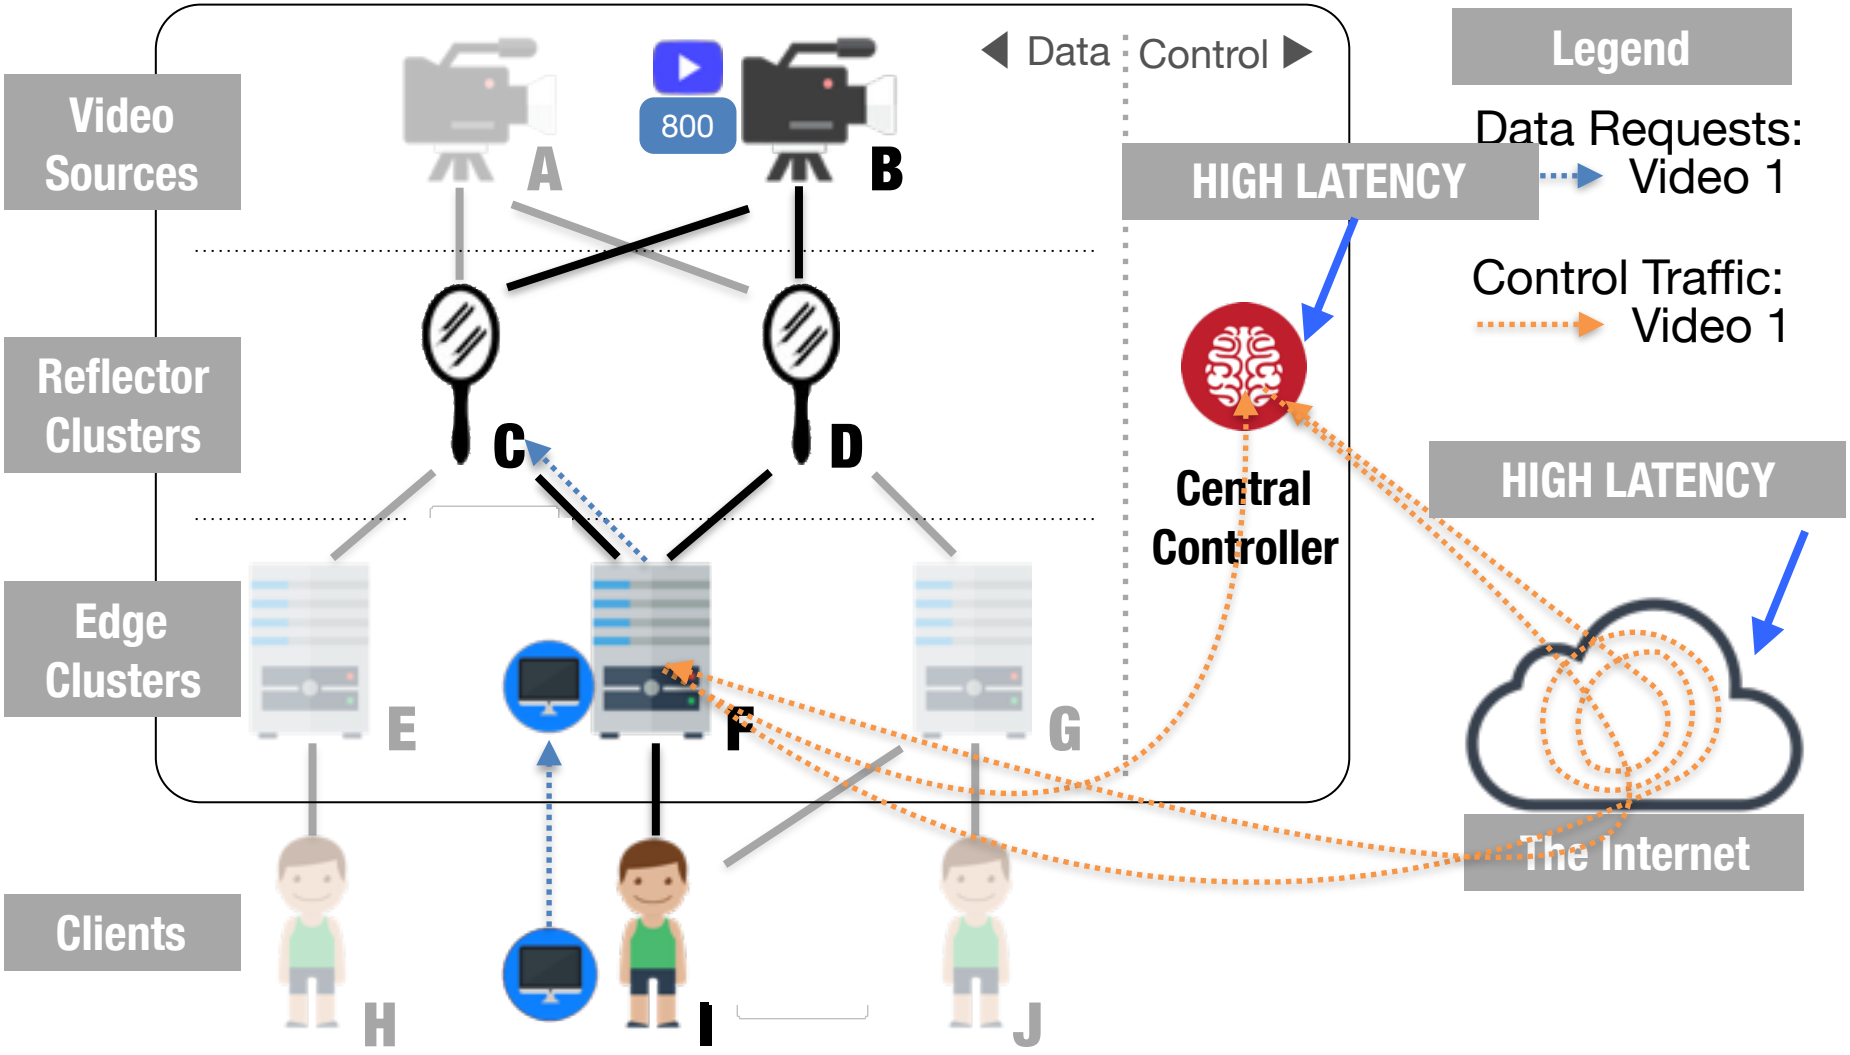

Outline

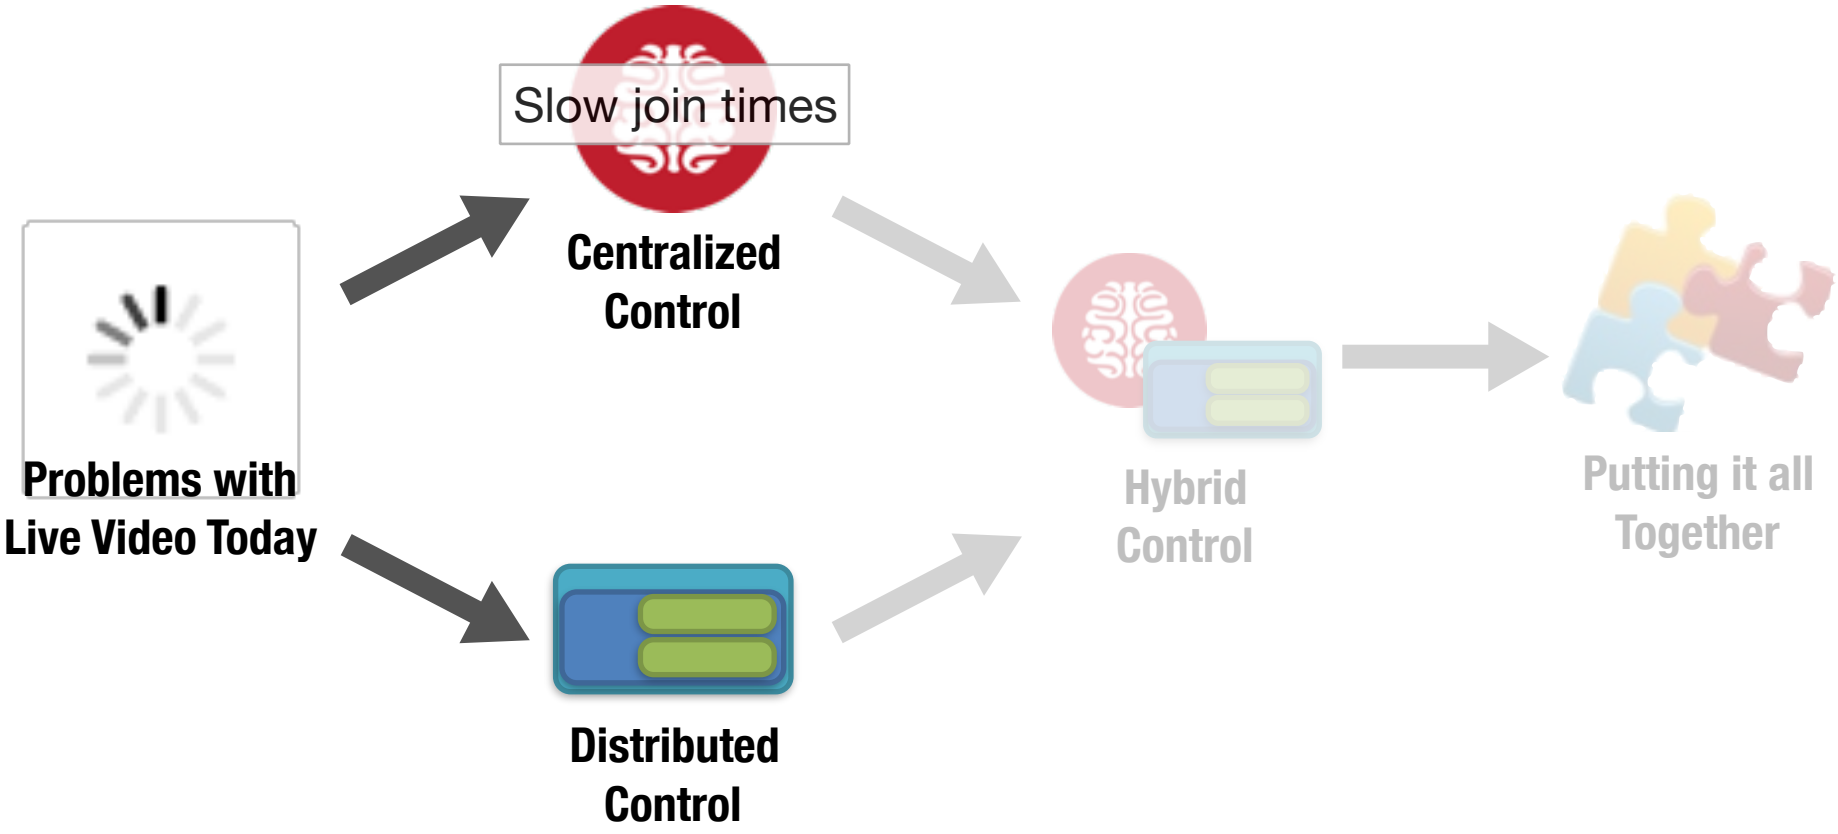
Slow join times
Centralized Control
Problems with Live Video Today
Distributed Control
Hybrid Control
Putting it all Together

Alternate Approach: Distributed

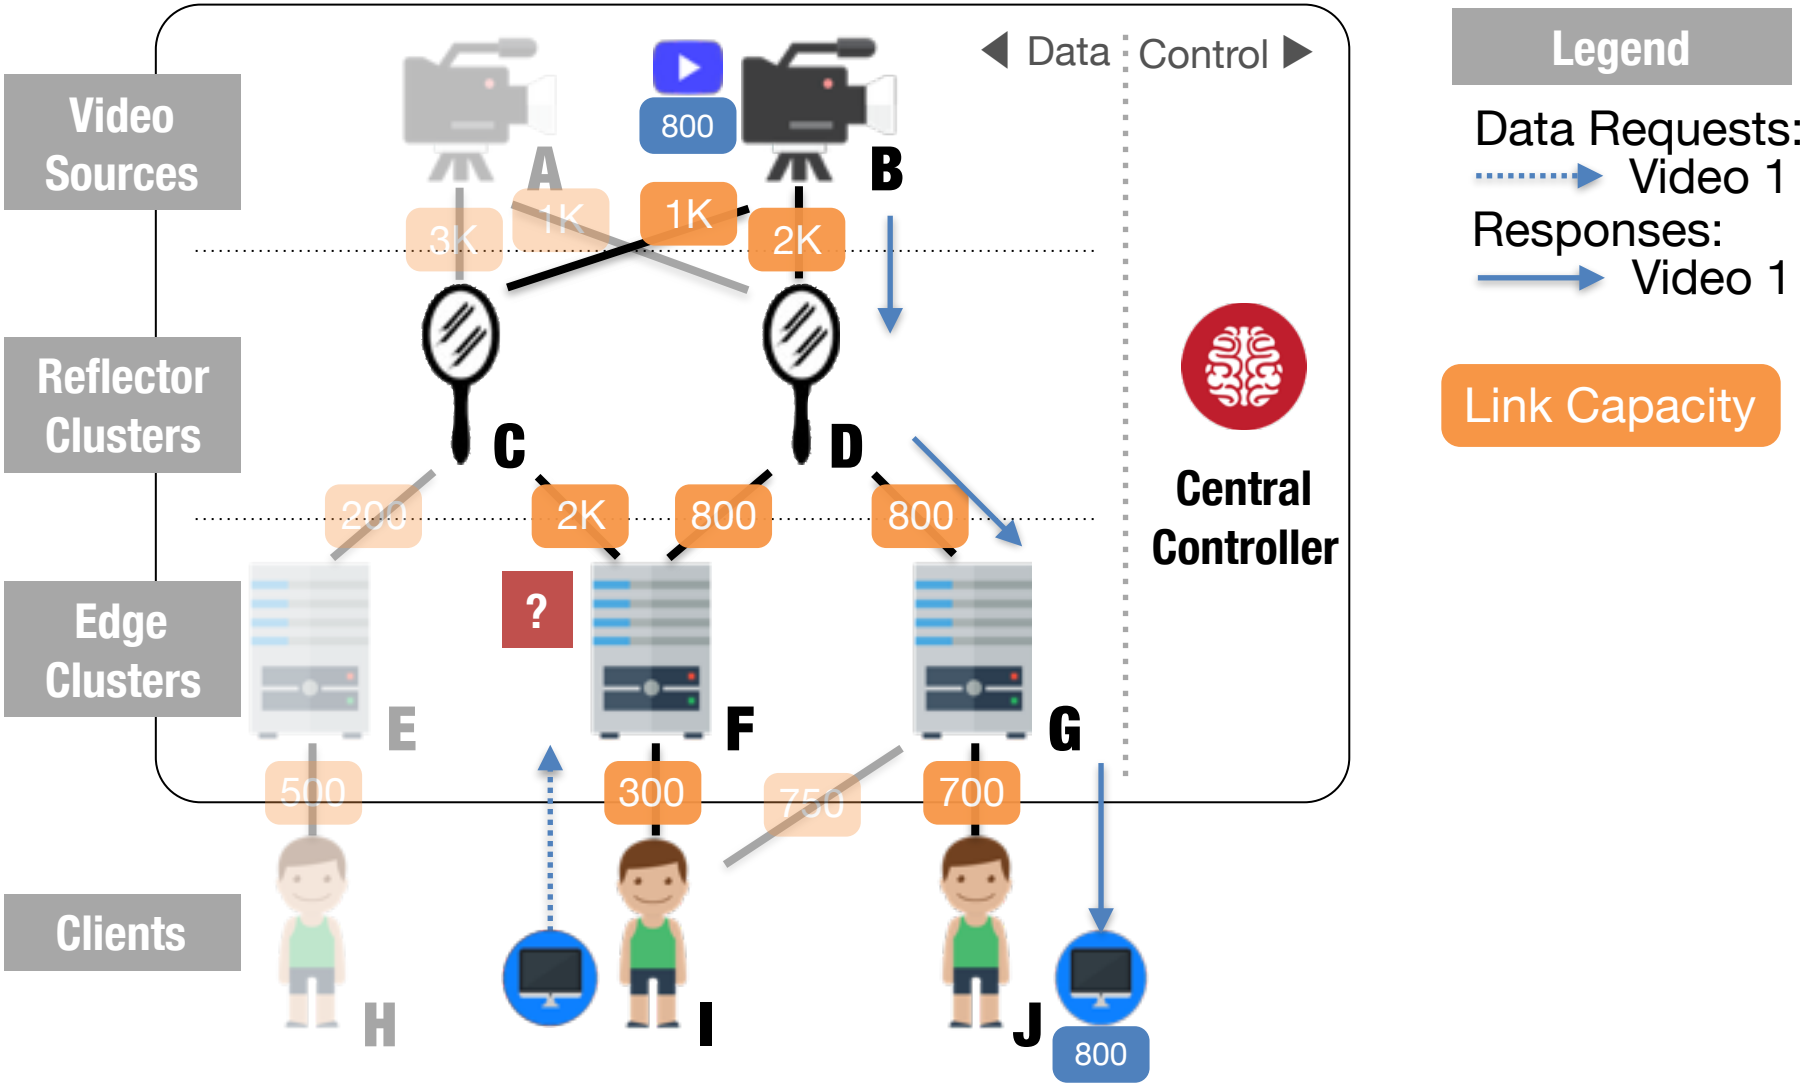

Legend

Data Requests:
Video 1

Responses:
Video 1

Link Capacity

Alternate Approach: Distributed

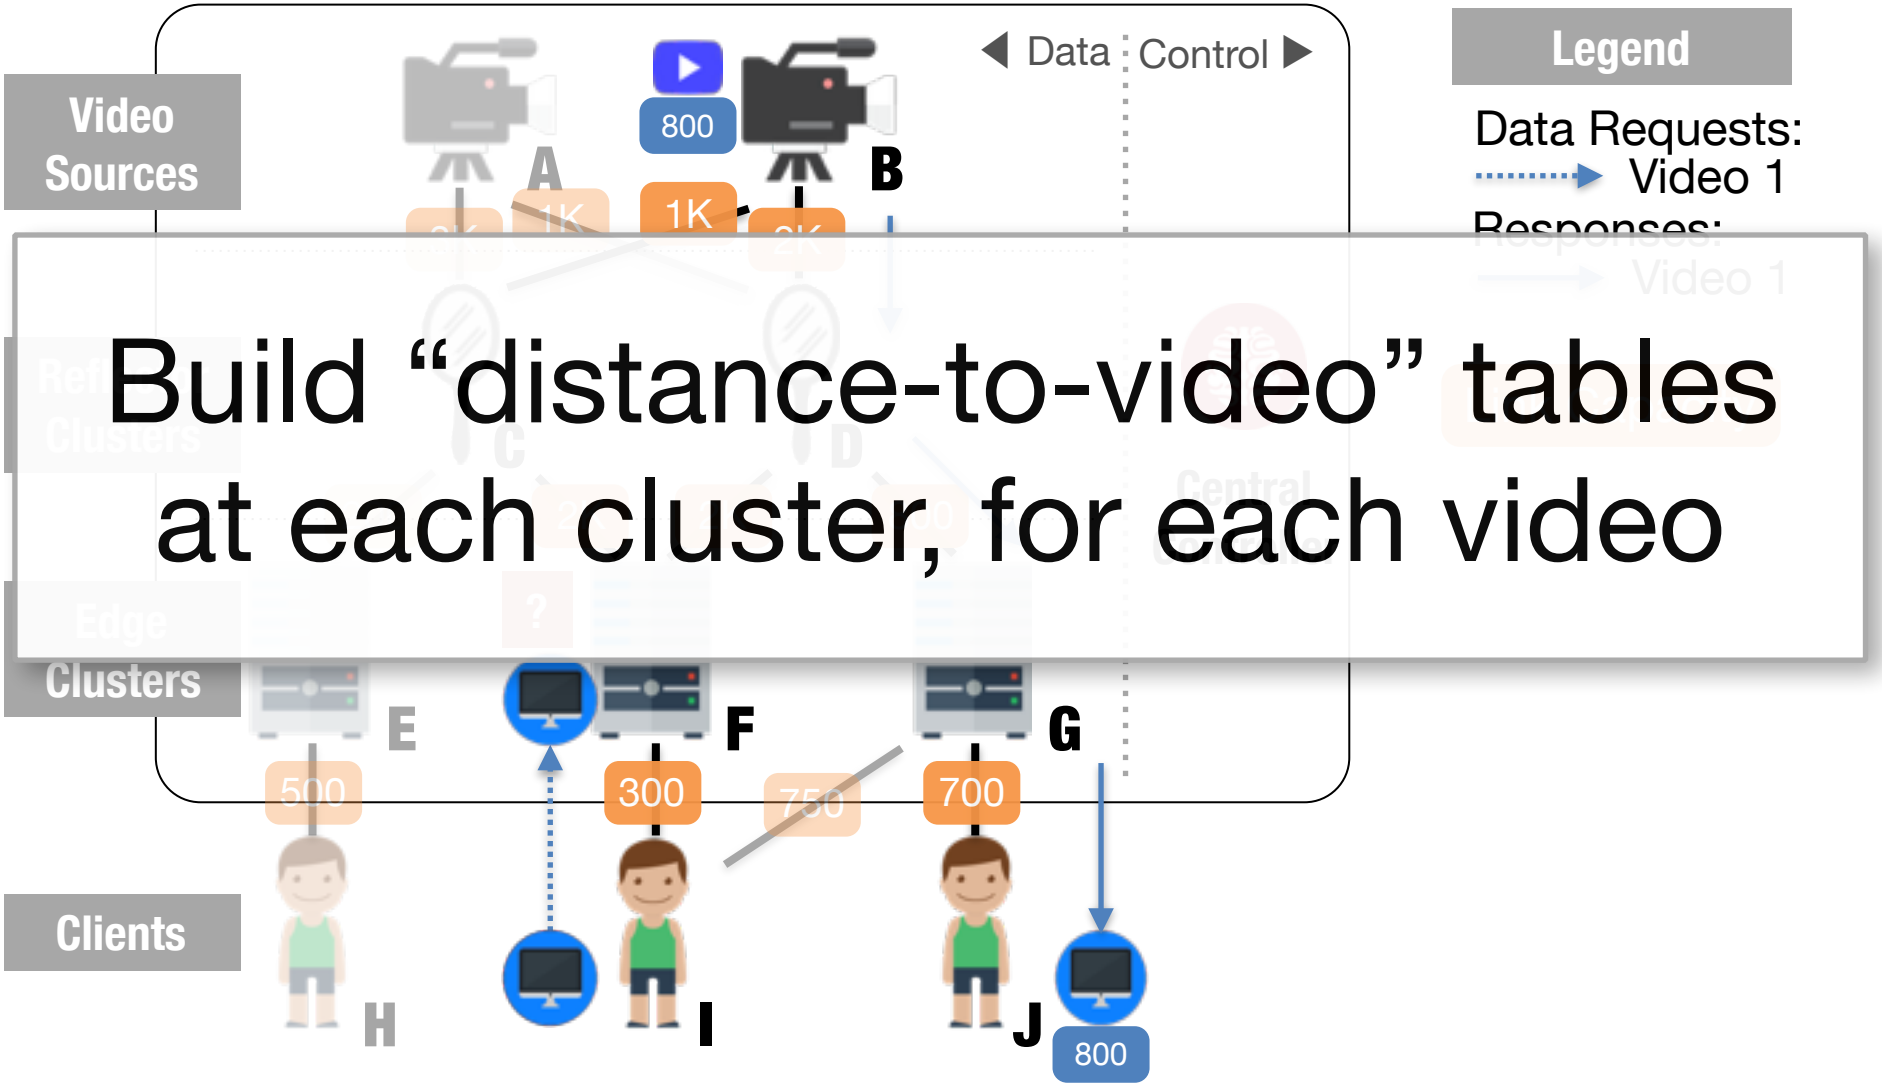

Build “distance-to-video” tables at each cluster, for each video

Alternate Approach: Distributed

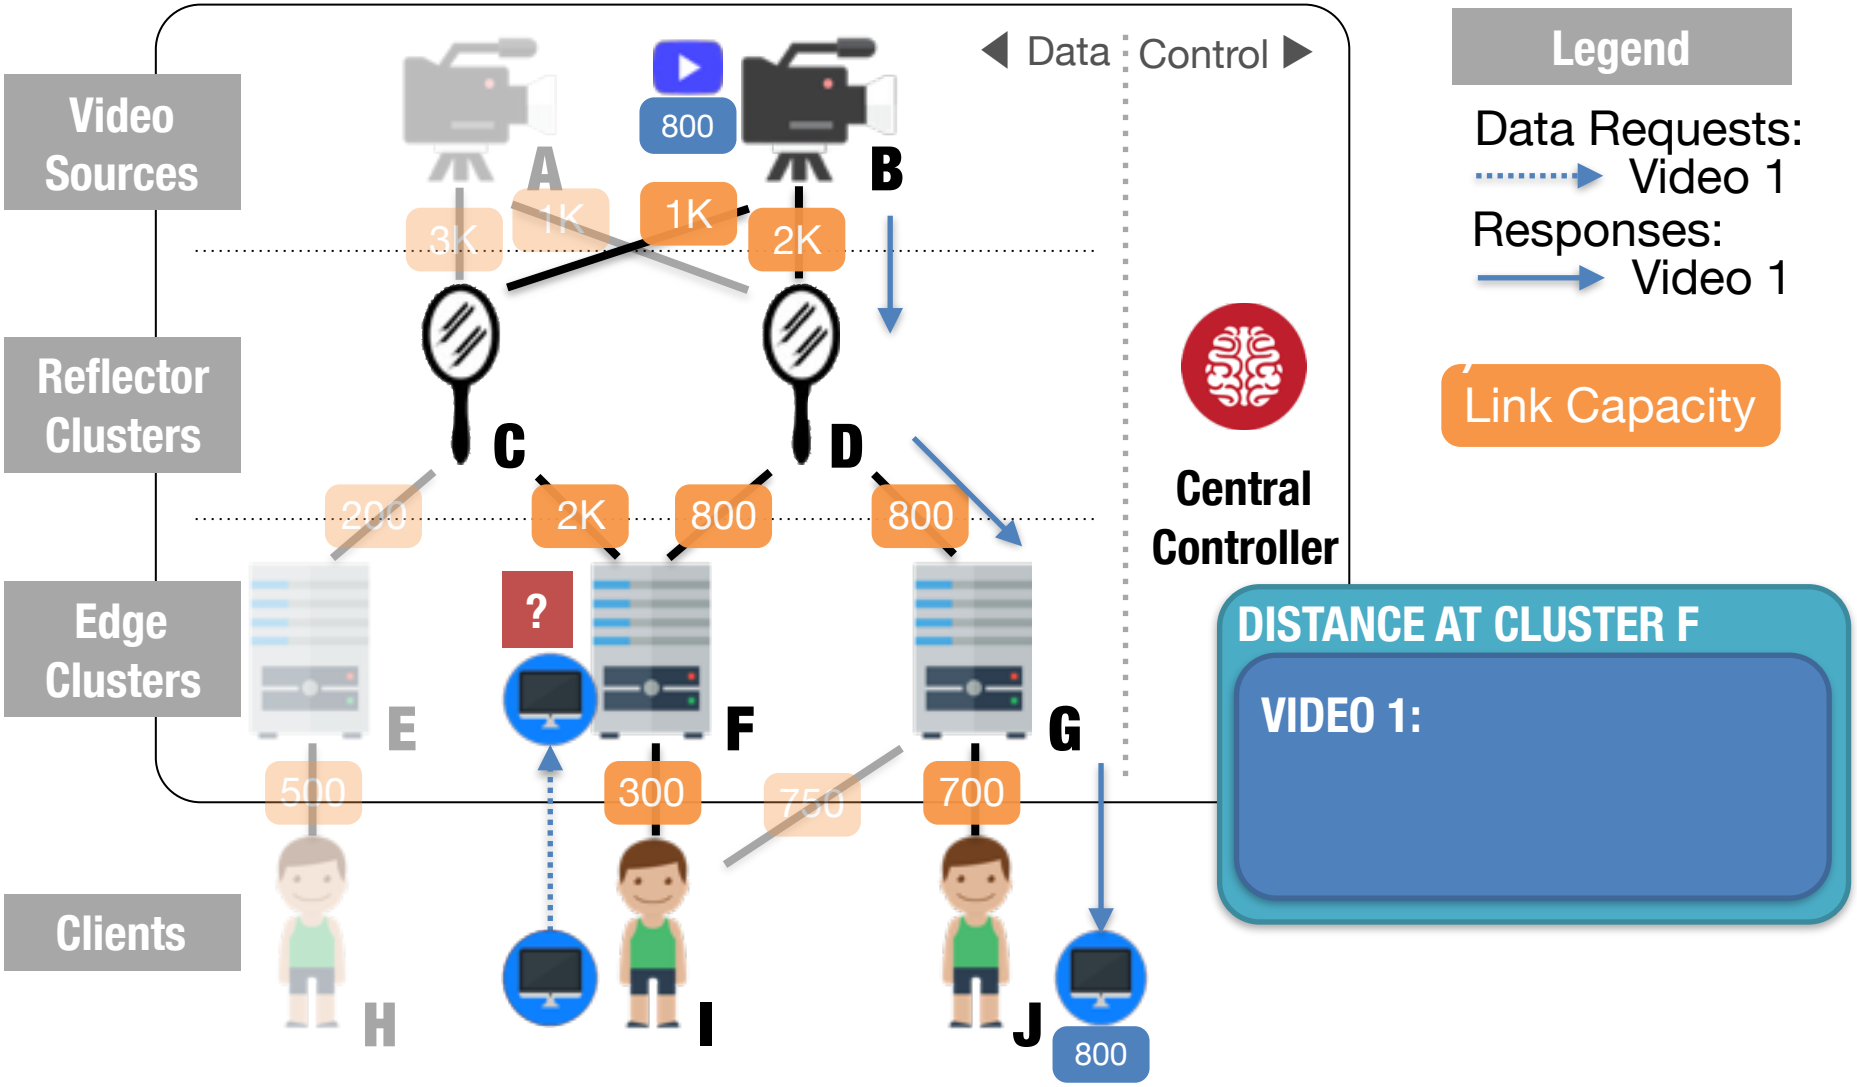

Alternate Approach: Distributed

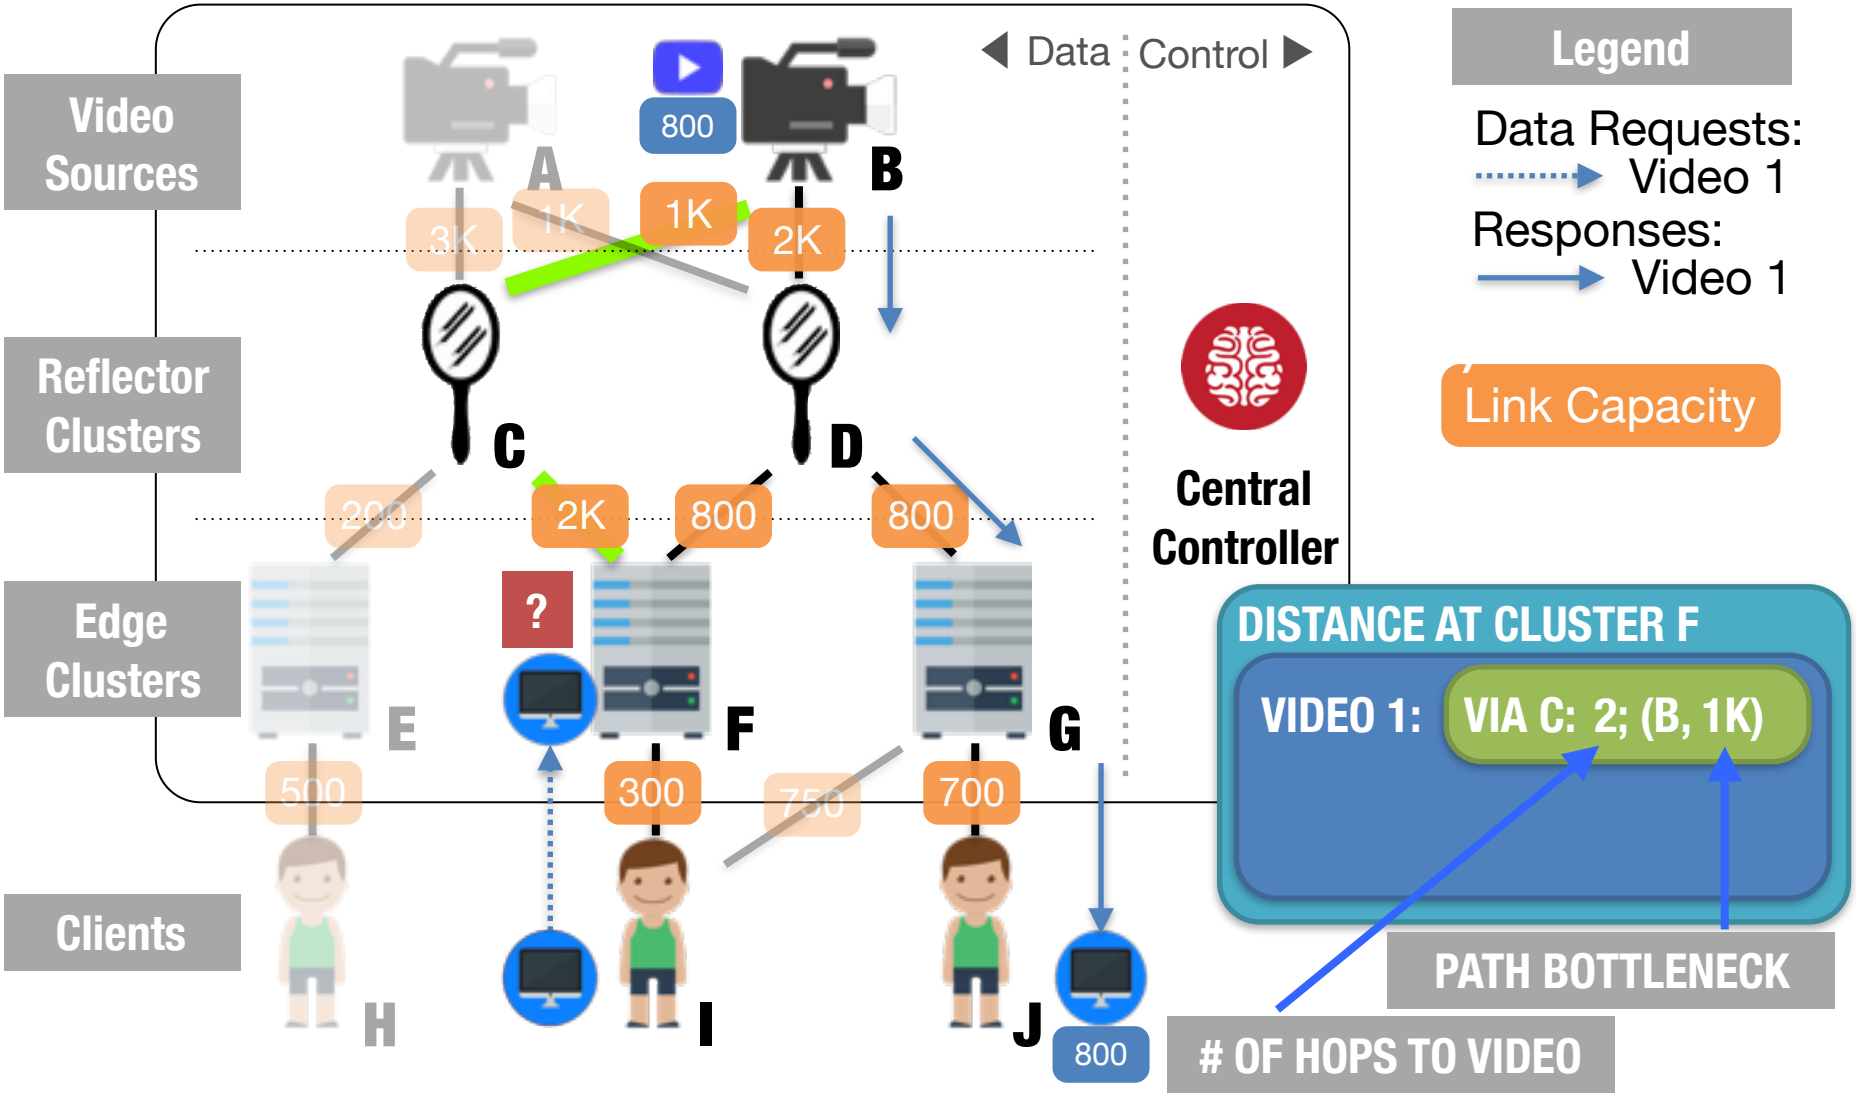

Alternate Approach: Distributed

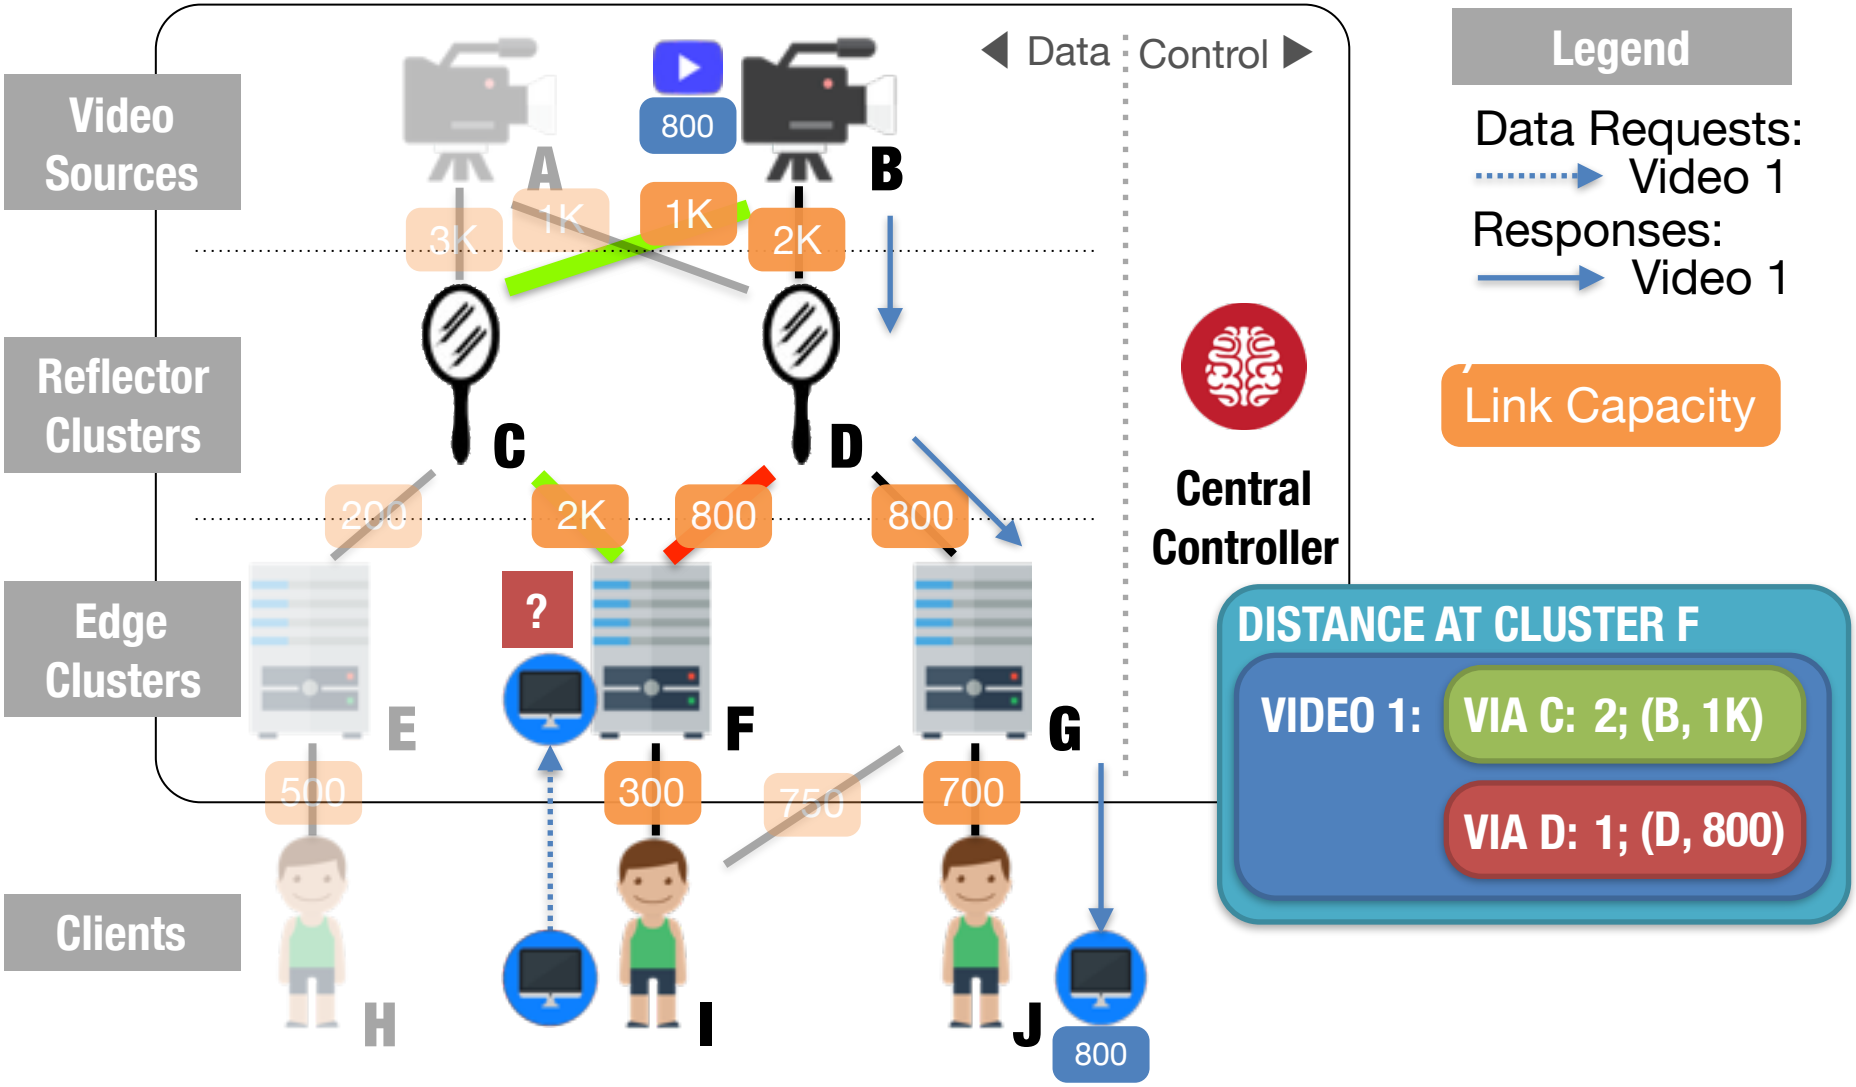

Alternate Approach: Distributed

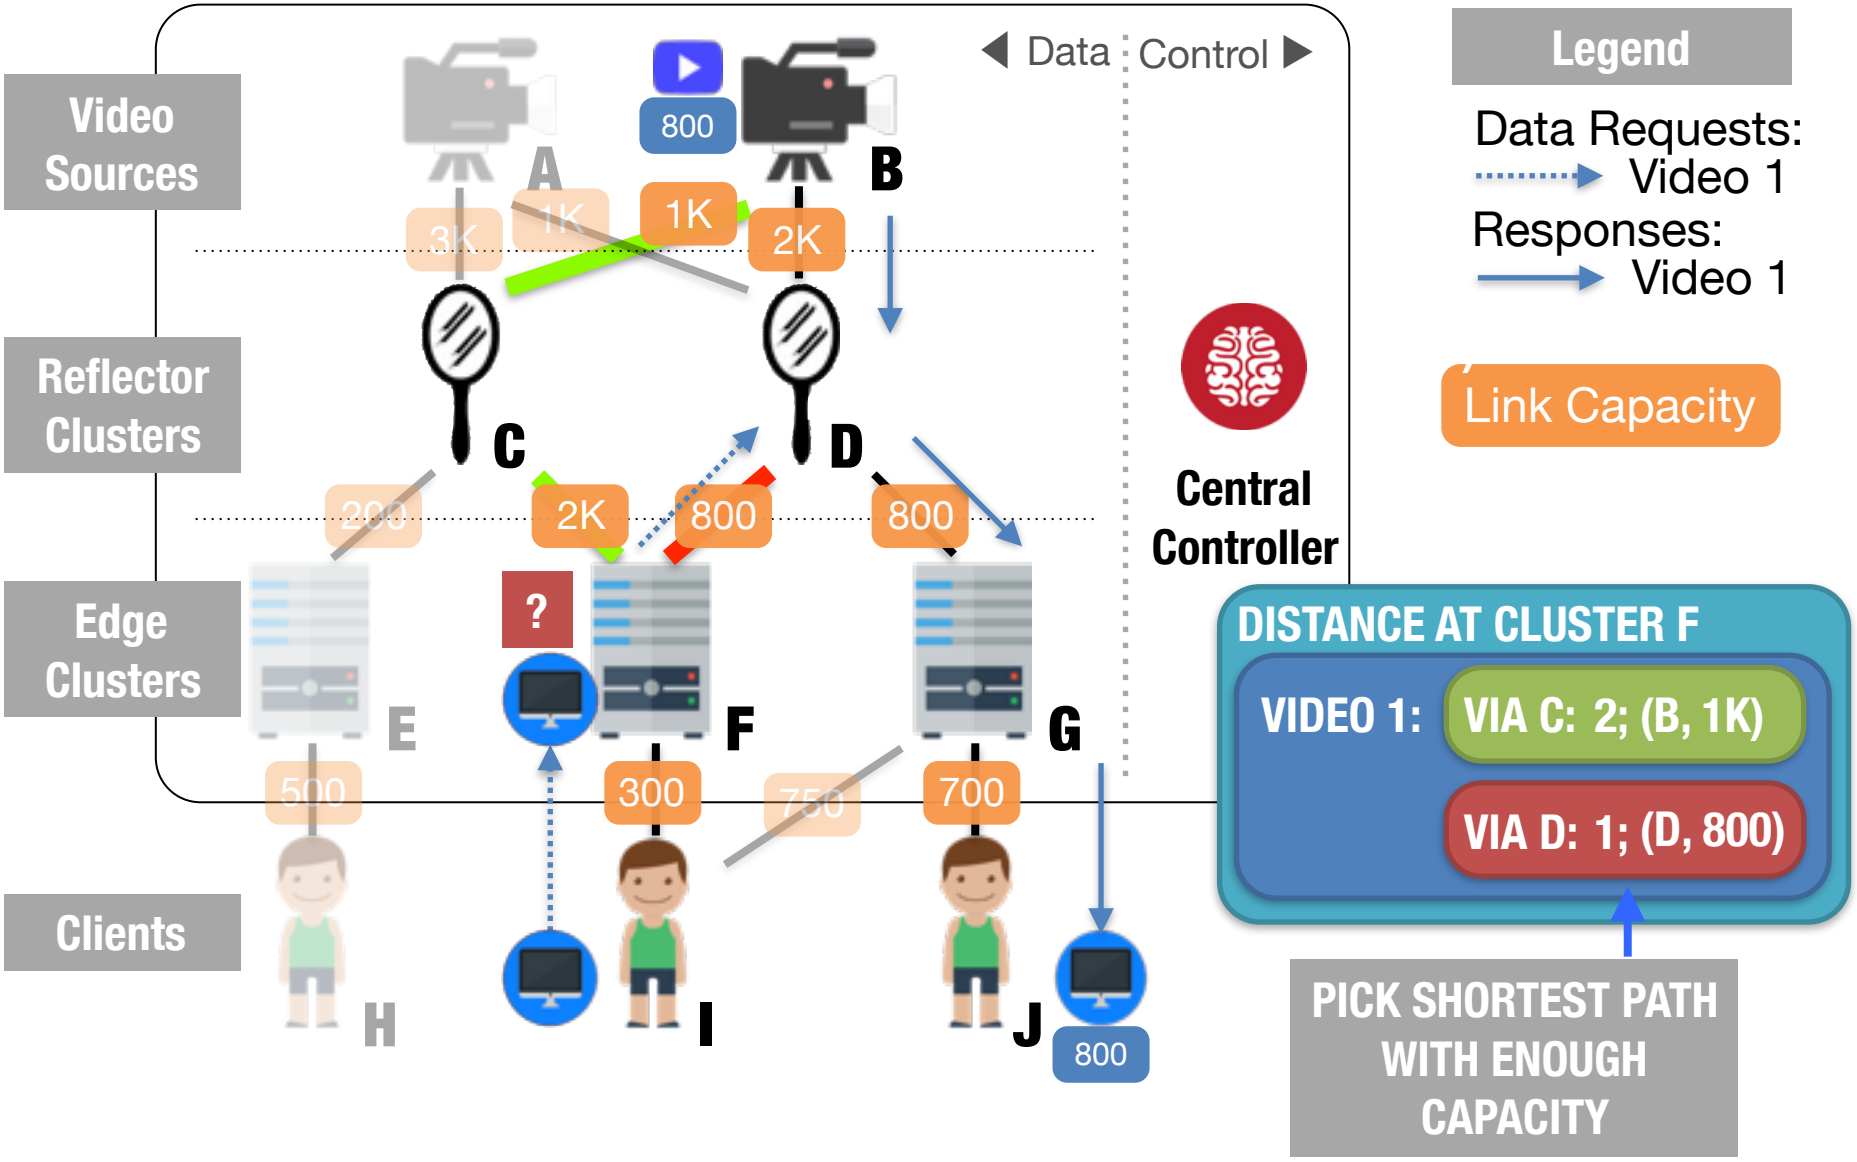

Alternate Approach: Distributed

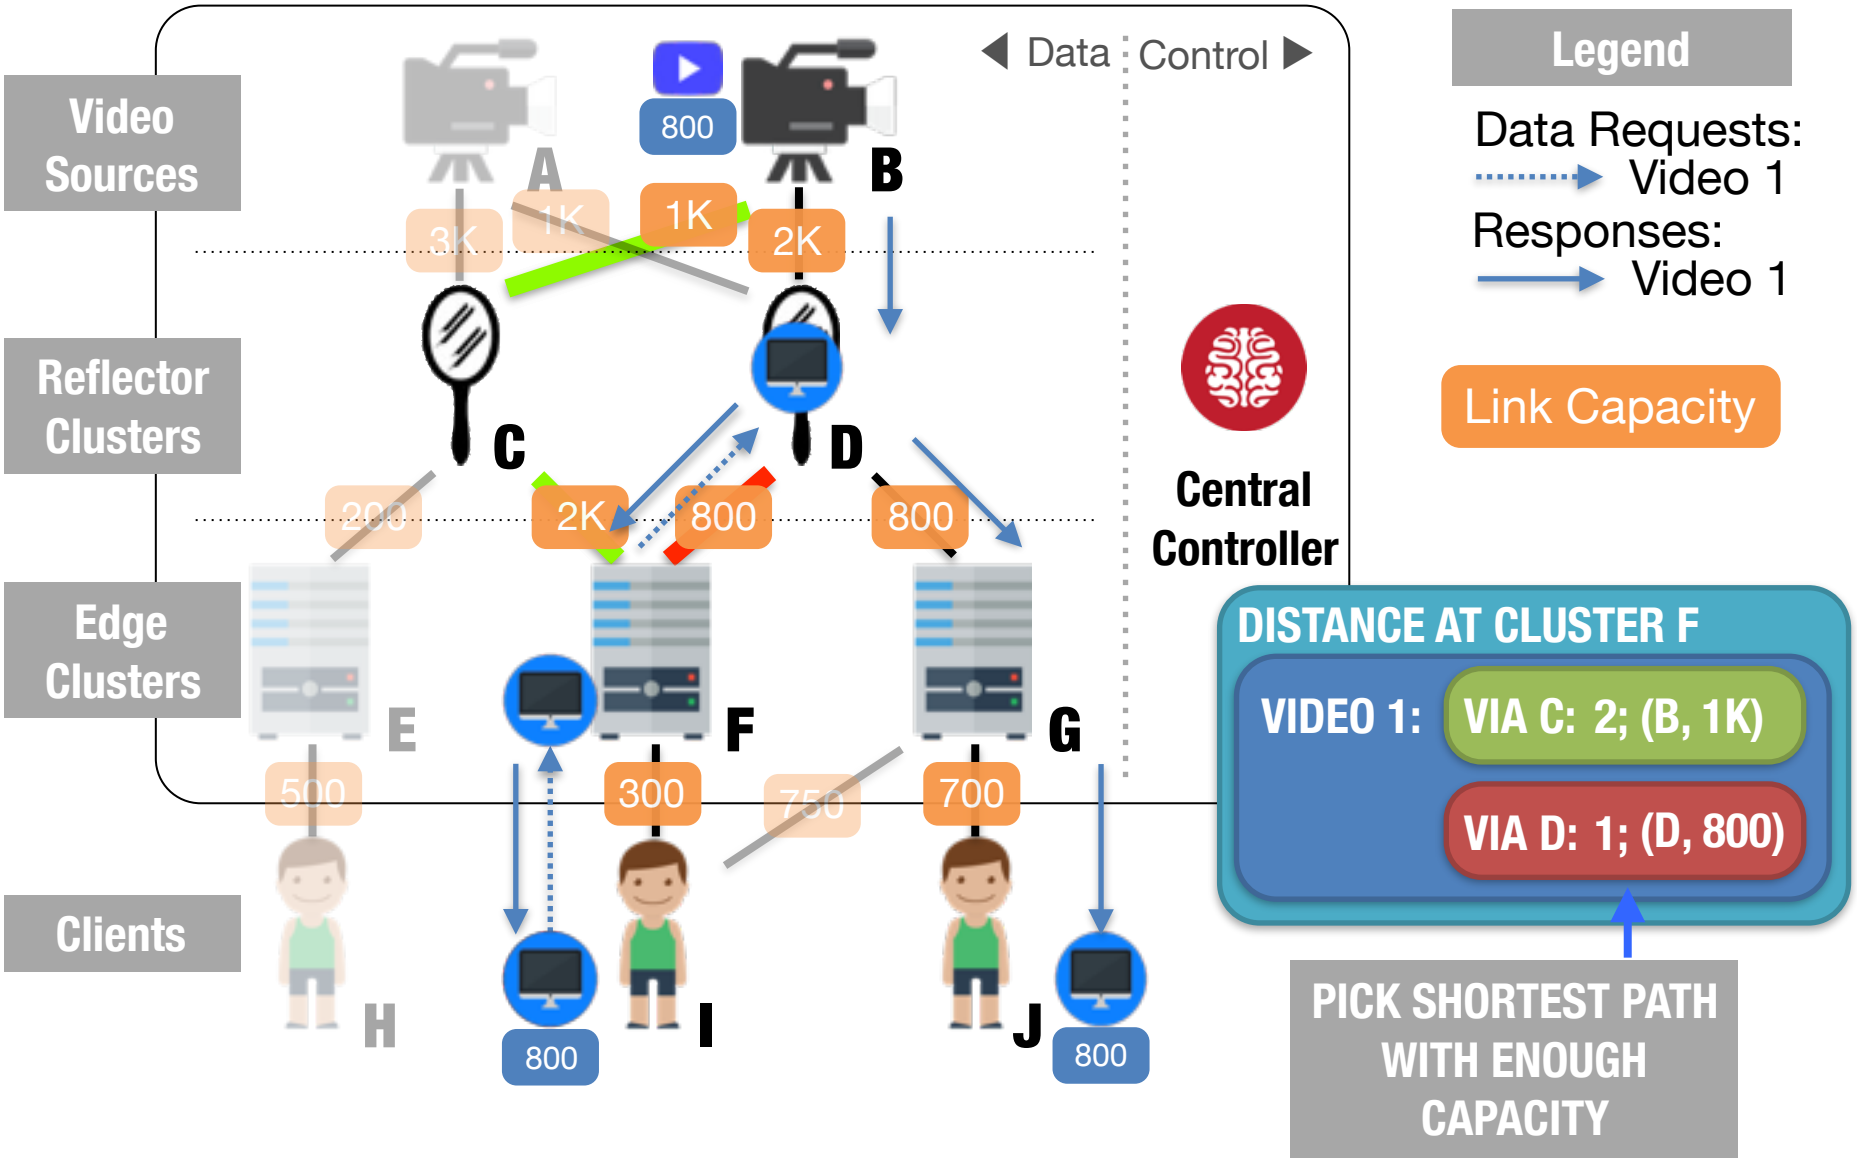

Alternate Approach: Distributed

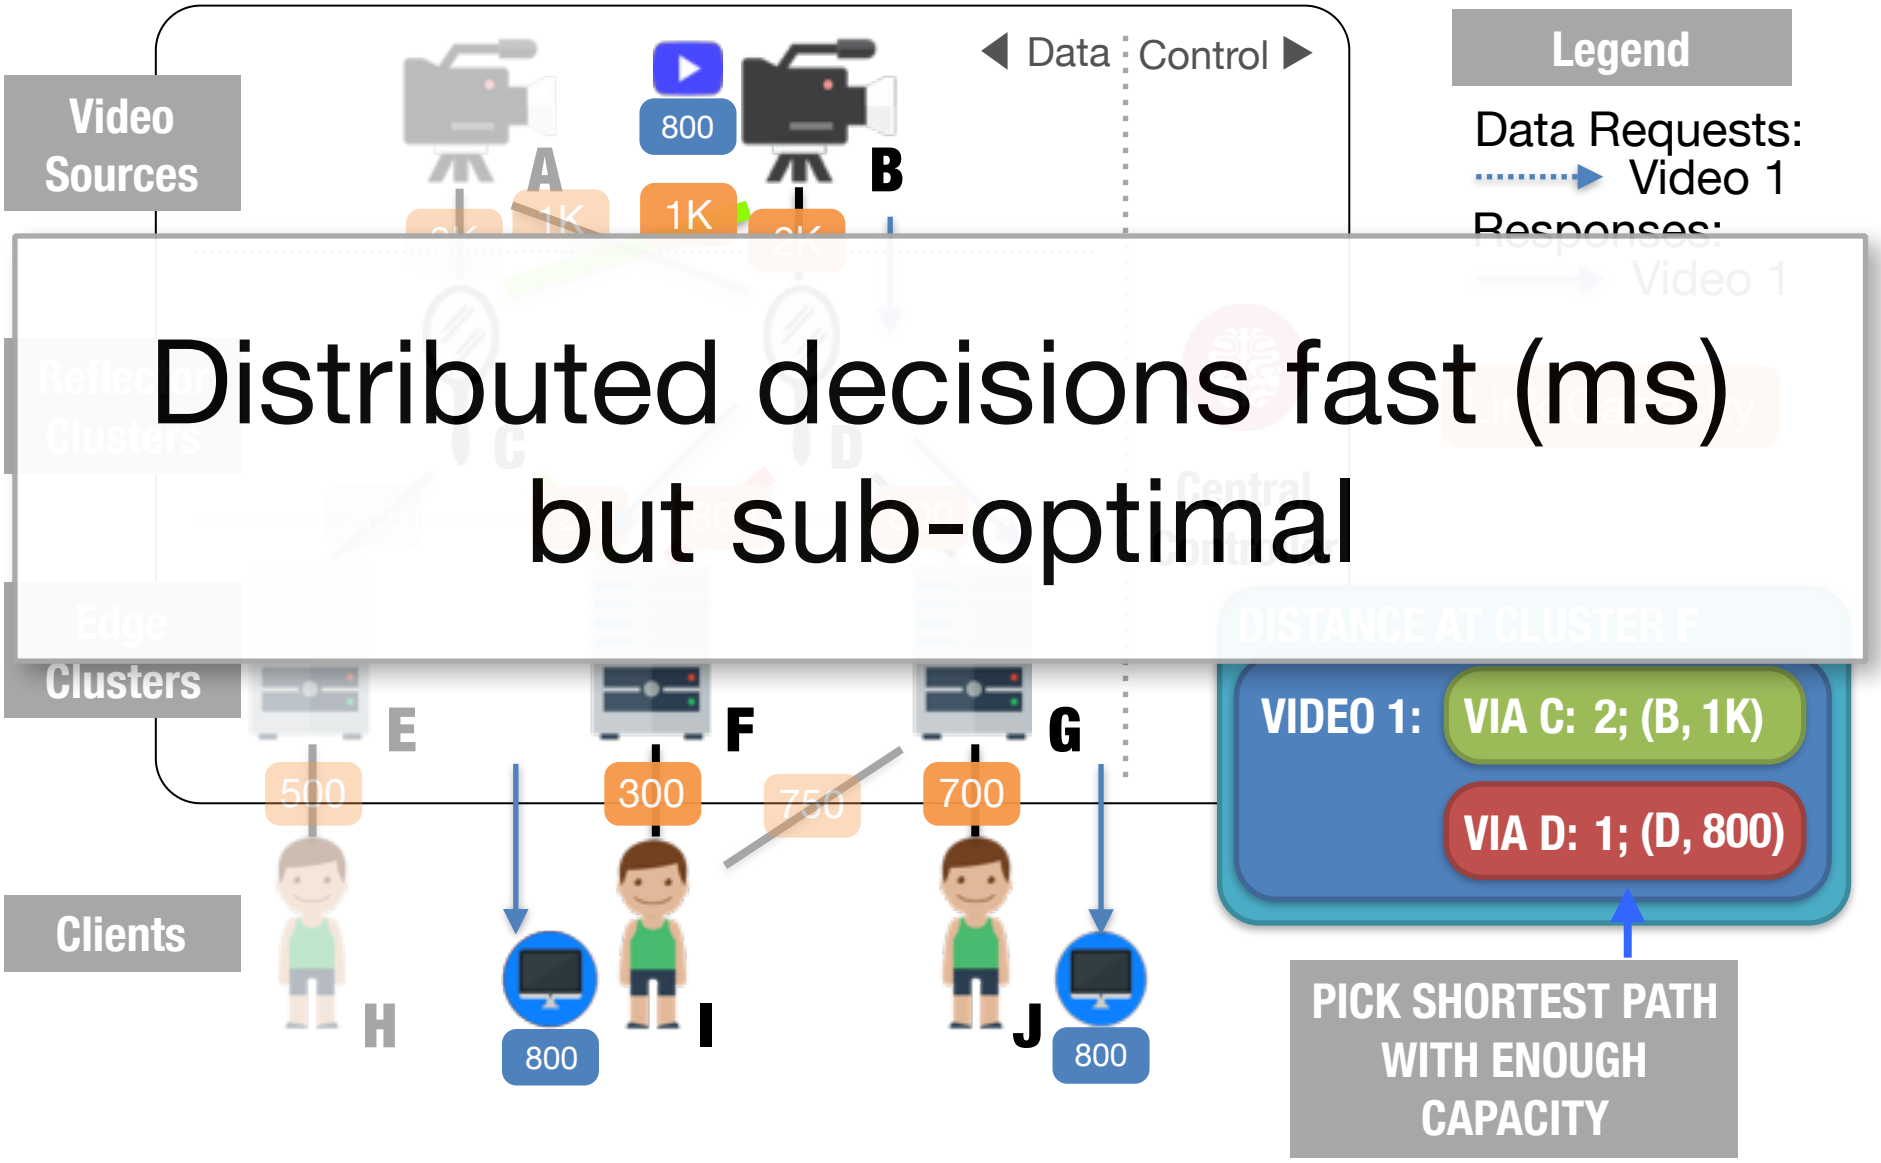

Distributed decisions fast (ms)
but sub-optimal

Alternate Approach: Distributed

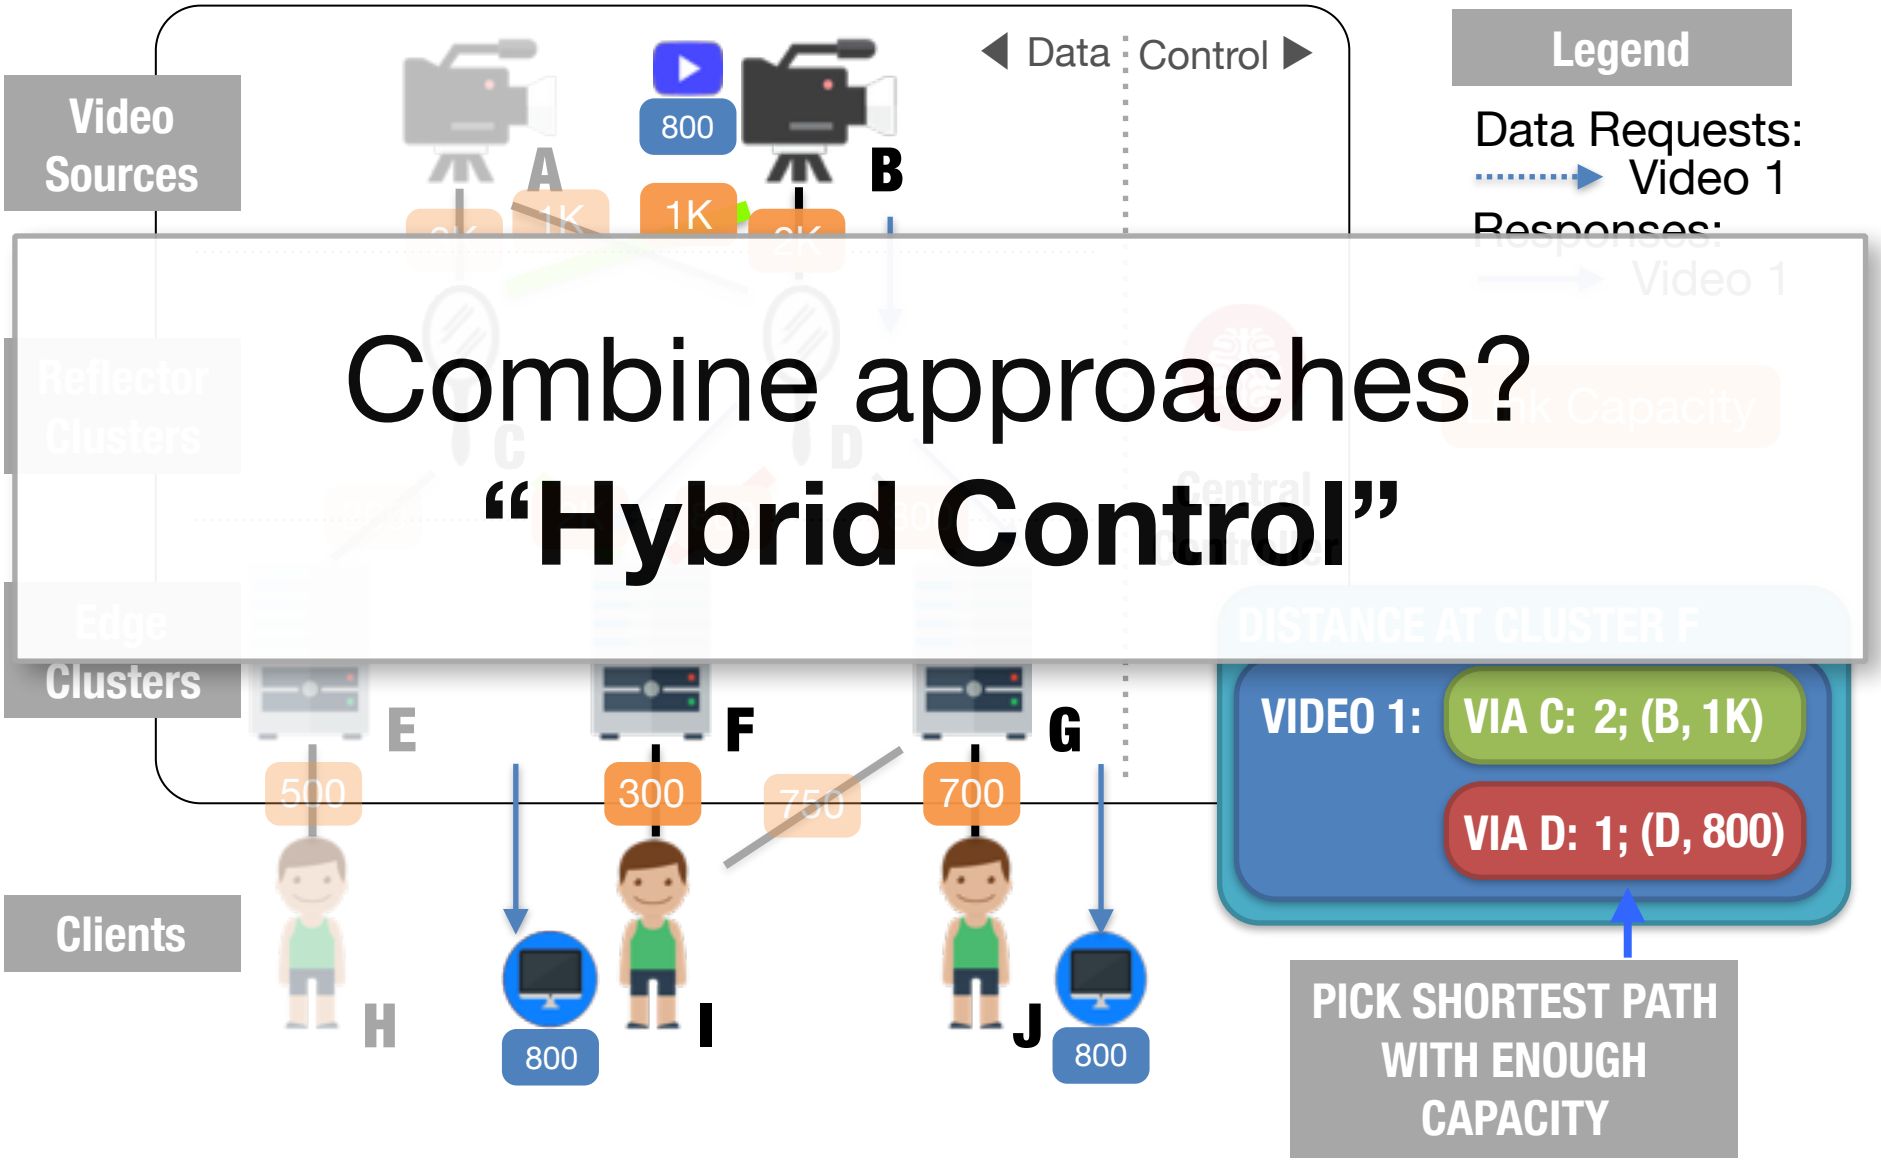

Combine approaches?

“Hybrid Control”

Outline

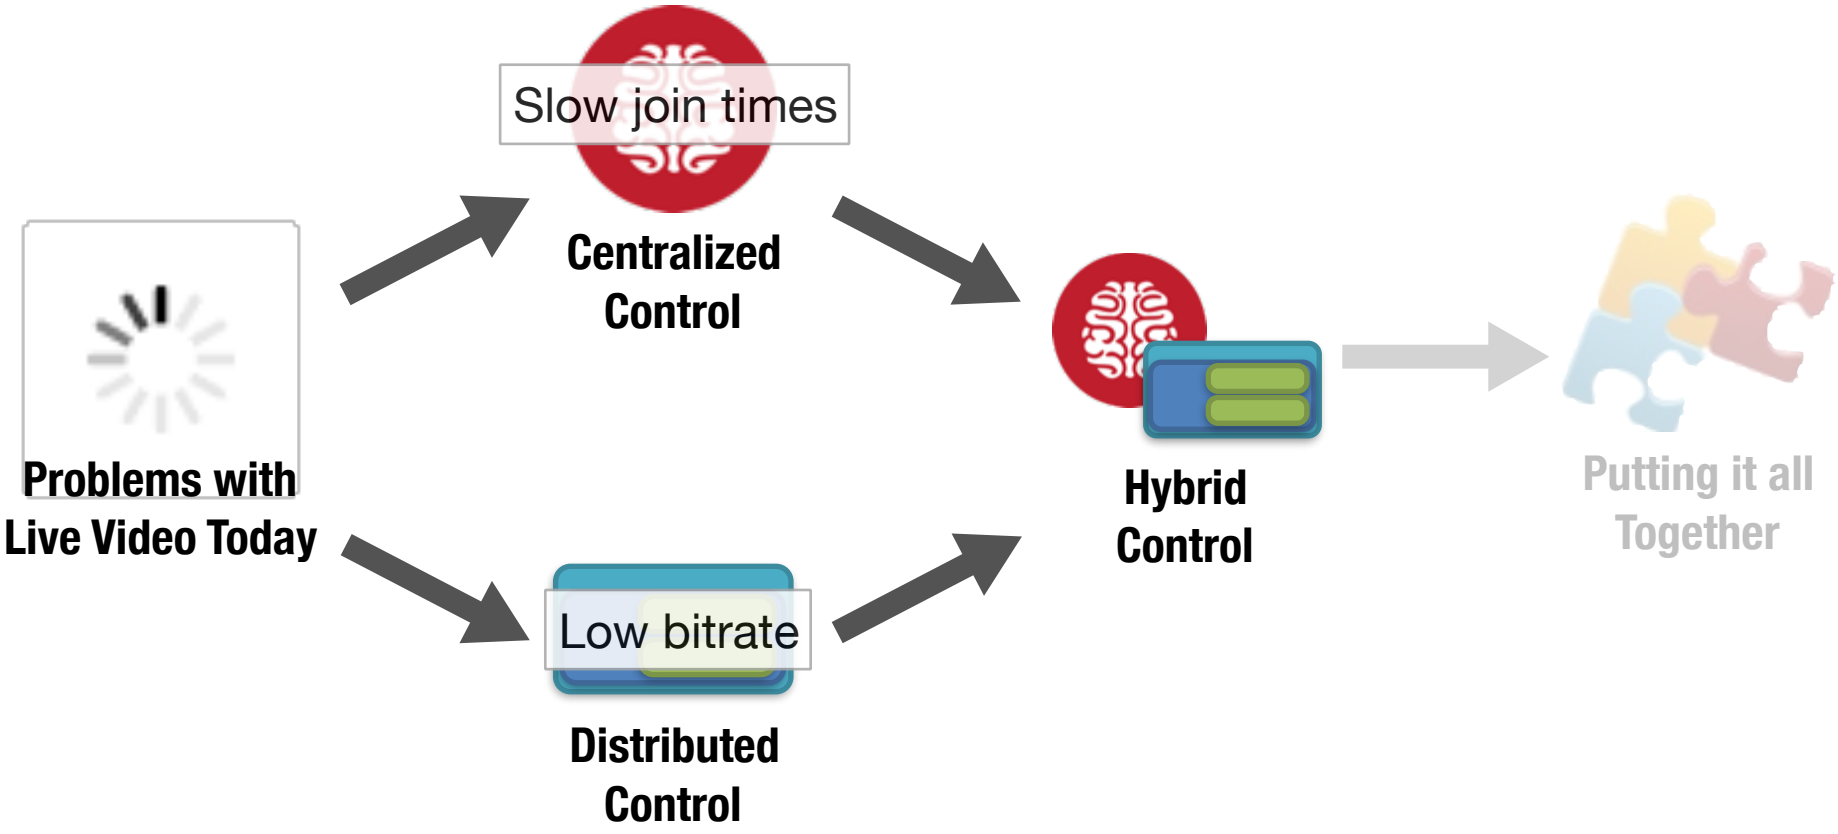
Slow join times
Centralized Control
Problems with Live Video Today
Low bitrate
Distributed Control
Hybrid Control
Putting it all Together

Combining Approaches: **Hybrid**

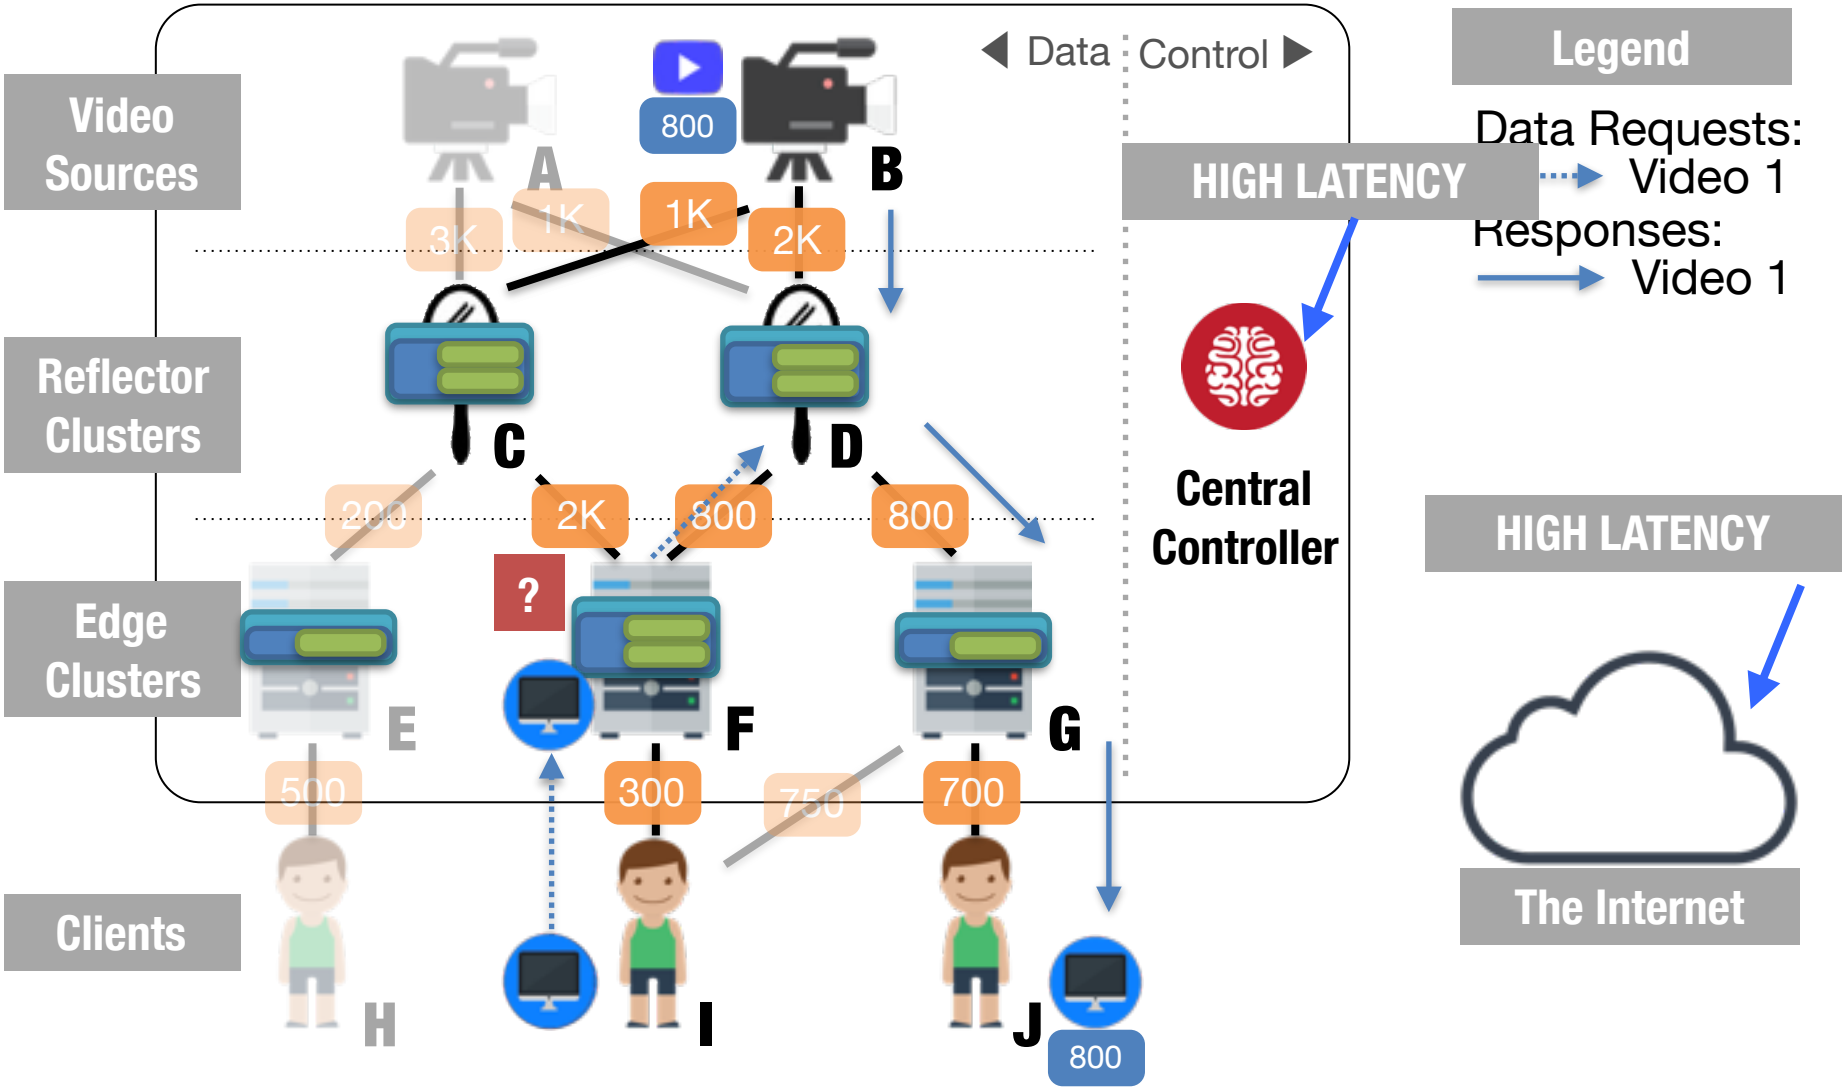

Combining Approaches: **Hybrid**

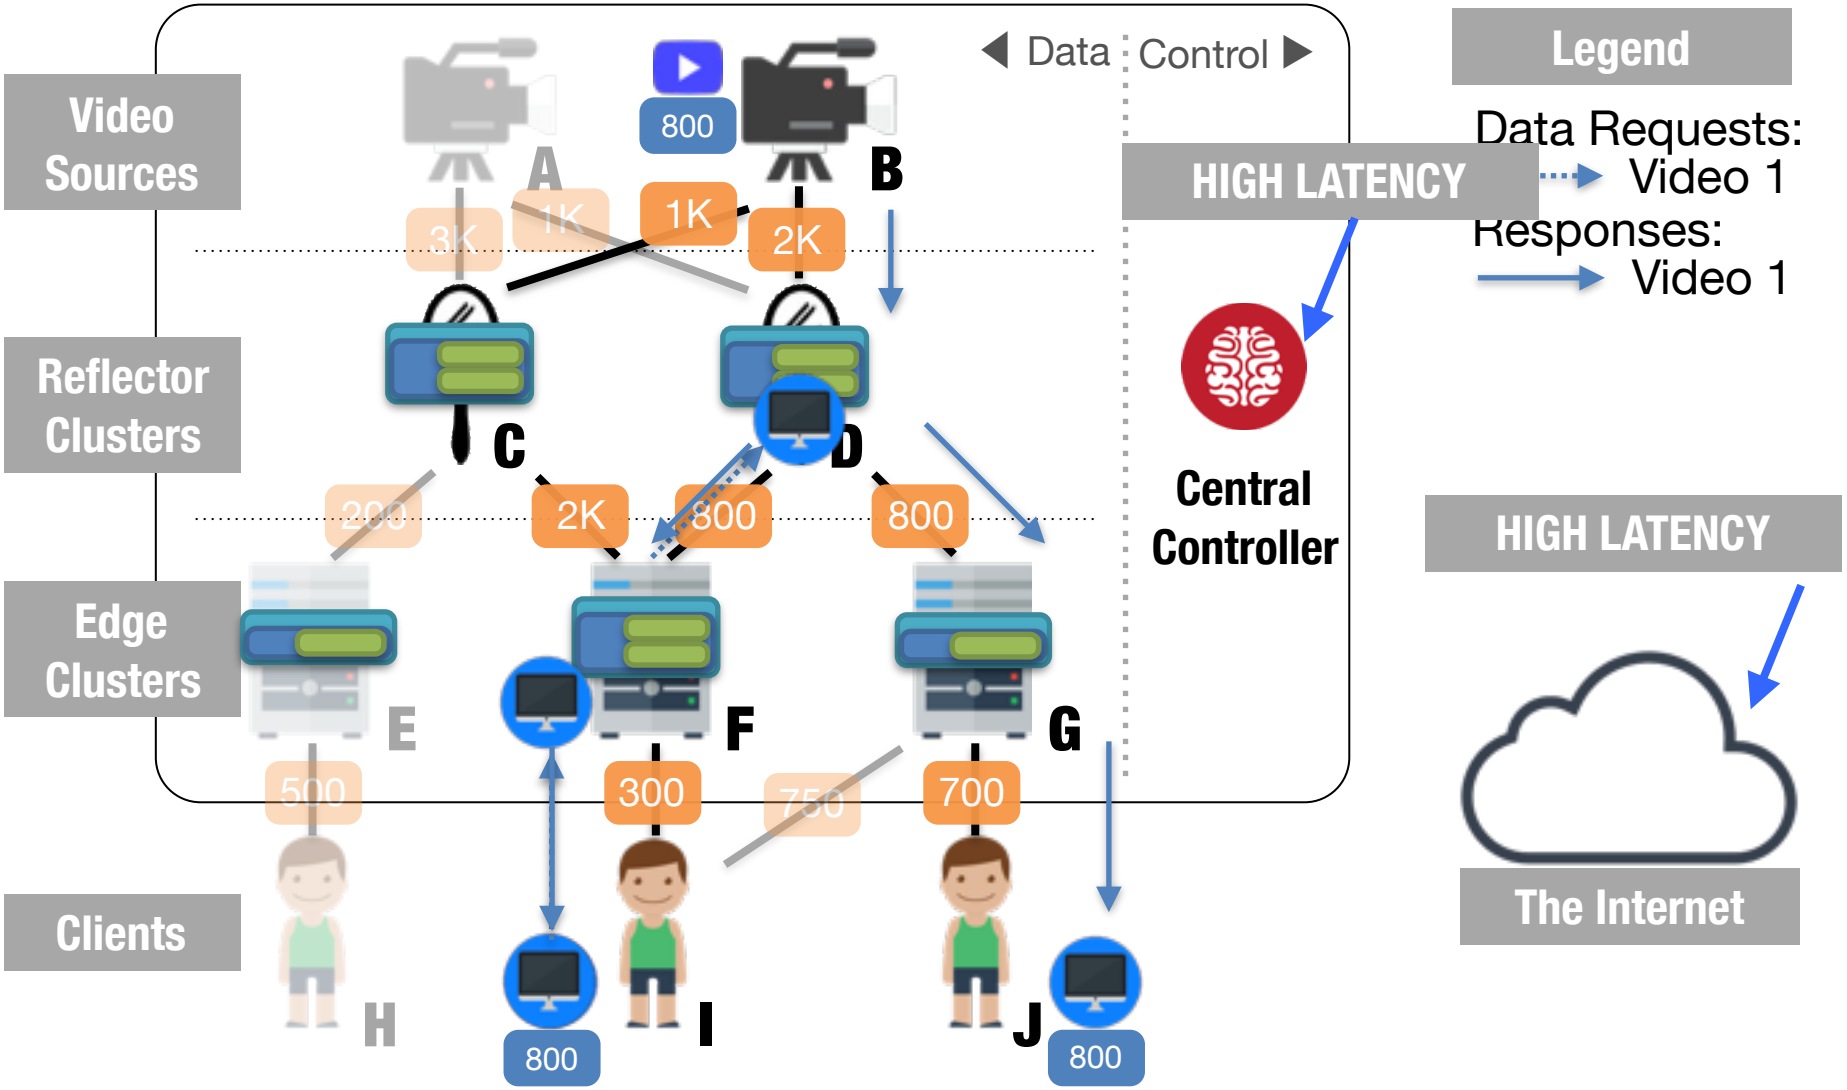

Combining Approaches: **Hybrid**

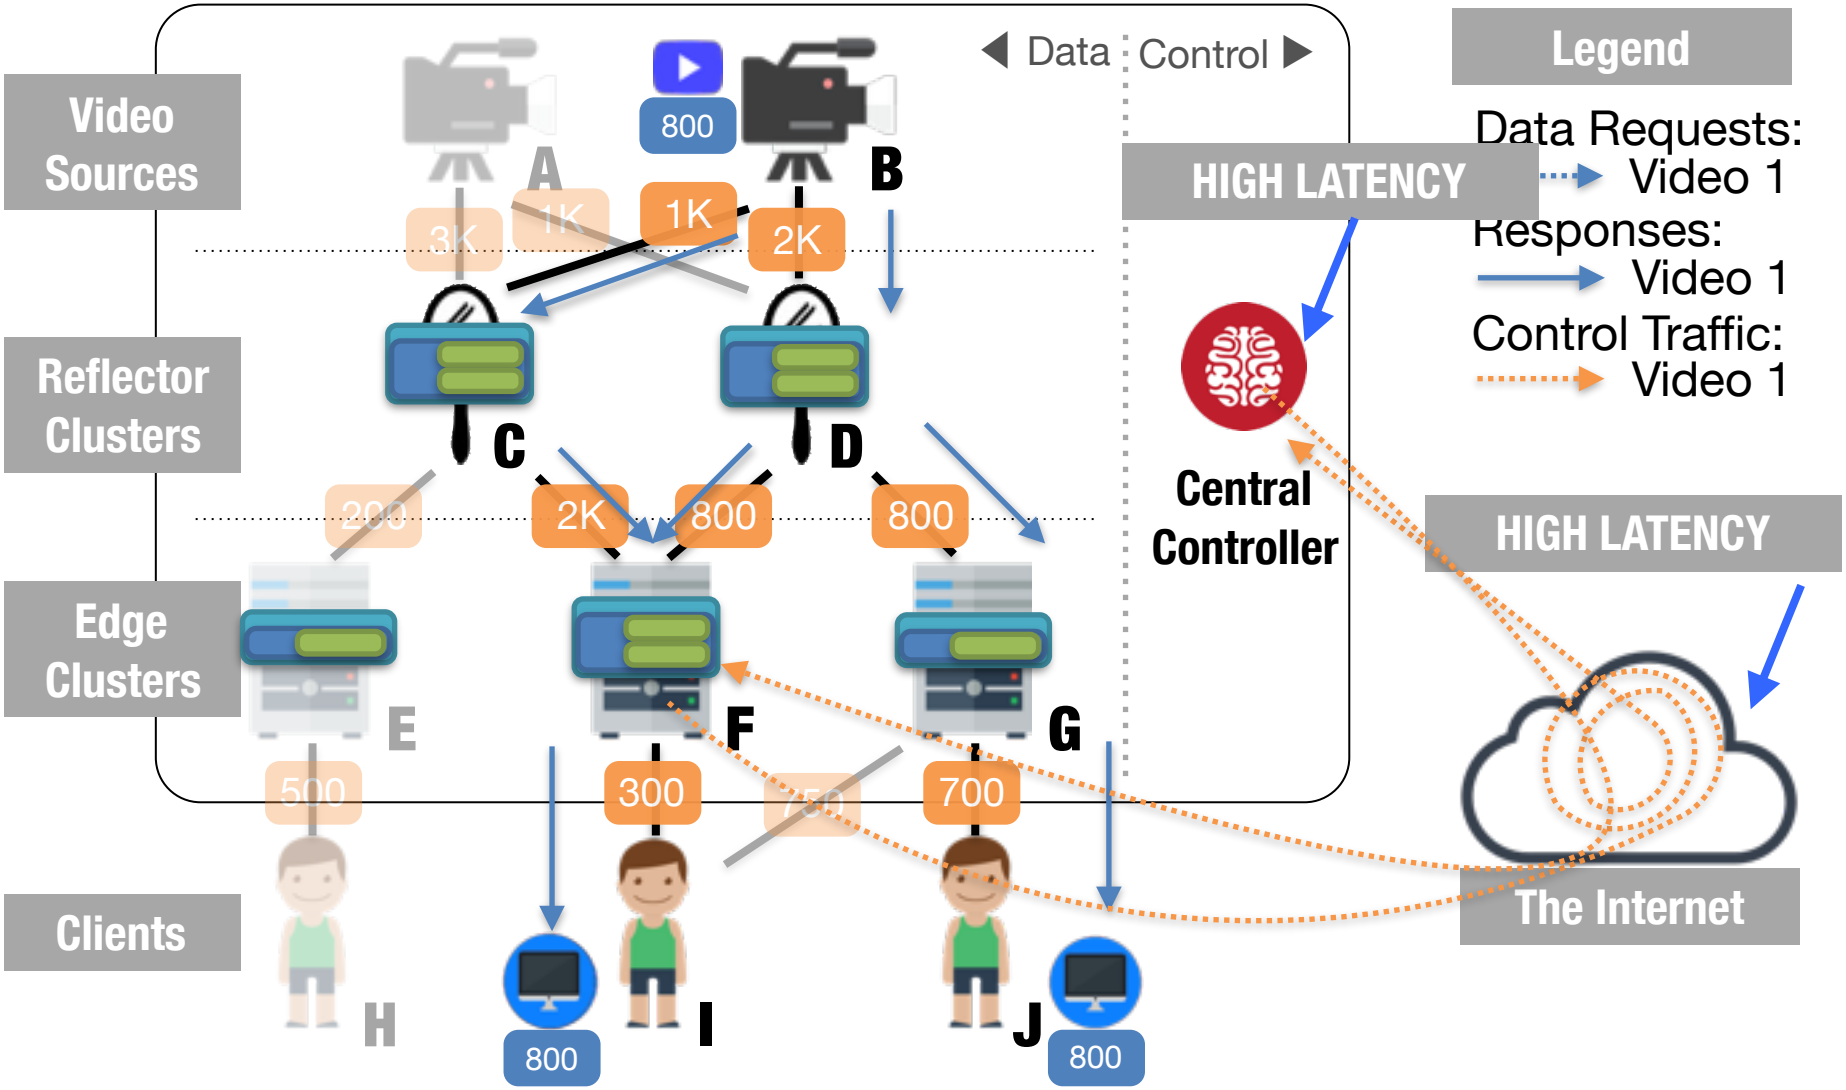

Challenges of Hybrid Control

- Forwarding loops
  - Always forward requests upwards

TRIVIAL

- State transitions
  - Versioning and “shadow FIBS”

PRIOR WORK

- Avoid bad control loop interactions

CHALLENGING

Challenges of Hybrid Control

- Avoid bad control loop interactions

CHALLENGING

1. Centralized decision has priority
2. Distributed uses residual after centralized
3. Distributed has no impact on current/future centralized decisions
4. Distributed's changes don't propagate

Hybrid Control and Responsiveness

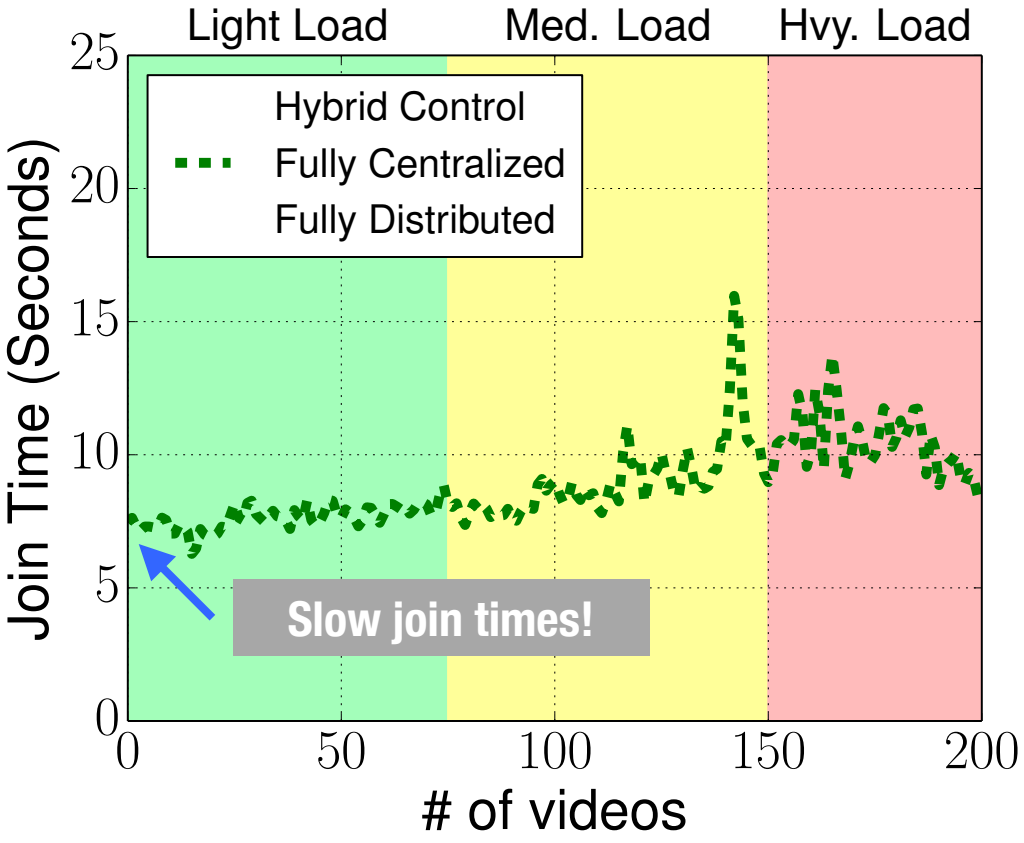

Experiments on EC2 nodes with a centralized controller at CMU across the Internet

Hybrid Control and Responsiveness

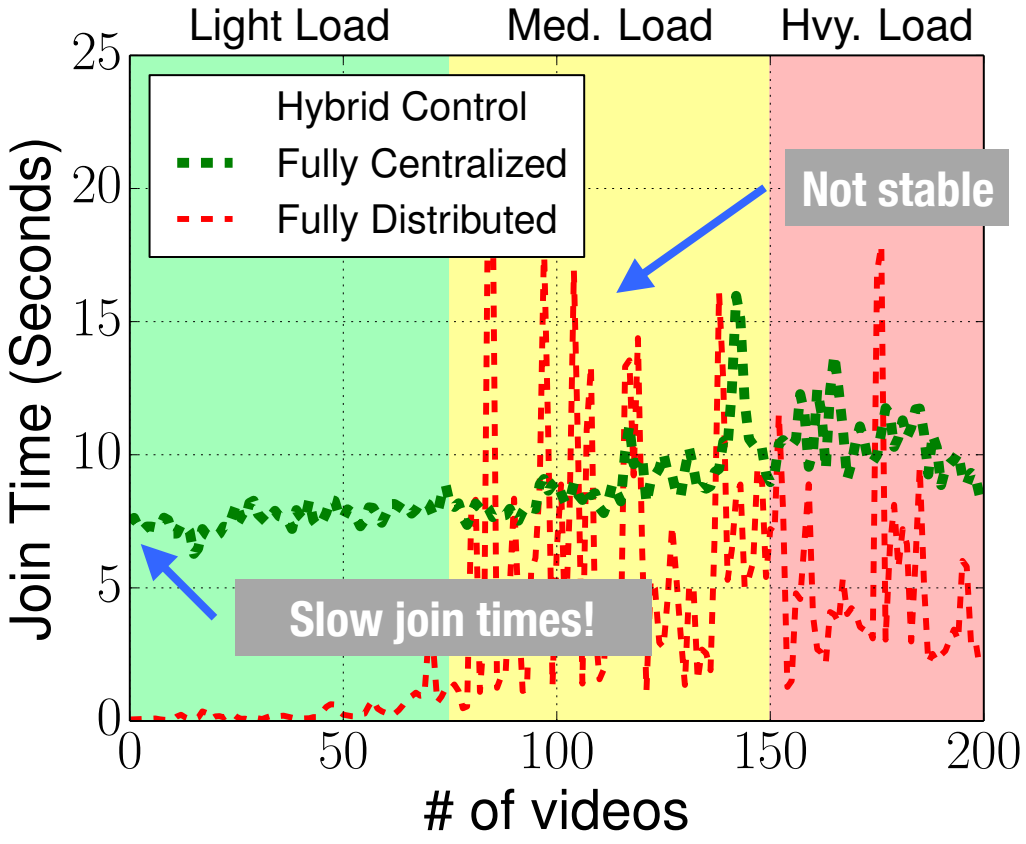

Experiments on EC2 nodes with a centralized controller at CMU across the Internet

Hybrid Control and Responsiveness

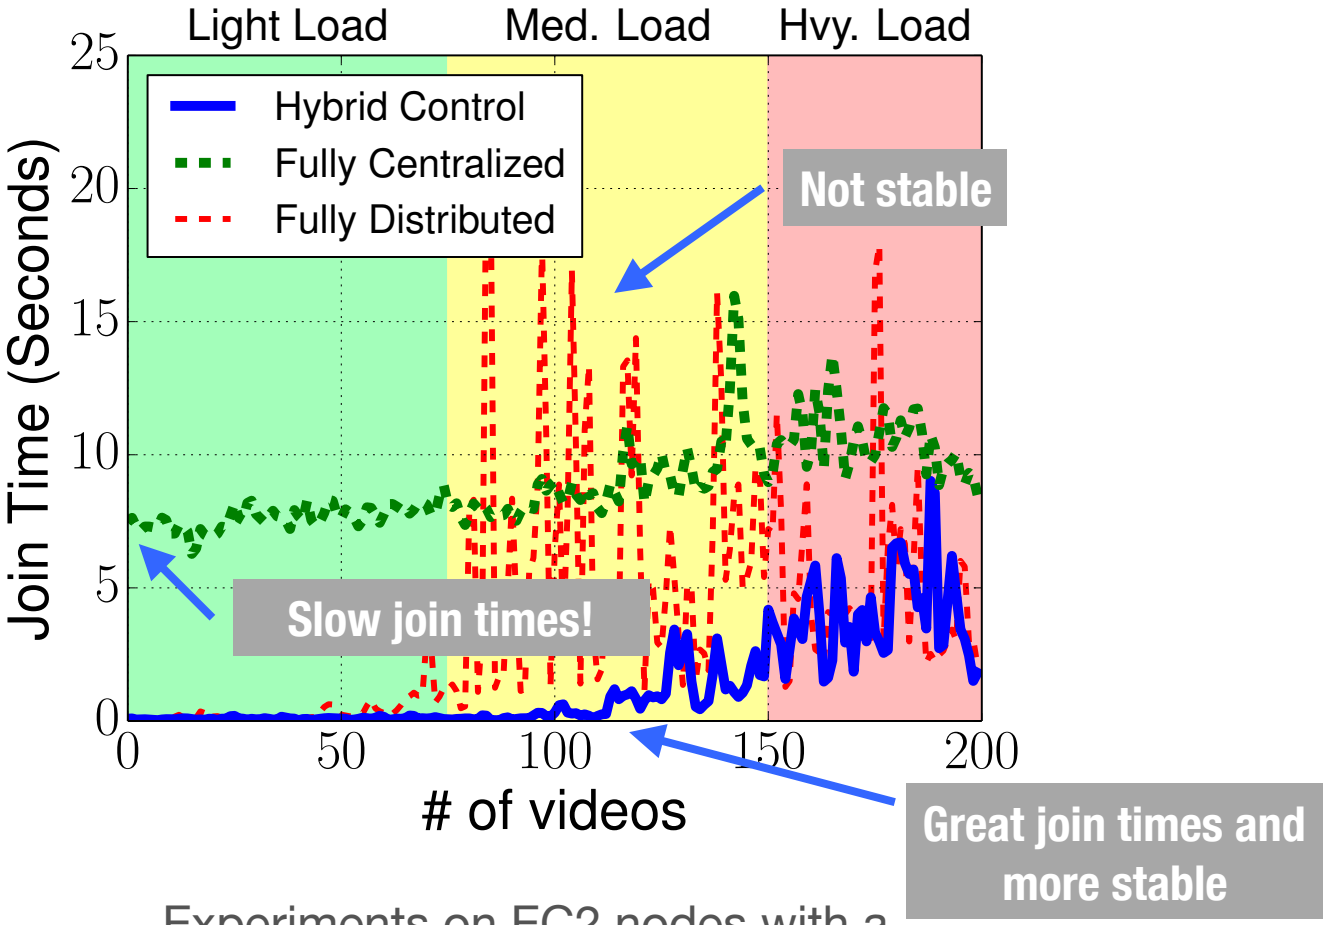

Experiments on EC2 nodes with a centralized controller at CMU across the Internet

Conclusion

- We present a *possible* solution for combating decision plane latency

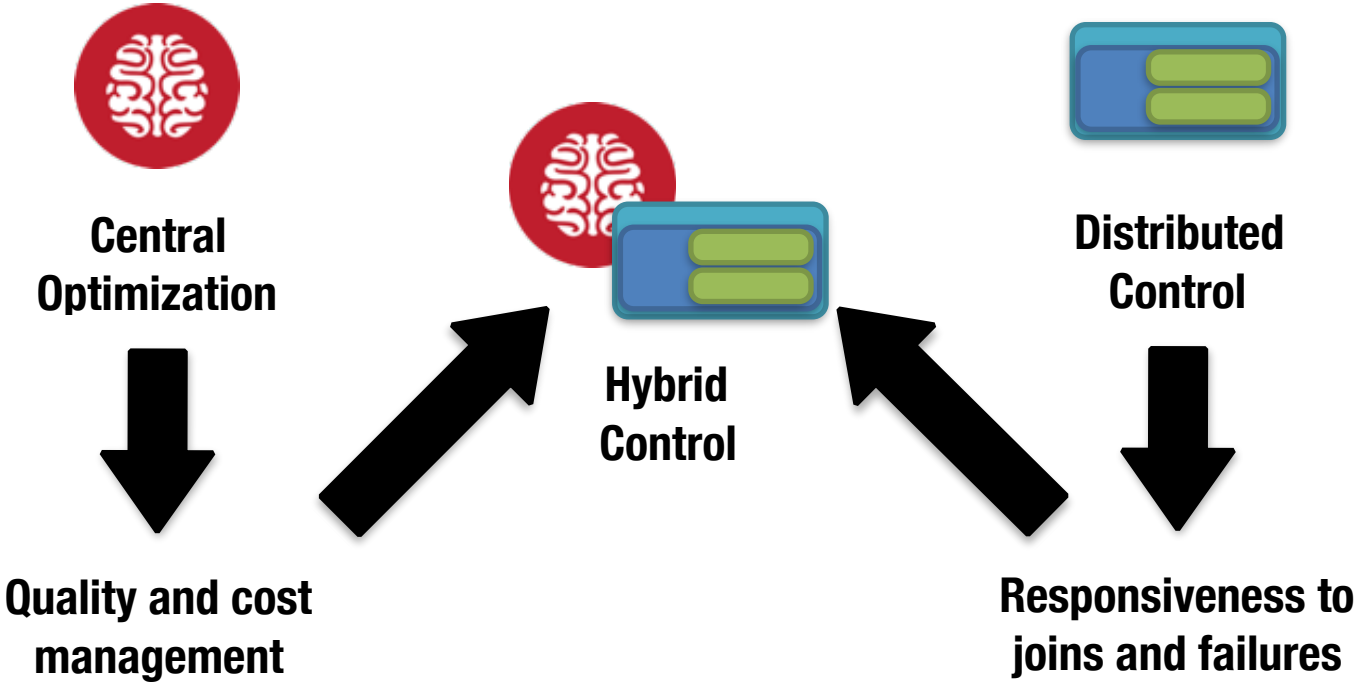

Conclusion

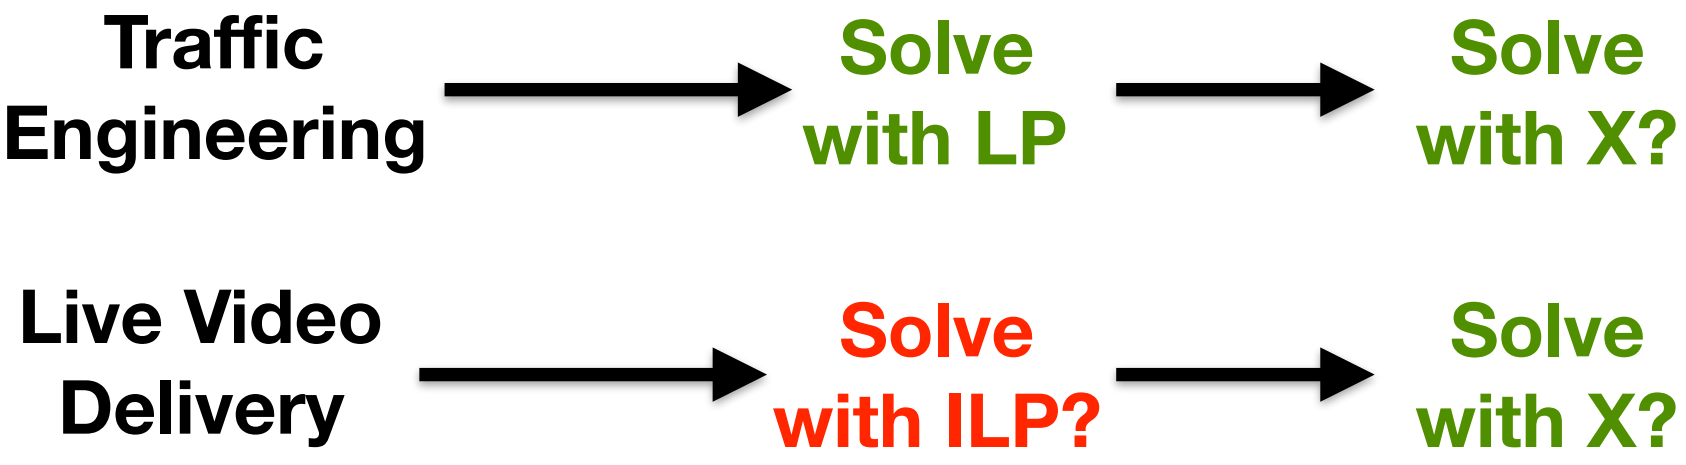
Traffic Engineering
Solve with LP
Solve with X?
Live Video Delivery
Solve with ILP?
Solve with X?

Practical, Real-time Centralized Control for CDN-based Live Video Delivery

Matt Mukerjee, David Naylor,
Junchen Jiang, Dongsu Han,
Srini Seshan, Hui Zhang

Carnegie Mellon University

CONVIVA

KAIST

Backup slides…

Problems with Traffic Engineering

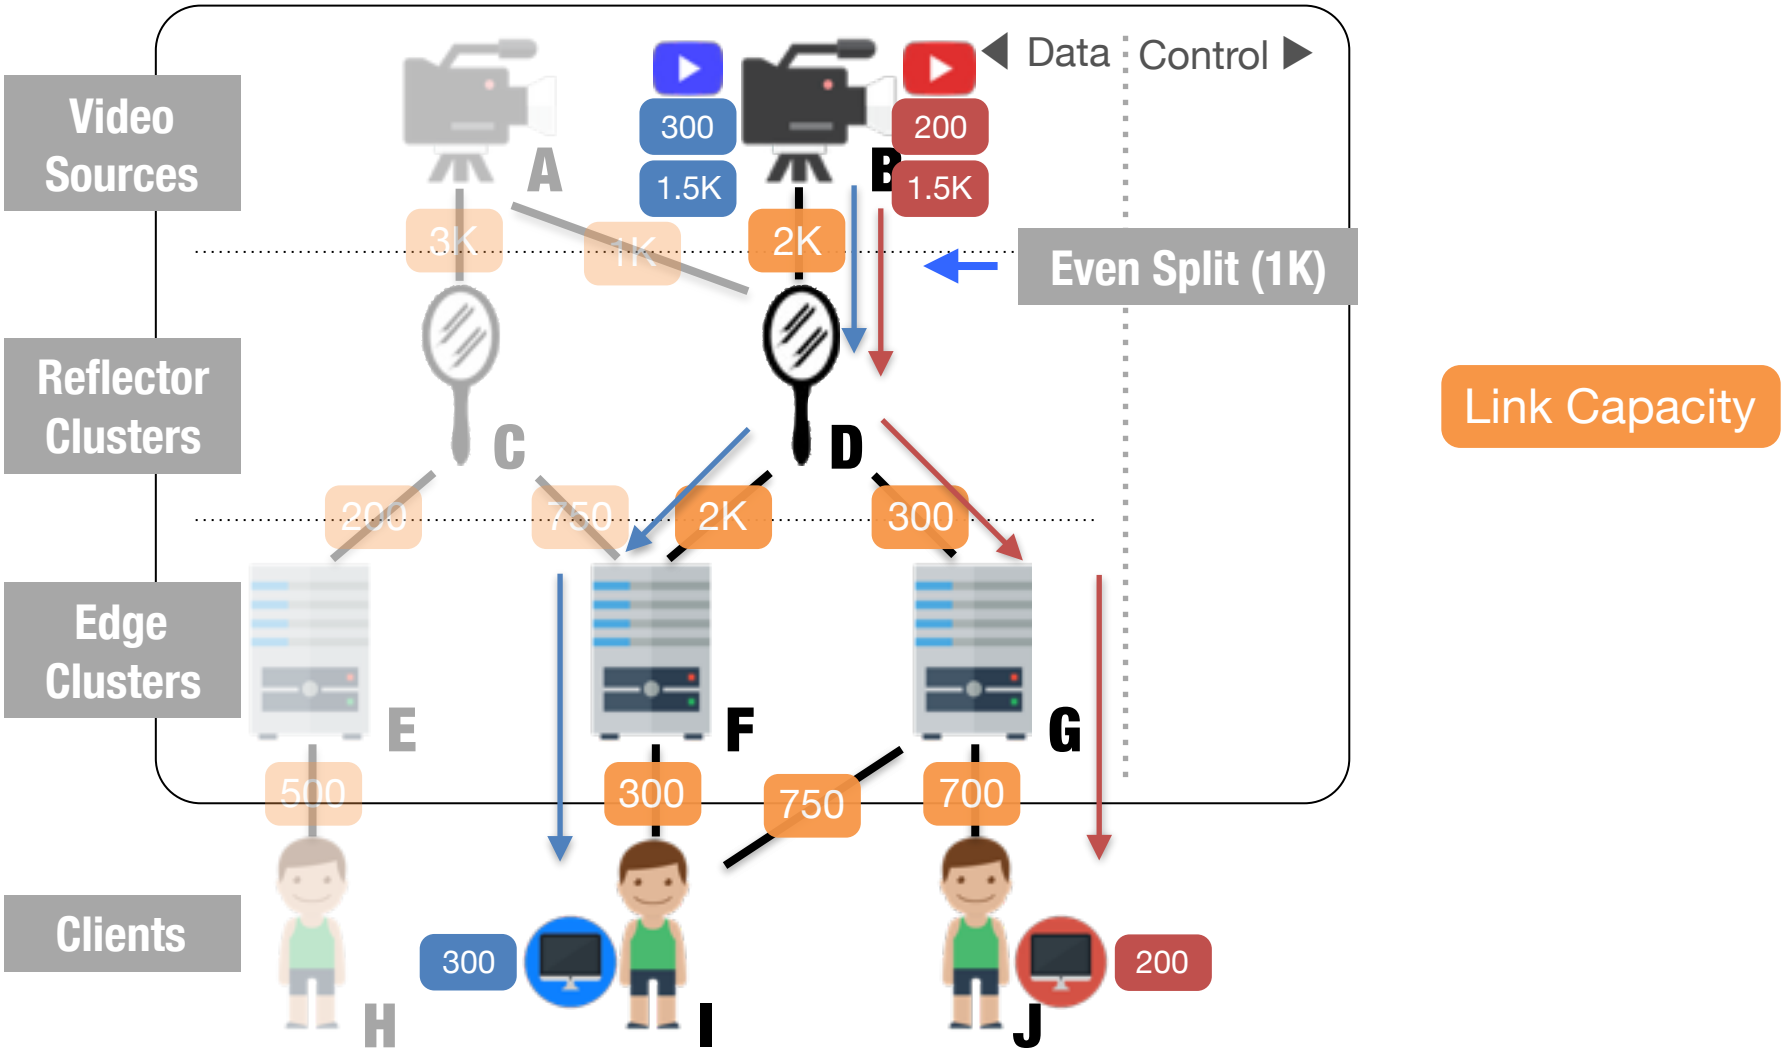

Link Capacity

Problems with Traffic Engineering

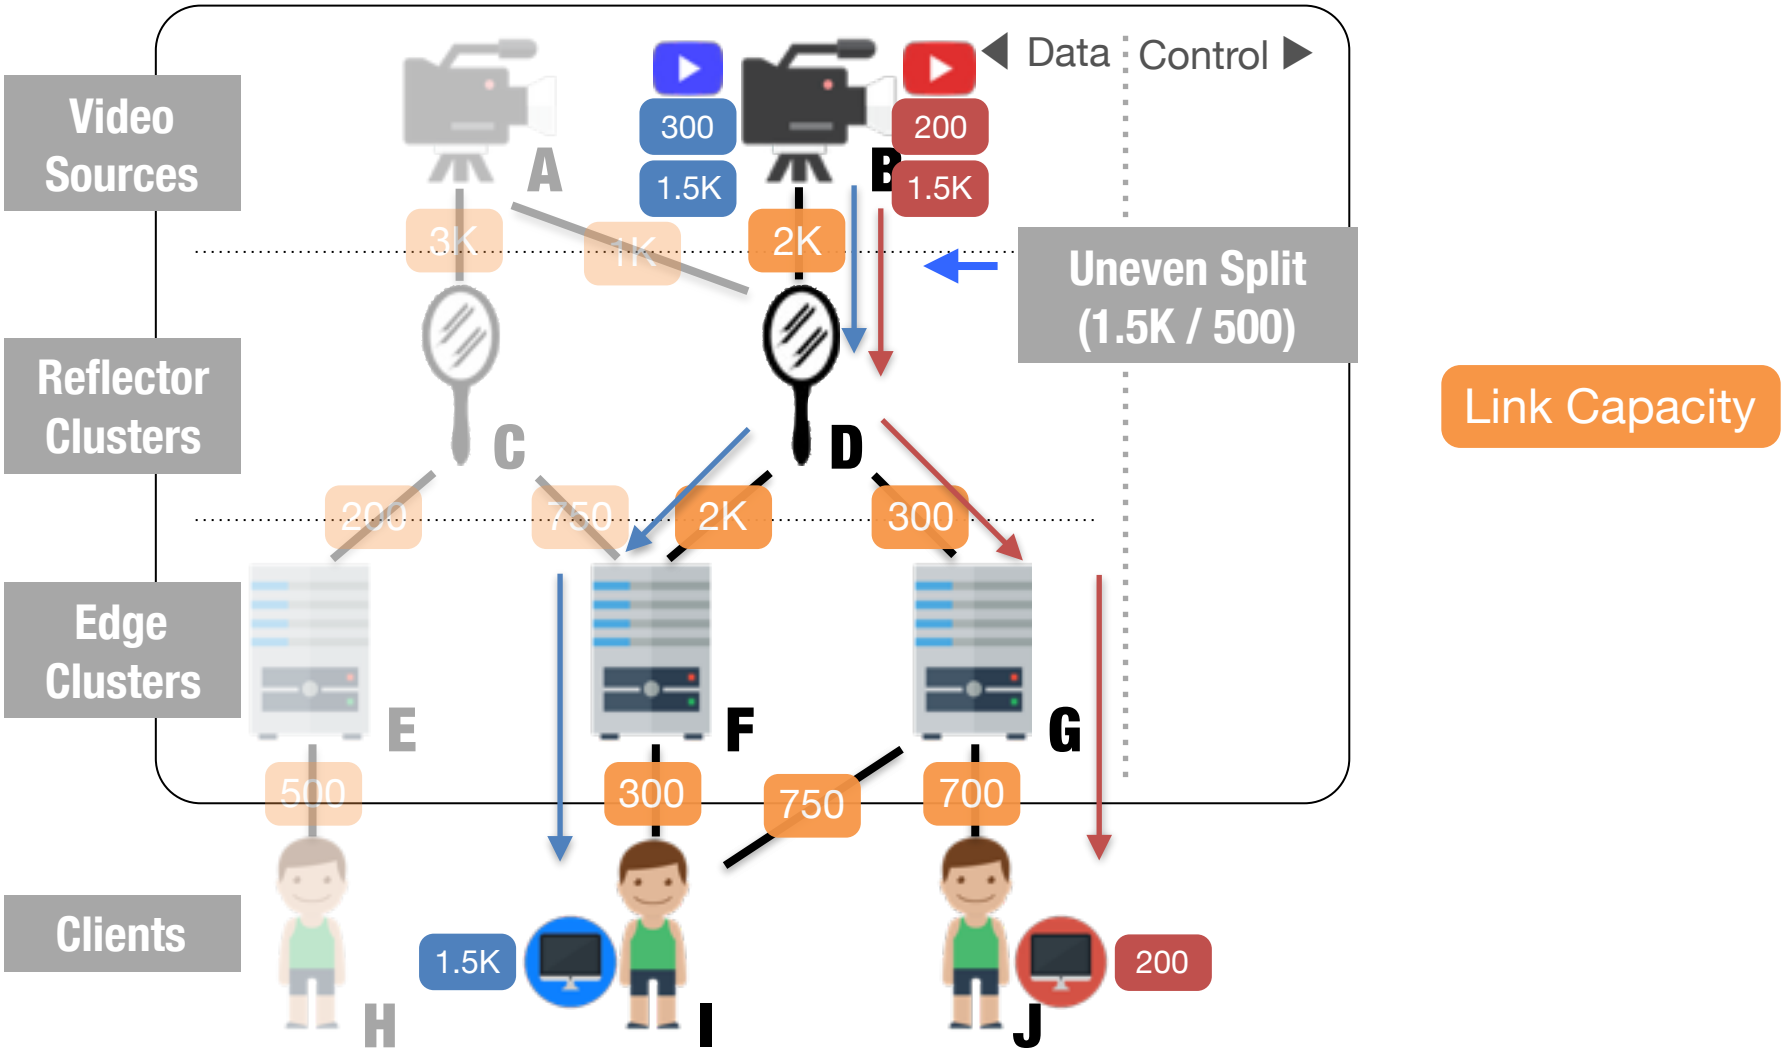

Distributed: Example of Sub-optimal

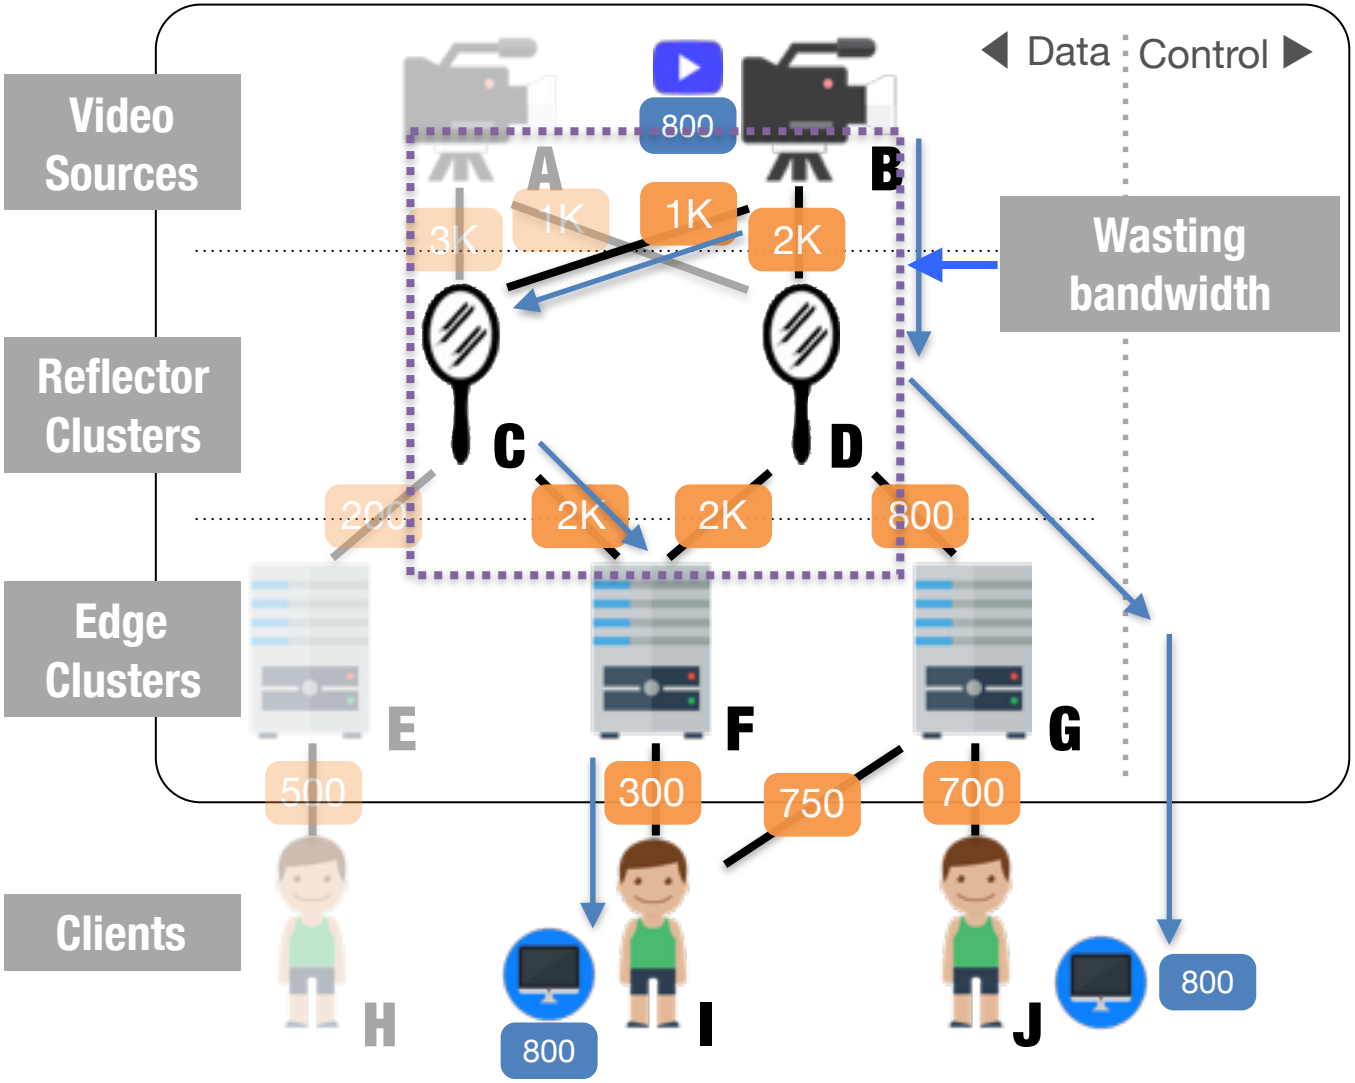

Legend

Data Requests: Video 1

Responses: Video 1

Link Capacity

Distributed: Example of Sub-optimal

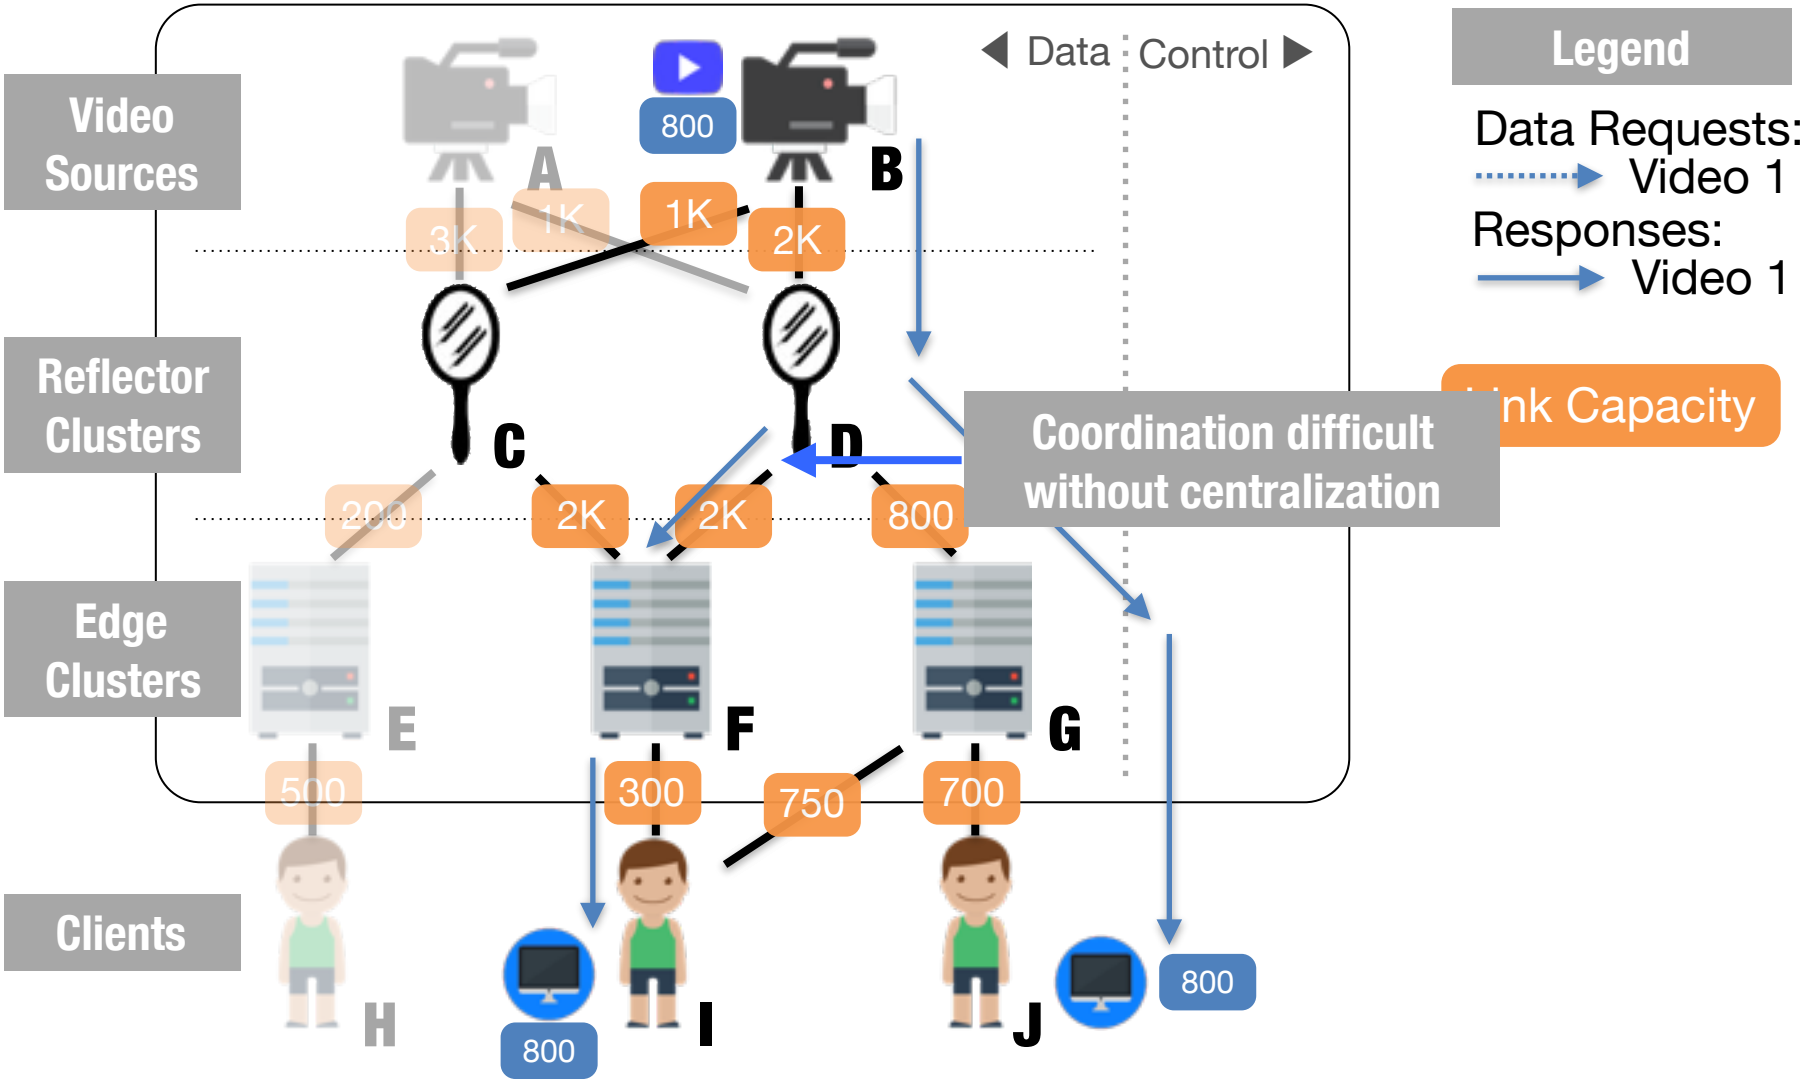

Trace-Driven Eval

- 3 Traces
  - Avg Day: raw trace of music video provider
  - Large Event: synthesized basketball game
  - Heavy Tail: synthesized twitch/ustream like workload
- 4 Systems
  - Everything Everywhere: all vids to all servers
  - Overlay Multicast: globally optimal; no coordination
  - CDN: greedy distribution scheme w/ DNS
  - VDN: our system

Trace-Driven Eval

	EE	CDN	VDN
Avg. Bitrate *(kbps)*	588	**2,725**	**2,716**
Cost / Sat. Req. *(norm.)*	103	1.5	**1**
Clients at Reqs. BR *(%)*	18.73%	**100%**	**99.83%**

Table 1: Results for Average Day trace.

	EE	CDN	VDN
Avg. Bitrate *(kbps)*	0.08	**2,725**	**2,725**
Cost / Sat. Req. *(norm.)*	178K	2.2	**1**
Clients at Reqs. BR *(%)*	0%	**100%**	**100%**

Table 2: Results for Large Event trace.

	EE	CDN	VDN
Avg. Bitrate *(kbps)*	685	1748	**3366**
Cost / Sat. Req. *(norm.)*	8.9	**1.21**	**1**
Clients at Reqs. BR *(%)*	22%	49%	**77%**

Table 3: Results for Heavy-Tail trace.

Existing Solutions

- Traffic Engineering (SWAN, B4, …)
  - Works on aggregates at coarse timescales
- Overlay Multicast (Overcast, Bullet, …)
  - Not designed for coordinating across streams
- Modern CDNs
  - Previous work shows a centralized system could greatly improve user experience but would be difficult to design over Internet