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

Live Video is Becoming Wildly Popular

- Commercial sports streams

- User-generated streams

Live Video is Becoming Wildly Popular

- Commercial sports streams
  - **Single** World Cup stream = **40%** global Internet traffic
- User-generated streams (e.g., Twitch)
  - Users watch **150b min of live video per month**
  - Amazon buys Twitch for ~**$1Billion**

Our Contributions

- We design a video delivery network (**VDN**) to efficiently manage quality and cost, with high responsiveness

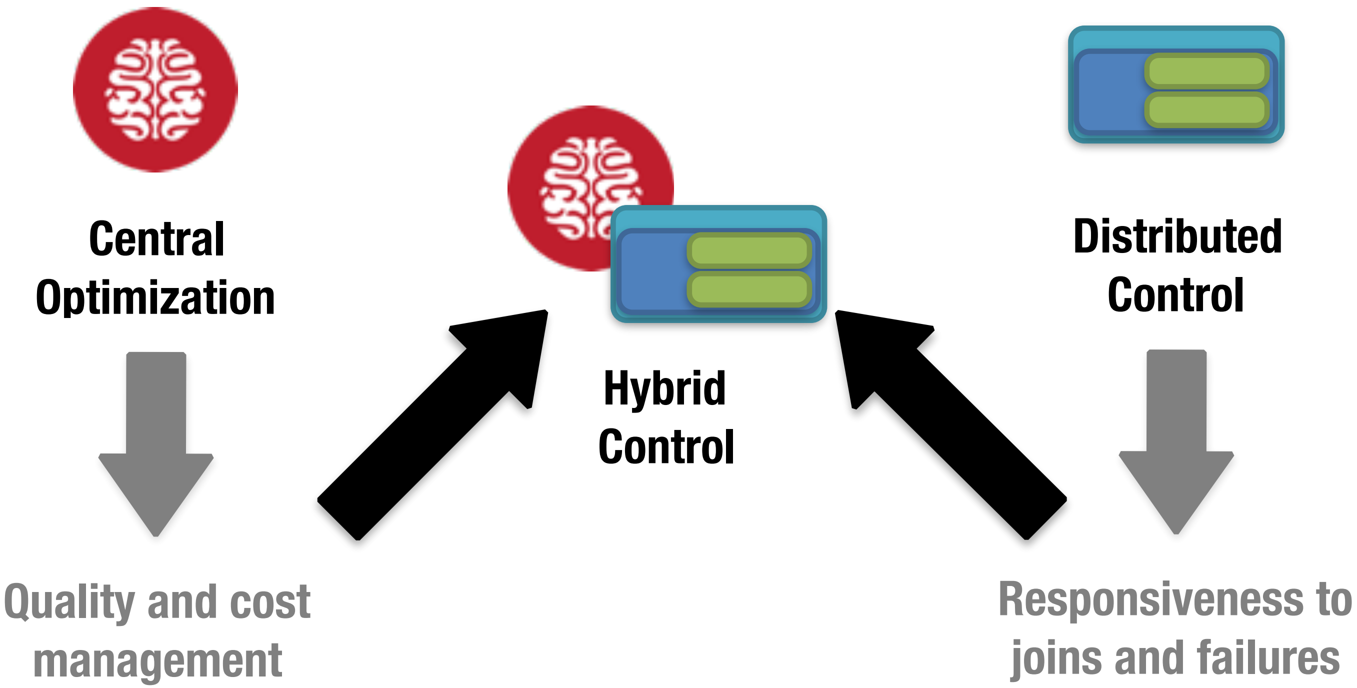

Outline

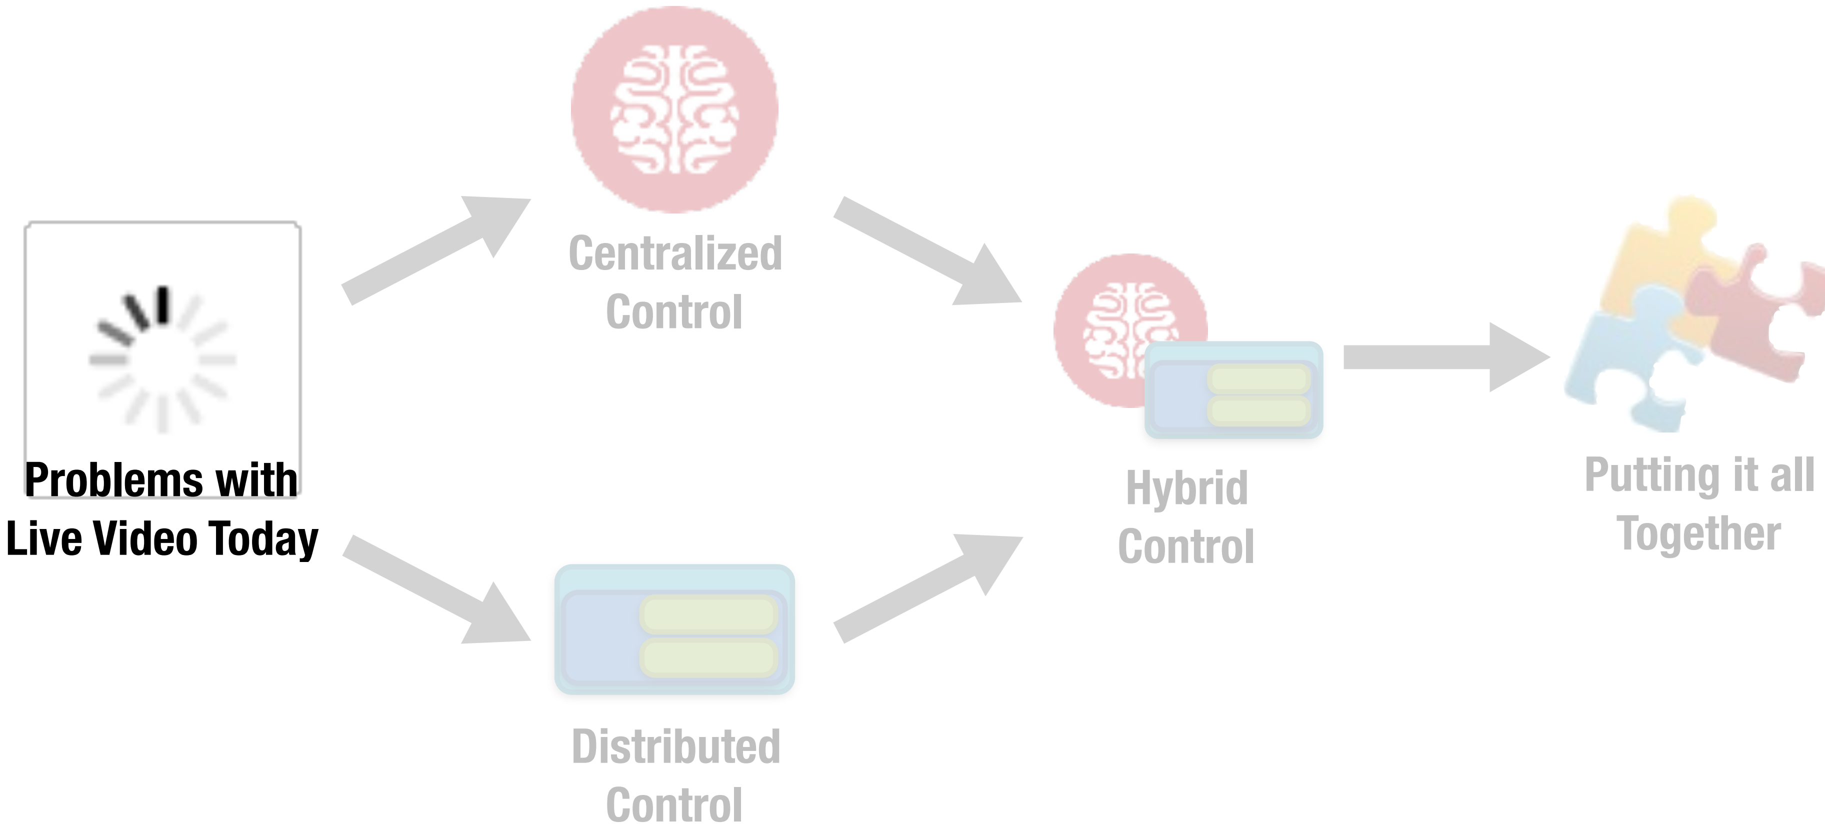
Problems with Live Video Today
Centralized Control
Distributed Control
Hybrid Control
Putting it all Together

CDN Live Video Delivery Background

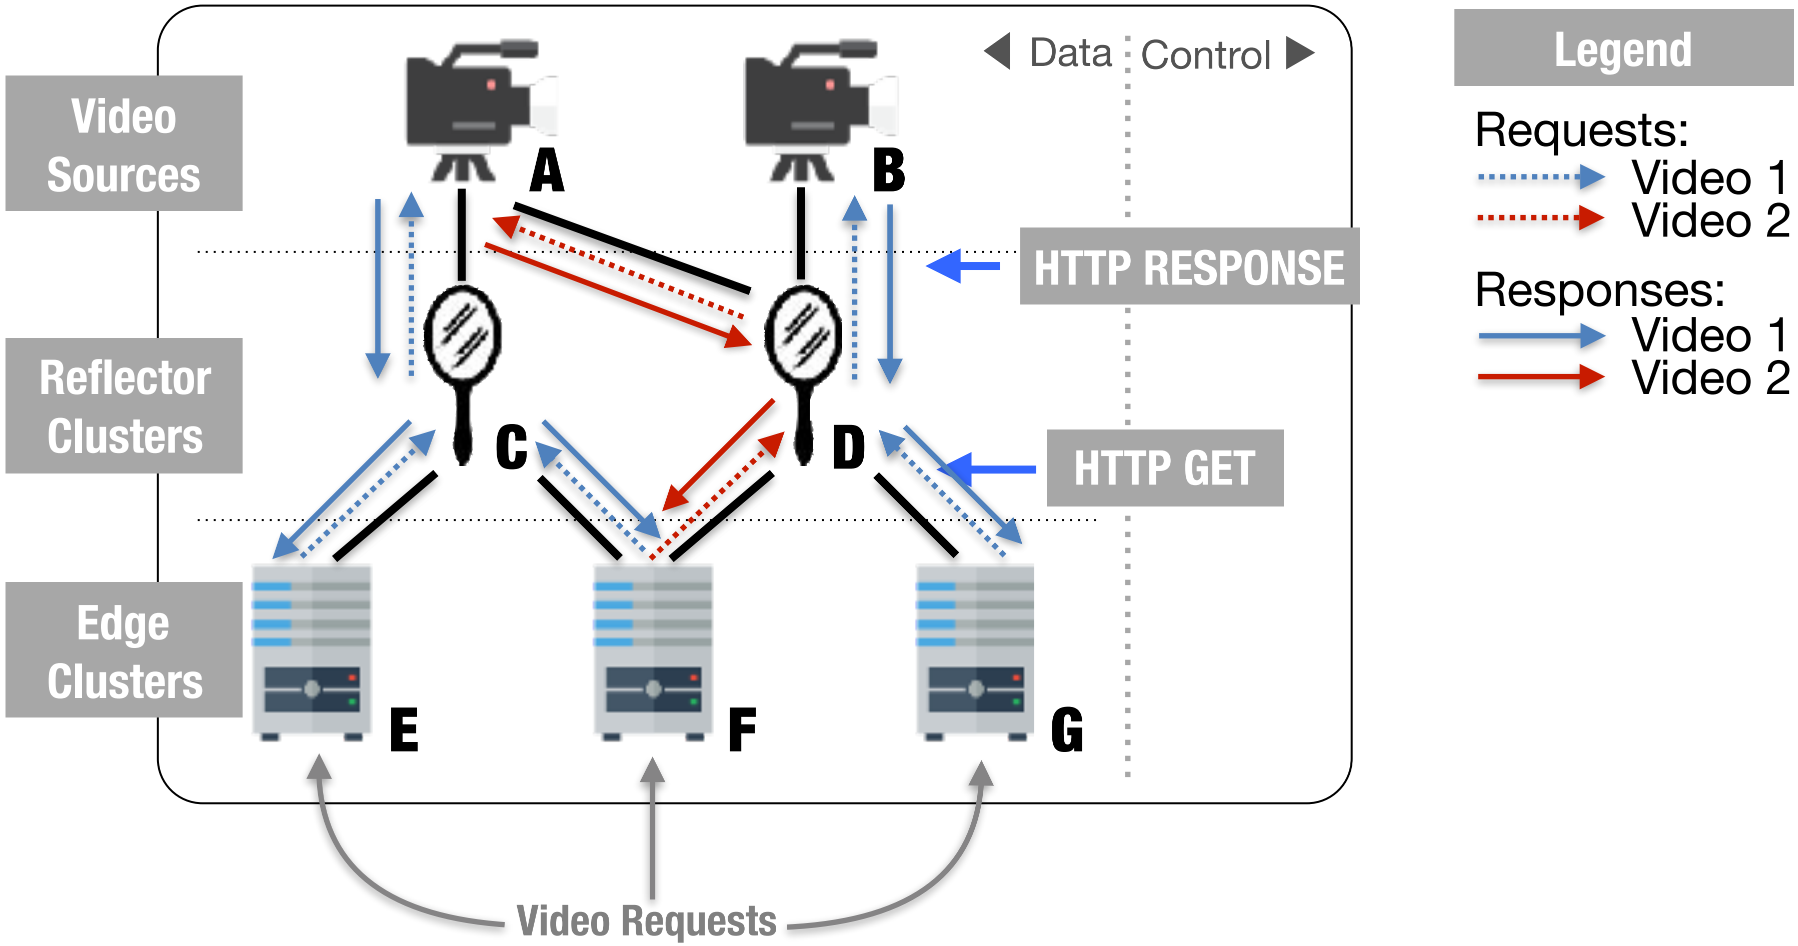

CDN Live Video Delivery Background

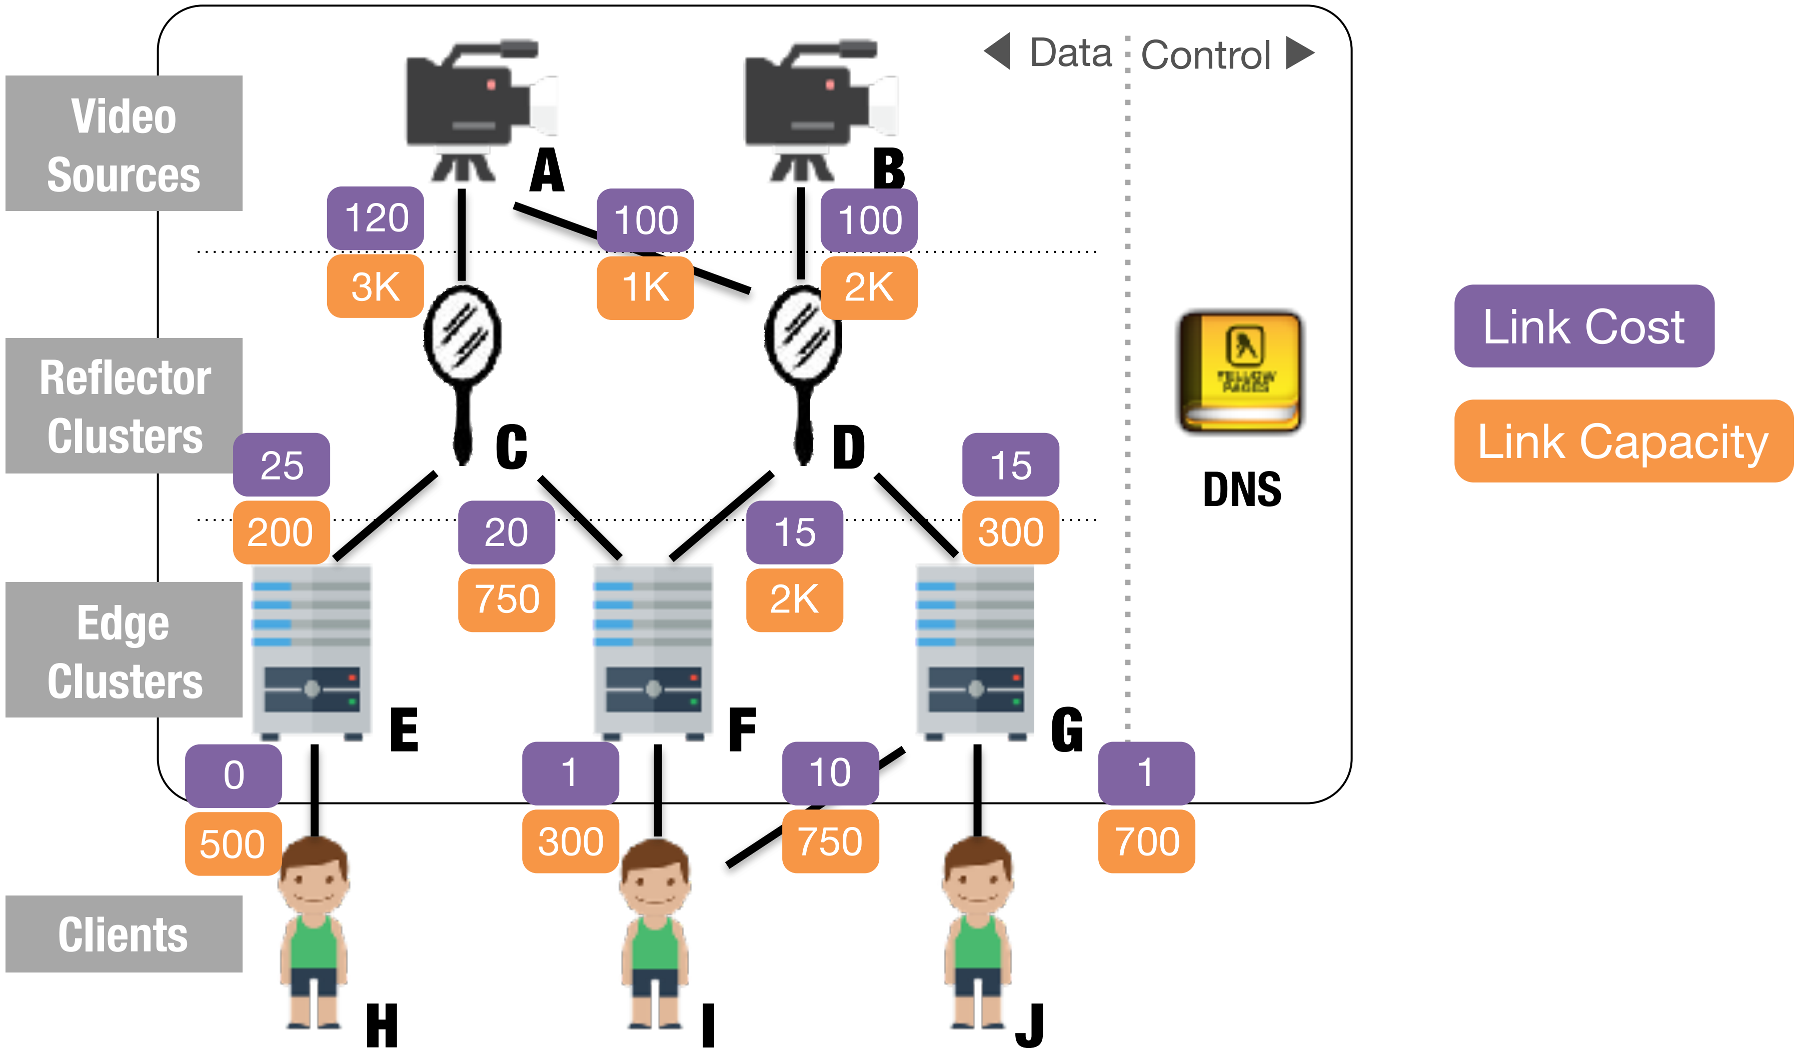

CDN Live Video Delivery Background

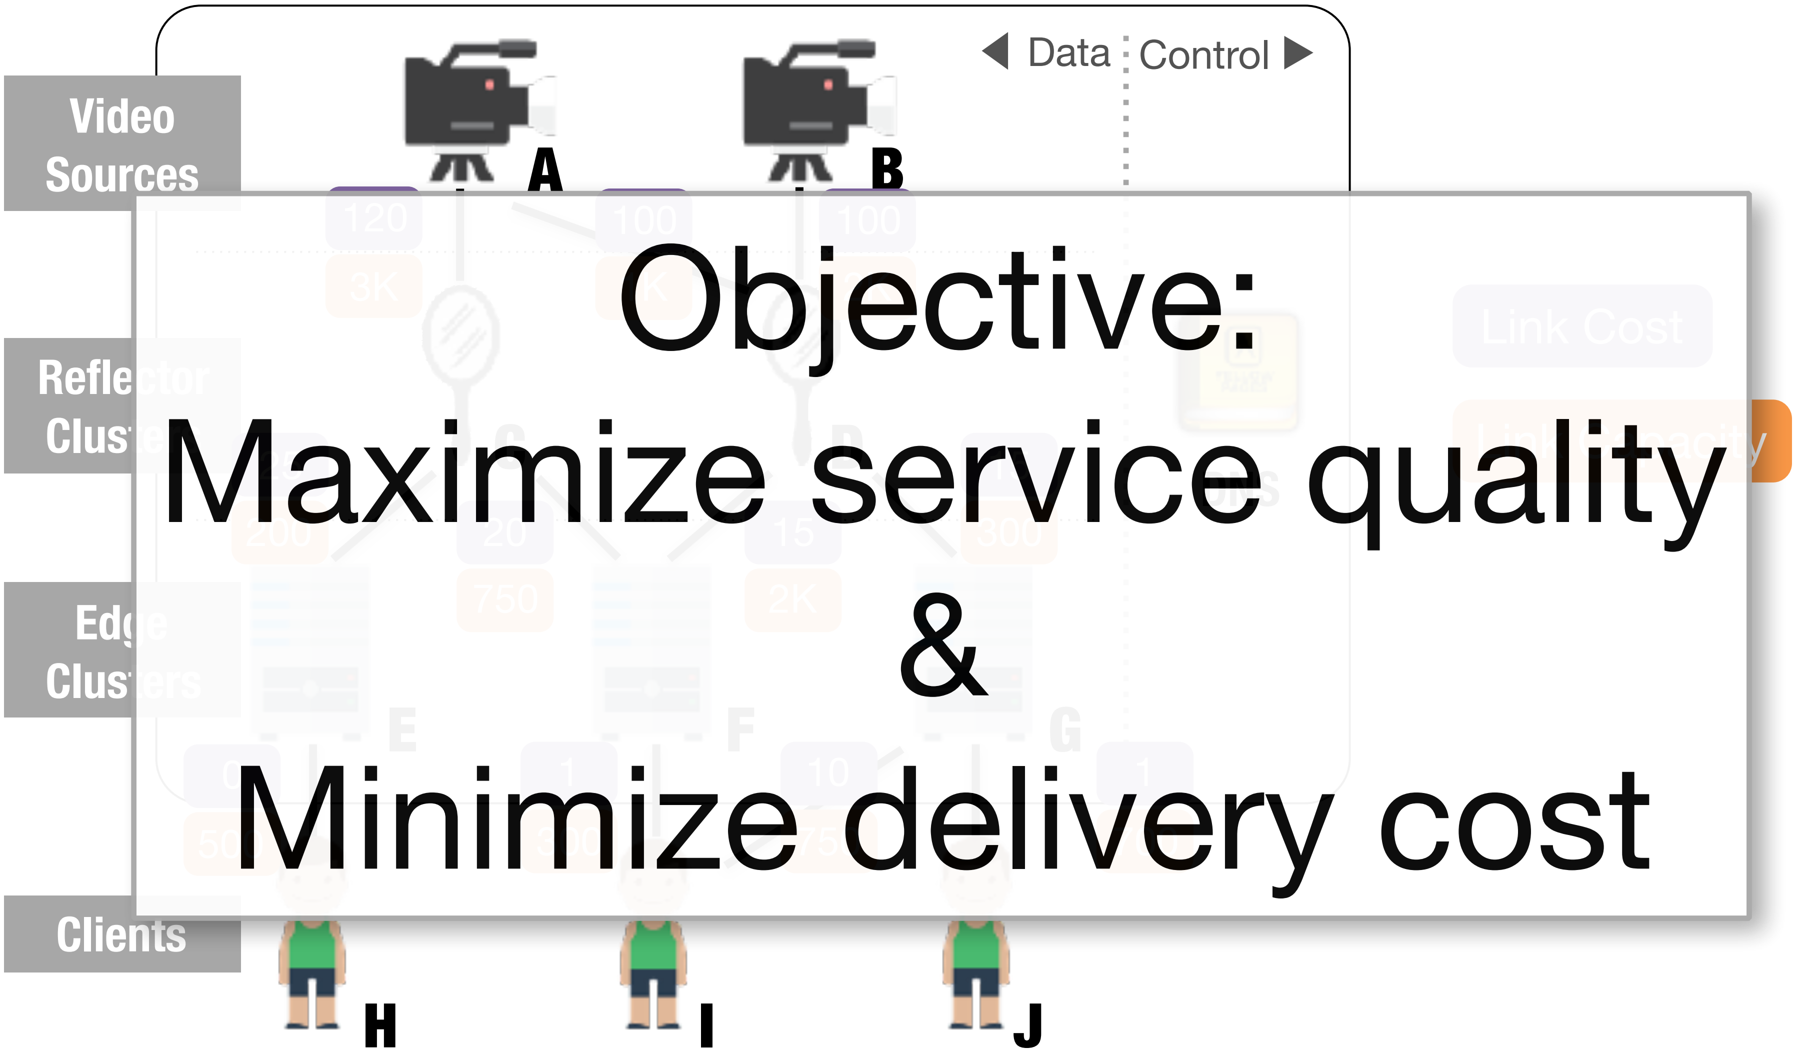

Objective:

Maximize service quality

&

Minimize delivery cost

Problems with CDNs Today

Service Quality

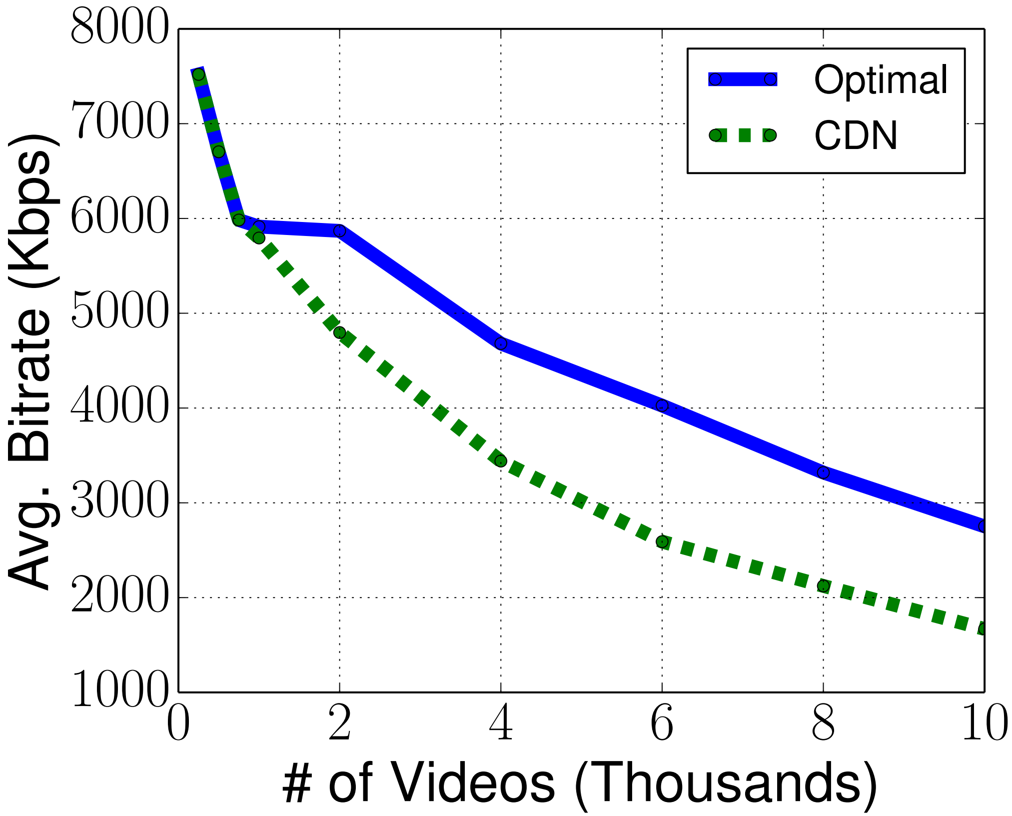

Simulation using Conviva traces, modeling user-generated content

Delivery Cost

(per request)

CDN

2.0x

OPTIMAL

1.0x

Simulation using Conviva traces, modeling large sports events

Problems with CDNs Today

Service Quality

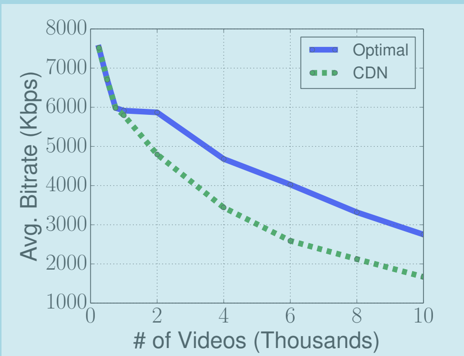

Delivery Cost

CDN 2.0x

OPTIMAL 1.0x

QUANTITATIVE

Not Fine-Grained

Videos aggregated into large groups

Slow DNS Updates

Can't push updates

DNS entries get cached

QUALITATIVE

Goals

Service Quality

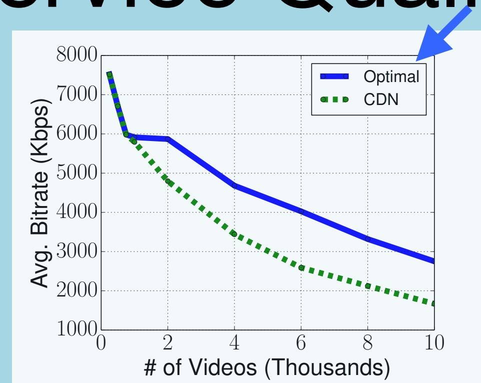

Delivery Cost

CDN
2.0x

OPTIMAL
1.0x

Fine-Grained Control

Per-video Control

Real-time Response

Sub-second response
to failures and joins

Room for improvement, but Internet latency / loss

Goals

Service Quality

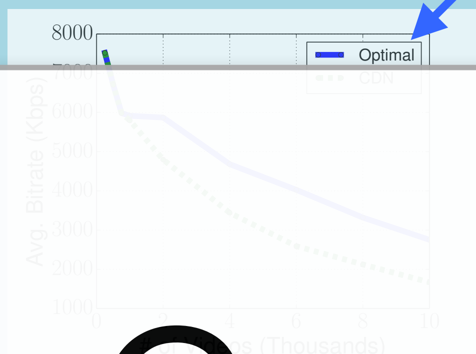

Fine-Grained Control

Per-video Control

Centralization!

Delivery Cost

Real-time Response

[Liu, Xi et. al. A Case for a Coordinated Video Control Plane. SIGCOMM 2012]

CDN 2.0x

OPTIMAL 1.0x

Sub-second response to failures and joins

Room for improvement, but Internet latency / loss

Outline

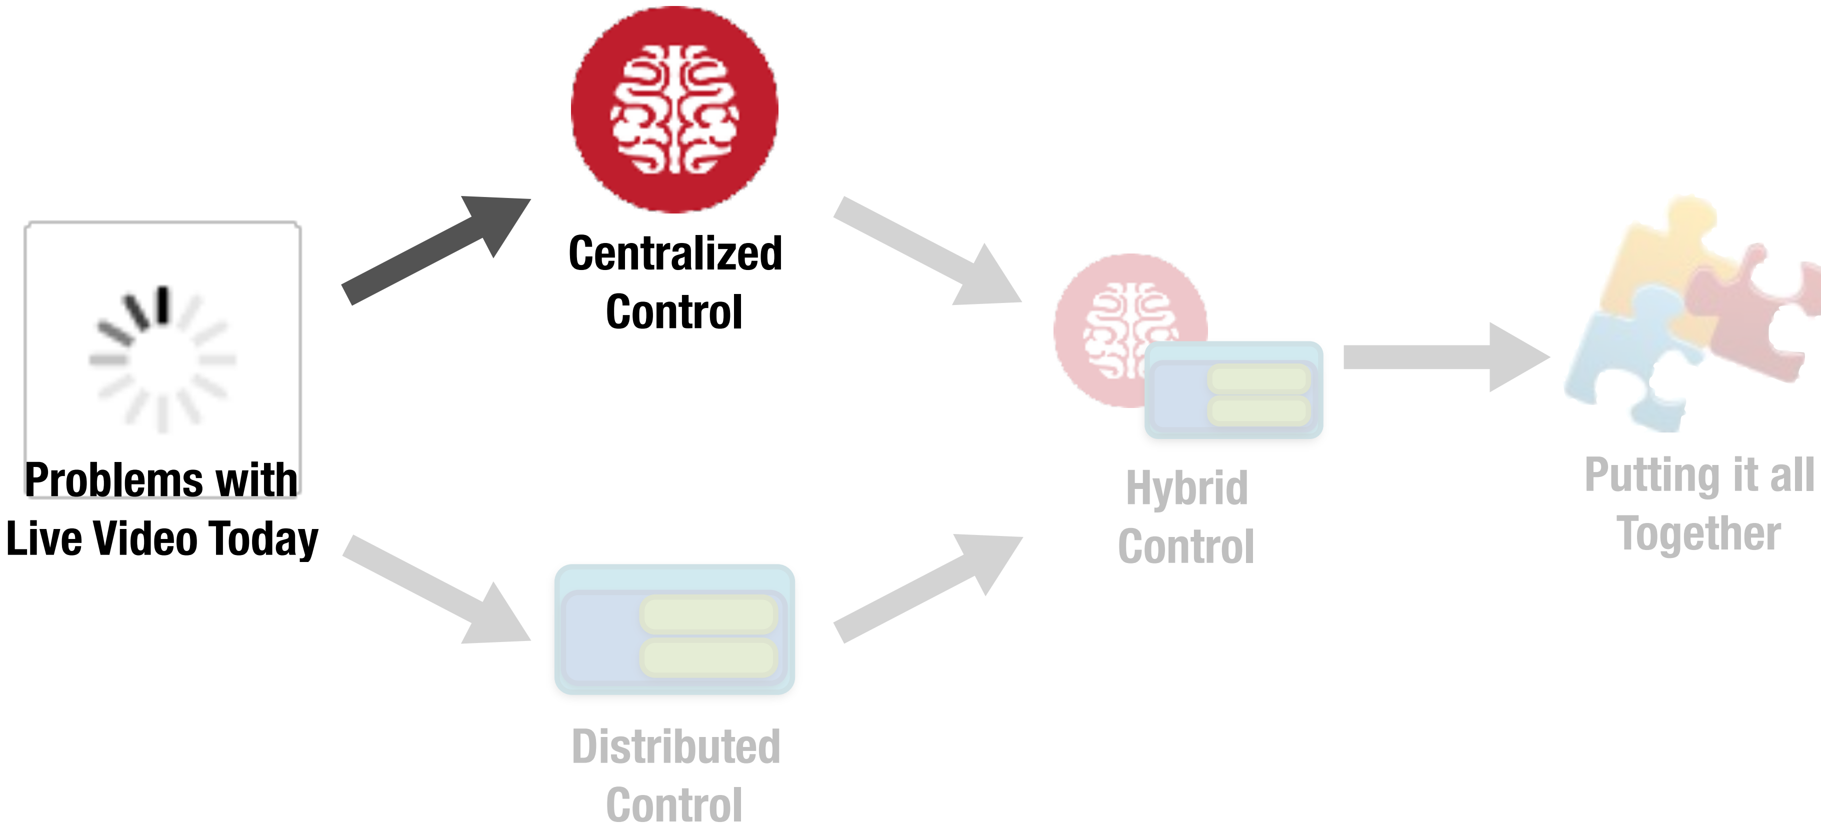
Problems with Live Video Today
Centralized Control
Distributed Control
Hybrid Control
Putting it all Together

Motivating Centralized Optimization

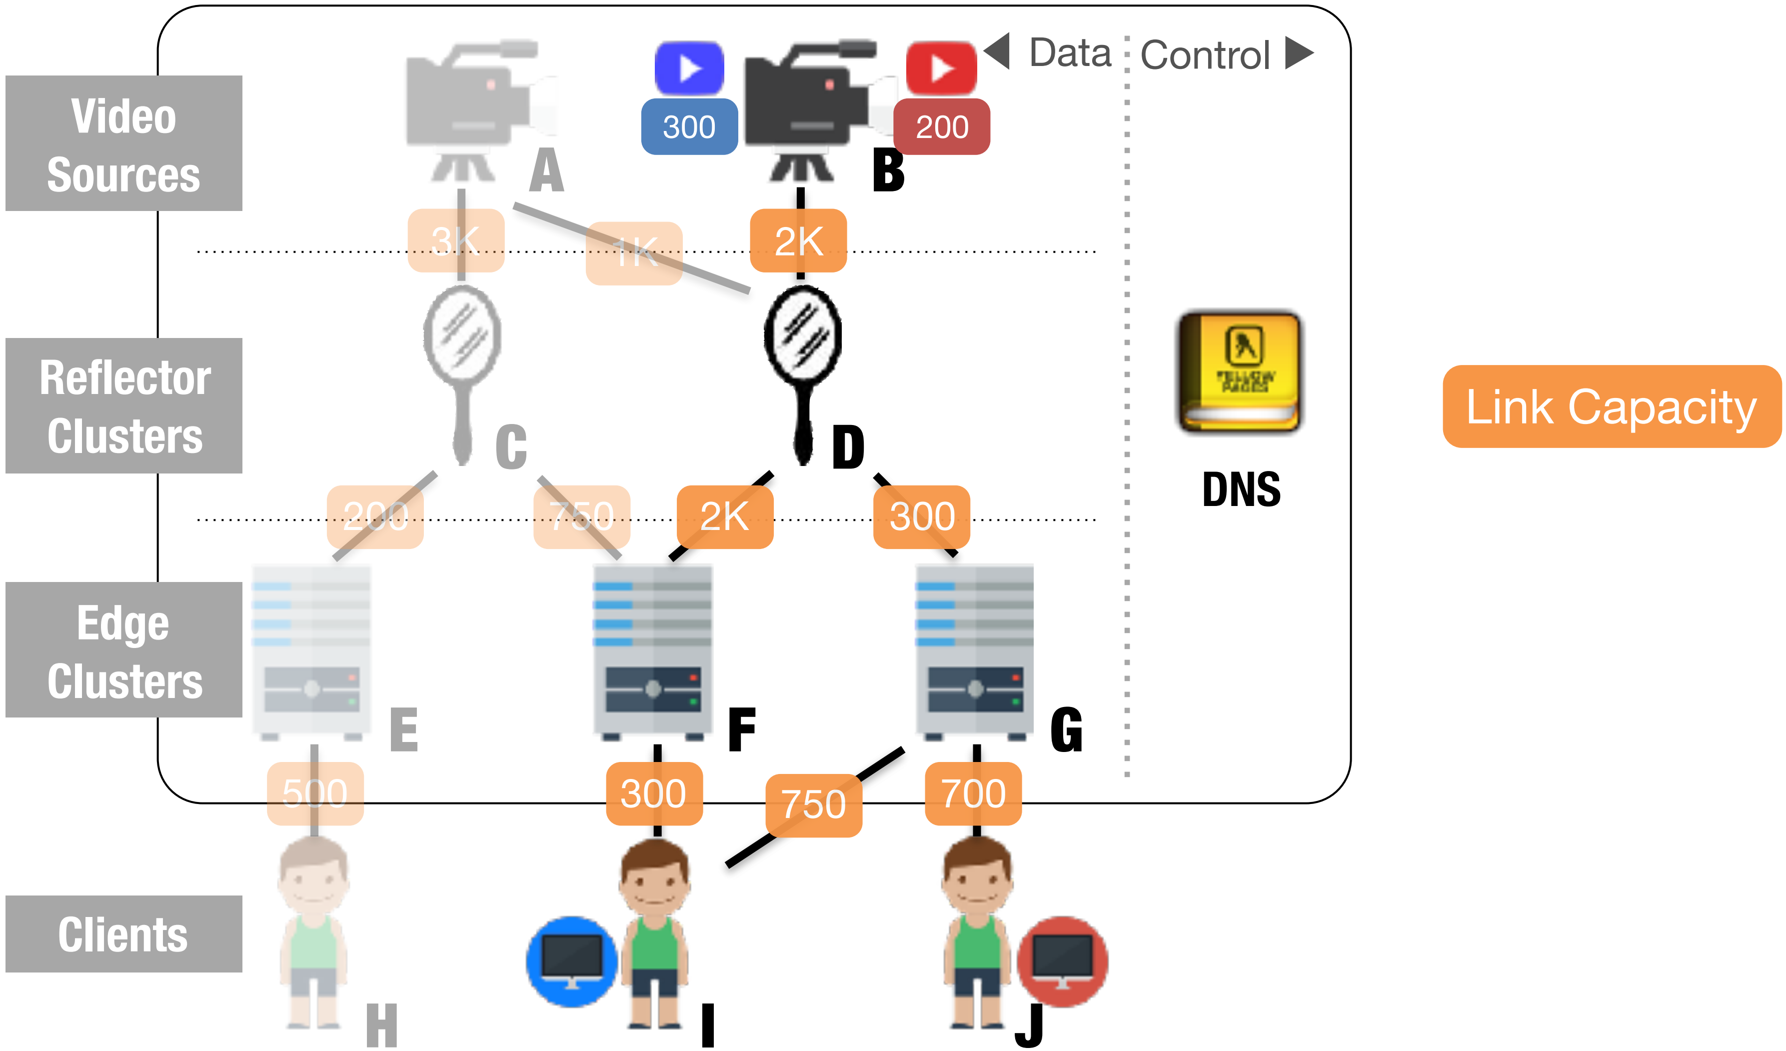

Motivating Centralized Optimization

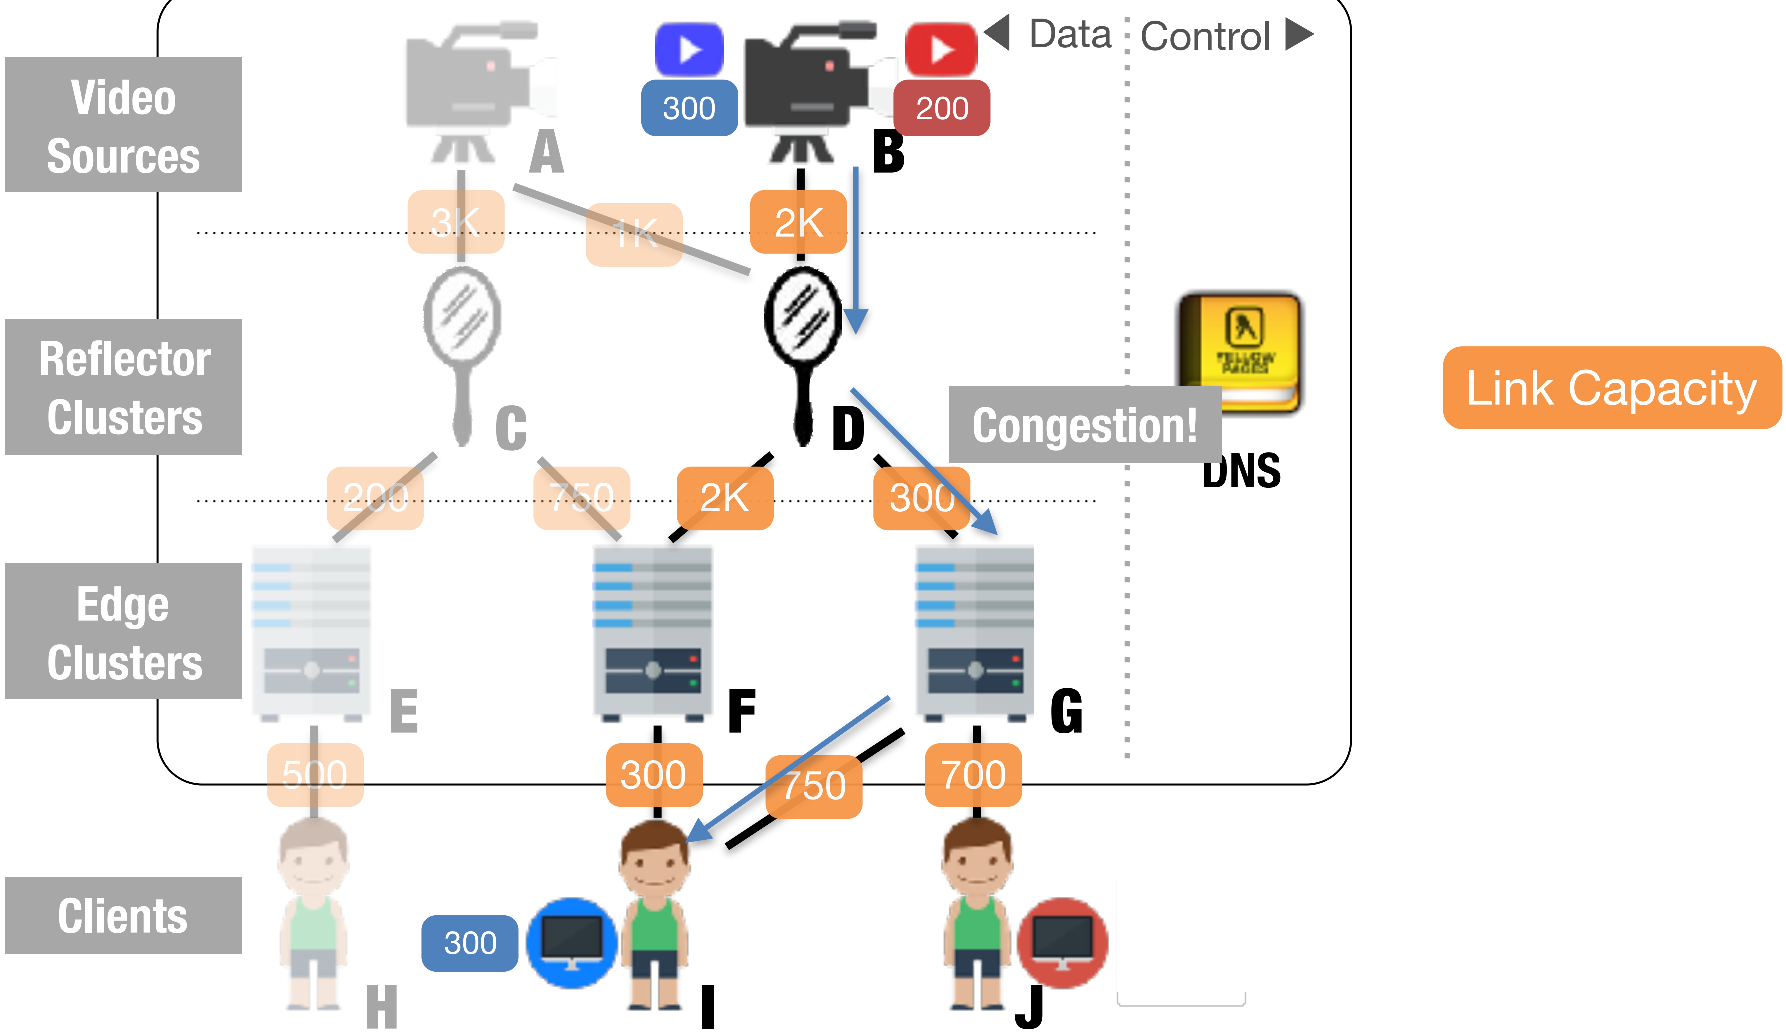

Motivating Centralized Optimization

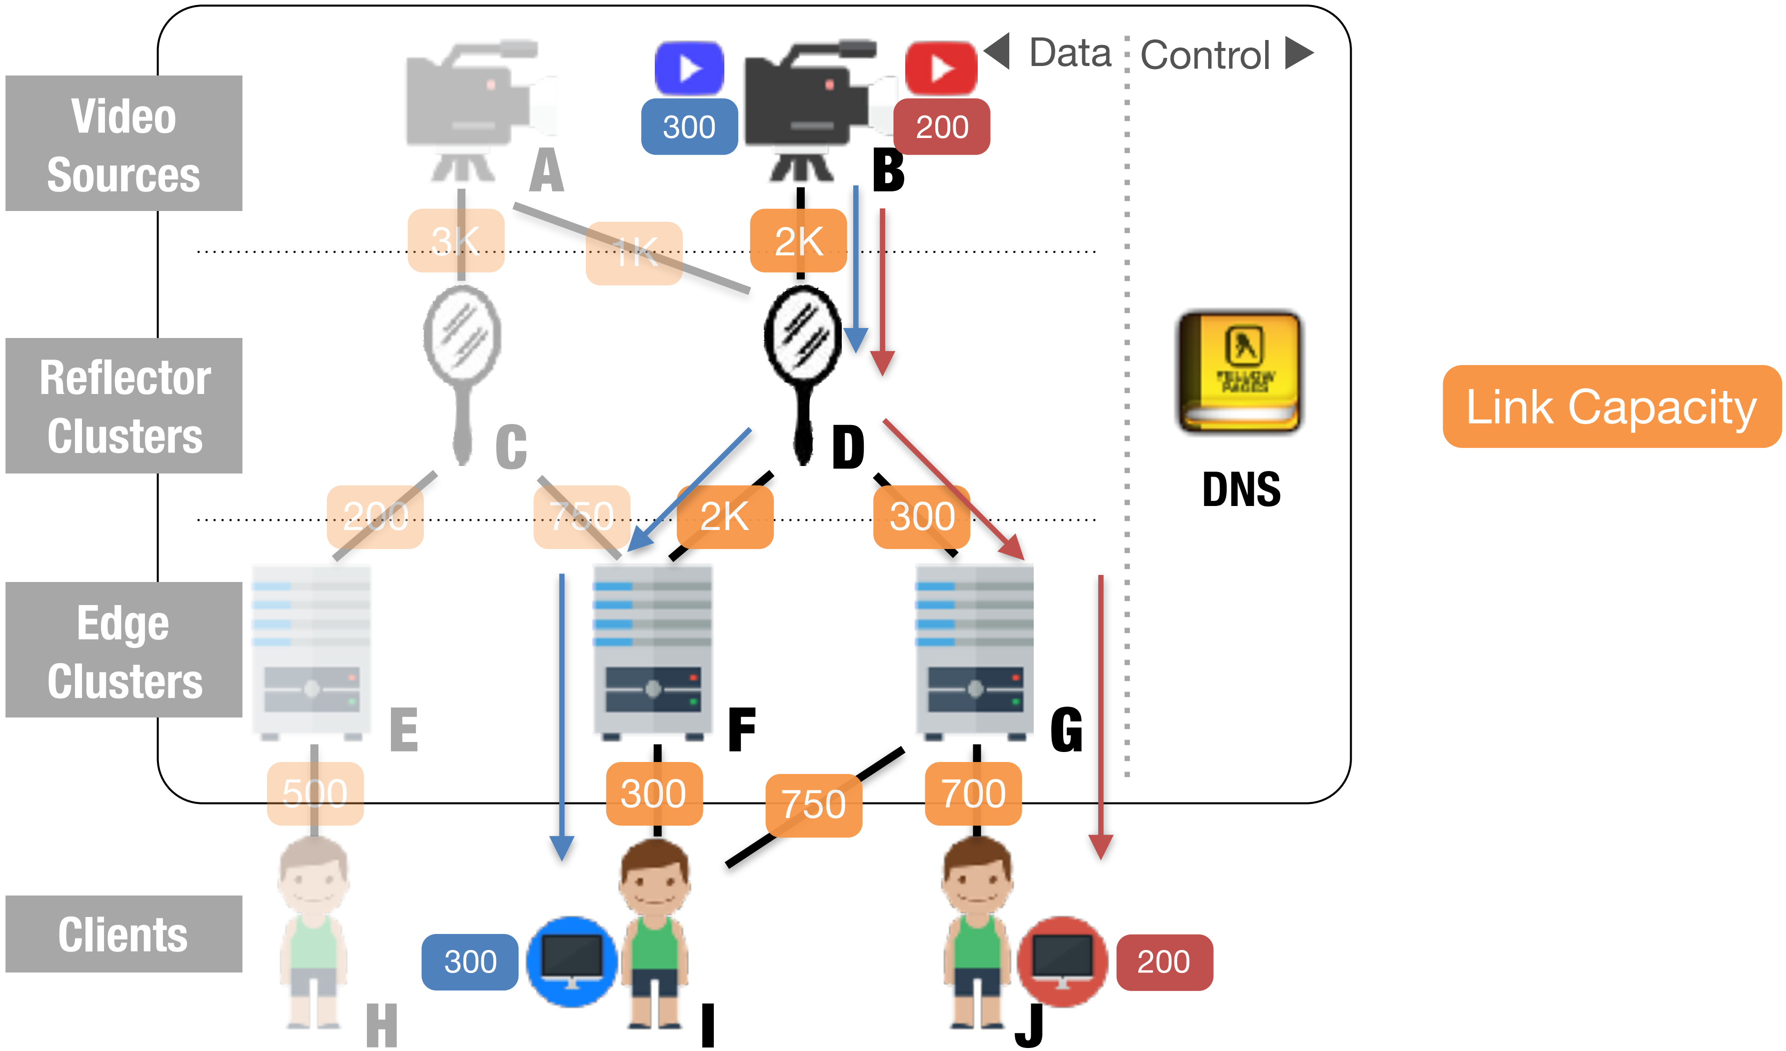

Motivating Centralized Optimization

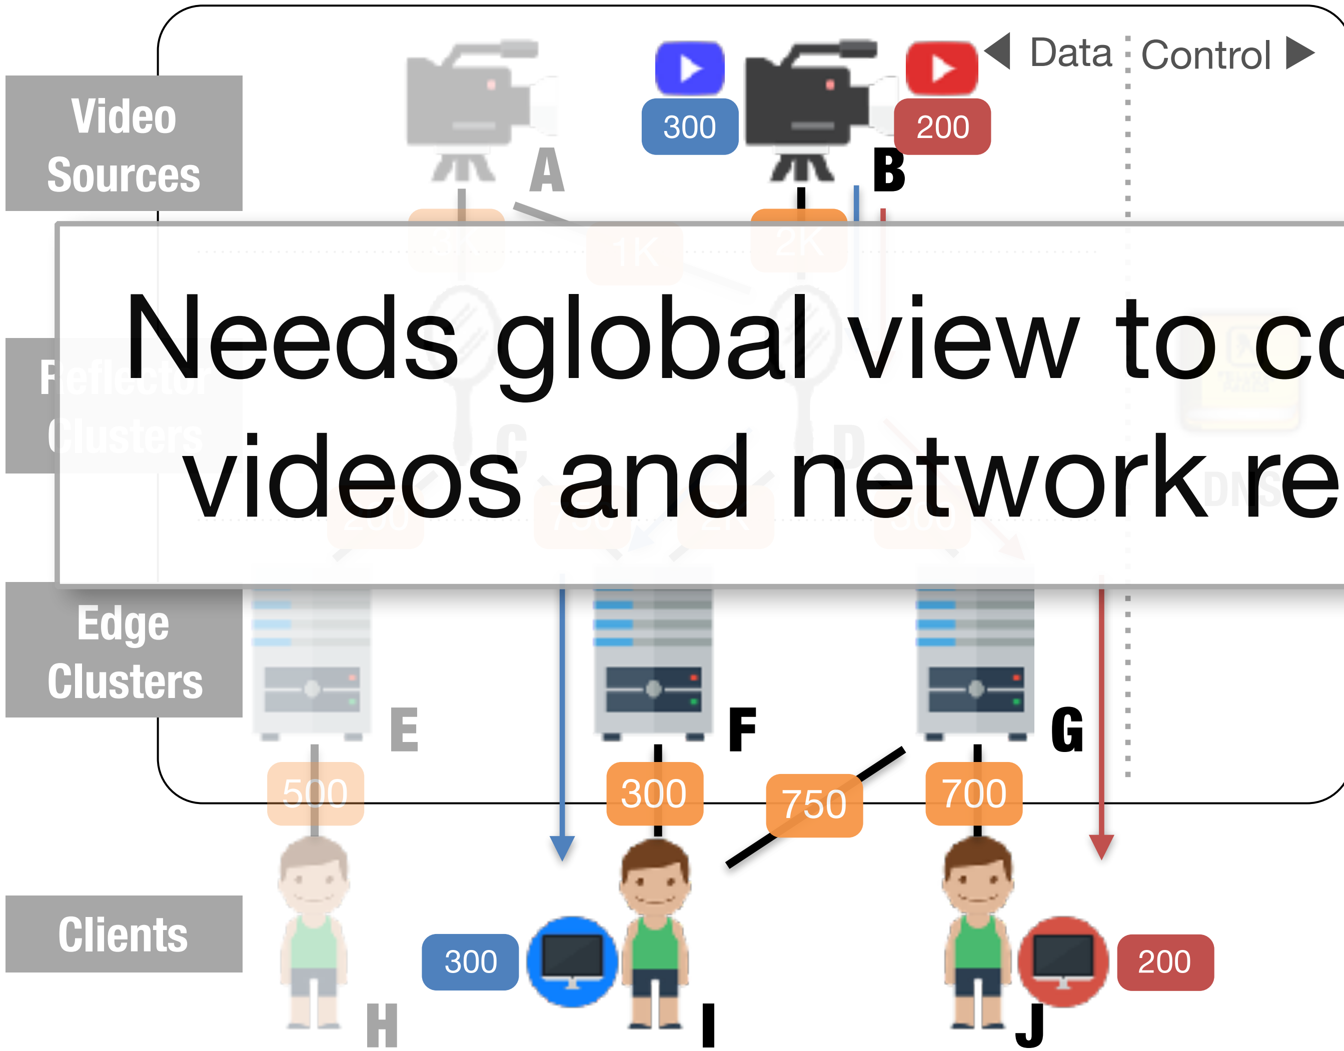

Needs global view to coordinate videos and network resources

Motivating Centralized Optimization

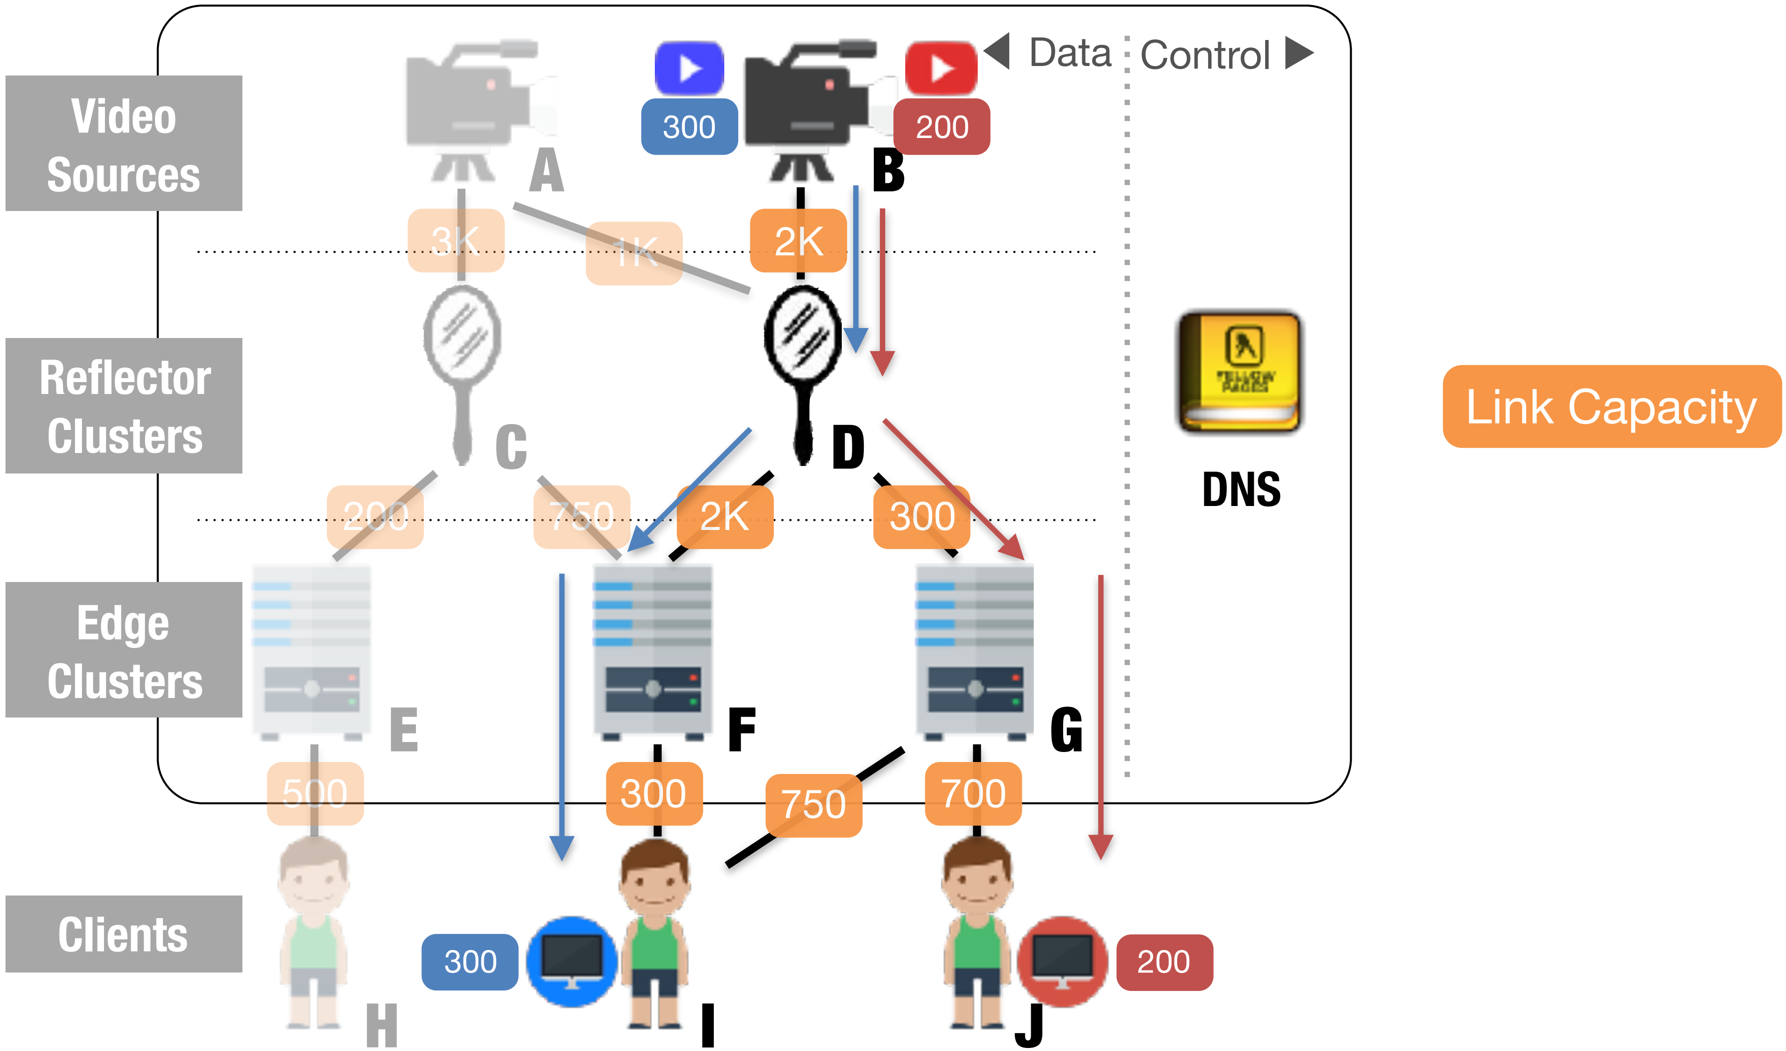

Motivating Centralized Optimization

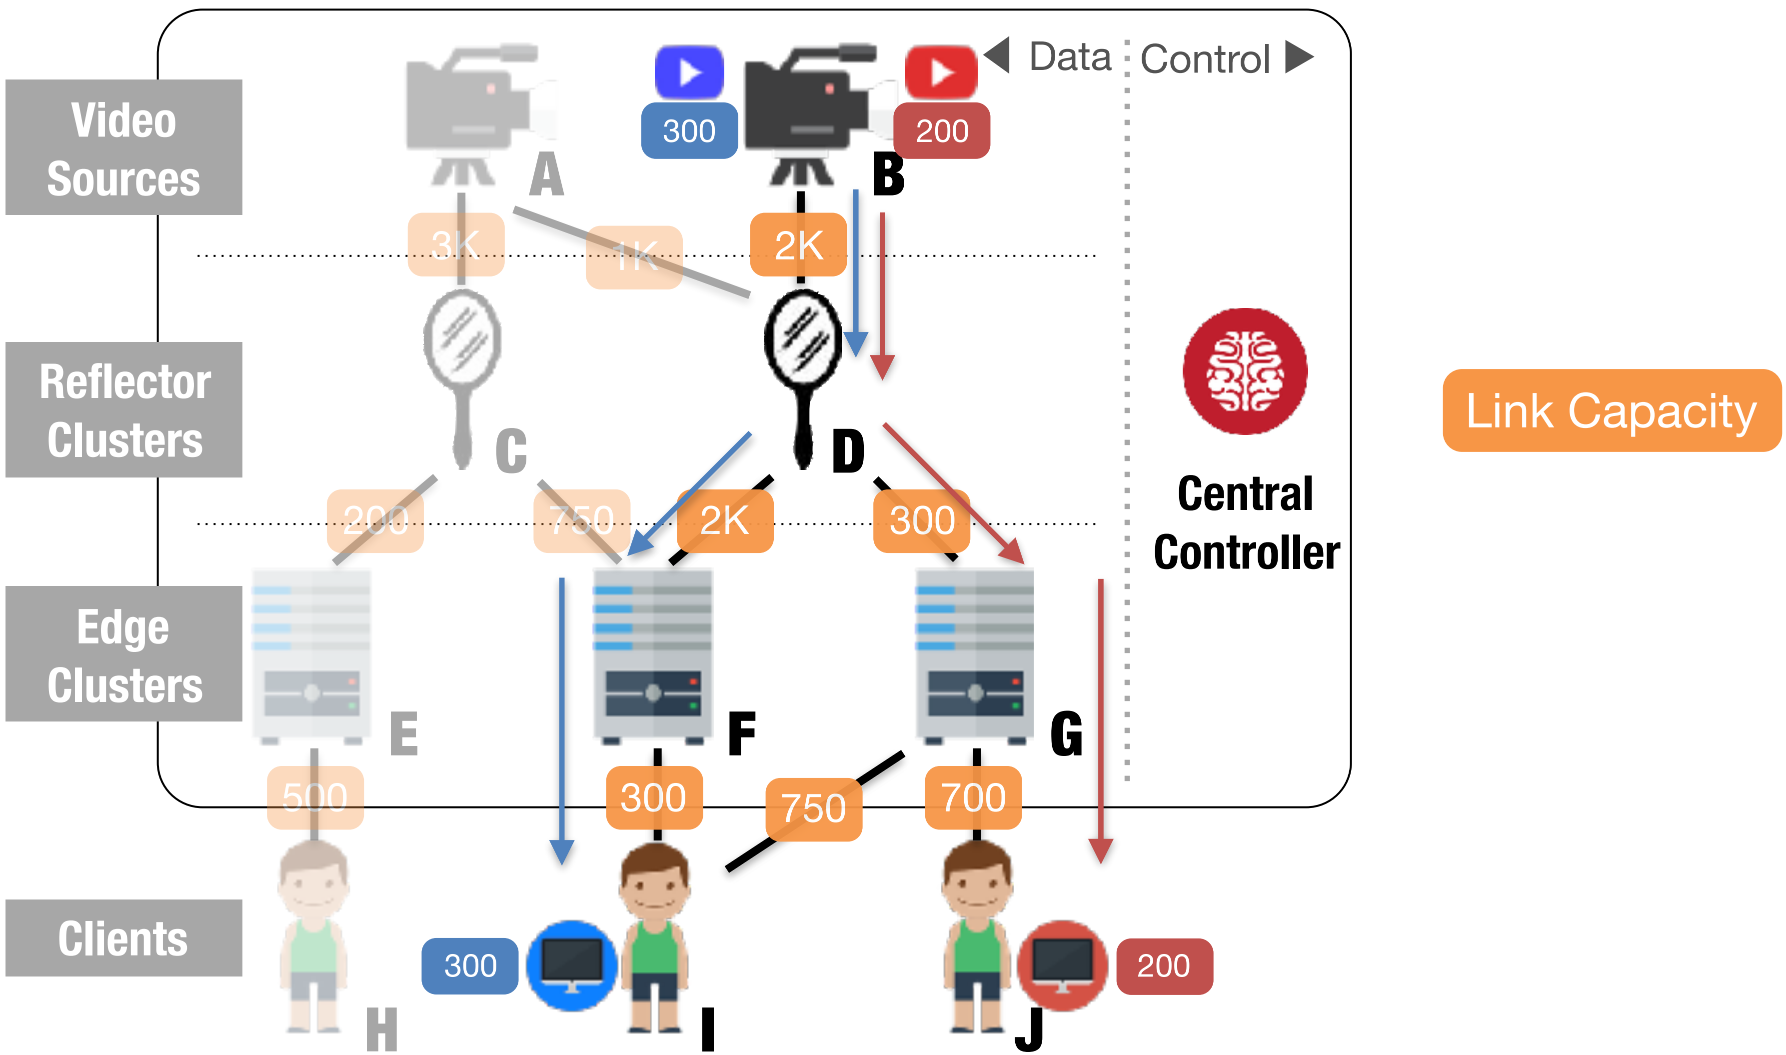

Solving Centralized Optimization

MAXIMIZE	SERVICE QUALITY
MINIMIZE	DELIVERY COST
SUBJECT TO	DON'T EXCEED LINK CAPACITY
	SENDER MUST HAVE RECEIVED VIDEO

Solving Centralized Optimization

SERVICE QUALITY

$$\max \; w_s \cdot \sum_{l \in L_{AS}, o \in O} \text{Priority}_o \cdot \text{Request}_{l,o} \cdot \text{Serves}_{l,o}$$
$$- \; w_c \cdot \sum_{l \in L, o \in O} \text{Cost}(l) \cdot \text{Bitrate}(o) \cdot \text{Serves}_{l,o}$$

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

Flexibility of Centralized Optimization

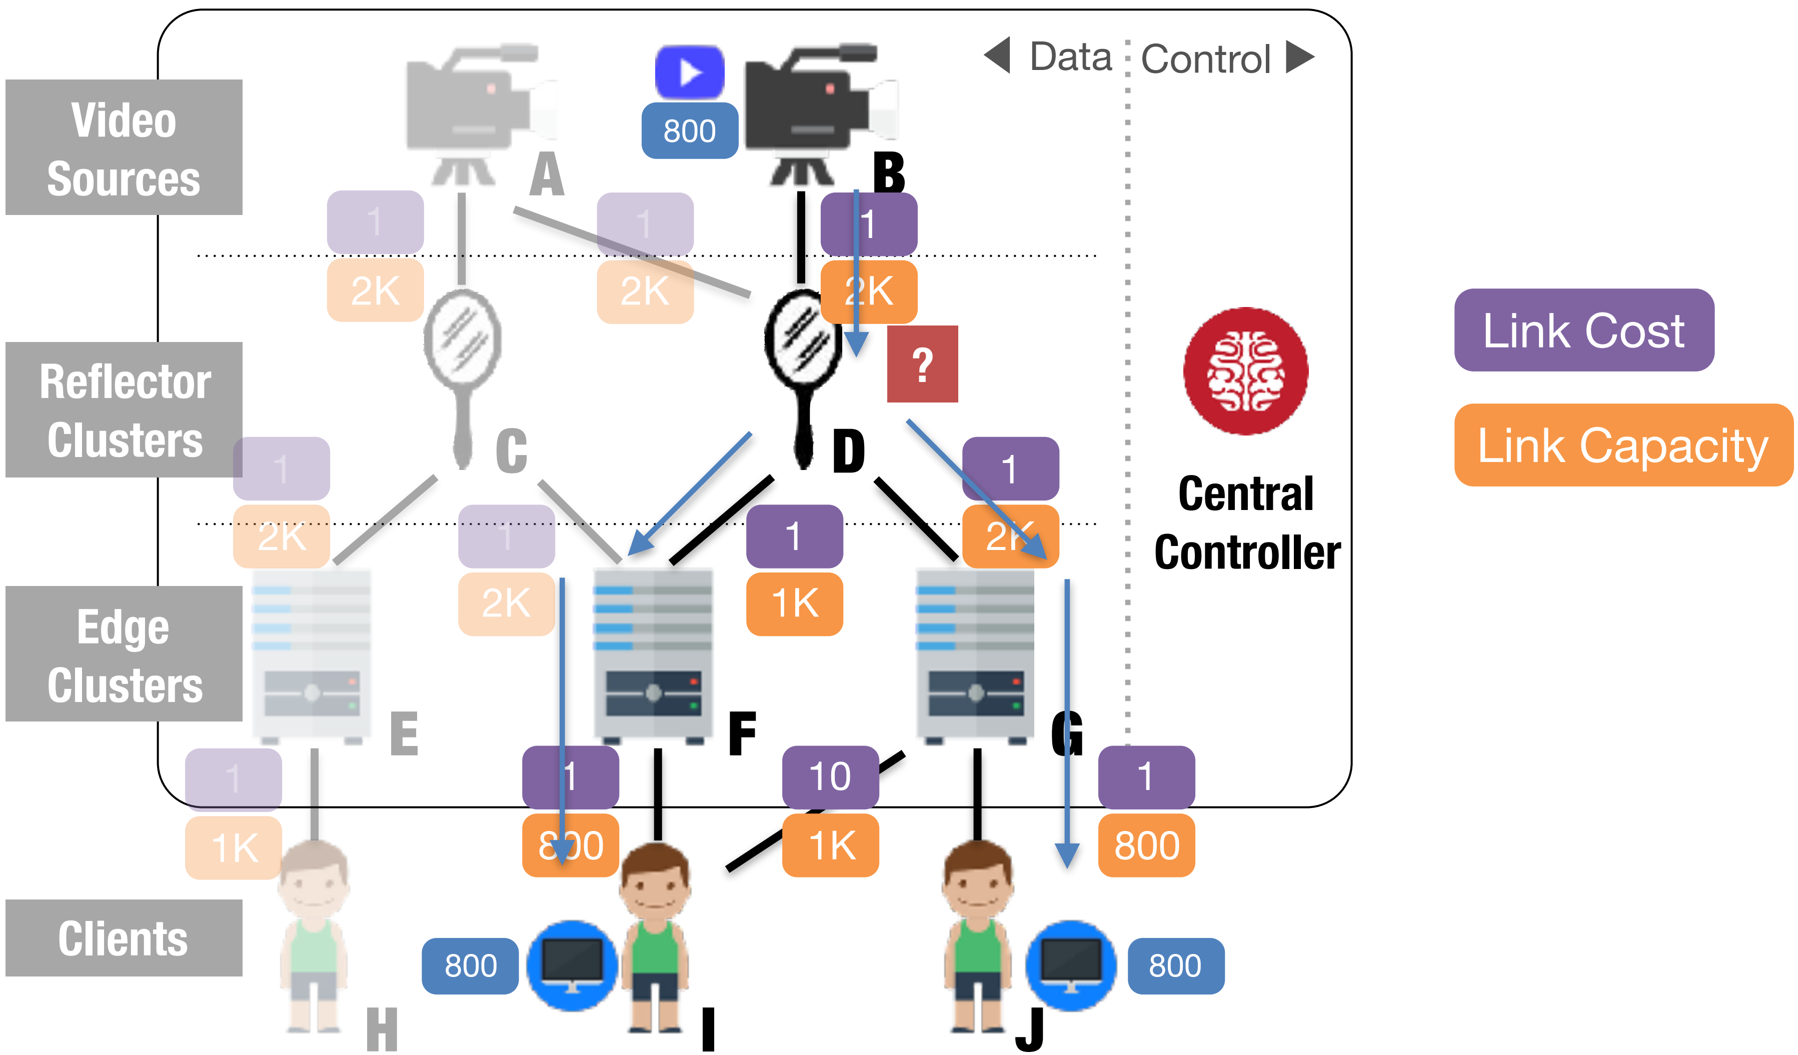

Flexibility of Centralized Optimization

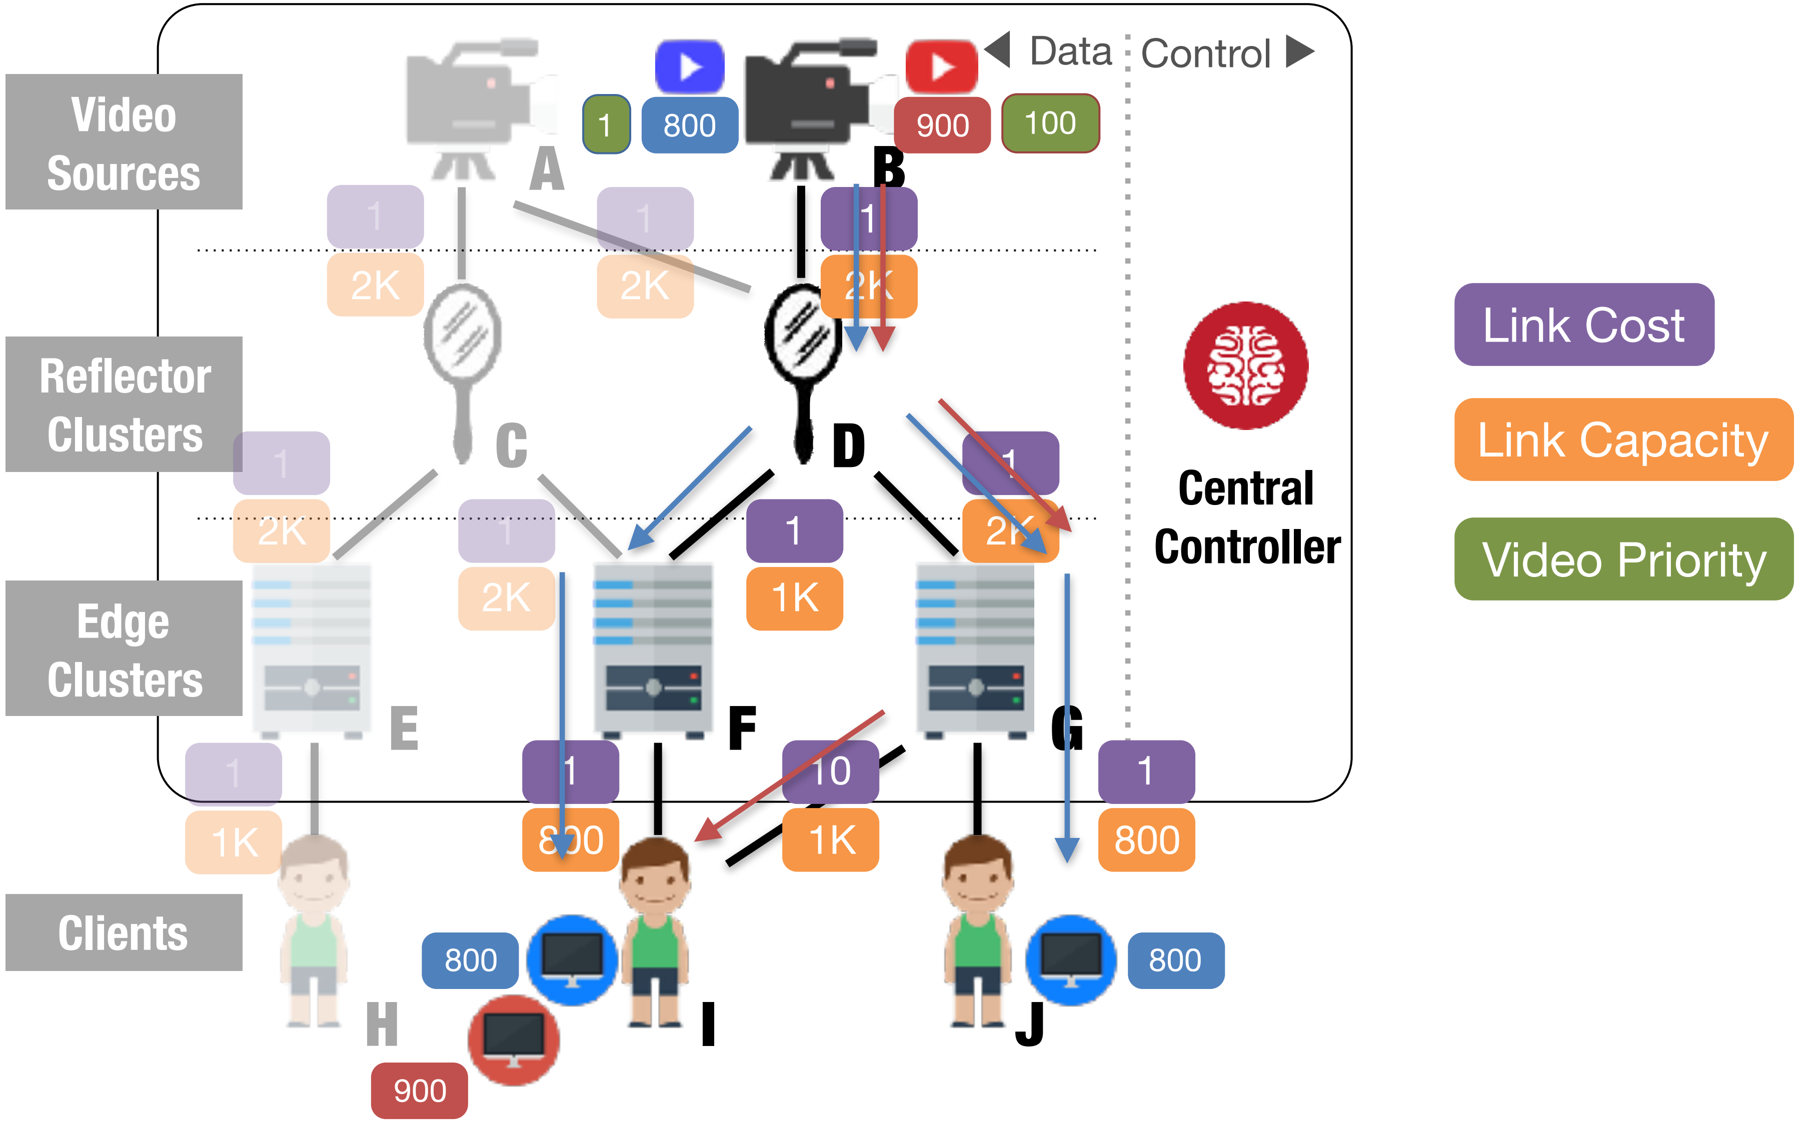

Centralized Optimization

Service Quality

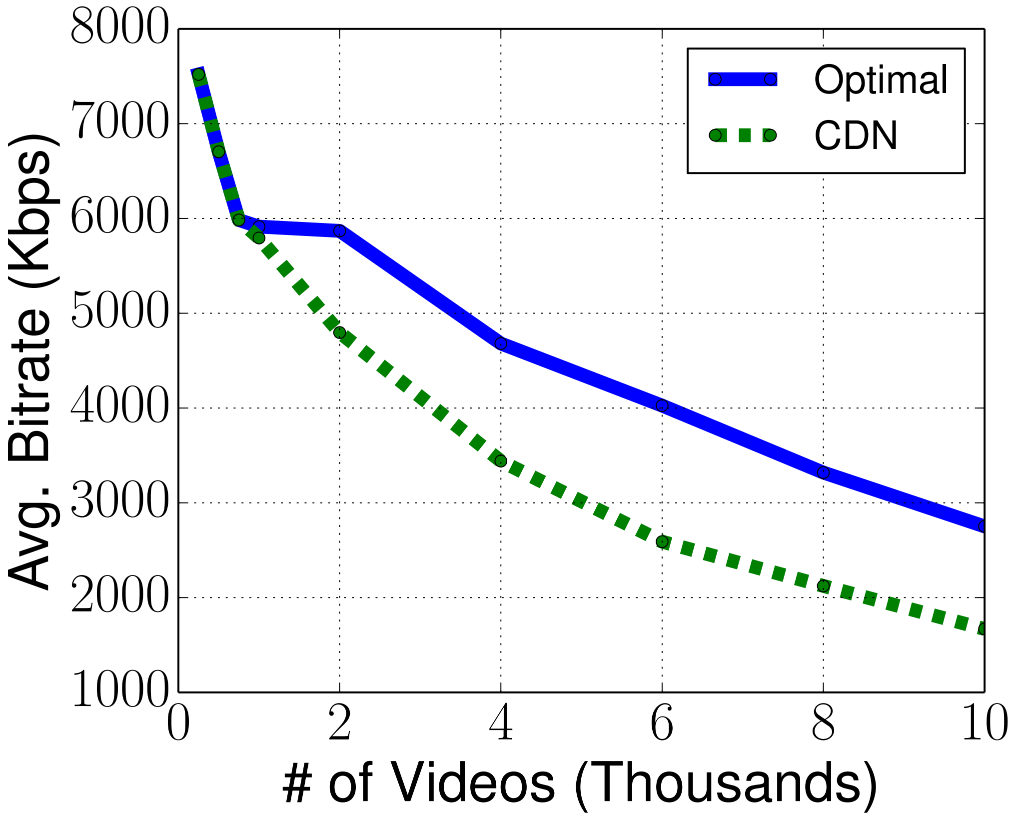

Simulation using Conviva traces, modeling user-generated content

Delivery Cost

(per request)

Simulation using Conviva traces, modeling large sports events

Centralized Optimization

Service Quality

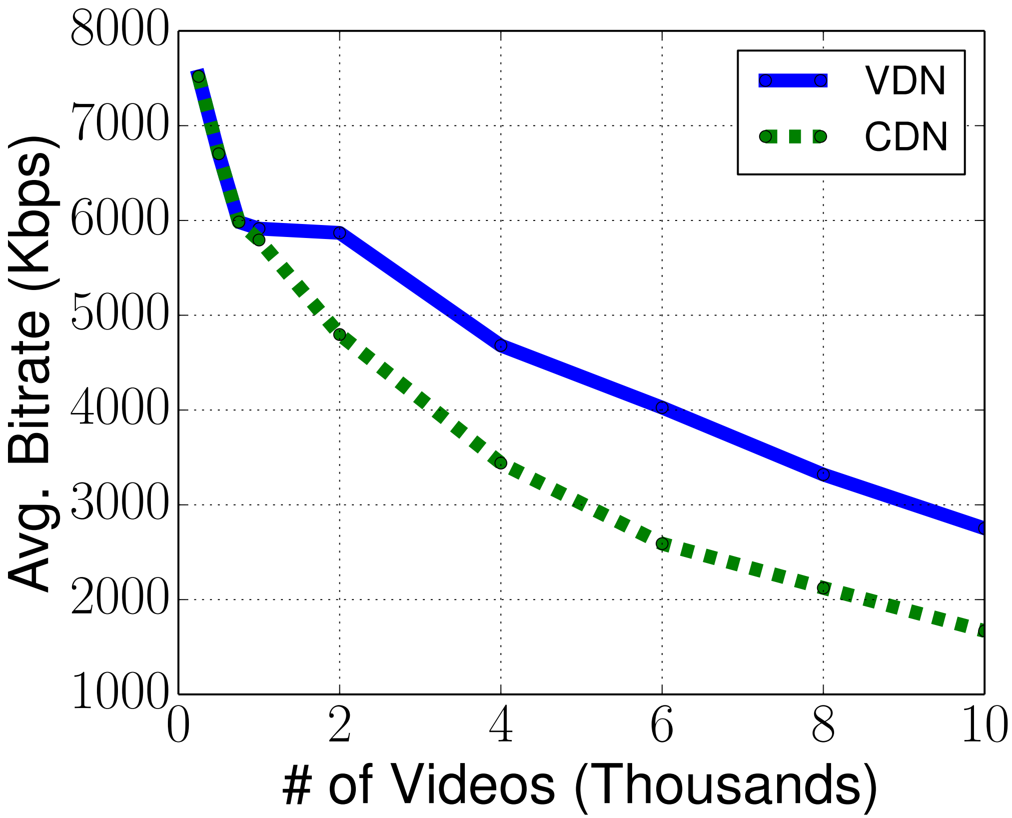

Simulation using Conviva traces, modeling user-generated content

Delivery Cost

(per request)

CDN
2.0x

VDN
1.0x

Simulation using Conviva traces, modeling large sports events

Unfortunately… No Free Lunch

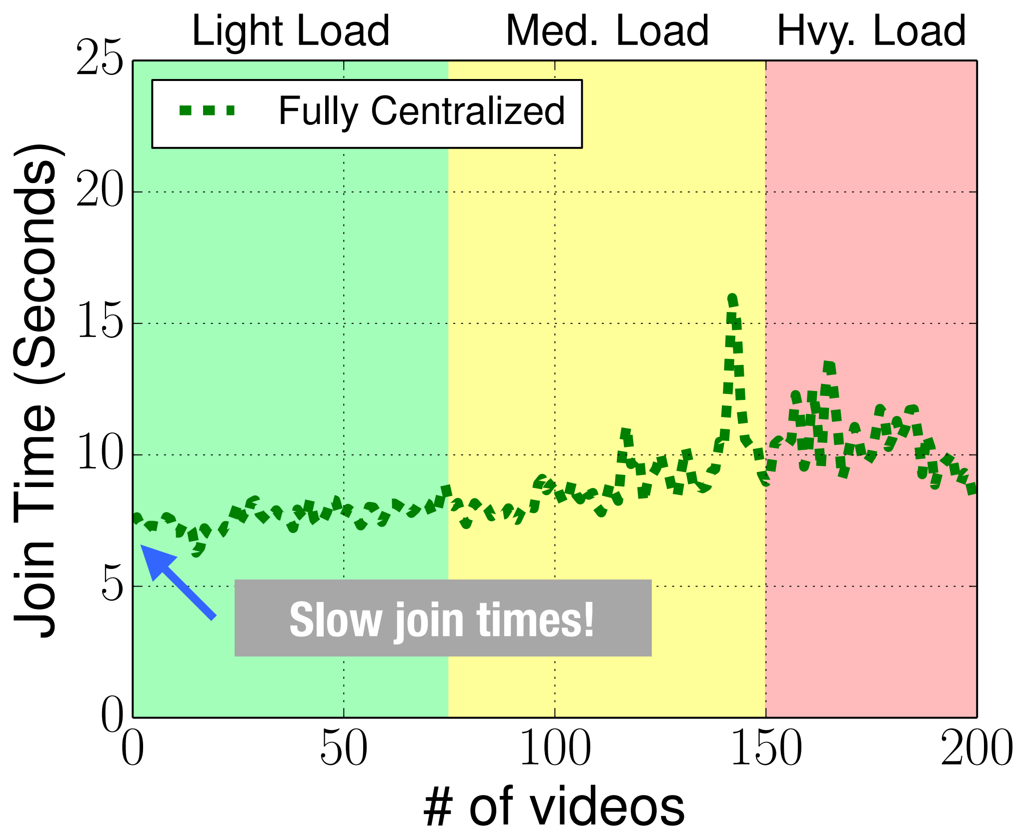

Experiments on EC2 nodes with a centralized controller at CMU across the Internet

Problems with Centralization

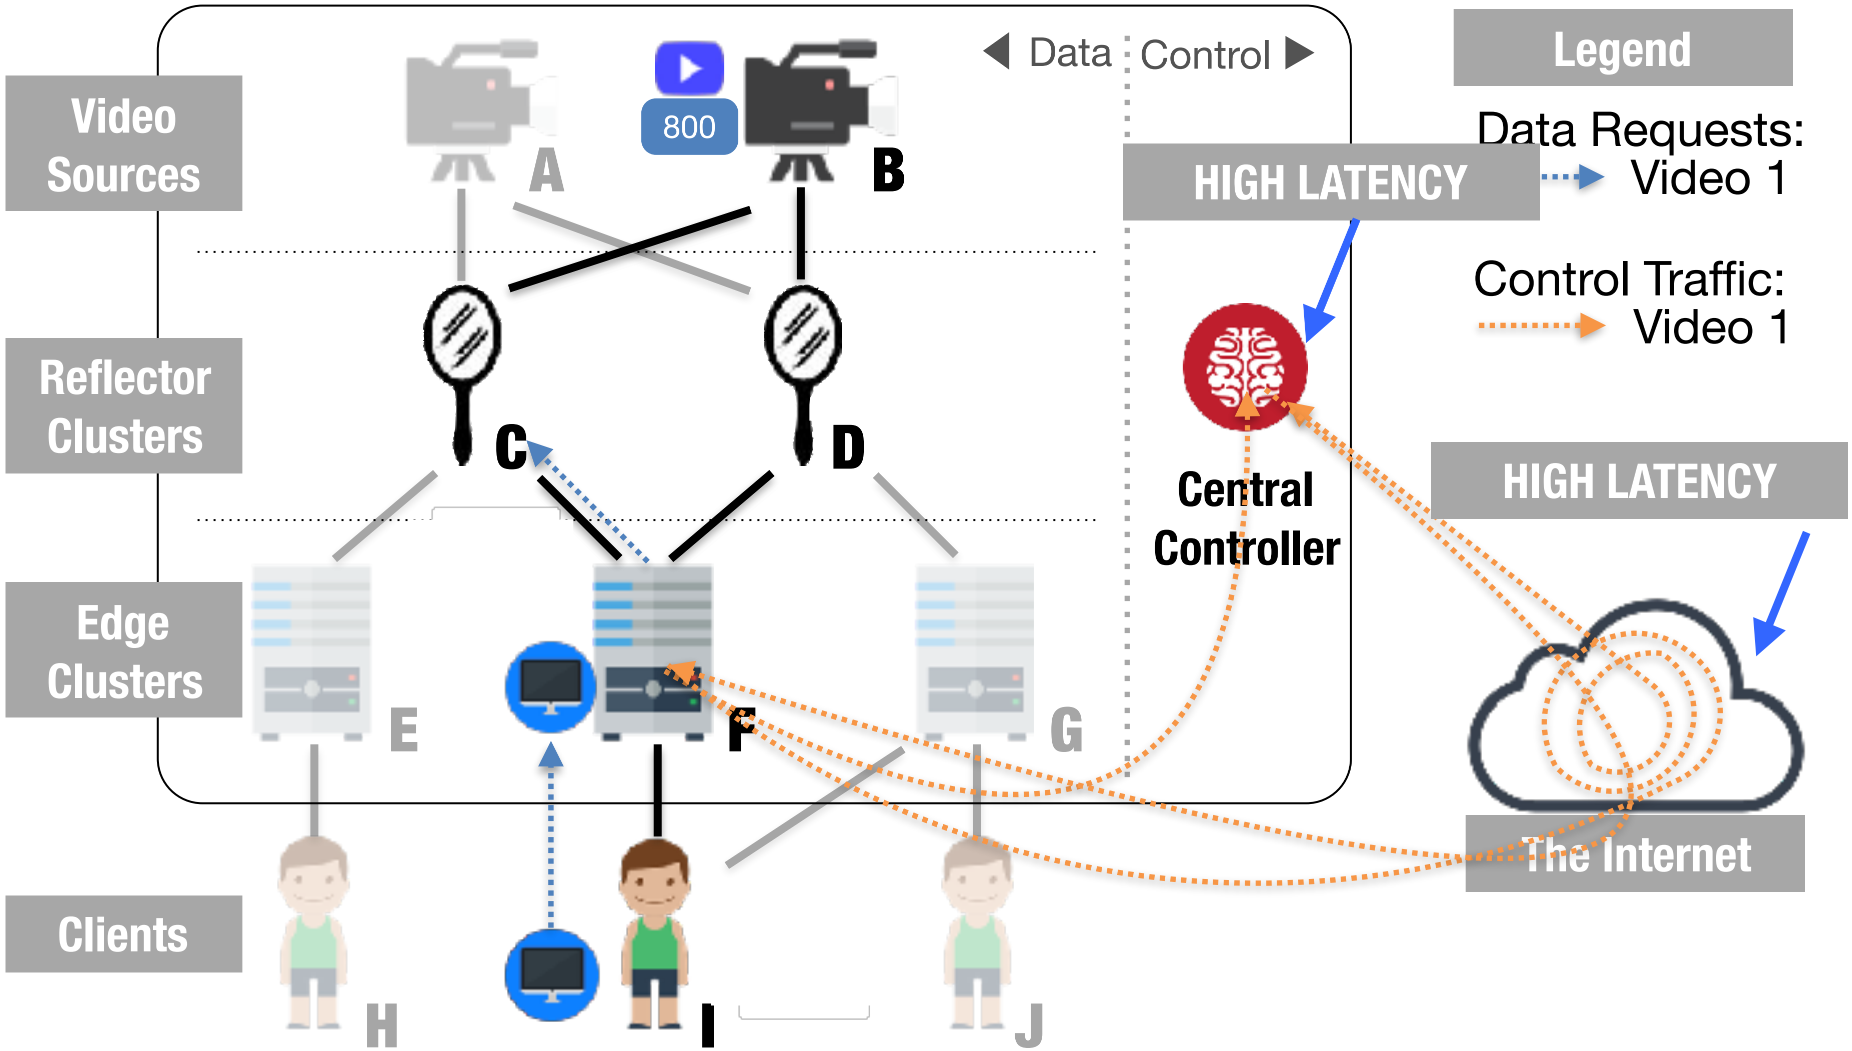

Outline

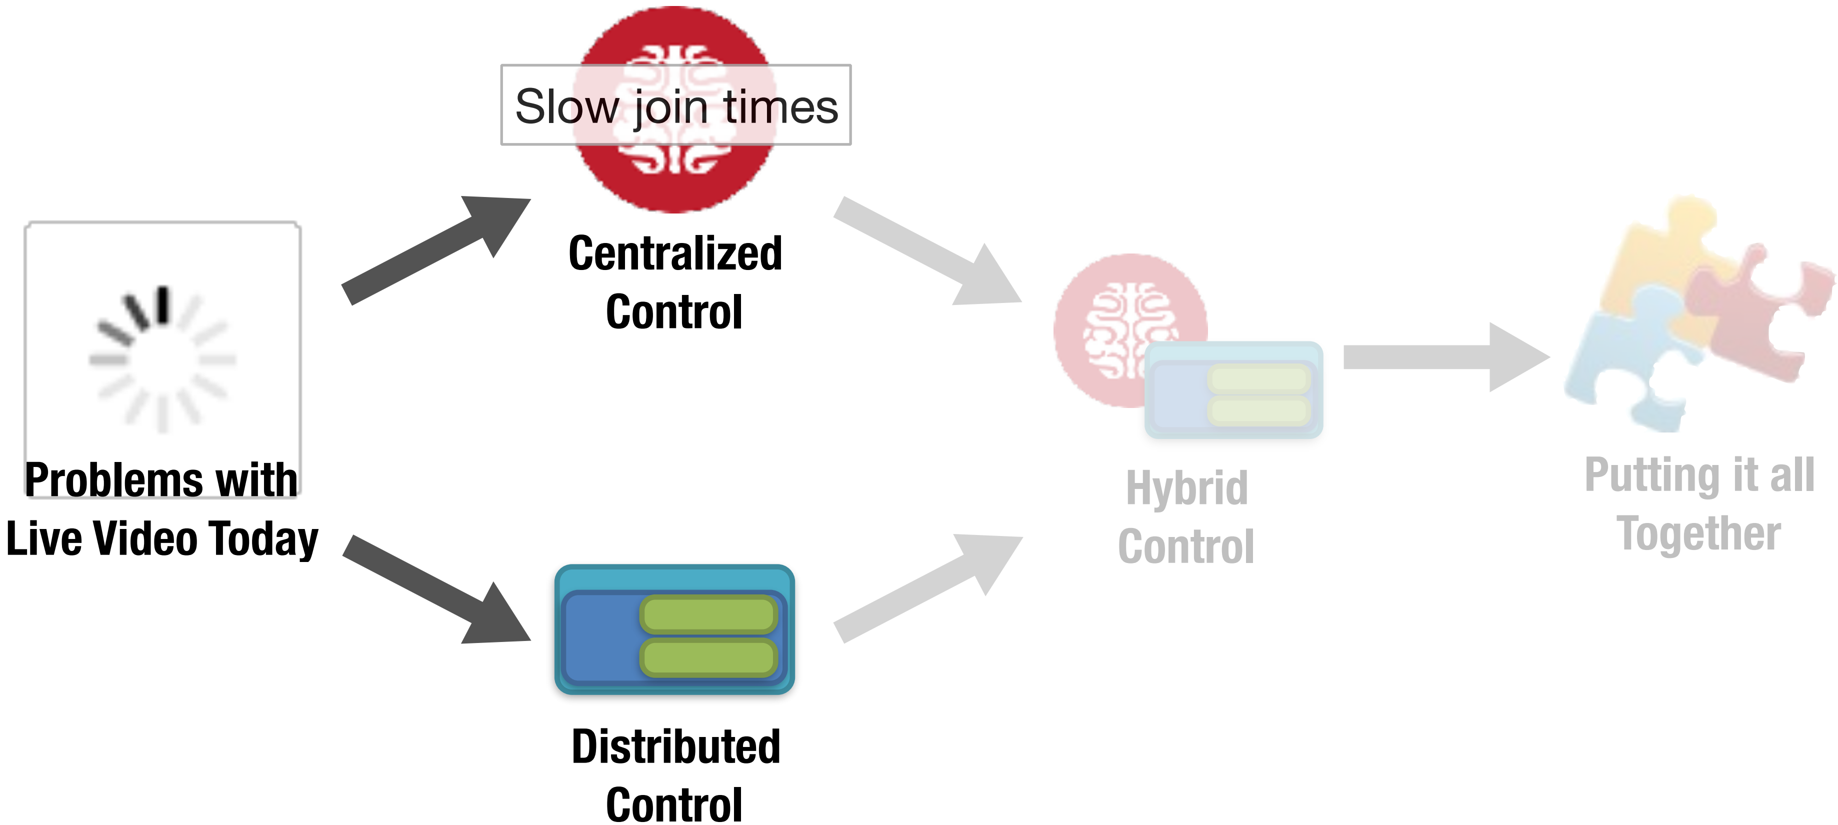
Slow join times
Centralized Control
Problems with Live Video Today
Hybrid Control
Putting it all Together
Distributed Control

Alternate Approach: Distributed

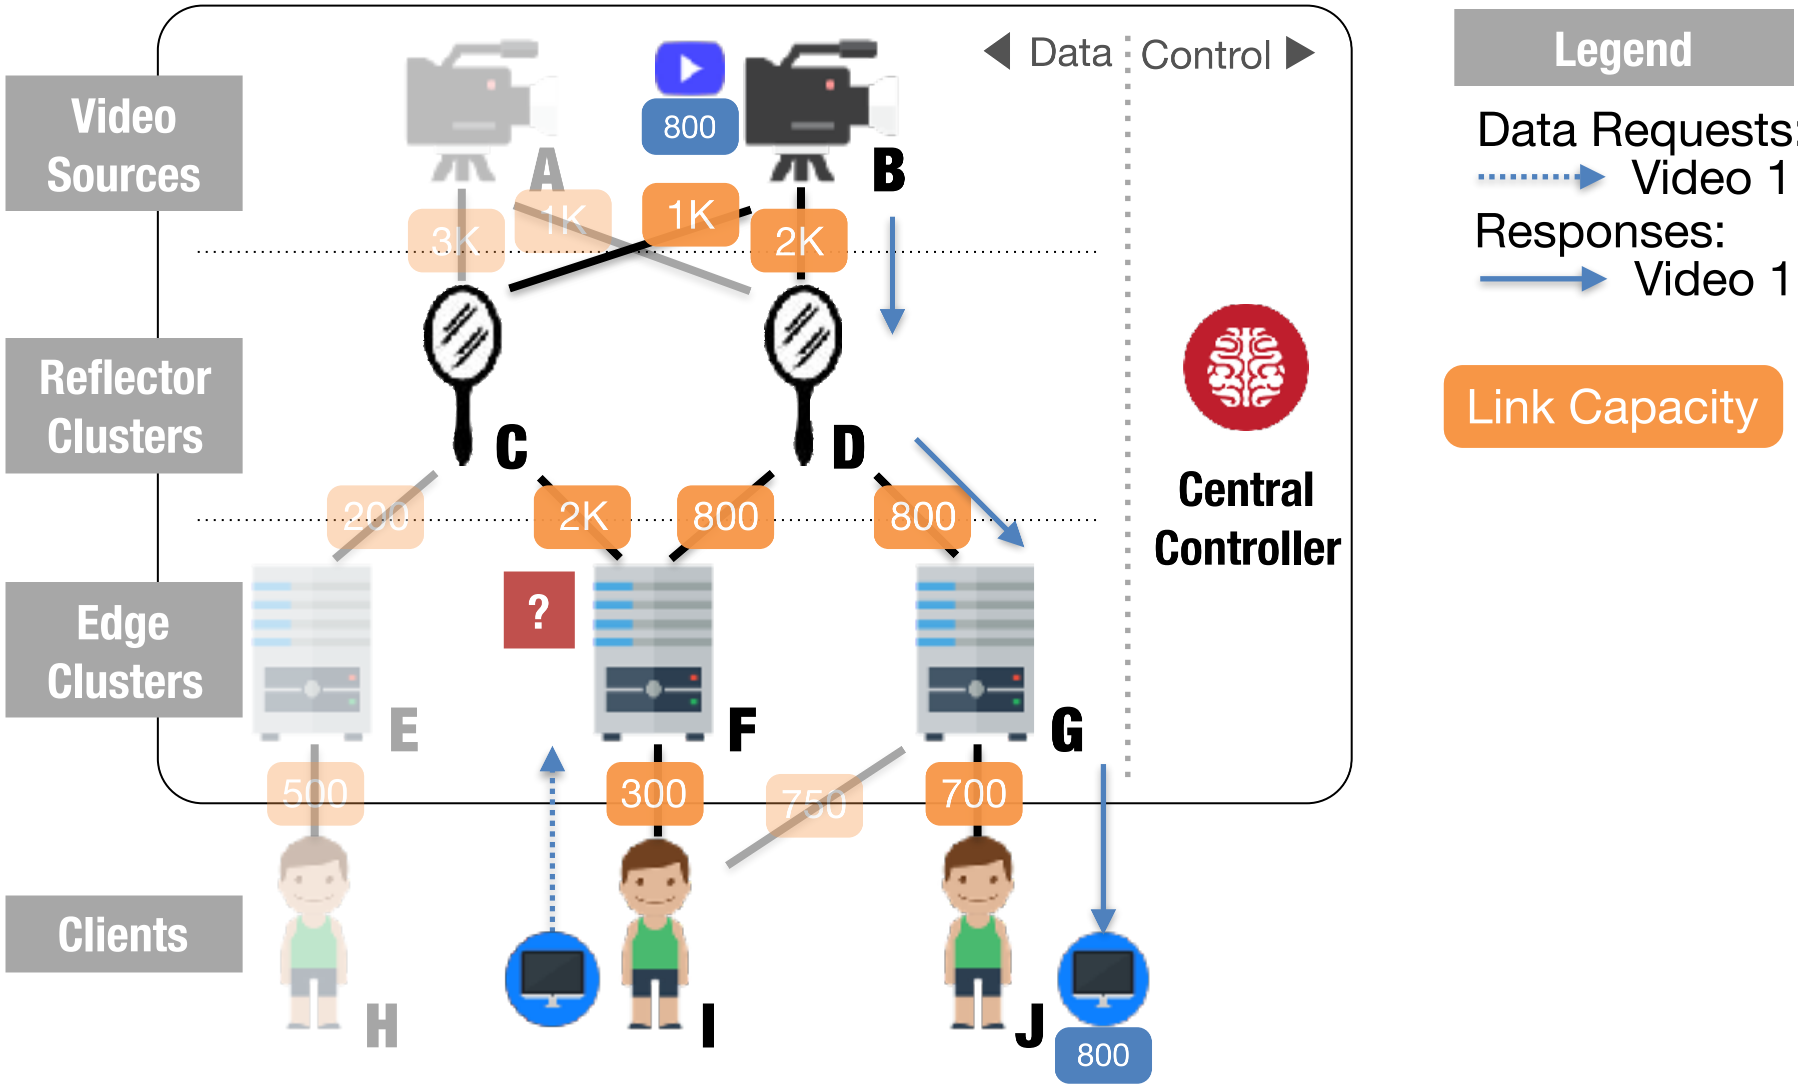

Alternate Approach: Distributed

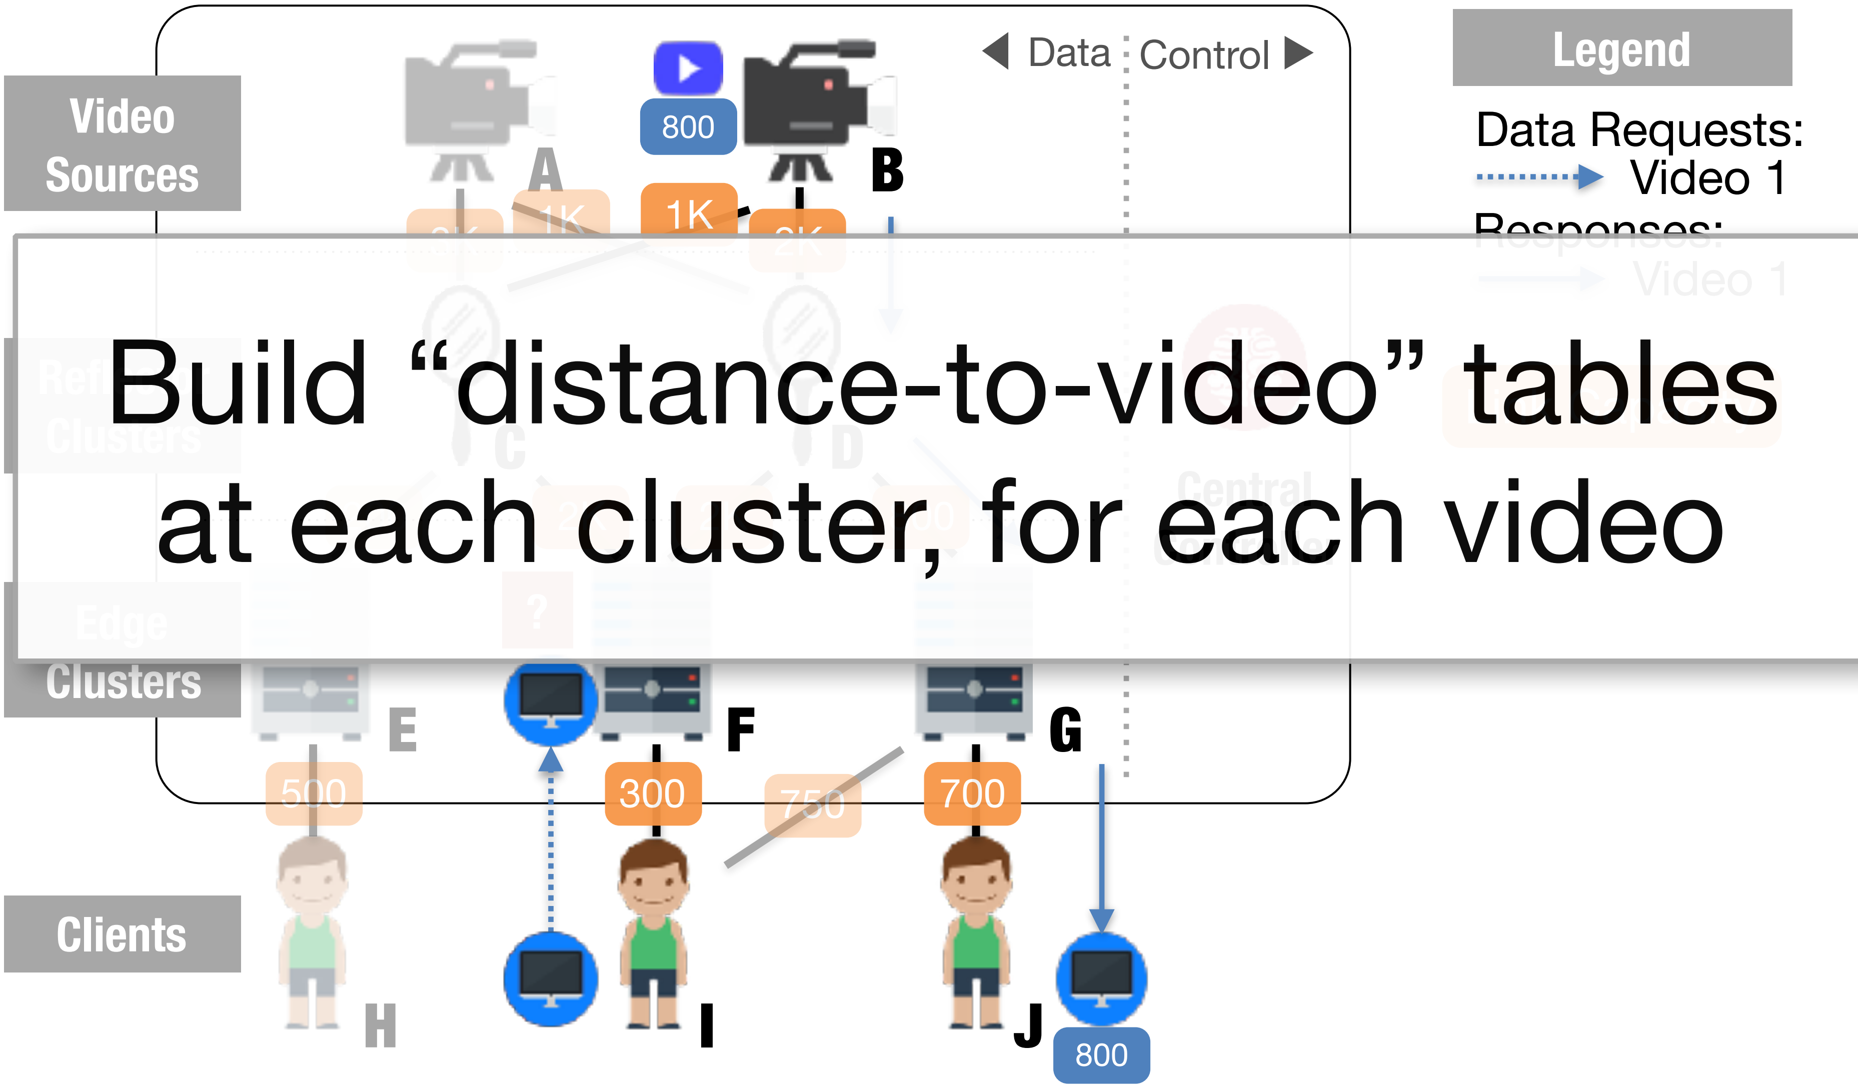

Build “distance-to-video” tables at each cluster, for each video

Alternate Approach: Distributed

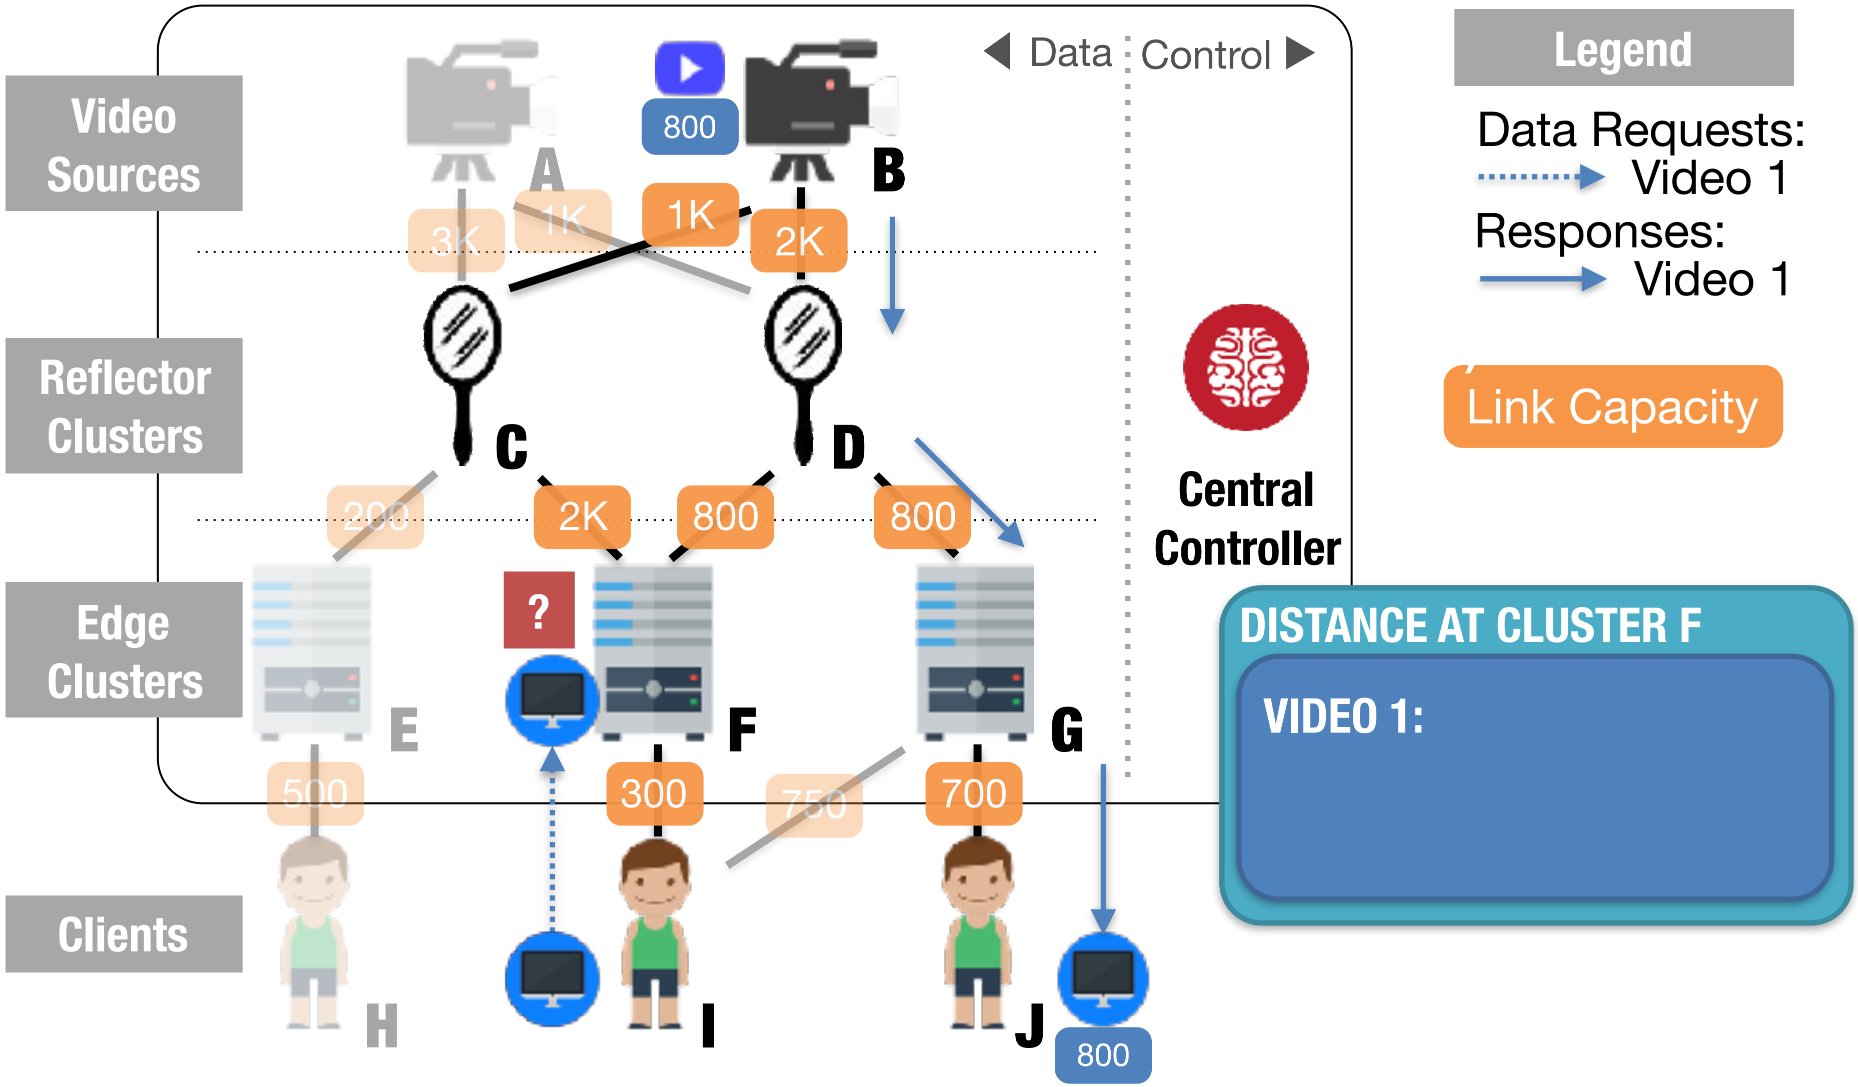

Alternate Approach: Distributed

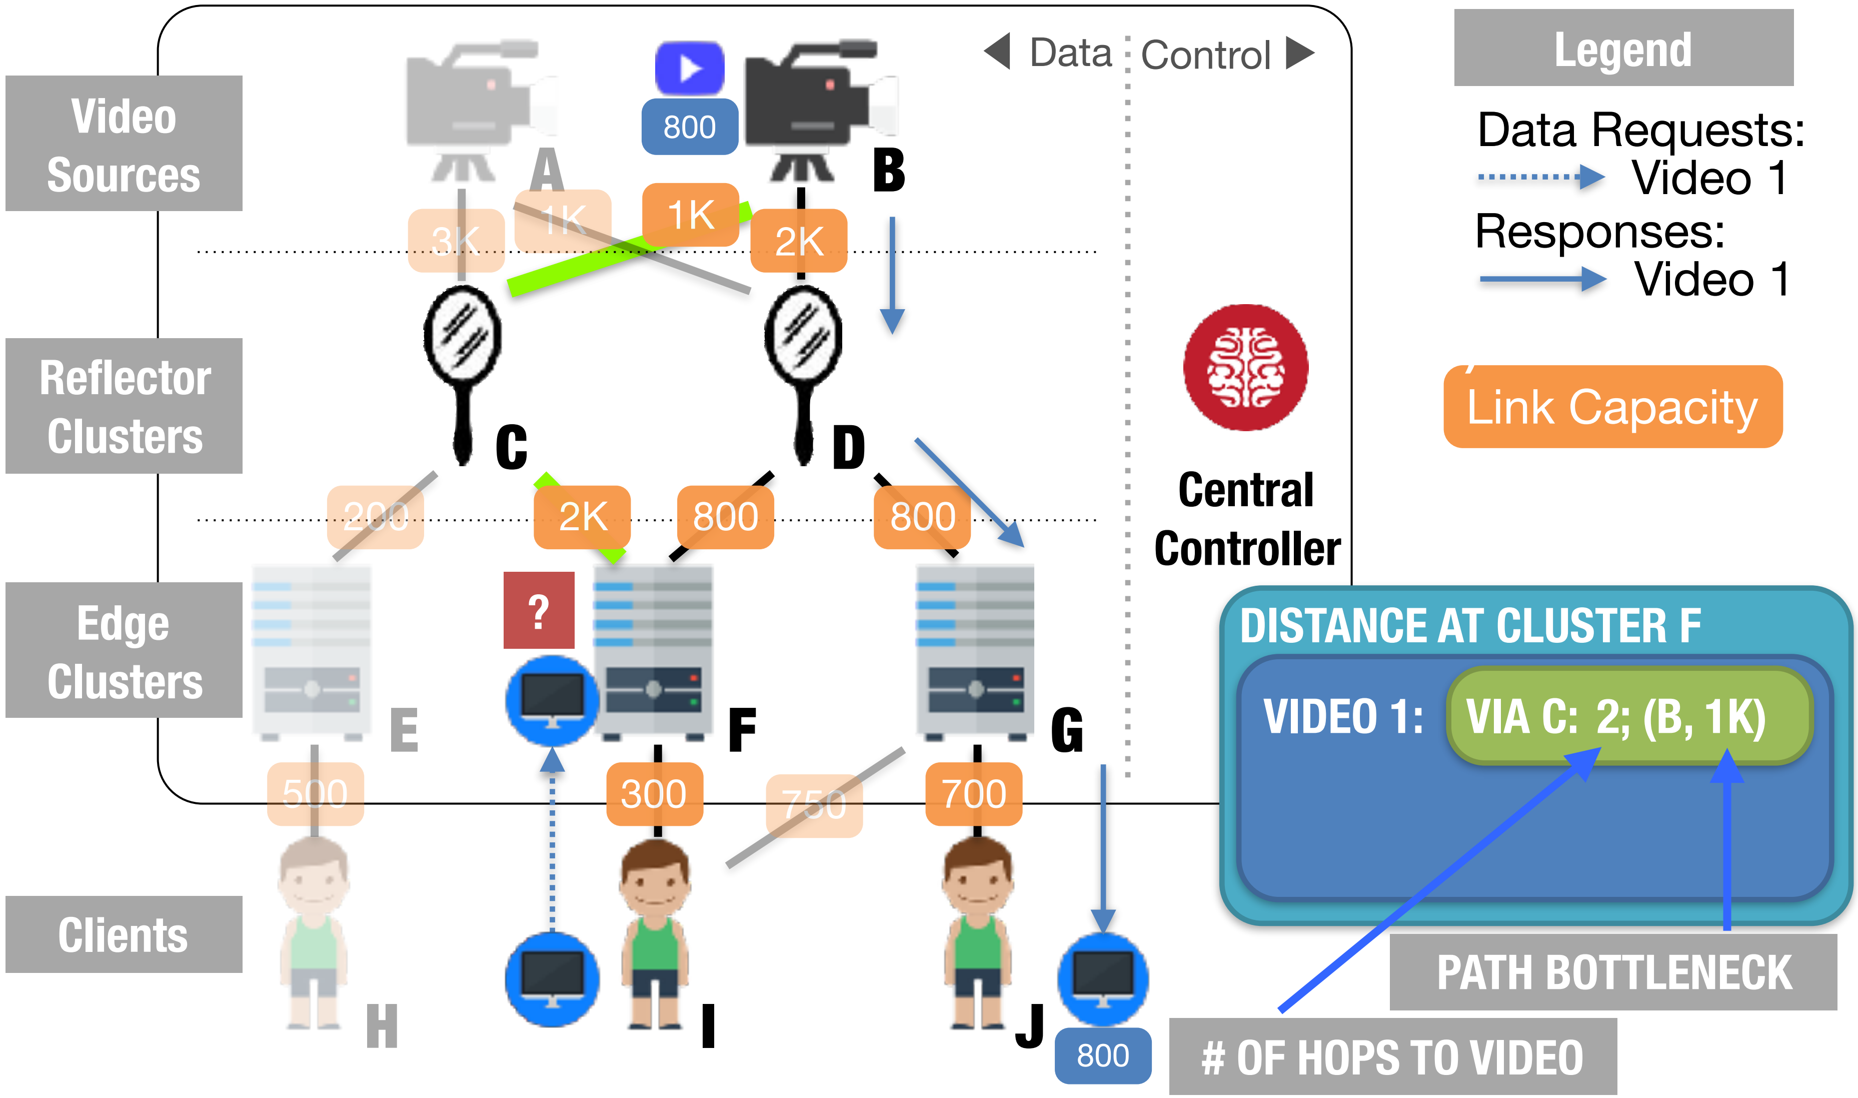

Alternate Approach: Distributed

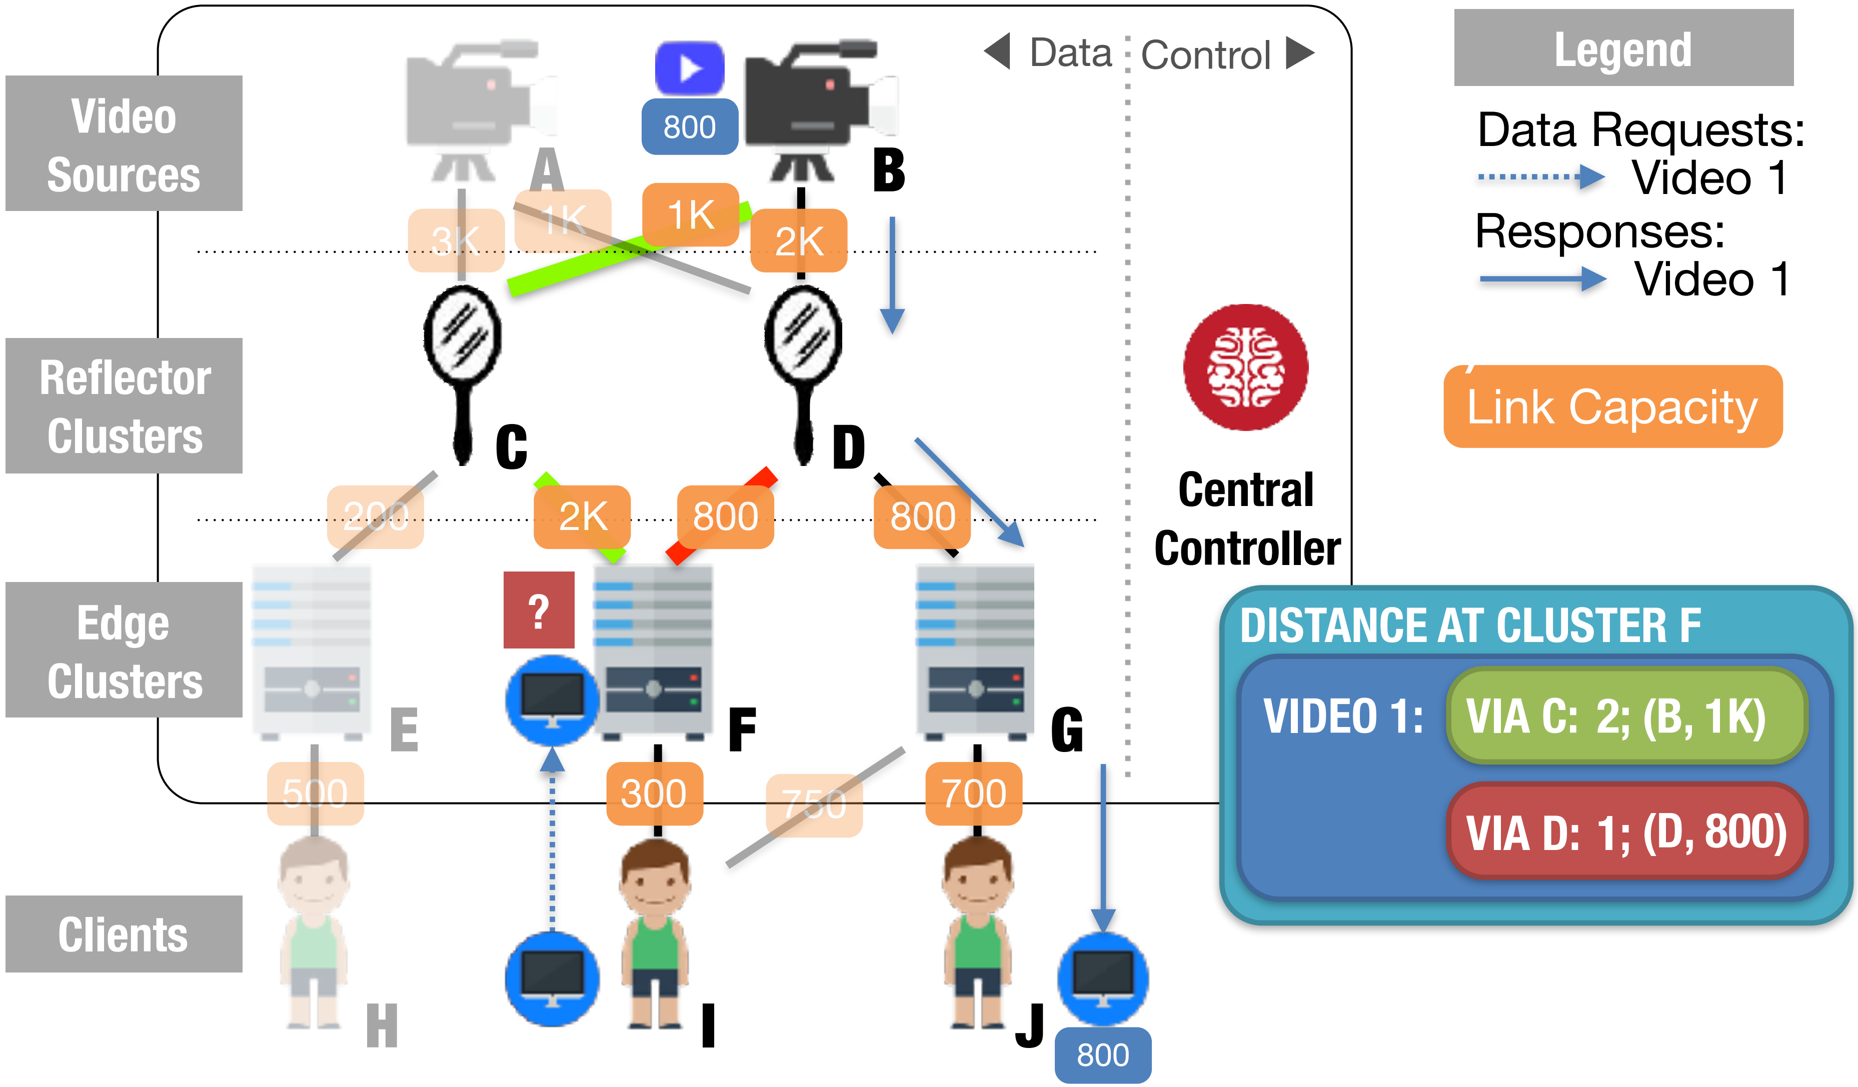

Alternate Approach: Distributed

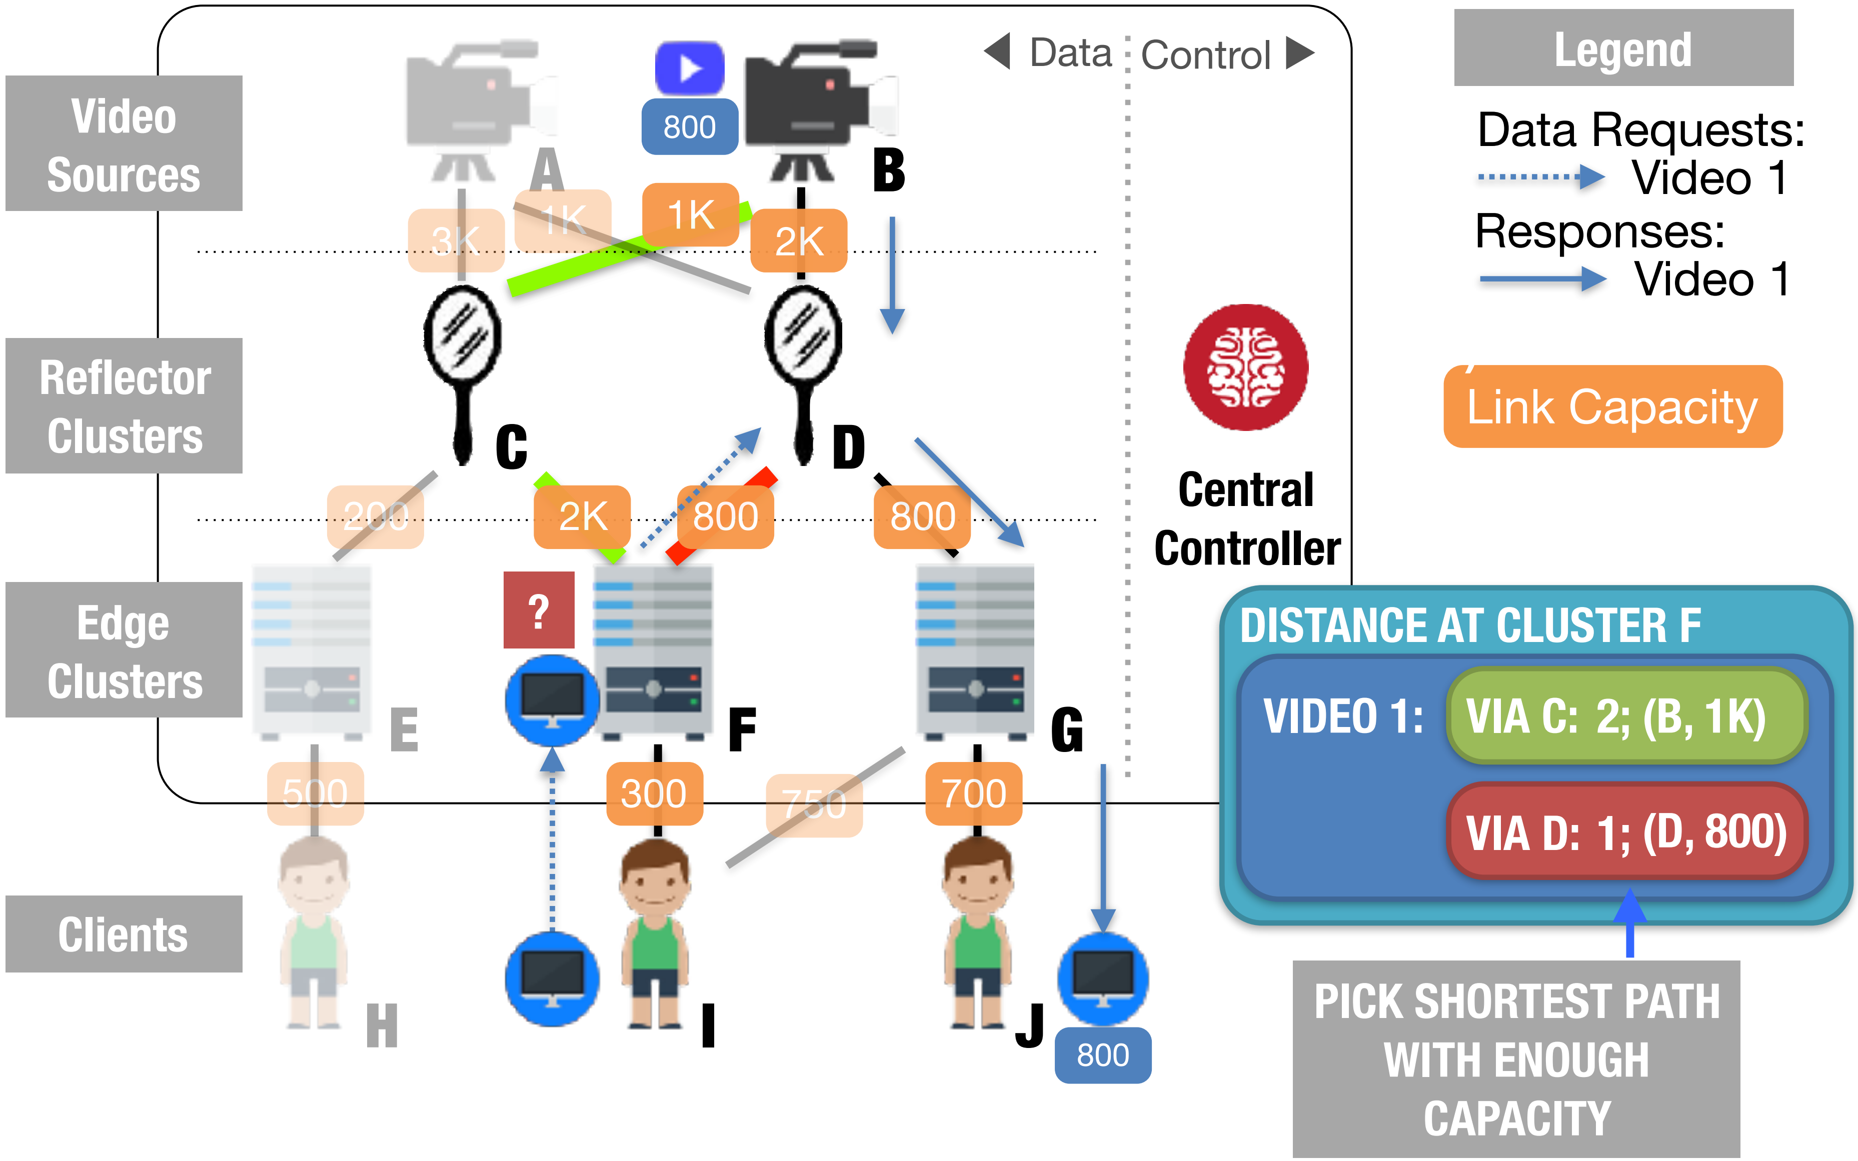

Alternate Approach: Distributed

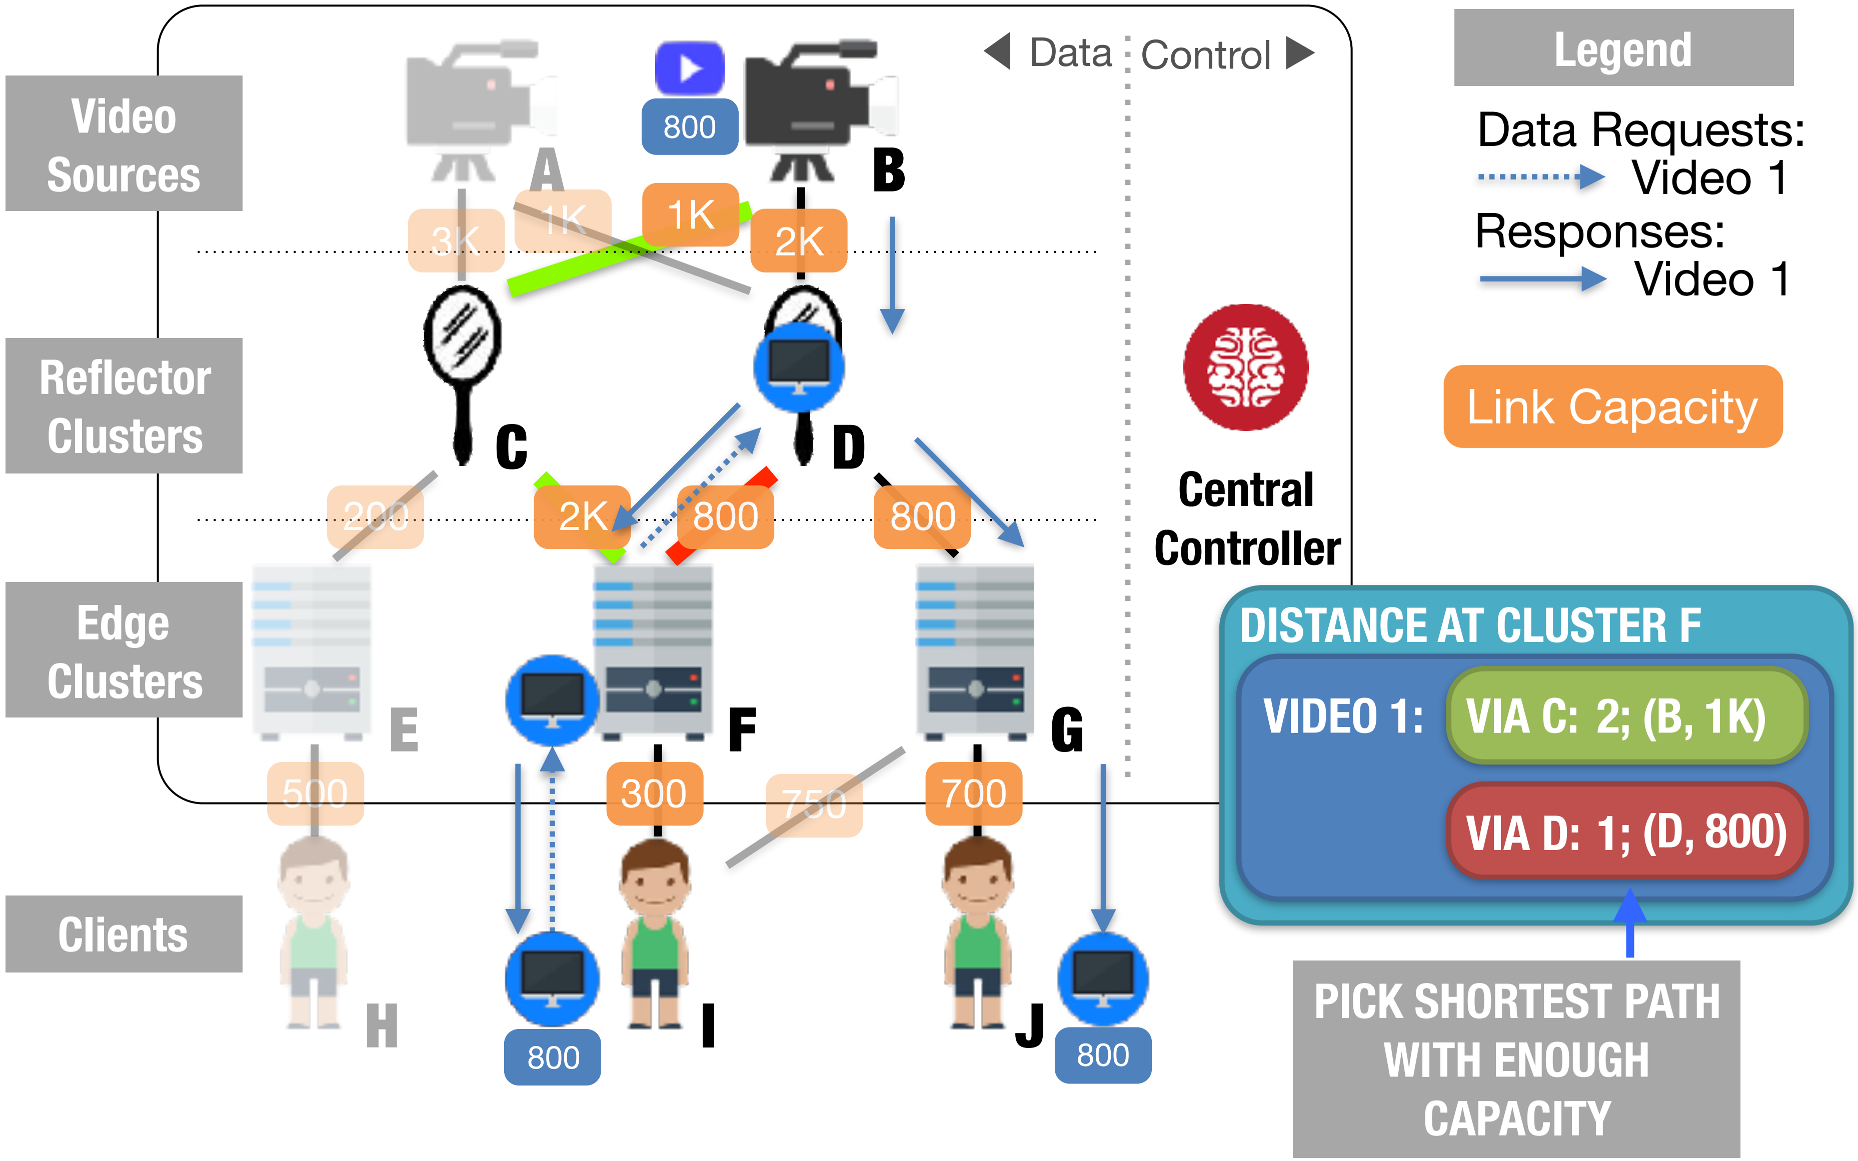

Alternate Approach: Distributed

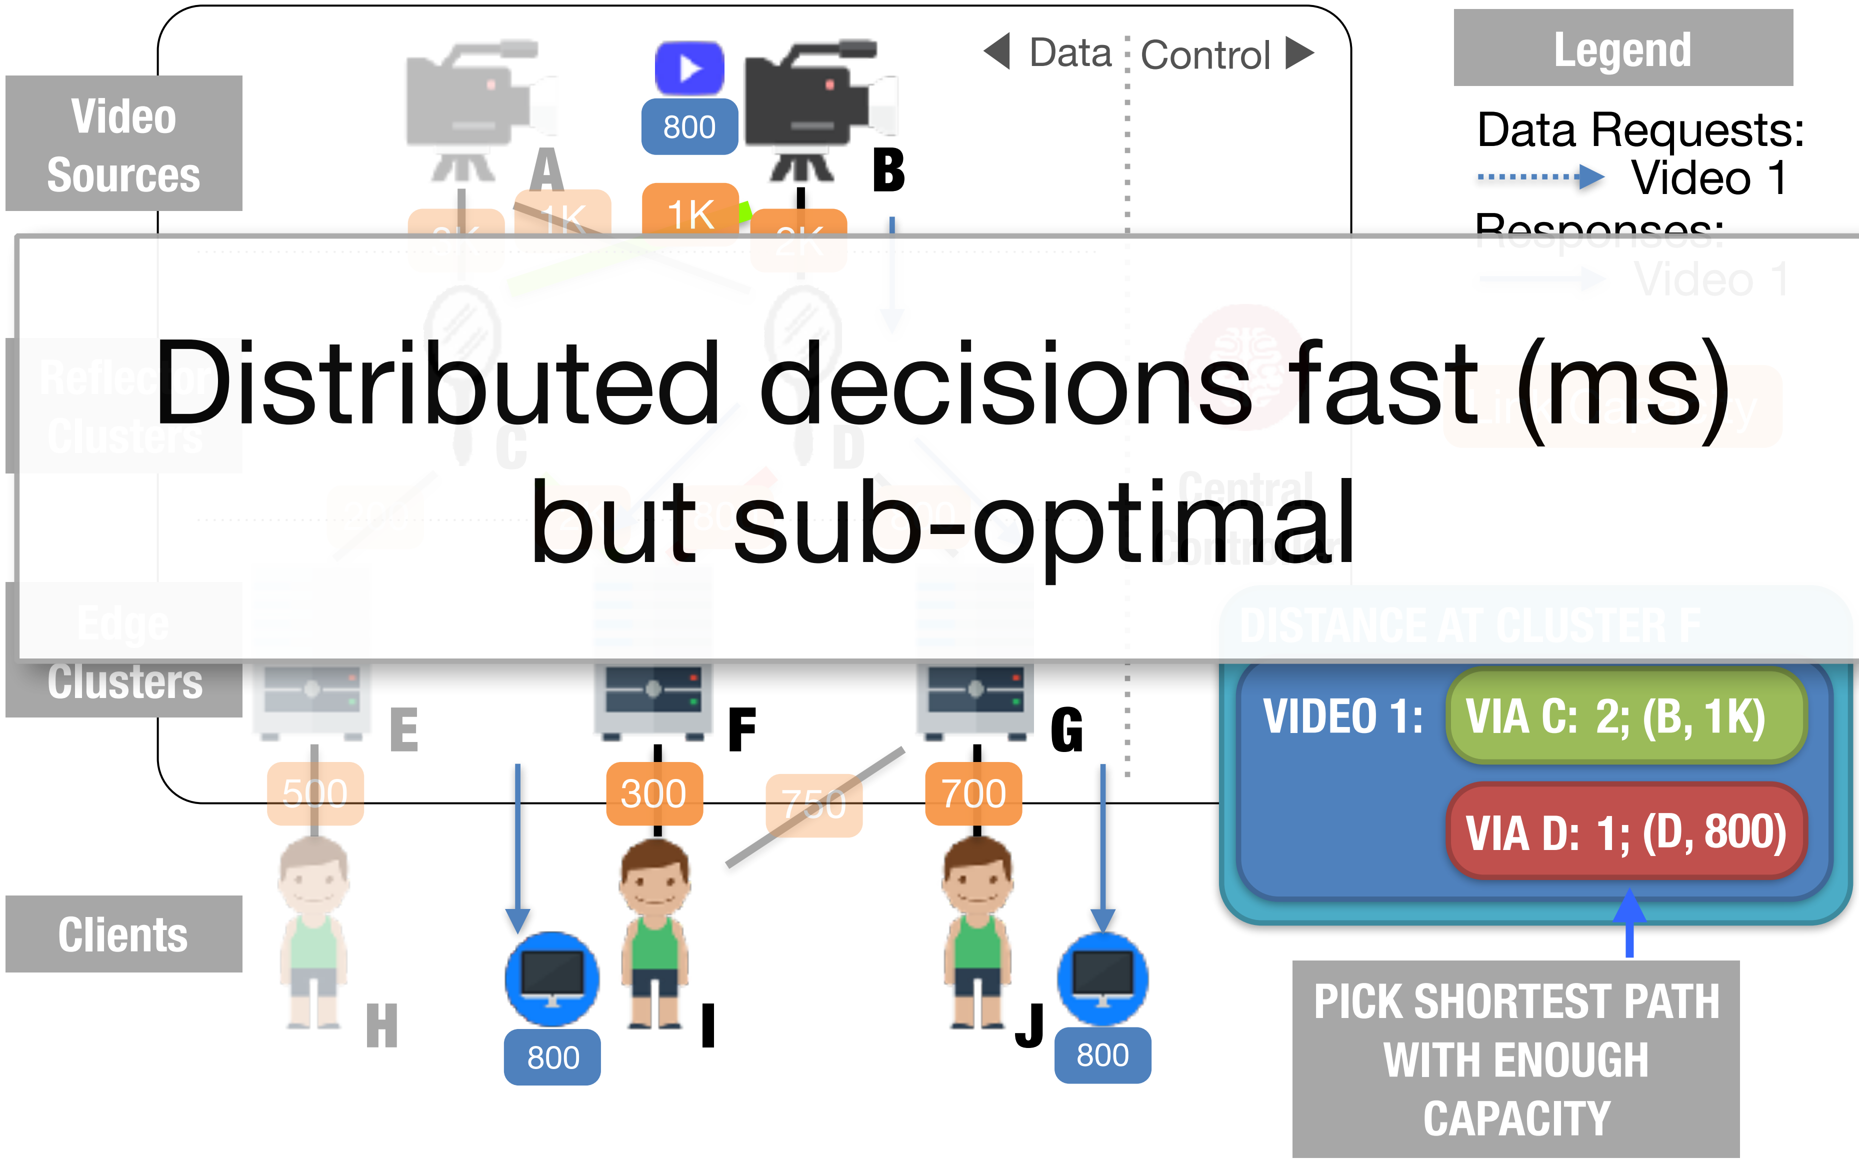

Distributed decisions fast (ms) but sub-optimal

Alternate Approach: Distributed

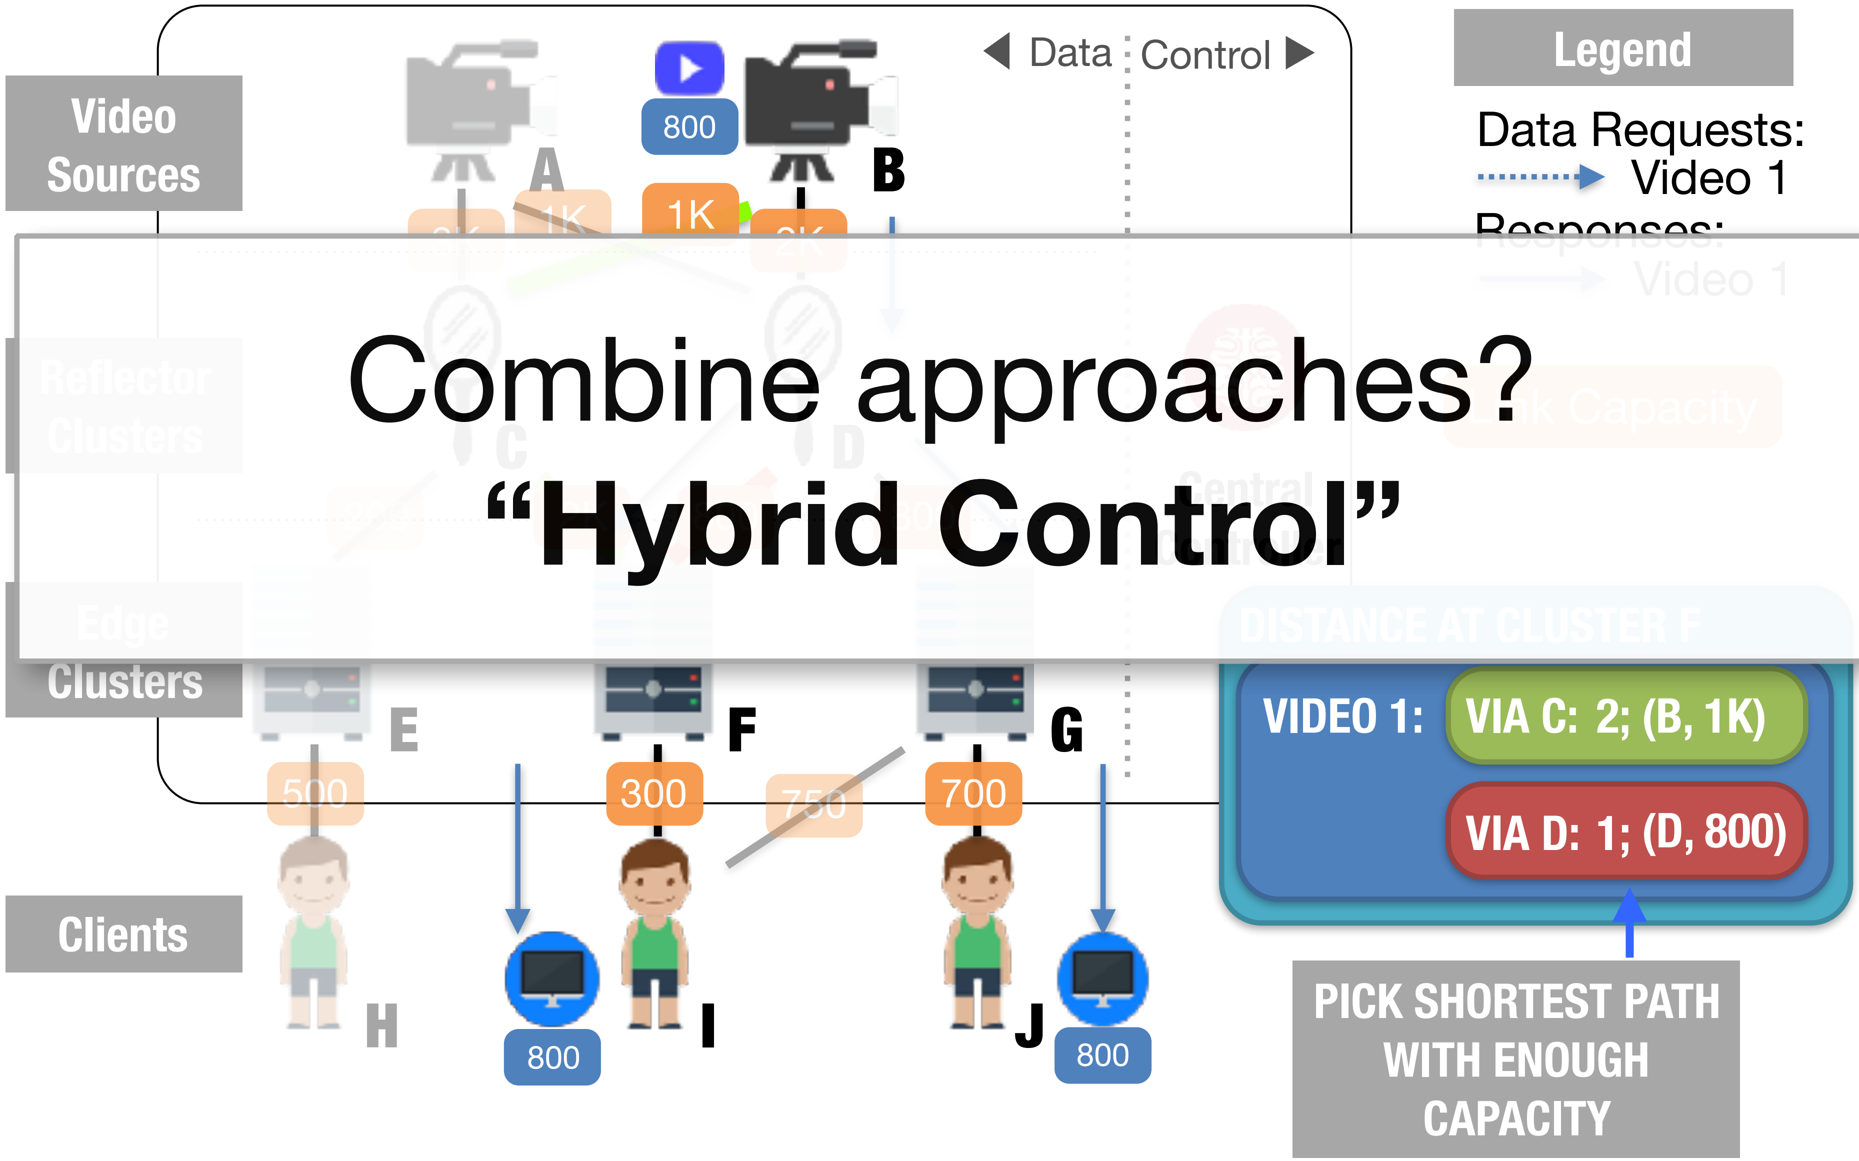

Combine approaches?

"Hybrid Control"

Outline

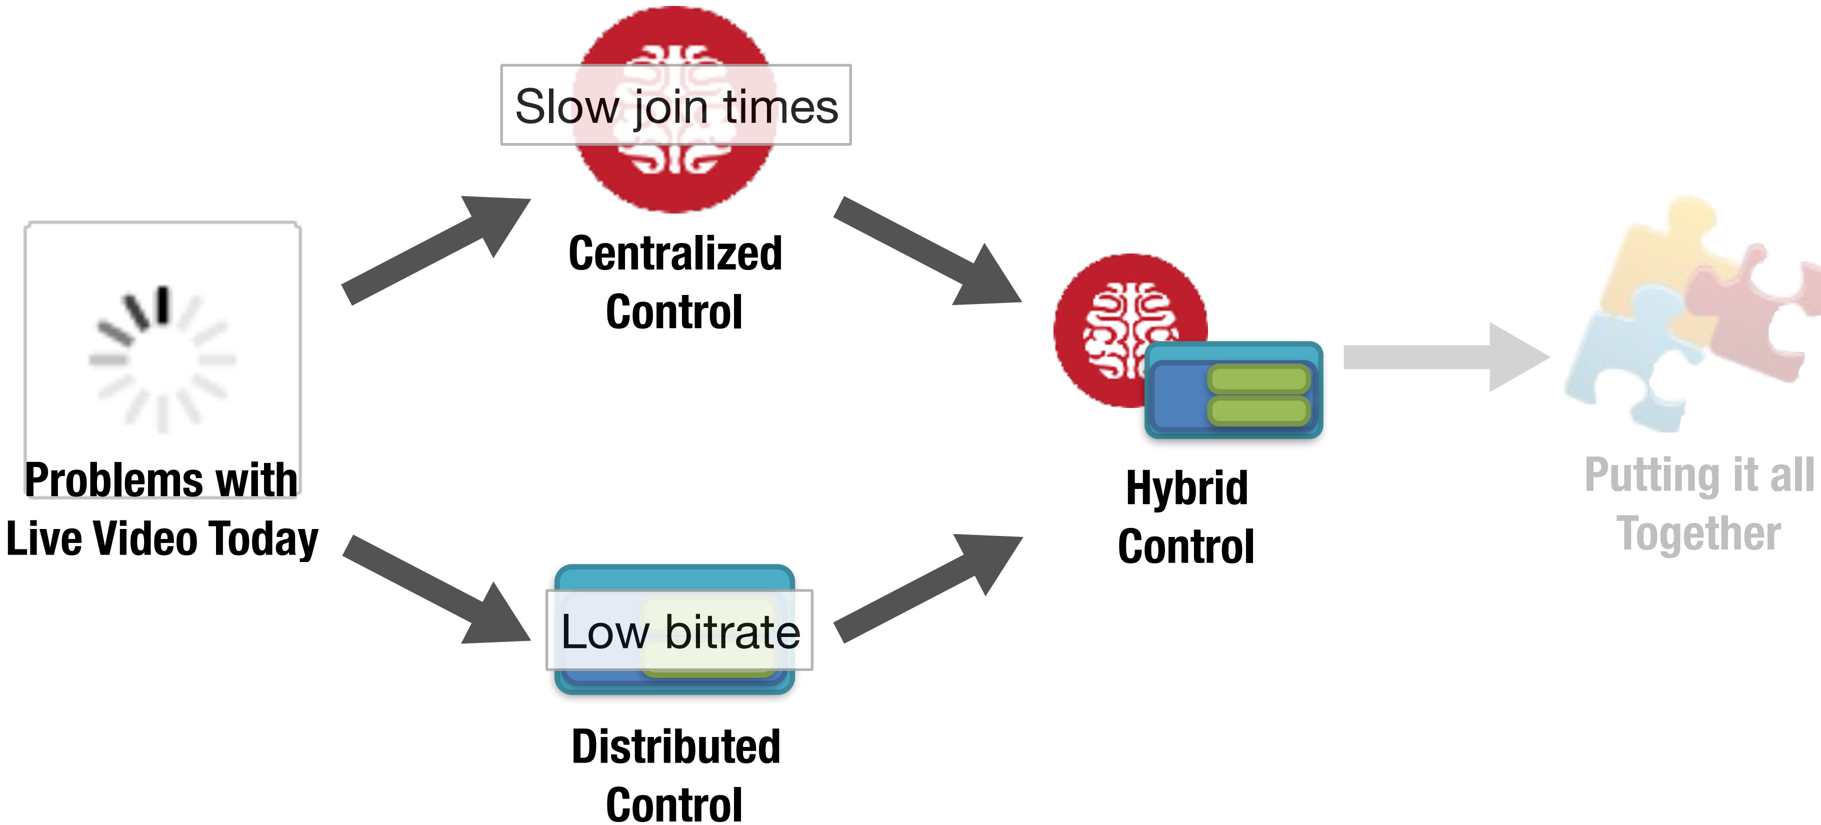

Slow join times
Centralized Control
Problems with Live Video Today
Hybrid Control
Putting it all Together
Low bitrate
Distributed Control

Hybrid Control

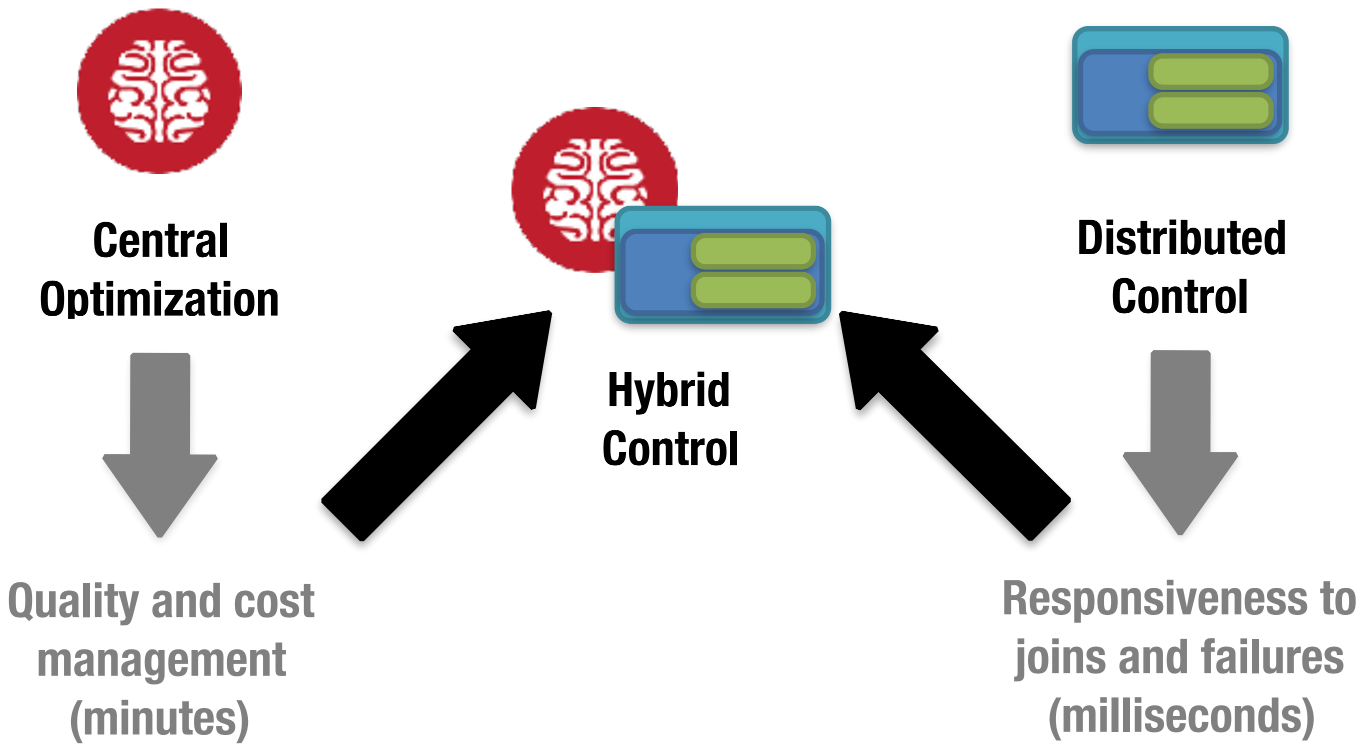

Challenges of Hybrid Control

- Forwarding loops — TRIVIAL
  - Always forward requests upwards
- State transitions — PRIOR WORK
  - Versioning and “shadow FIBS”
- Avoid bad control loop interactions — CHALLENGING

Combining Approaches: **Hybrid**

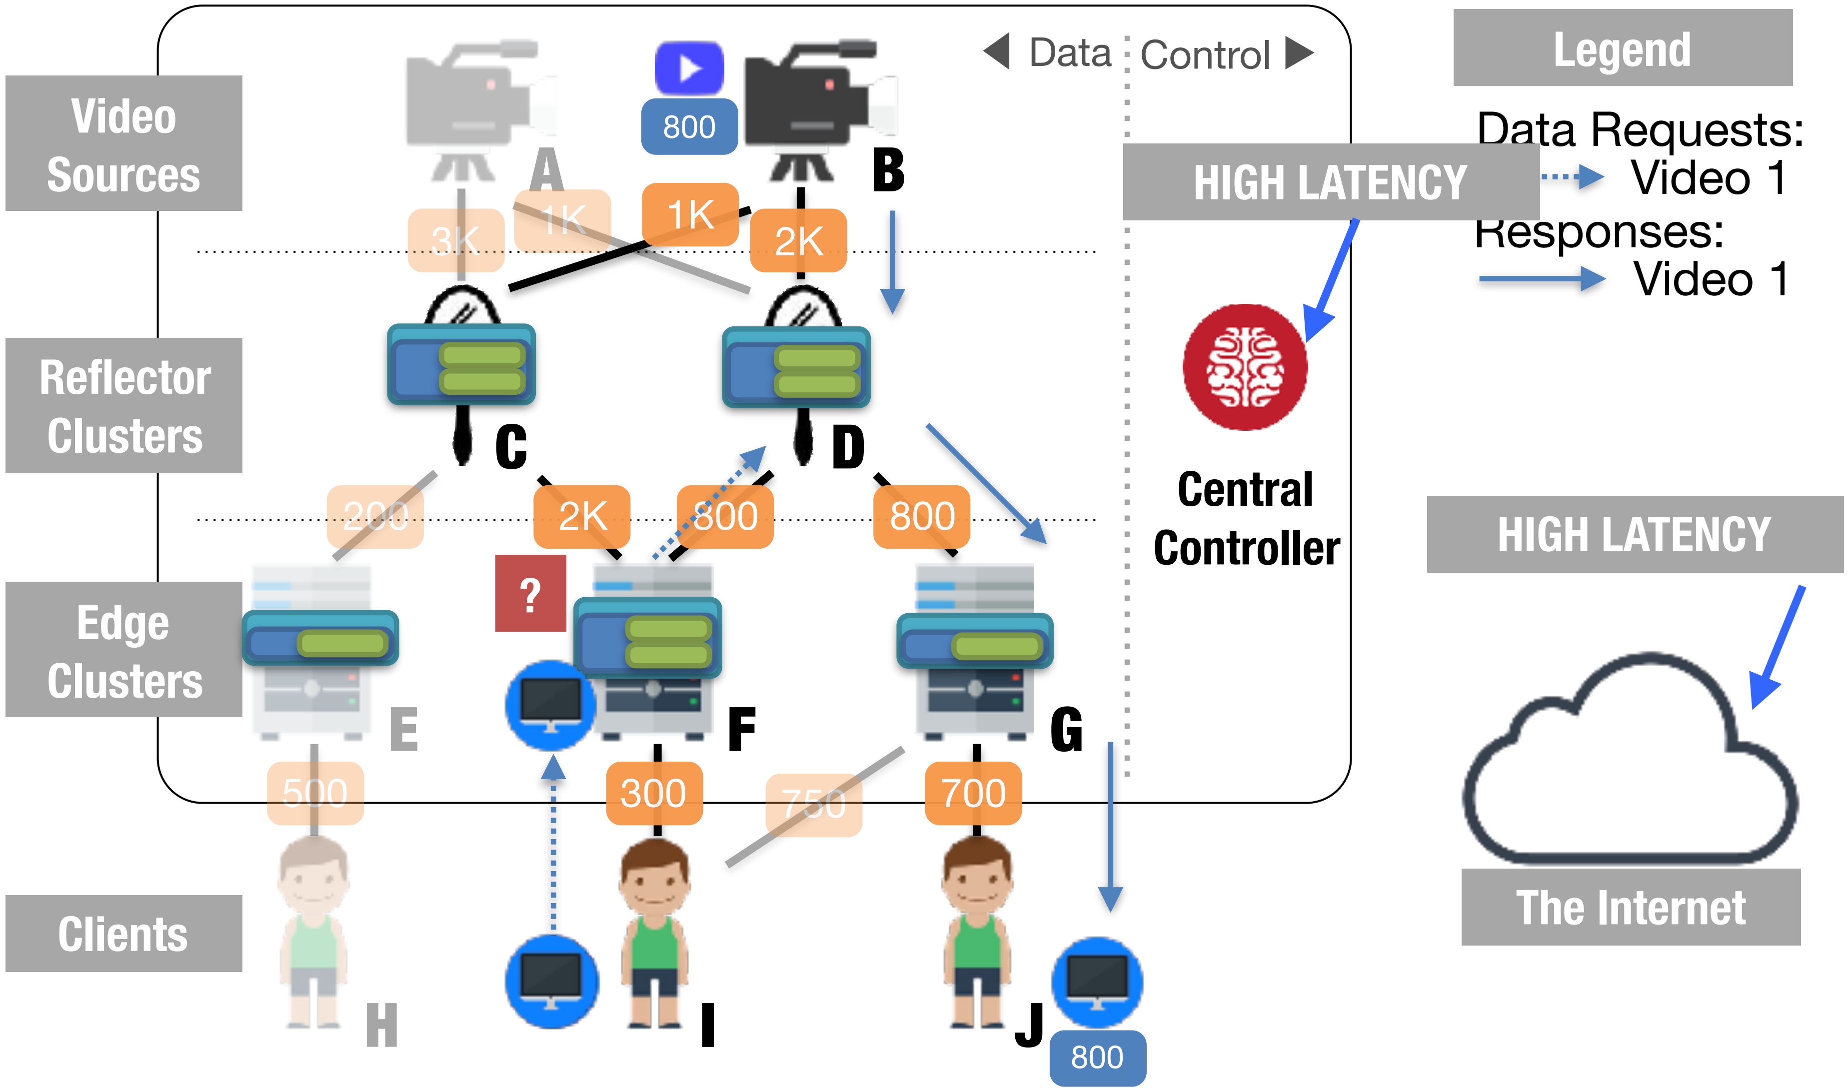

Combining Approaches: **Hybrid**

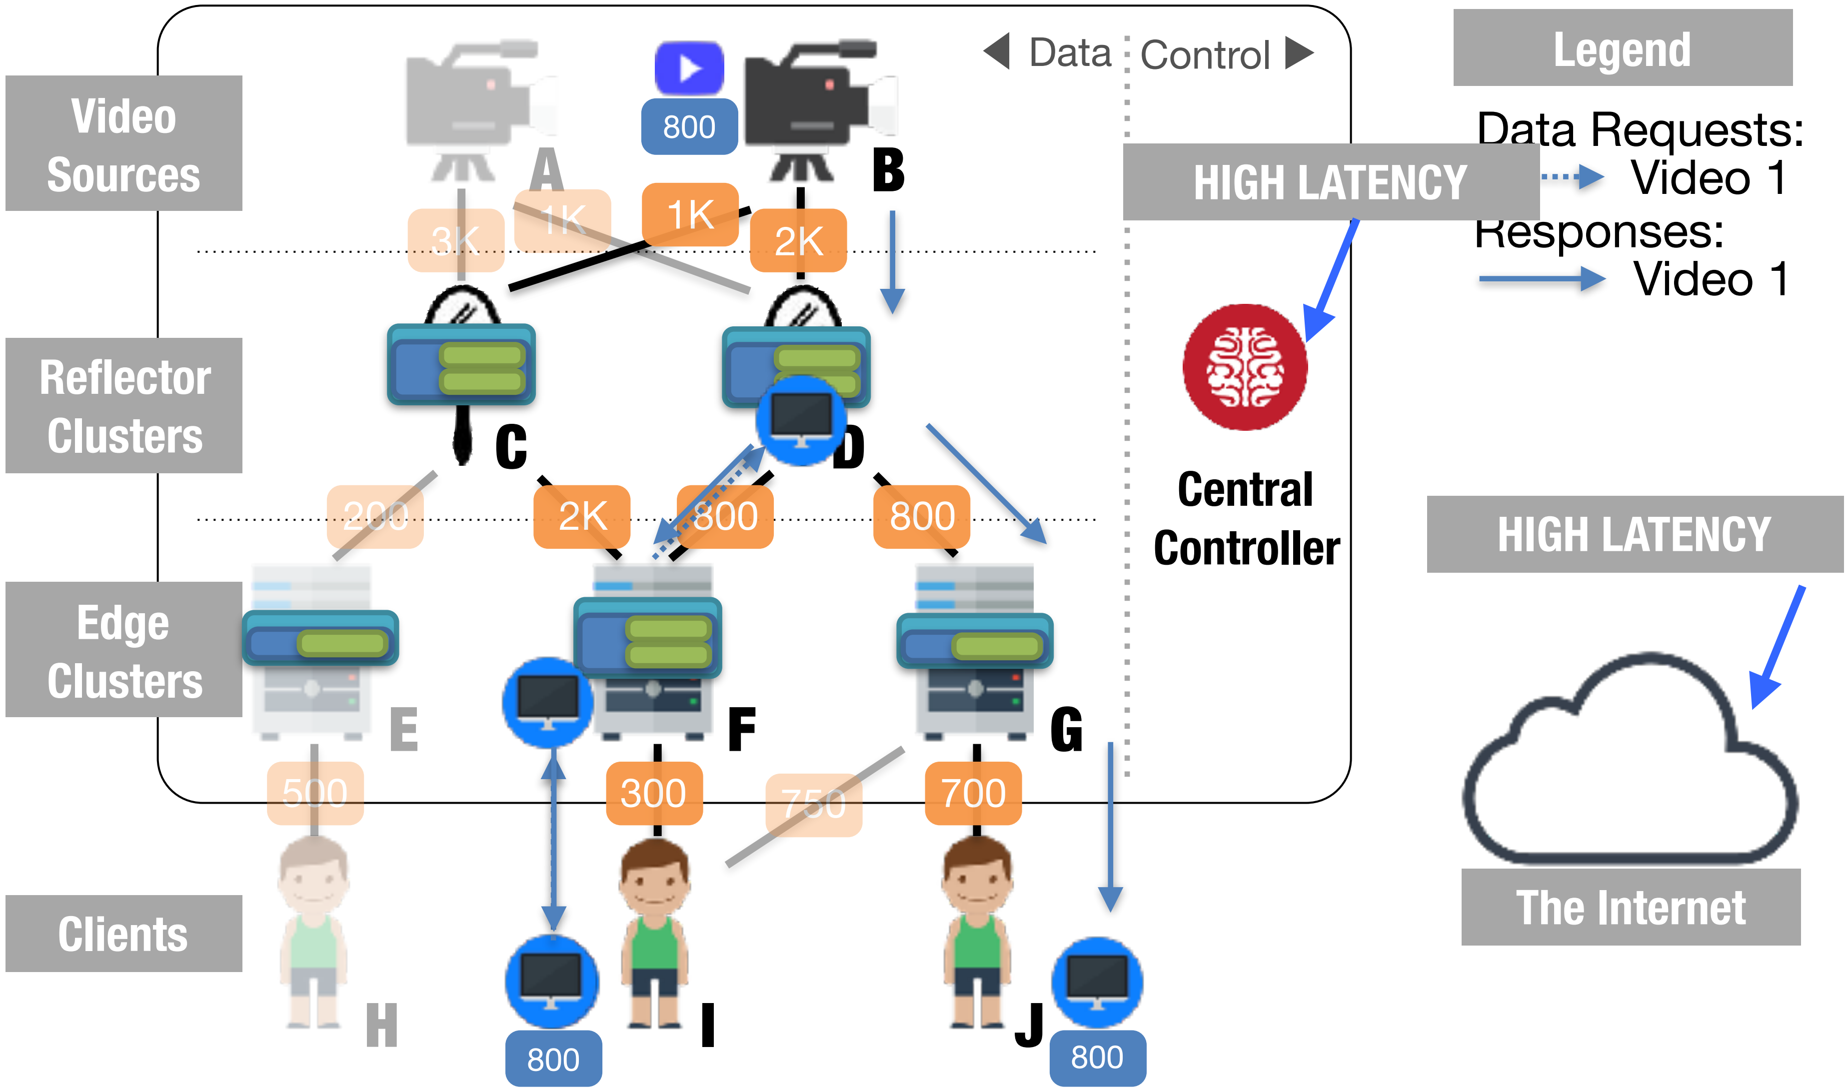

Combining Approaches: **Hybrid**

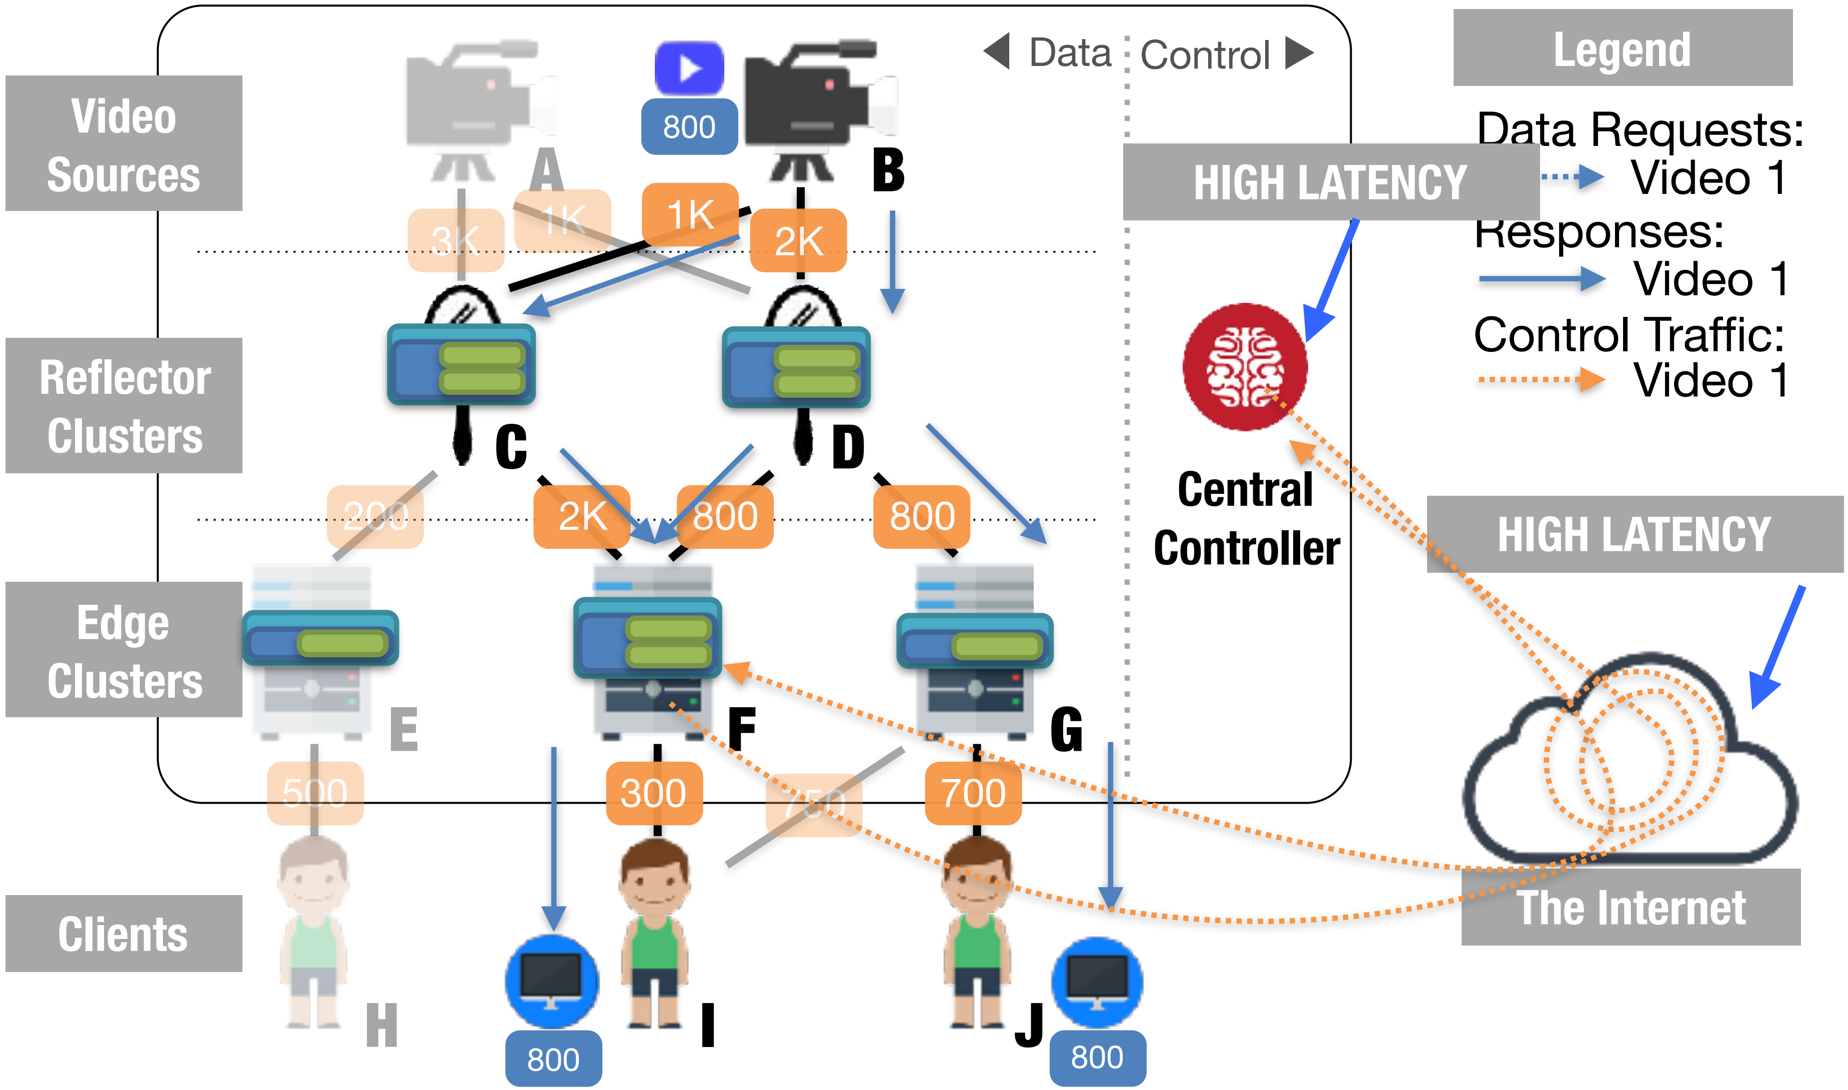

Challenges of Hybrid Control

- Forwarding loops — **TRIVIAL**
  - Always forward requests upwards
- State transitions — **PRIOR WORK**
  - Versioning and “shadow FIBS”
- 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

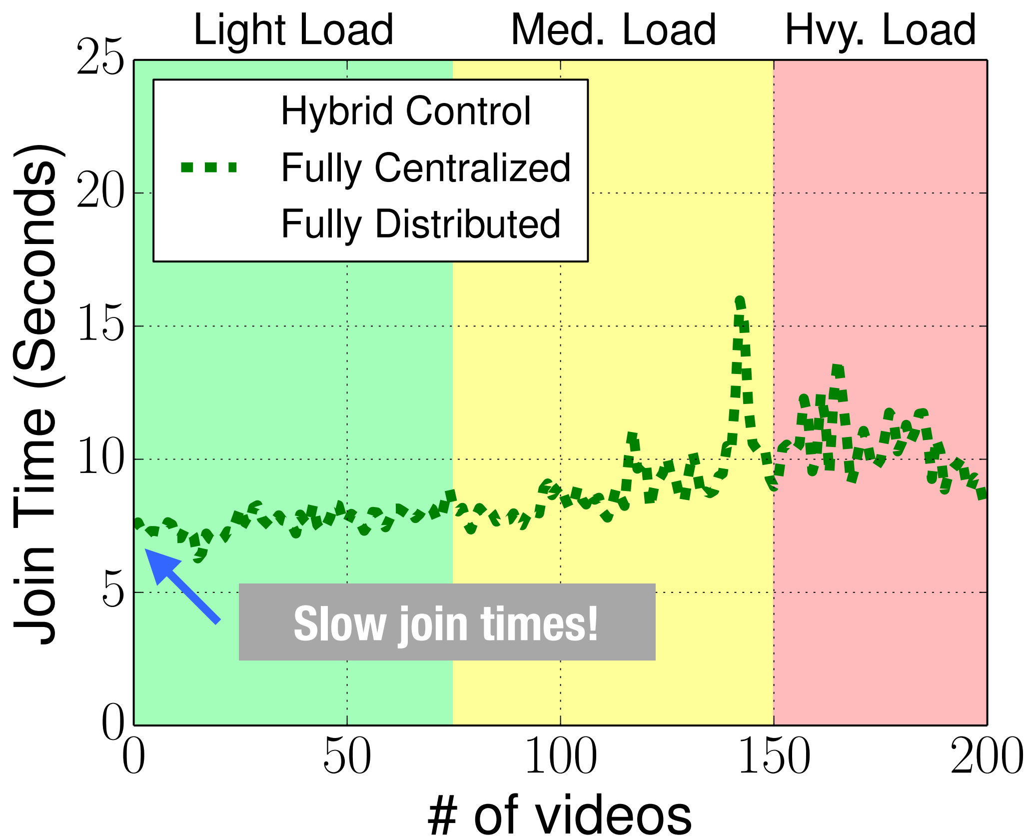

Experiments on EC2 nodes with a
centralized controller at CMU across the Internet

Hybrid Control and Responsiveness

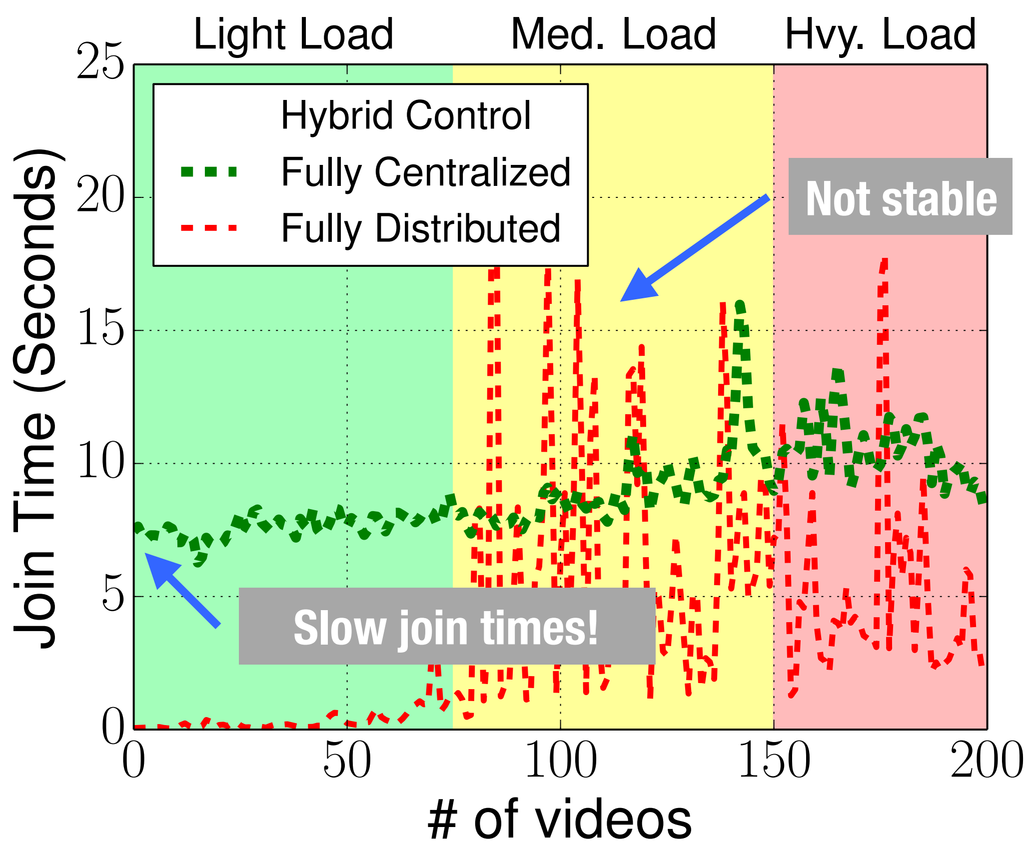

Experiments on EC2 nodes with a
centralized controller at CMU across the Internet

Hybrid Control and Responsiveness

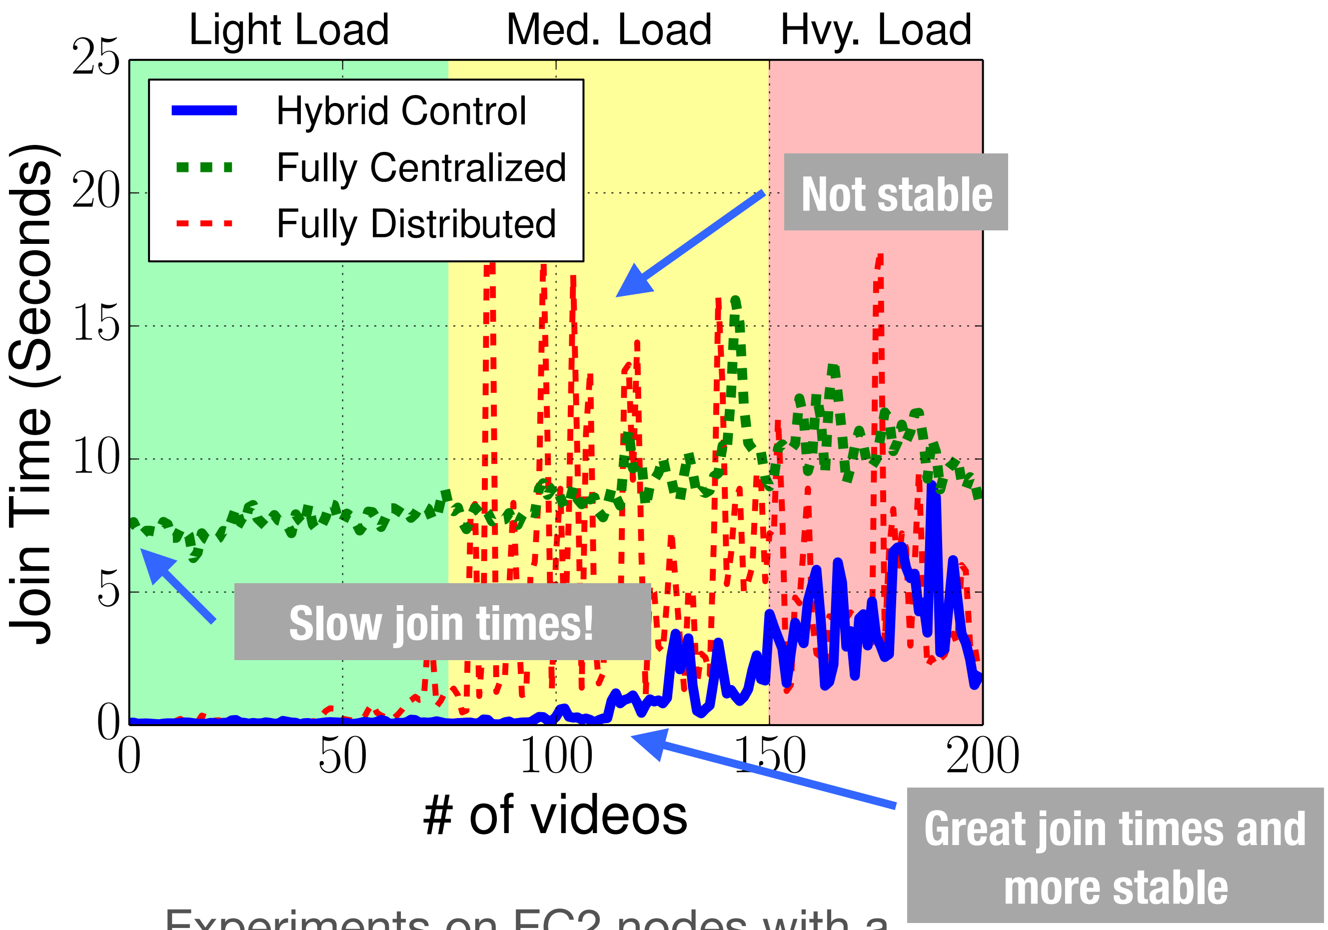

Experiments on EC2 nodes with a centralized controller at CMU across the Internet

Outline

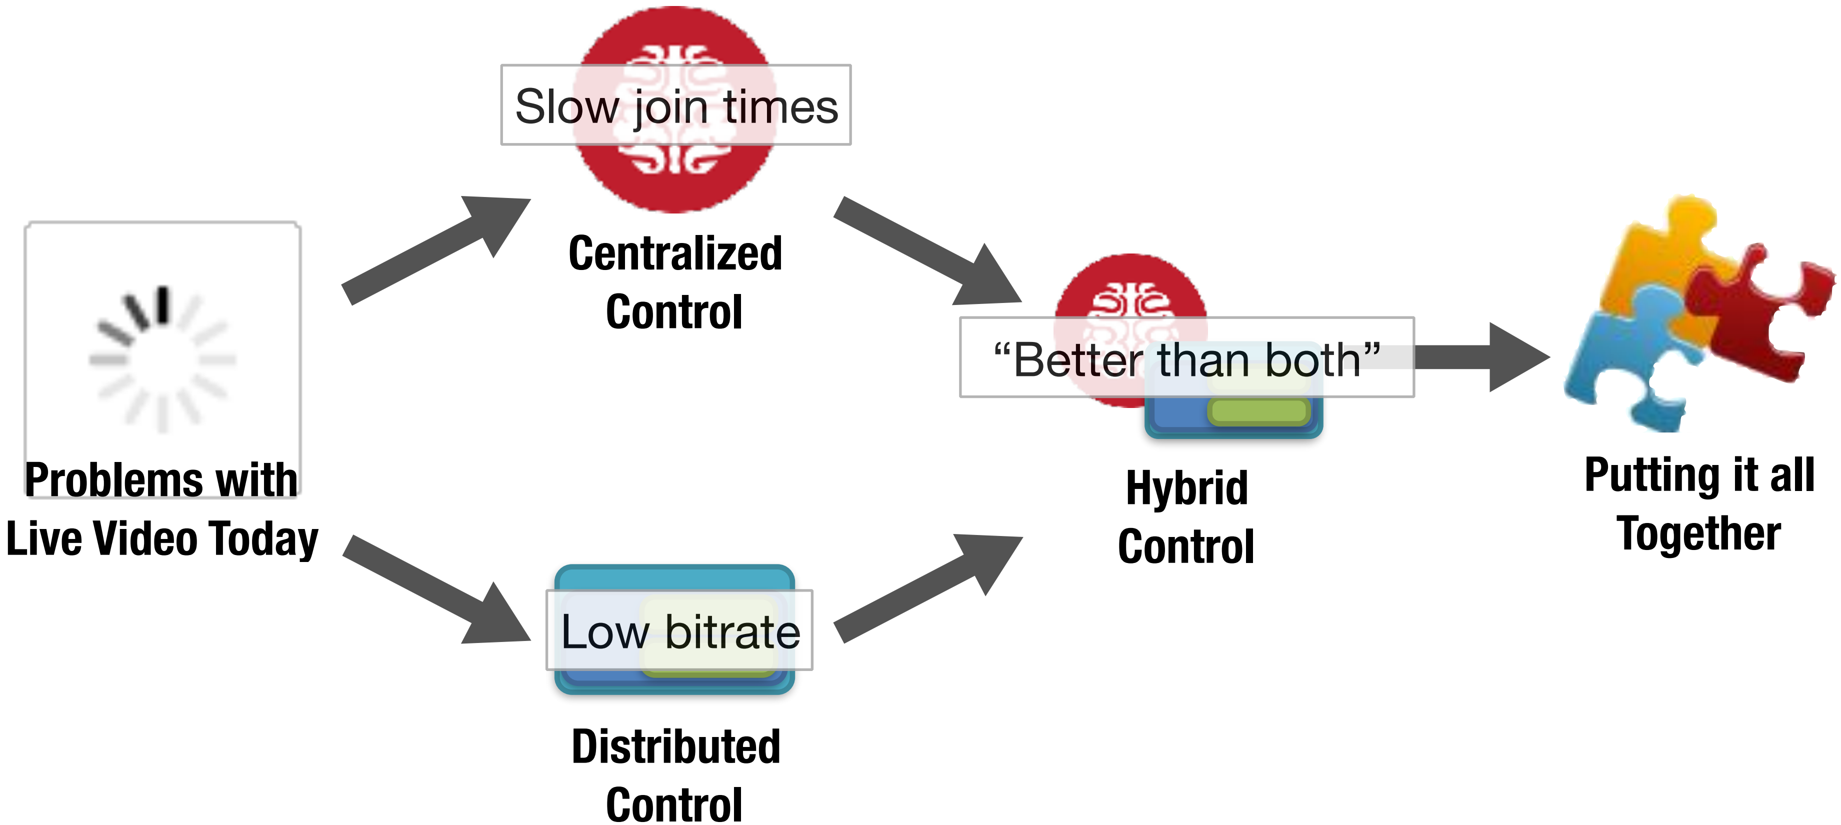
Slow join times
Centralized Control
Problems with Live Video Today
"Better than both"
Hybrid Control
Putting it all Together
Low bitrate
Distributed Control

Putting it all Together

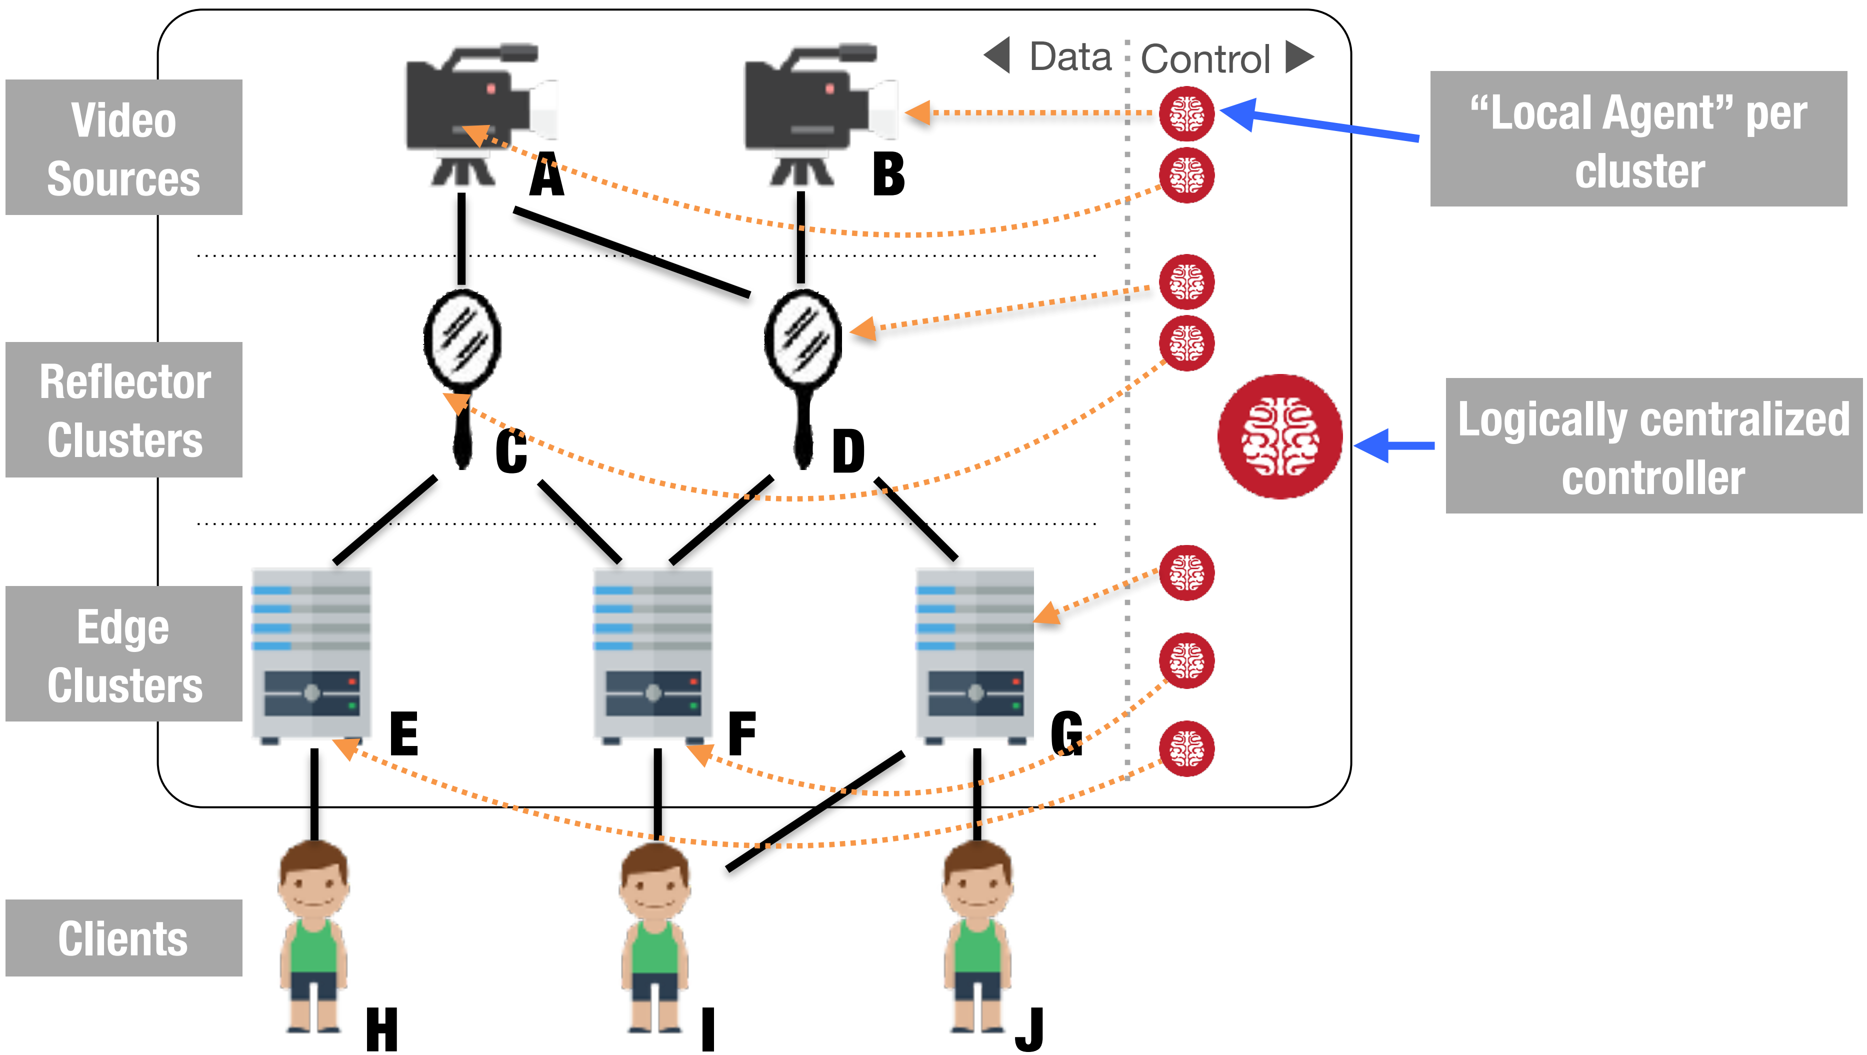

Putting it all Together

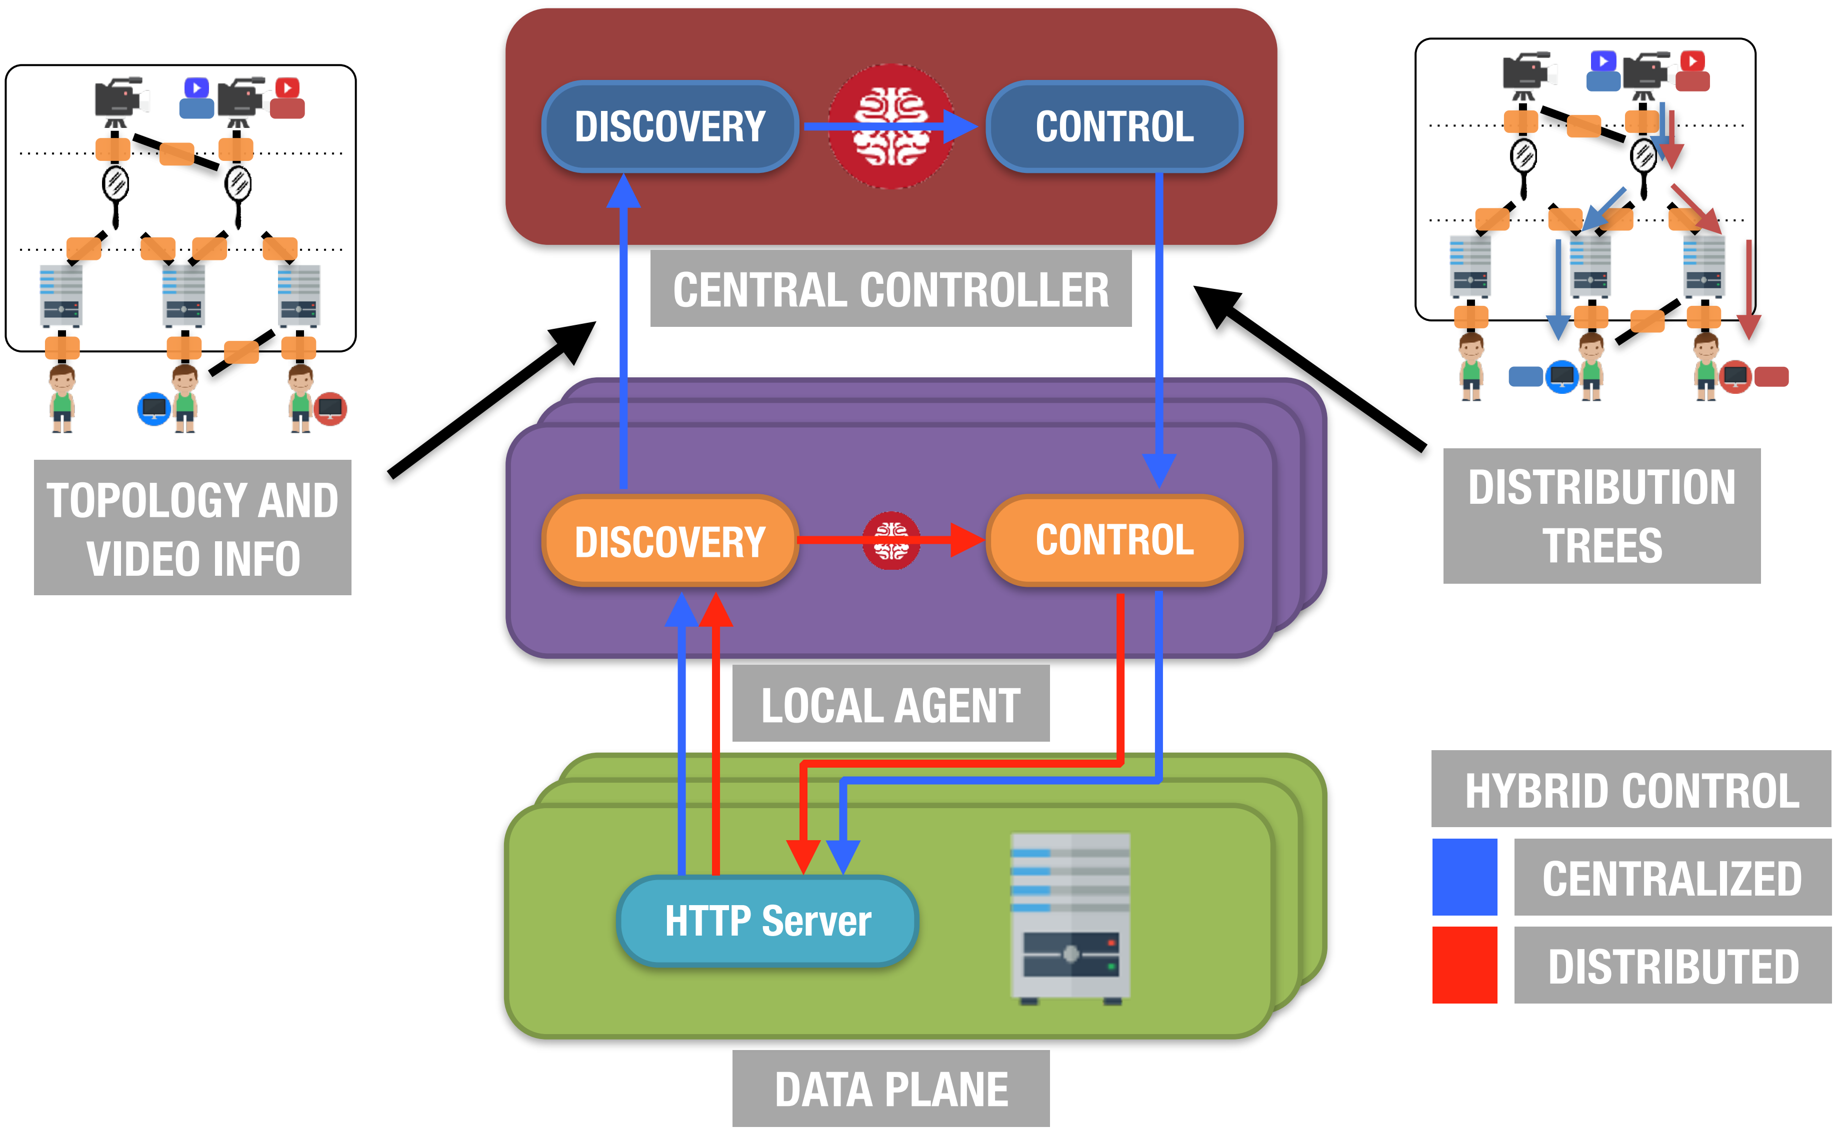

Key Results

- Trace-driven eval - **centralized optimization**
  - High quality & low delivery cost? **1.7x / 2x**
  - Scalable / fine grain? **10K videos**; **2K clusters**
- End-to-end eval - **hybrid control**
  - Responsive? **200ms**
- More results in paper
  - Operator Control? Failures? Partitions?

Conclusion

- VDN presents a new approach for CDN-based live video delivery

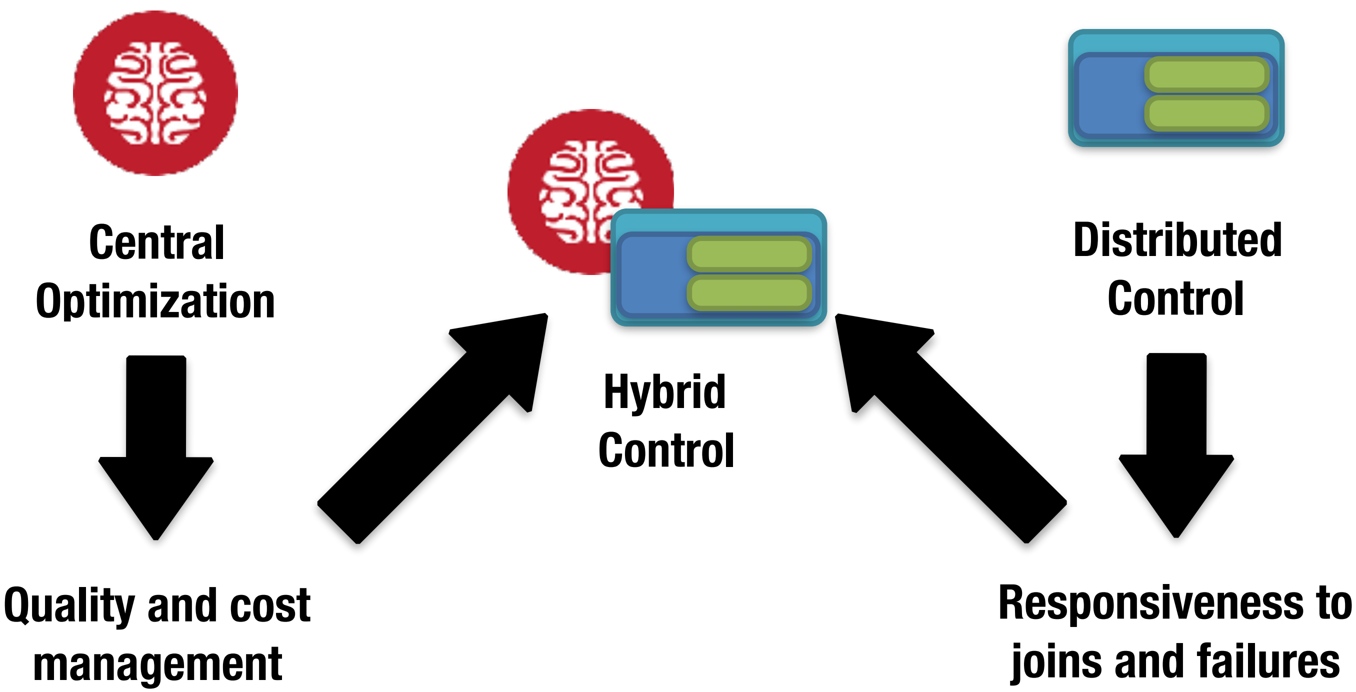

Practical, Real-time Centralized Control for CDN-based Live Video Delivery

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

Backup slides…

Problems with Traffic Engineering

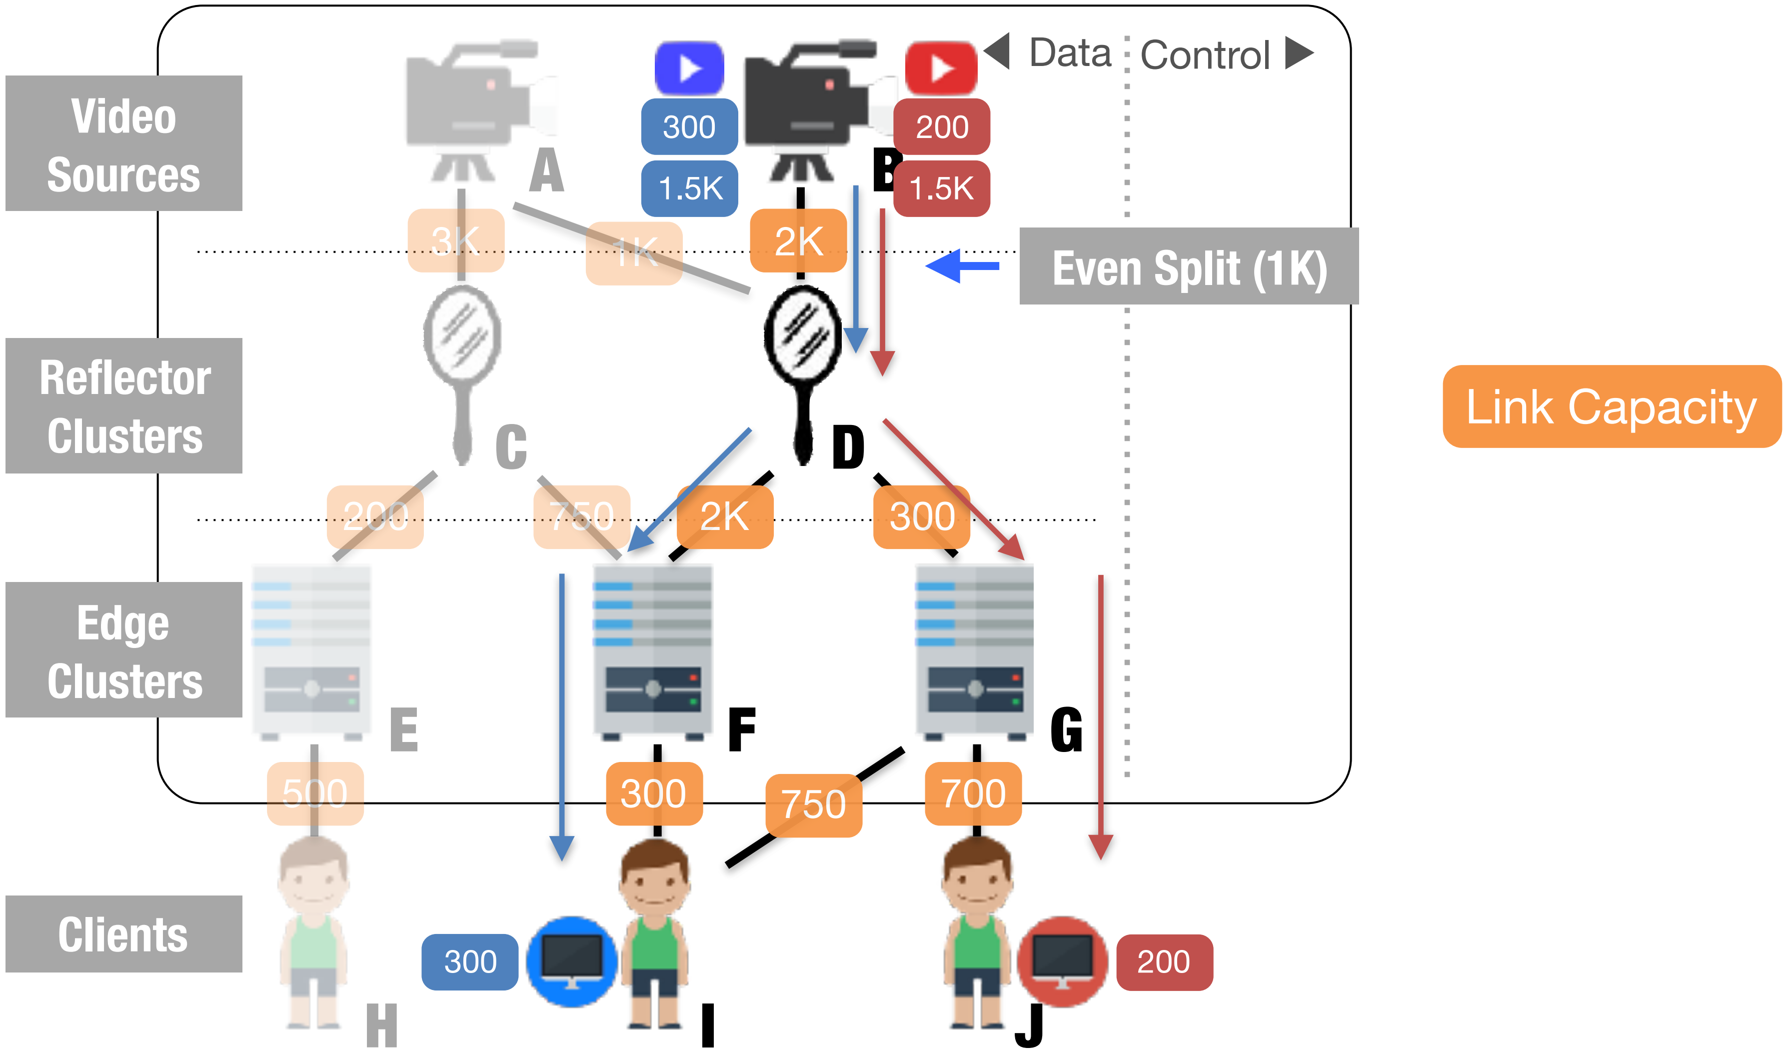

Problems with Traffic Engineering

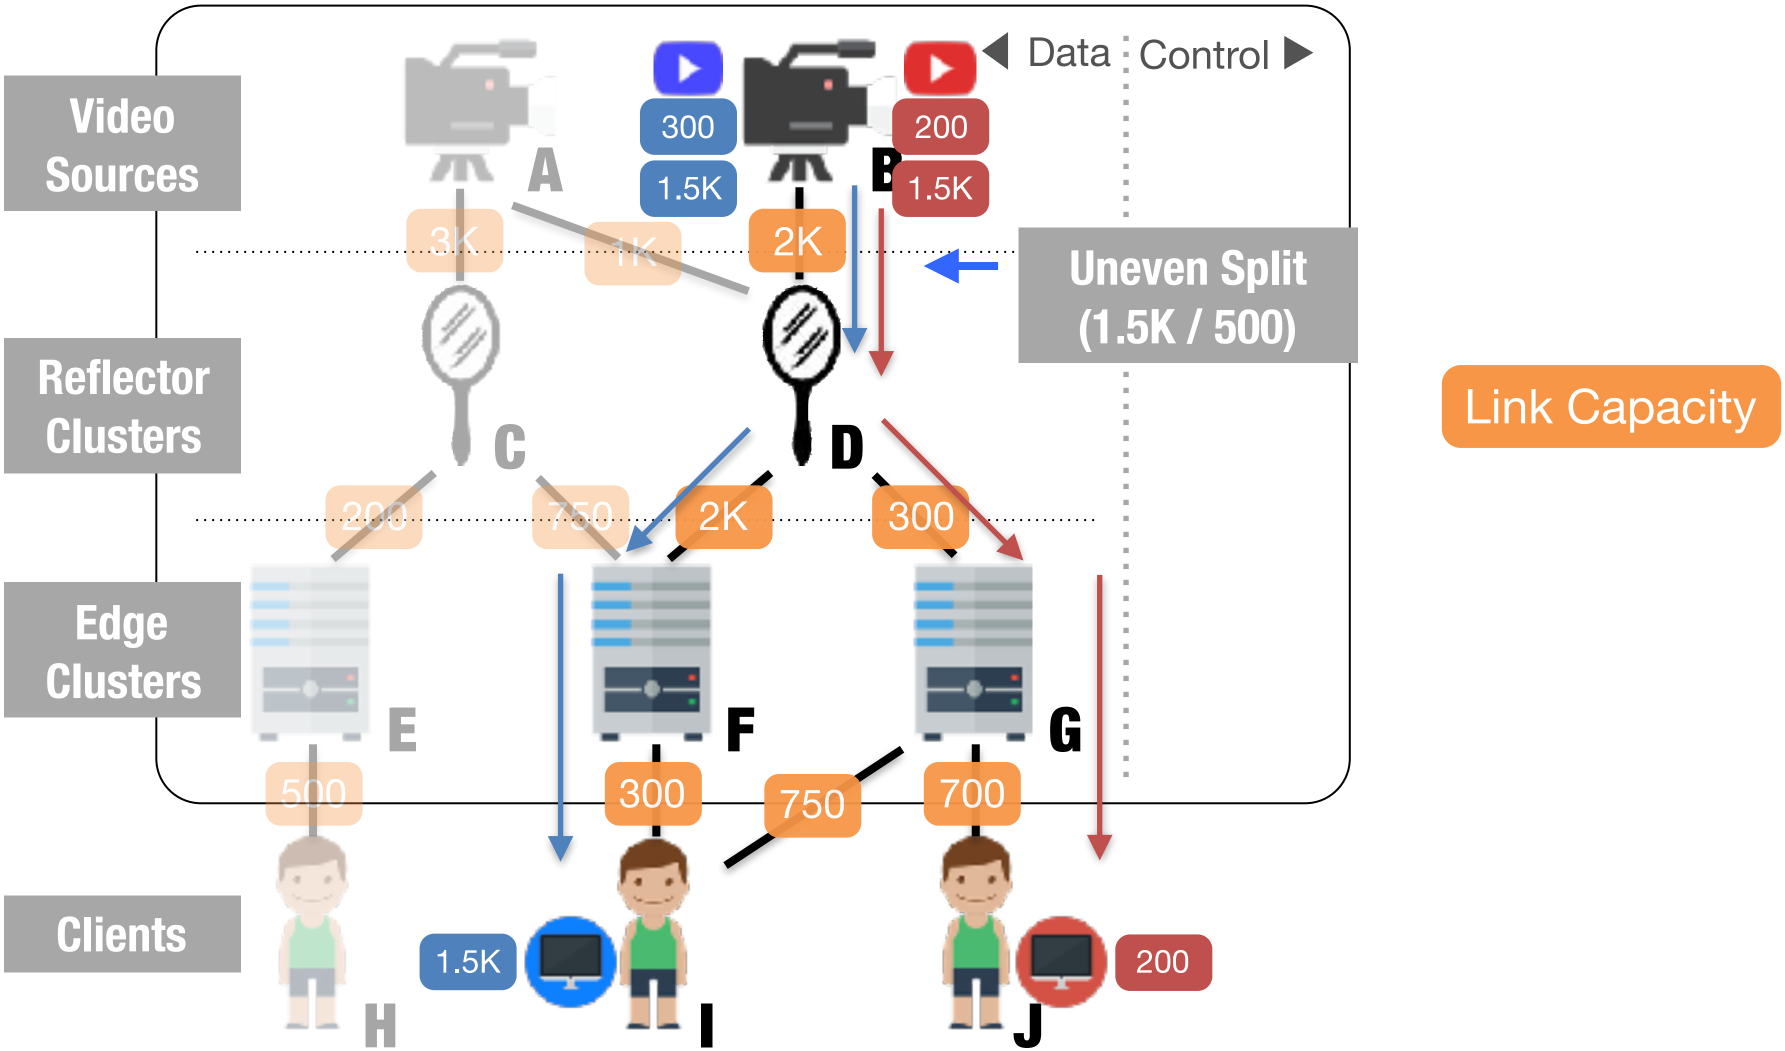

Distributed: Example of Sub-optimal

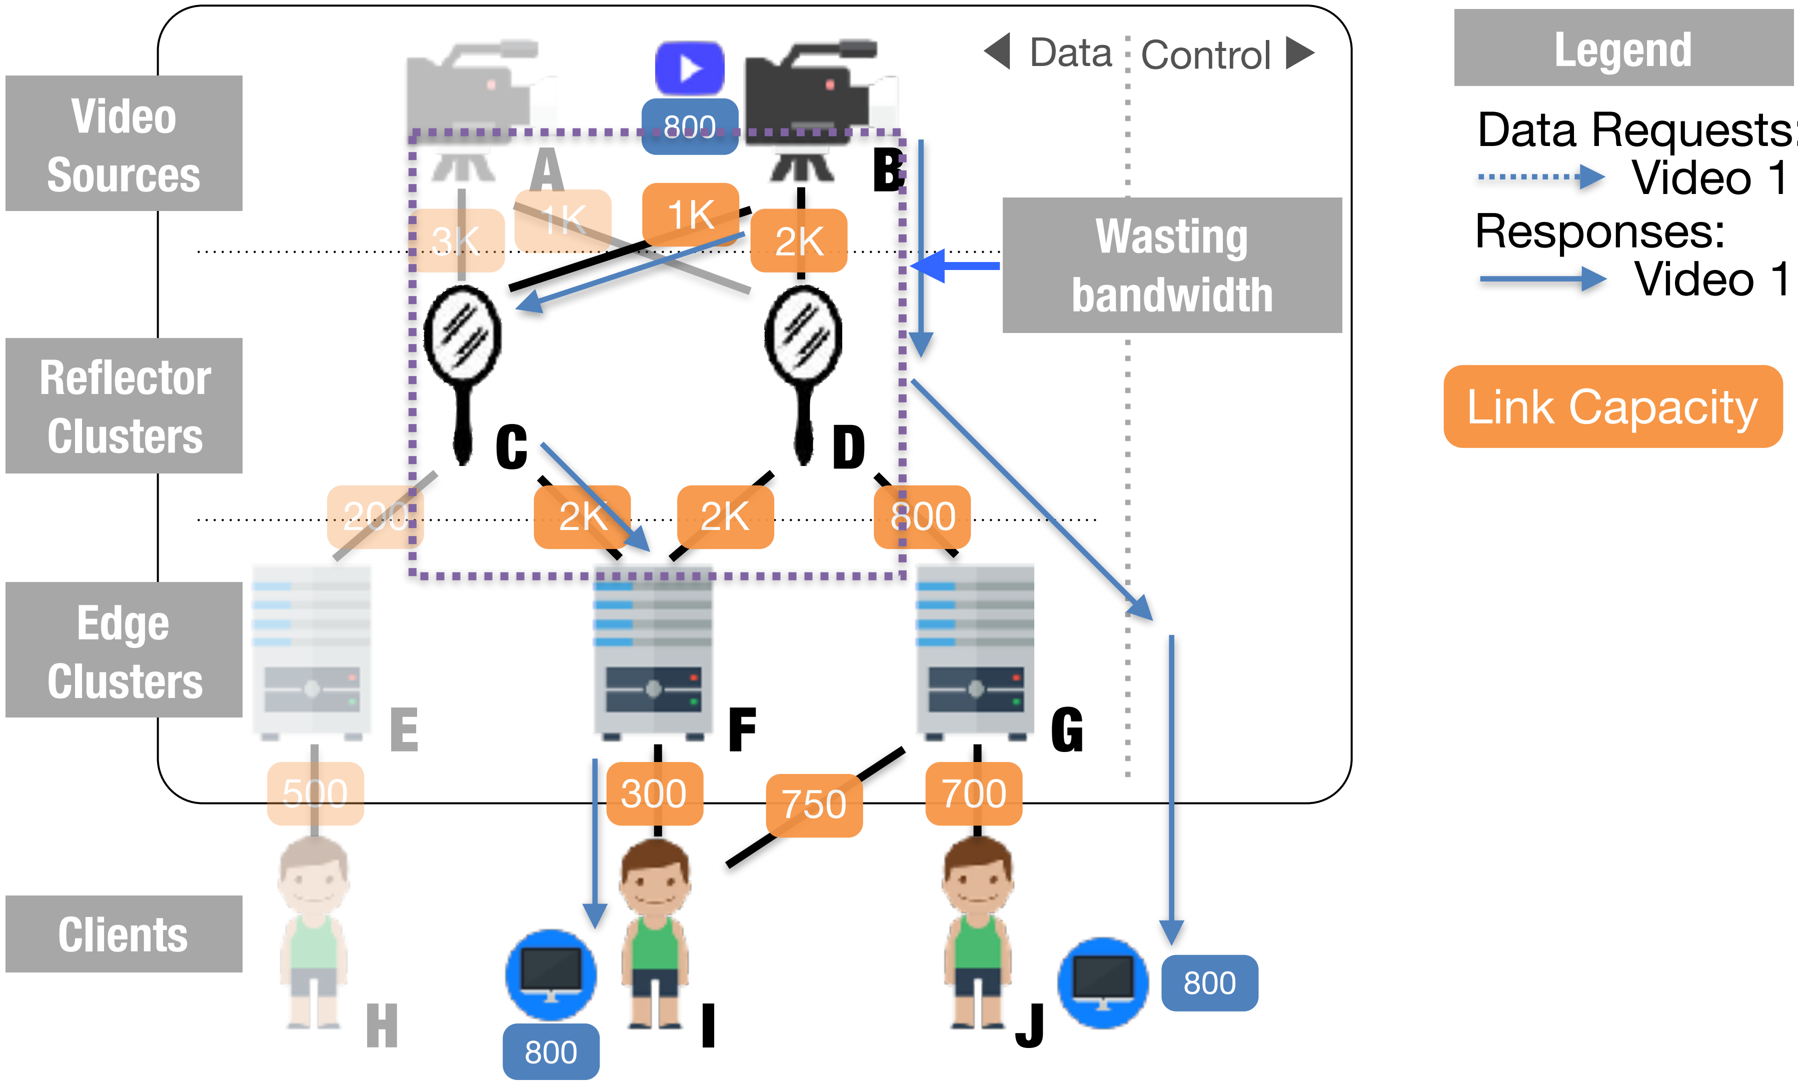

Distributed: Example of Sub-optimal

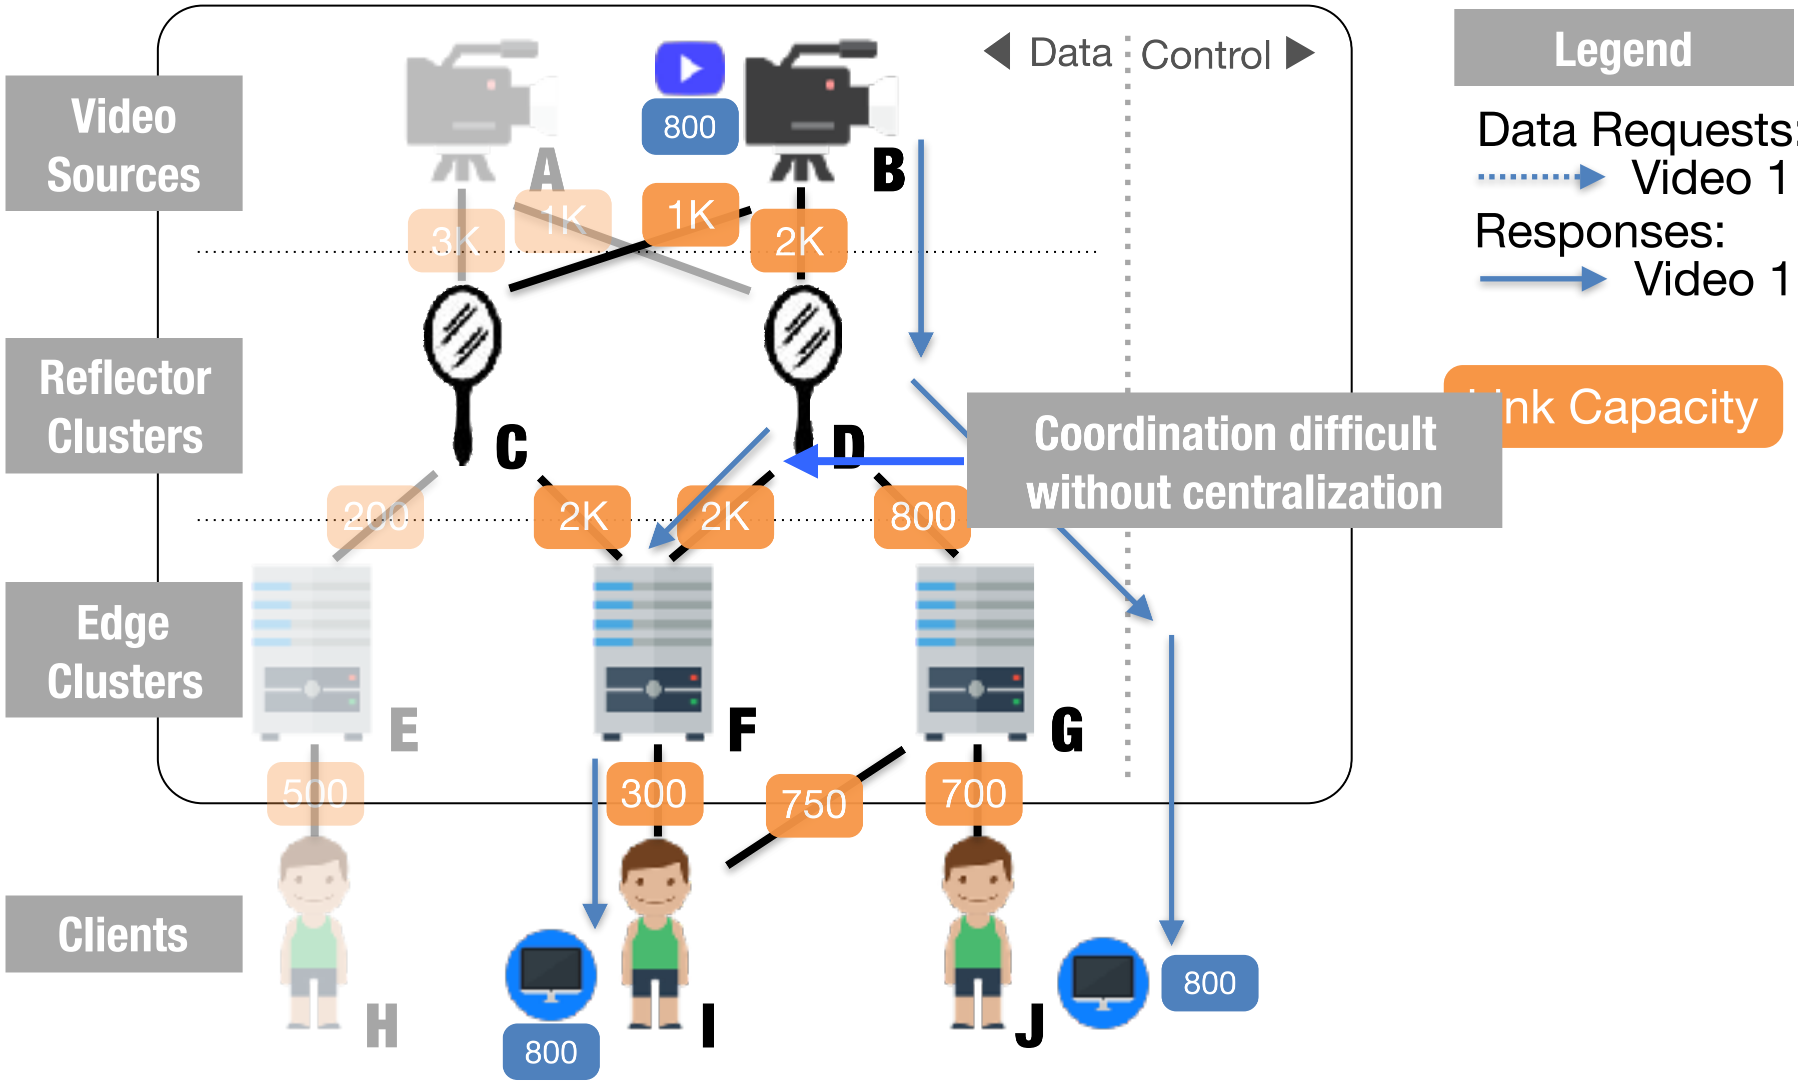

Trace-Driven Eval

- 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)*	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.

	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.

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