C3: INTERNET-SCALE CONTROL PLANE FOR VIDEO QUALITY OPTIMIZATION

Aditya Ganjam, Jibin Zhan, Xi Liu, Faisal Siddiqi, *Conviva*
Junchen Jiang, Vyas Sekar, *Carnegie Mellon University*
Ion Stoica, *University of California, Berkeley, Conviva*
Hui Zhang, *Carnegie Mellon University, Conviva*

HIGH EXPECTATIONS ON VIDEO QUALITY

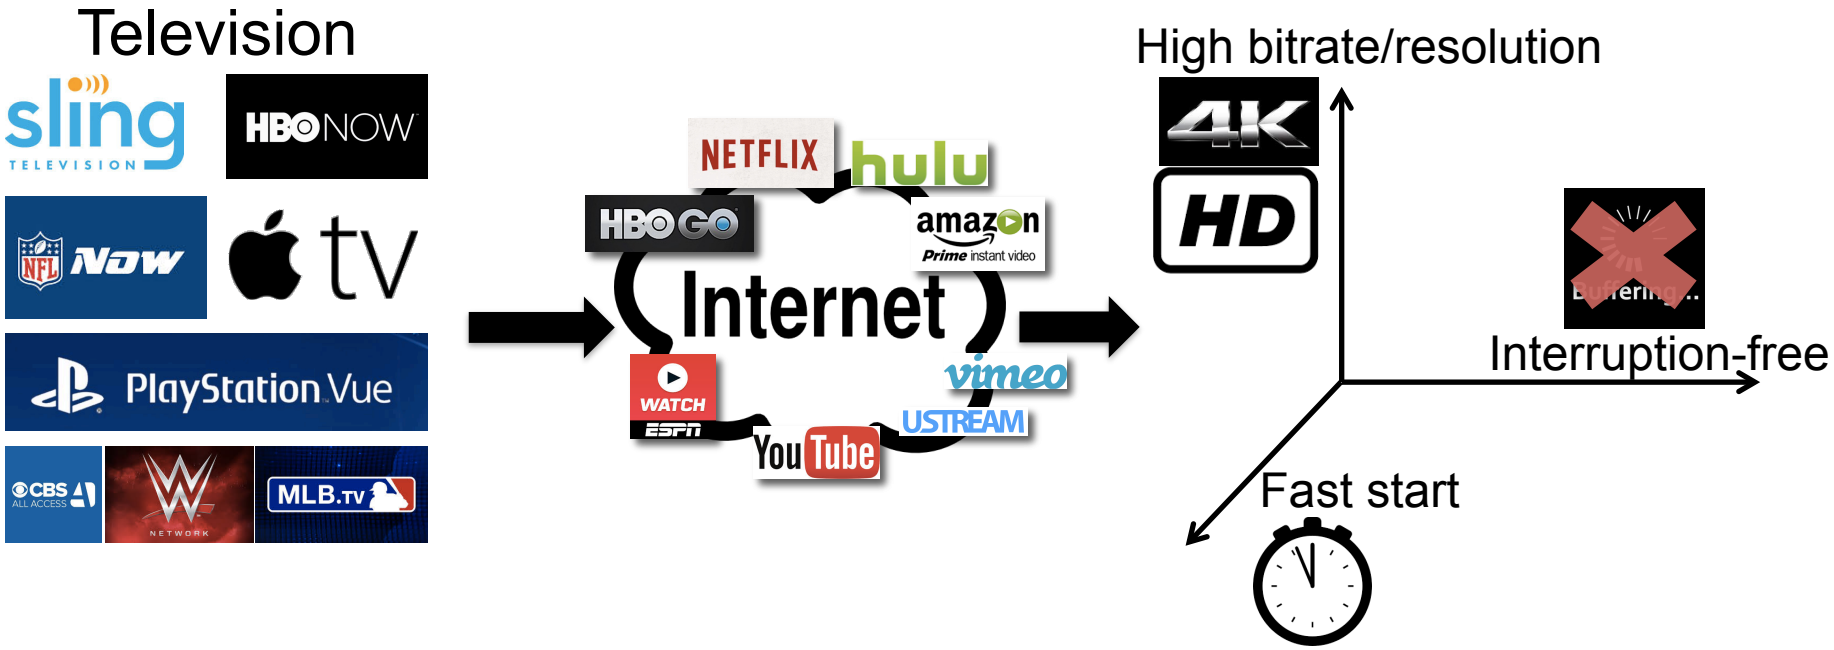

Need to optimize video quality through the lifetime of a session

WHERE TO IMPLEMENT OPTIMIZATION?

- Multiple choices for “actuation”
- Client-side actuation most favorable to incremental deployment

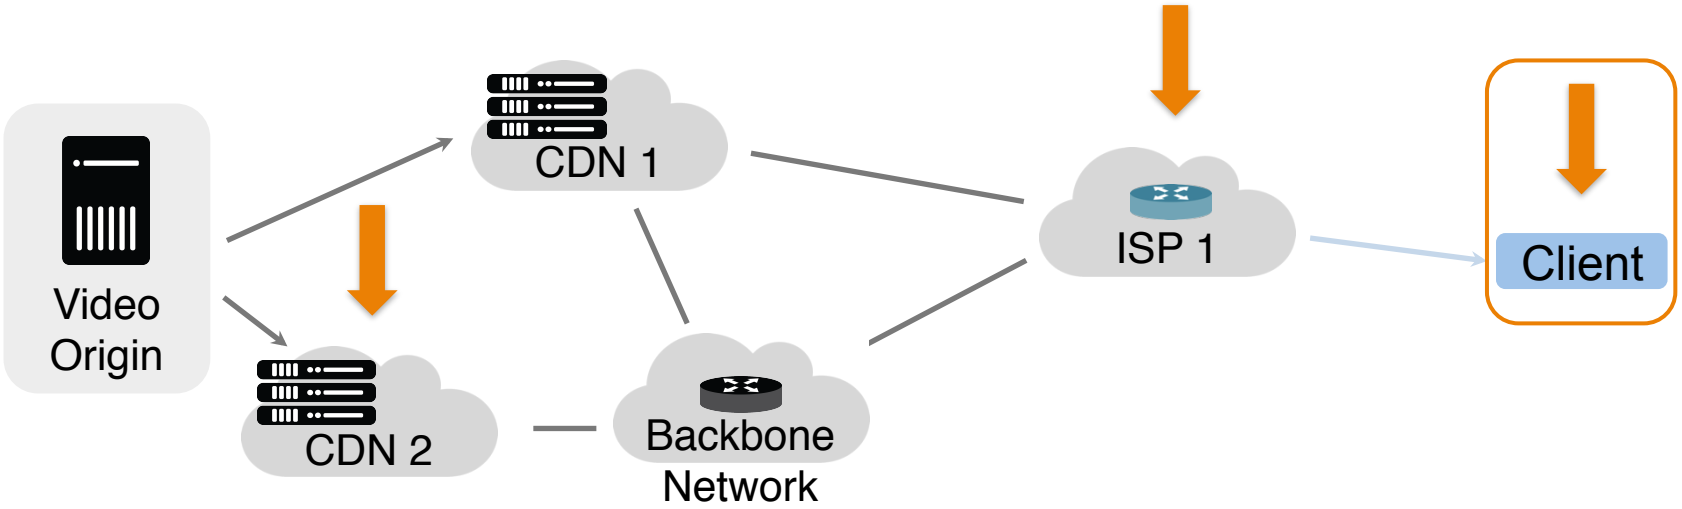

MANY CHOICES & NEED QUICK DECISION

Optimization parameters

- Bitrate x CDN (over time)

Existing approaches are reactive, apply a probe and update method

- Too slow to find optimal choice

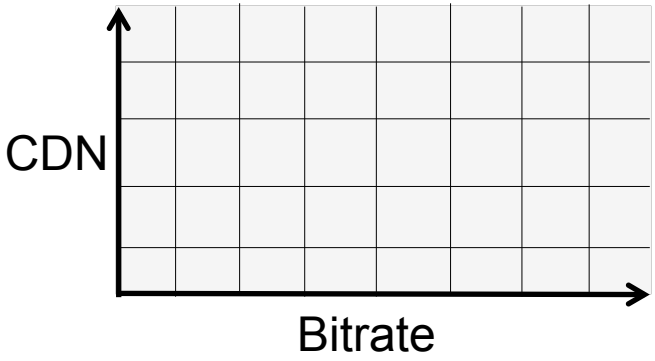

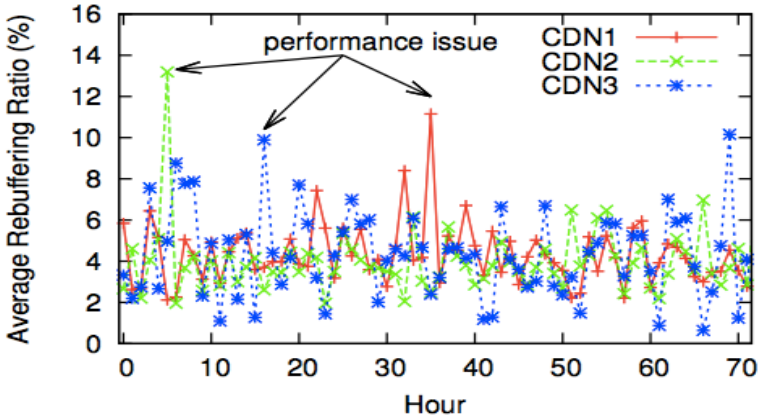

Rebuffering Ratio

CENTRALIZED PREDICTIVE CONTROL

Ideal solution

- Predict outcome of each parameter choice in real-time
- Continuously select optimal choice

Achieving ideal solution requires …

- Collect global quality information
- Build a prediction model
- Make per-client predictions and parameter decisions at Internet-scale

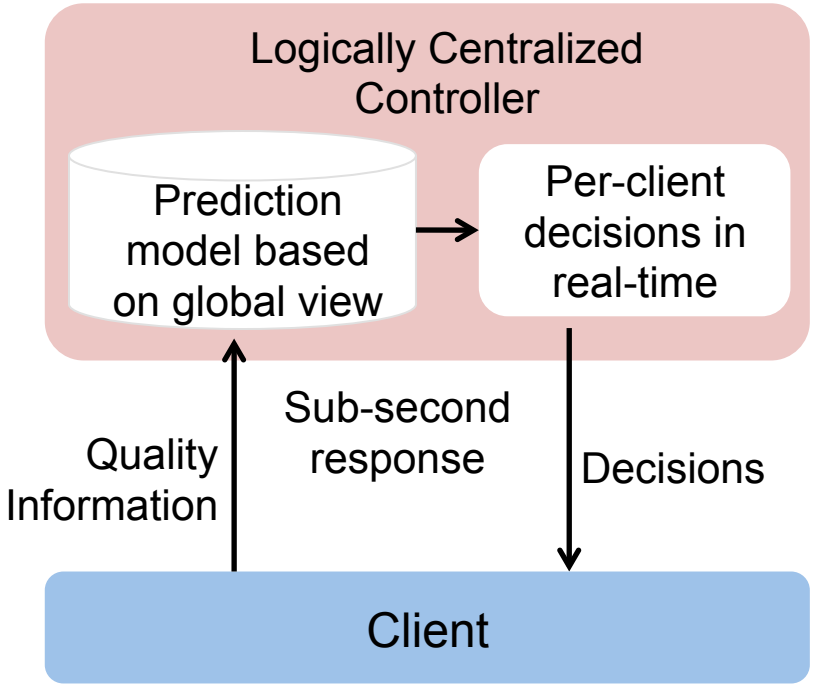

CHALLENGE:
CLIENT HETEROGENEITY

- Strawman: implement 100 times
- Heterogeneous software environments and interfaces
- Slow software update cycle

>100 unique client platforms

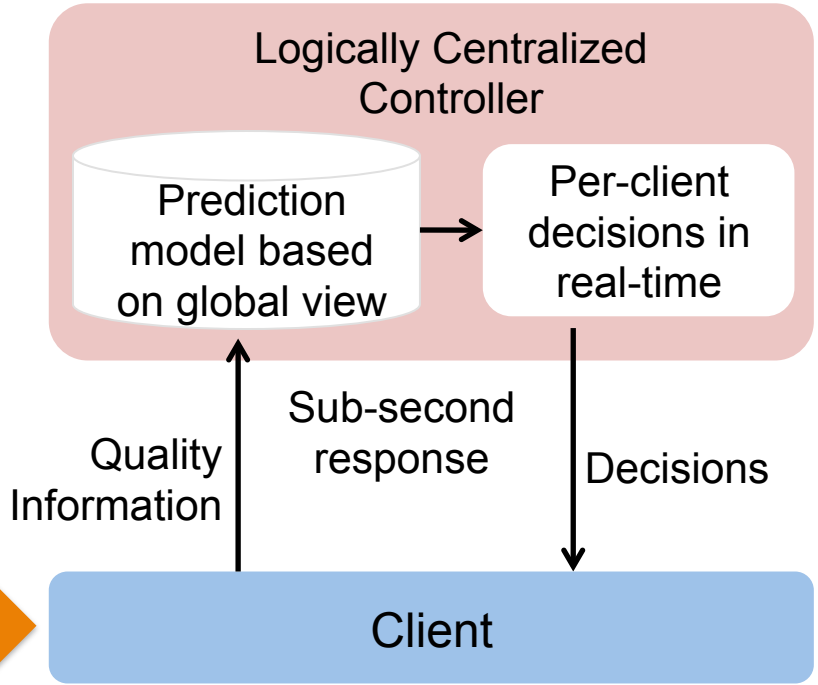

DESIGN CHOICE:
THIN SENSING/ACTUATION LAYER

- Minimal client footprint
  - ➢ Move computation to controller
- Define a “narrow waist” interface

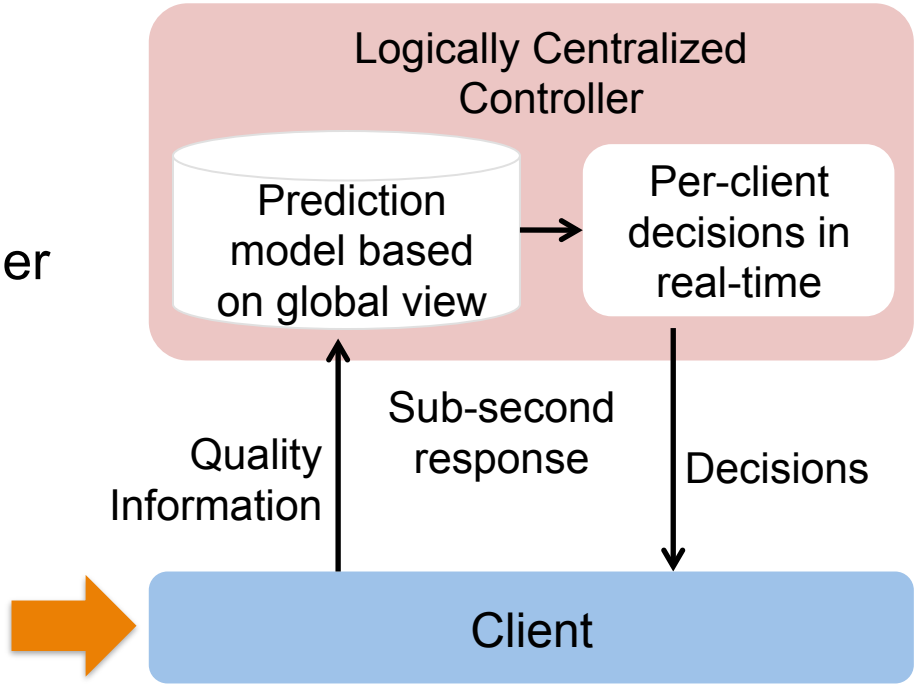

CHALLENGE:
REAL-TIME PREDICTIONS AT INTERNET SCALE

Internet scale?

- Geographic scale - all continents
- Network scale - all network types
- Client scale – 10s to 100s of millions concurrent

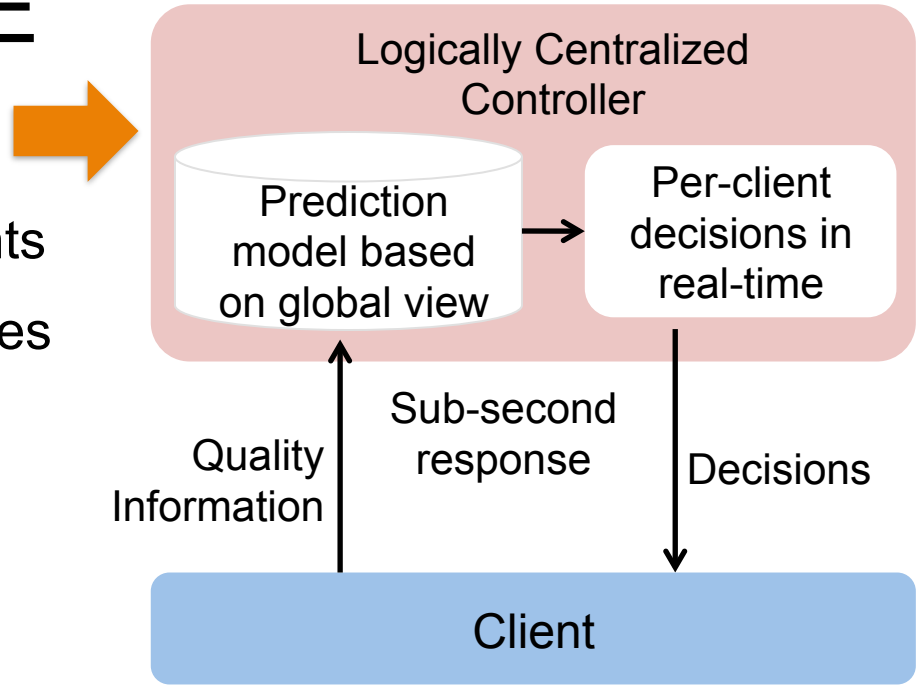

CHALLENGE: REAL-TIME PREDICTIONS AT INTERNET SCALE

- Achieving **real-time response**, with the most recent **global view** at **Internet-scale** is not feasible based on today's technologies
- Strawman solutions
  - Centralized global controller – not real-time
  - Partitioned controller – no global view

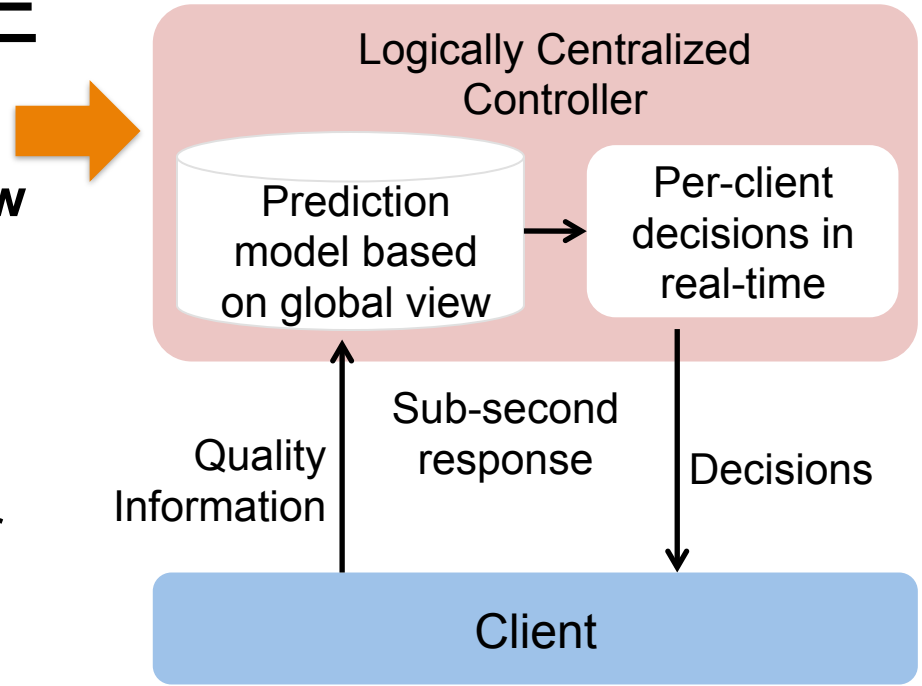

DESIGN CHOICE:
SPLIT CONTROL PLANE

- Create a global prediction model in near real-time (minutes)
- Make per-client decisions in real-time (sub-second) based on global model and most recent per-client local information

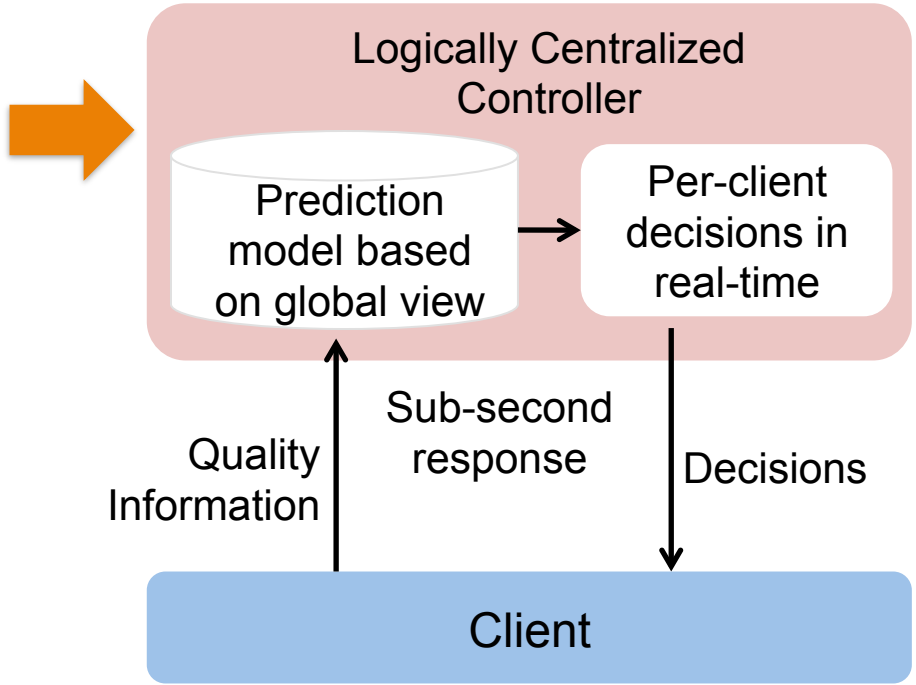

C3 ARCHITECTURE

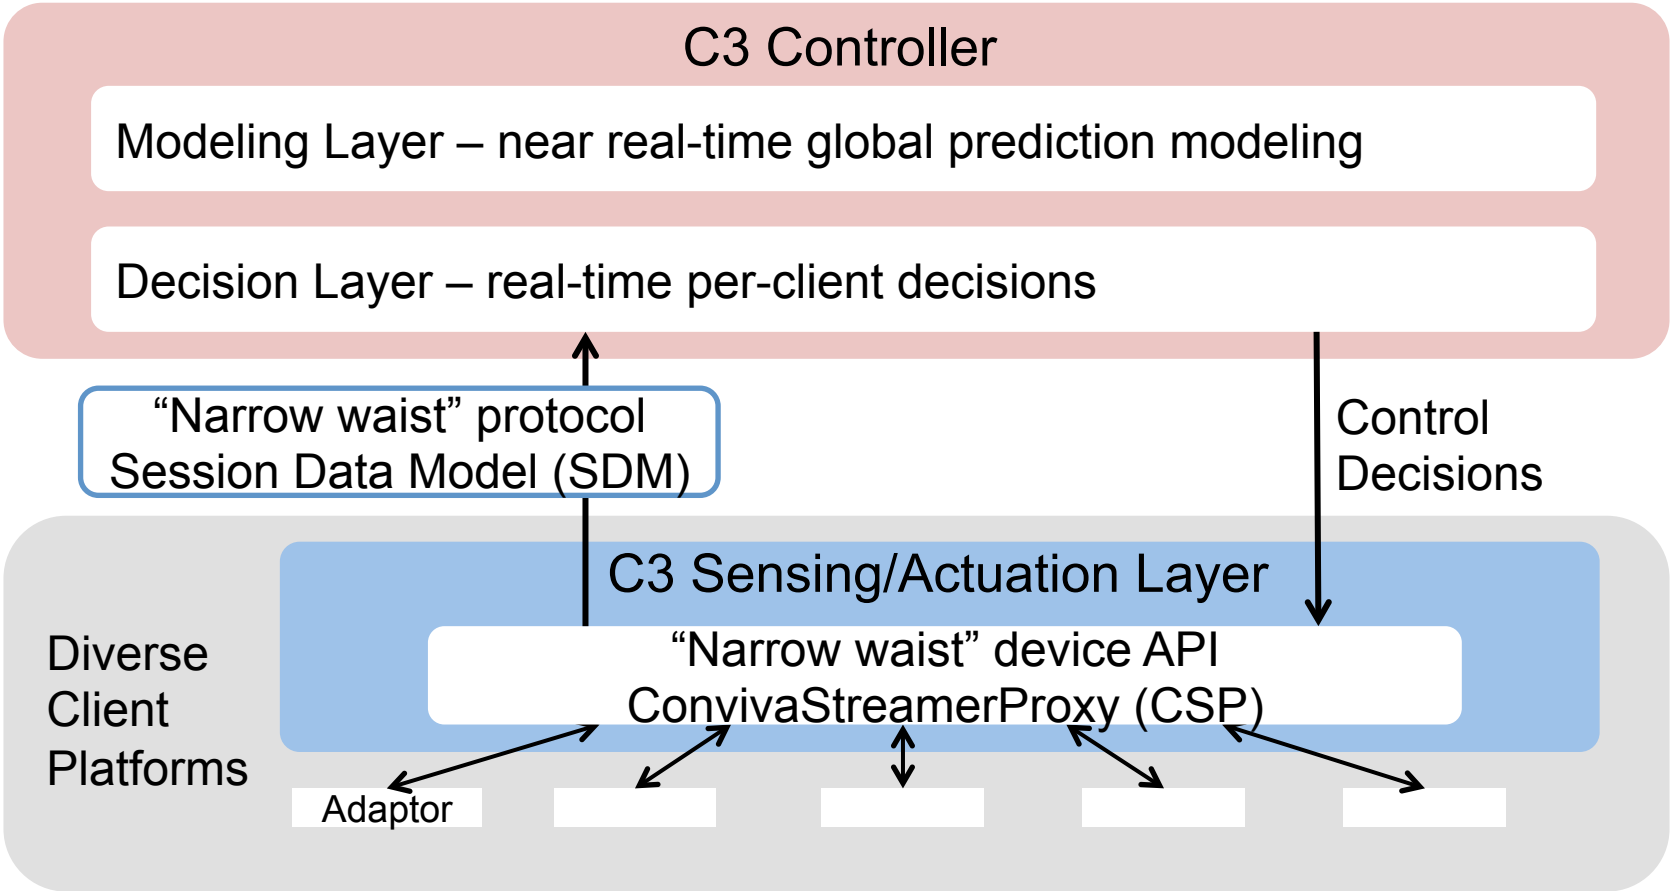

C3 ARCHITECTURE

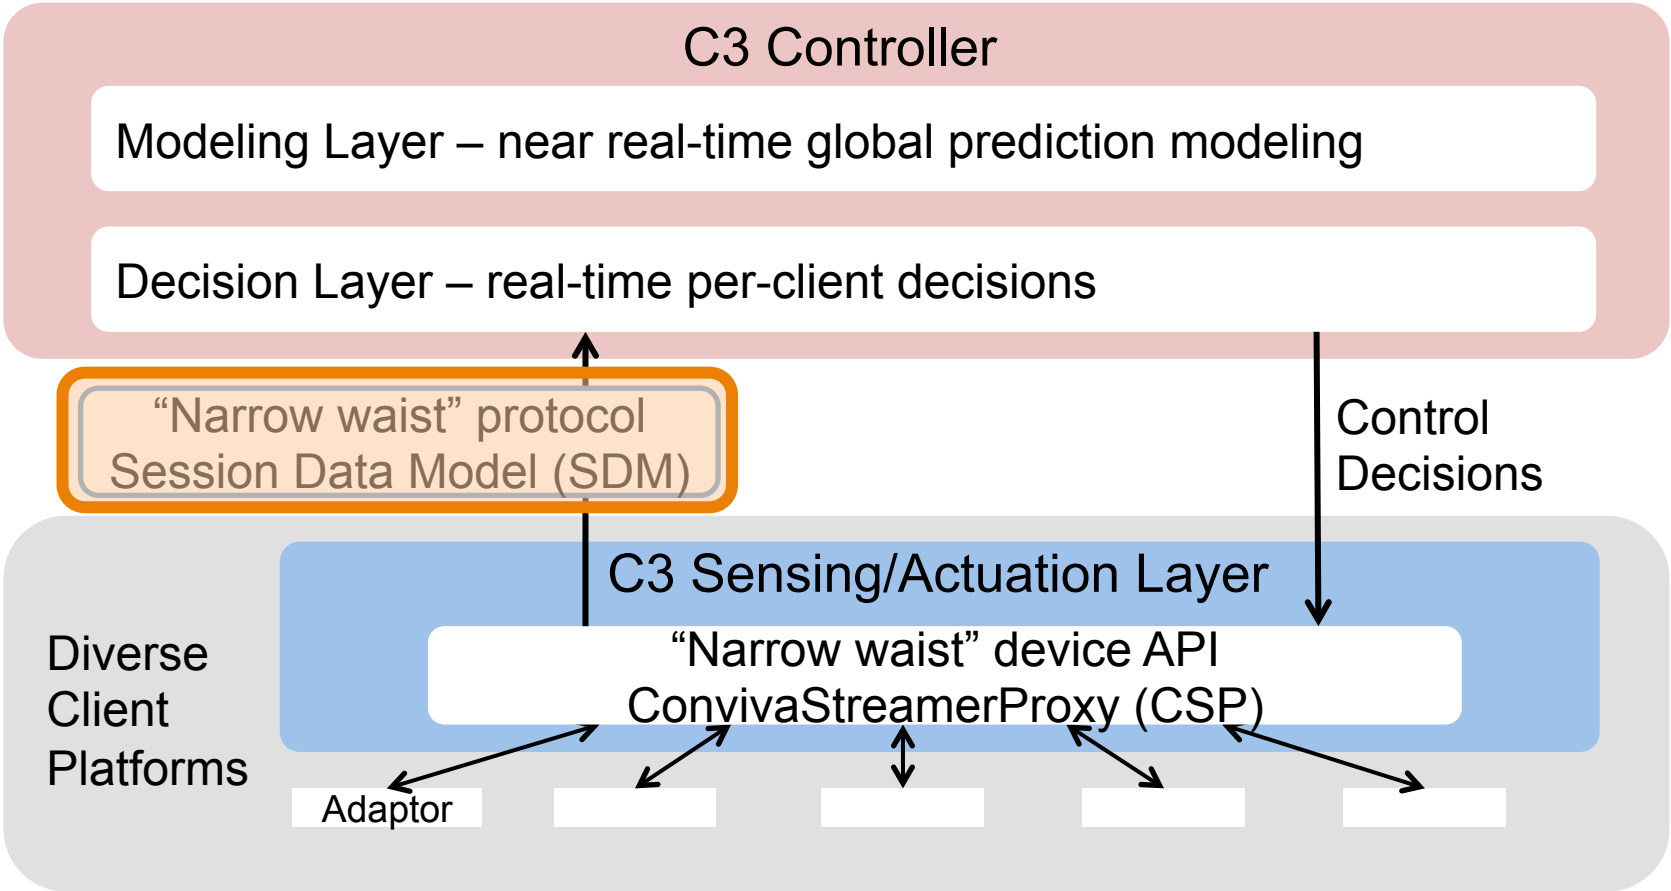

MOTIVATING EXAMPLE

Compute Video Startup Time metric

<u>Default Option:</u>

Compute Video Startup Time in the client, send to controller

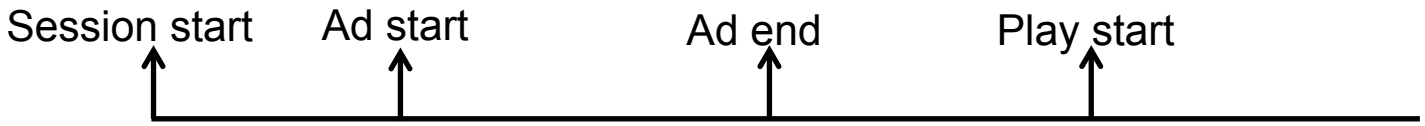

$\text{Startup Time}_1 = T(\text{play start}) - T(\text{session start})$

$\text{Startup Time}_2 = T(\text{play start}) - T(\text{session start}) - (T(\text{Ad end}) - T(\text{Ad start}))$

How do we adapt to changes?

IDEA: EXPOSE LOW-LEVEL ACTIONS

- Metric calculation pushed to the controller
- Session Data Model (SDM) – Events, States, Measurements
- States – robustness to message loss

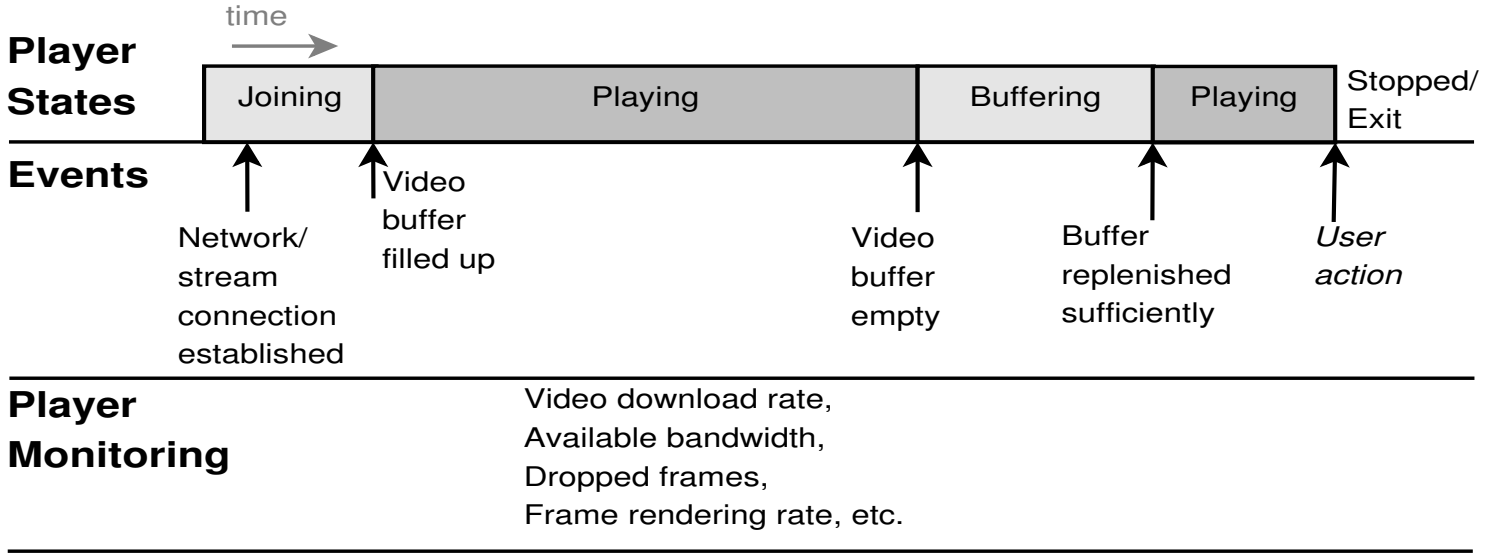

C3 ARCHITECTURE

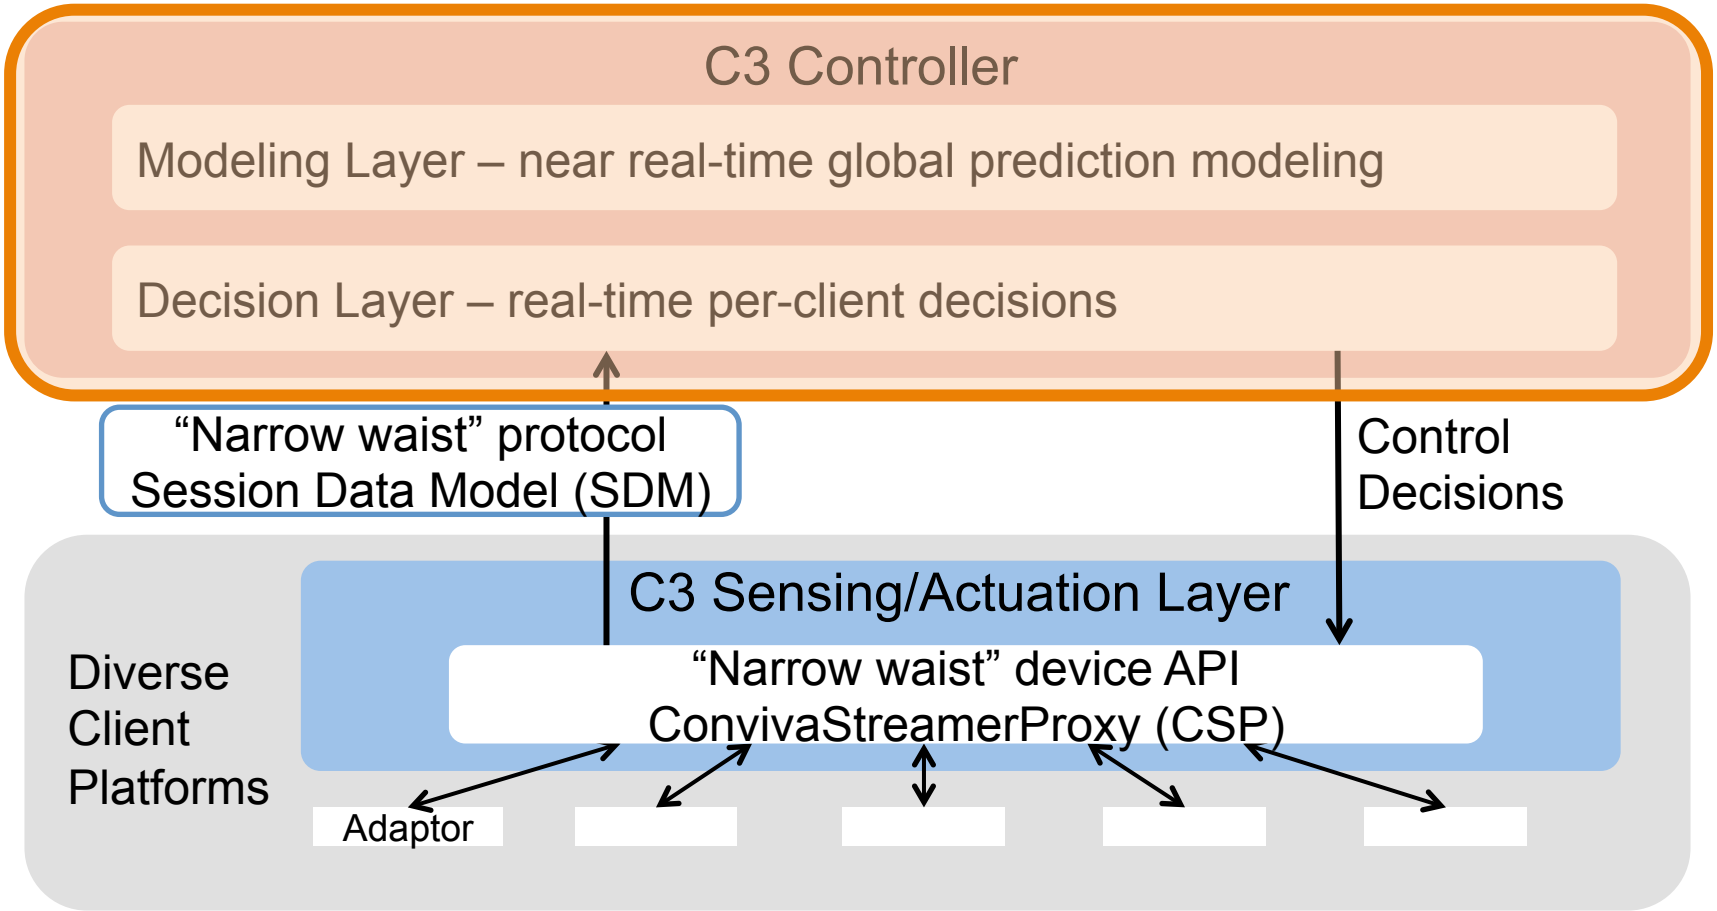

SPLIT CONTROL-PLANE

- Dual-loop control
  - Client-driven real-time loop (red)
  - Periodic global model learning and dissemination loop (blue)
- Intelligent Decision Layer
  - Per-client decisions based on global model and client quality info

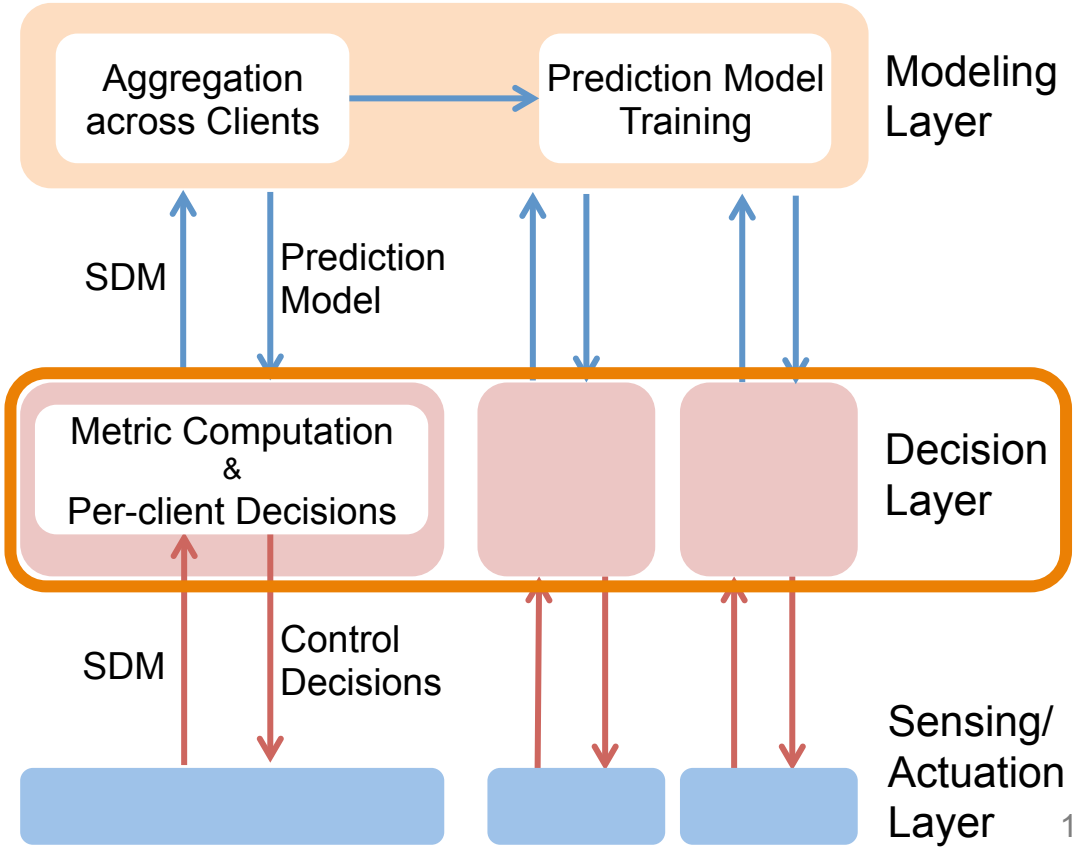

WHY SPLIT CONTROL CAN ACHIEVE NEAR OPTIMAL PERFORMANCE

- Domain specific insight – globally optimal decisions tend to be persistent on minute-level timescales
- Persistence results do not hold for individual clients - still need per-client up-to-date state

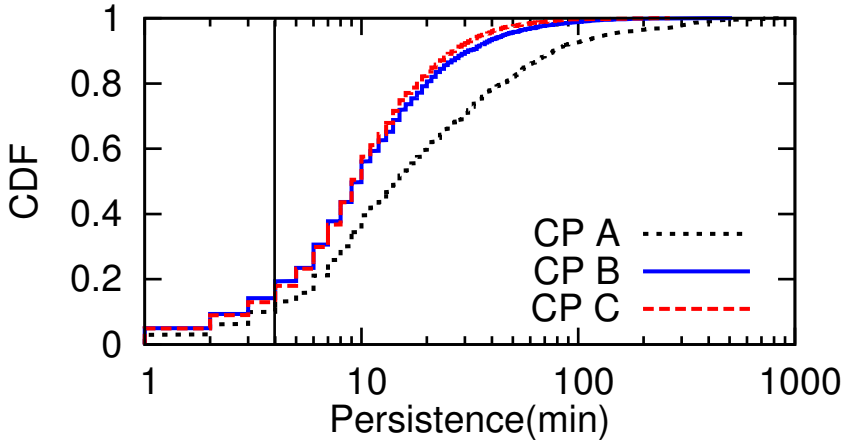

Figure shows the distribution of persistence of the best CDN for three content providers

RESPONSIVENESS & SCALE

- Sub-second response time
  - Geographically distributed decision layer
  - Geo-distribution using cloud
- Controller scale
  - Horizontally scaled decision layer
  - Burst scaling using cloud
  - Big-data technologies including Spark for modeling layer

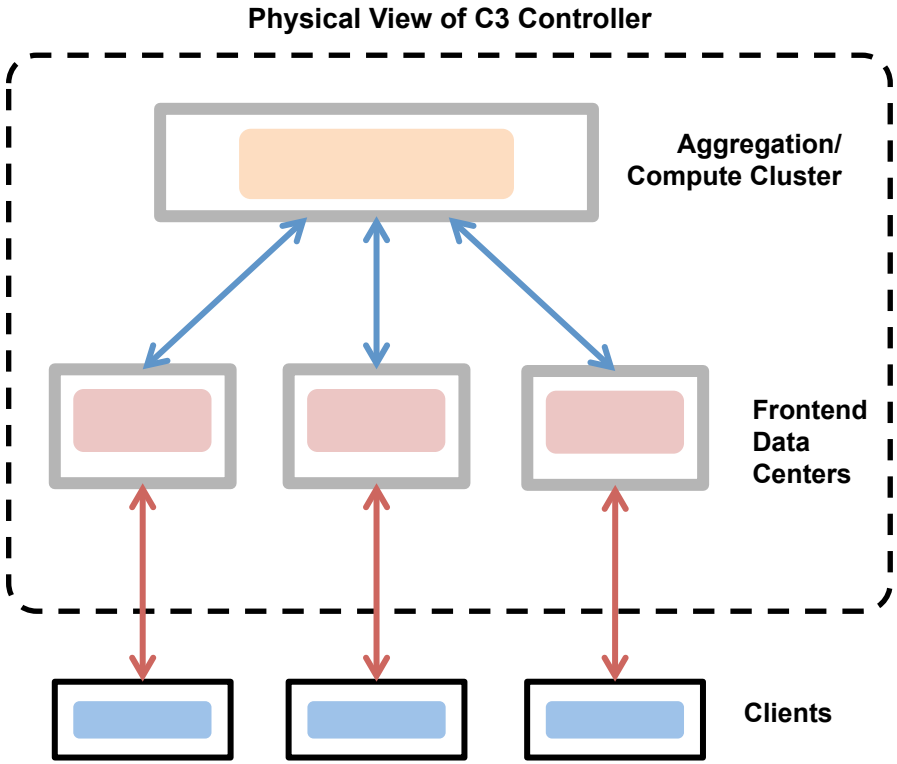

FAULT TOLERANCE

- Decision layer gracefully degrades due to stale or missing global model
- Client-side failover and decision layer state reconstruction to handle instance failure

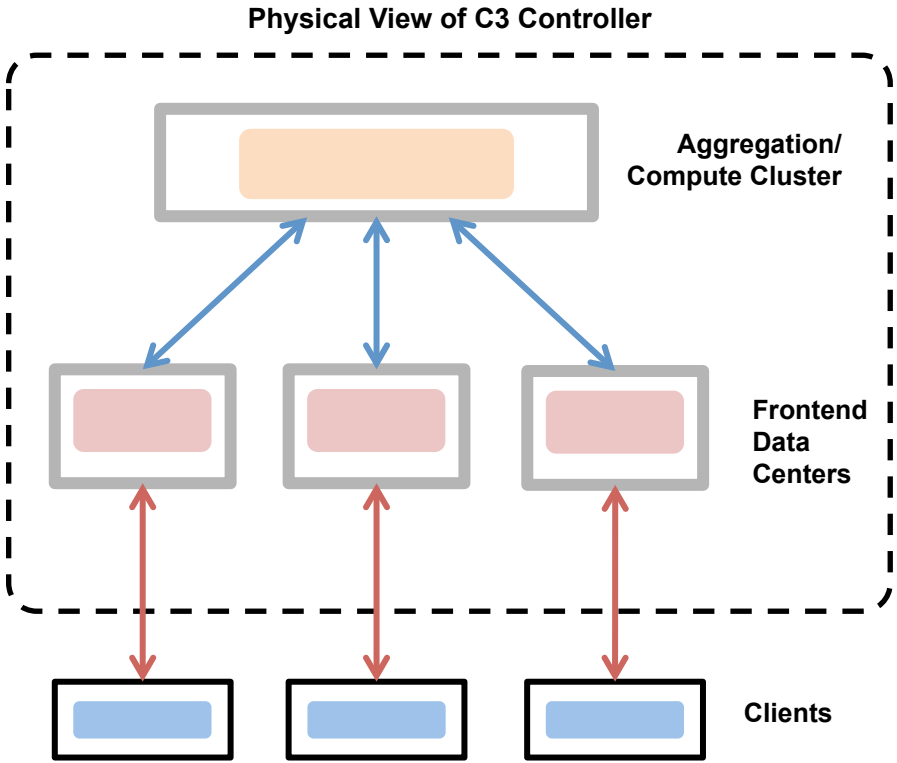

REAL-WORLD DEPLOYMENT & EVOLUTION

- Earlier phases explored alternate architectures
  - Partitioning
  - “Heavy client” (client offload of computation and decision logic)
- C3 platform used by premium video publishers over 8 years
  - Over 1 billion unique devices
  - Over 3 million concurrent devices during major events
    - E.g., World Cup, Super Bowl
- Significant quality improvement results

QUALITY IMPROVEMENT: REBUFFERING

- Premium content publisher
- One month A/B test
- Result: Greater then 50% reduction in Buffering Ratio

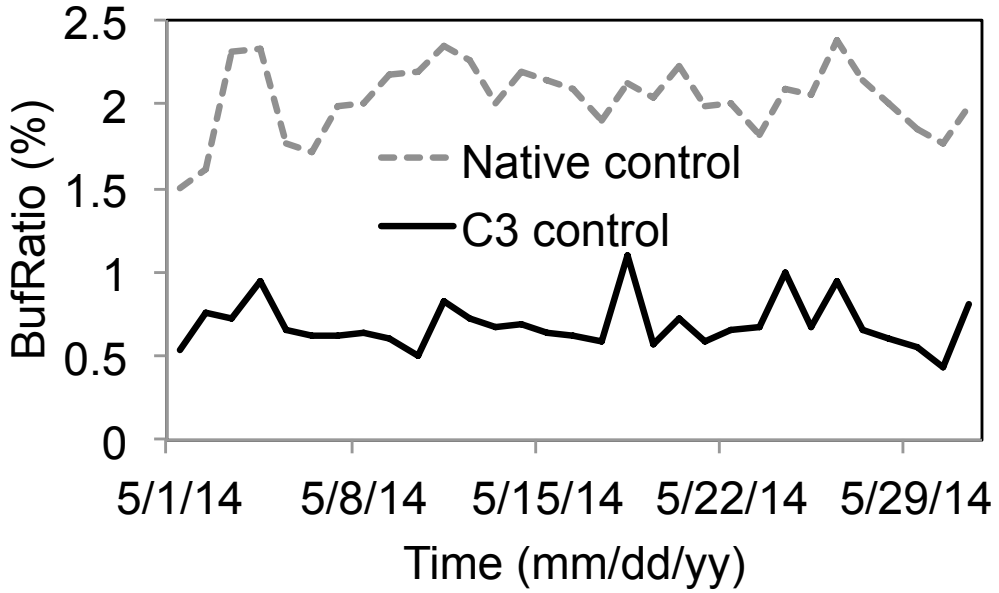

QUALITY IMPROVEMENT: START FAILURES

- Premium content publisher
- One month A/B test
- Result: Greater then 60% reduction in Video Start Failures

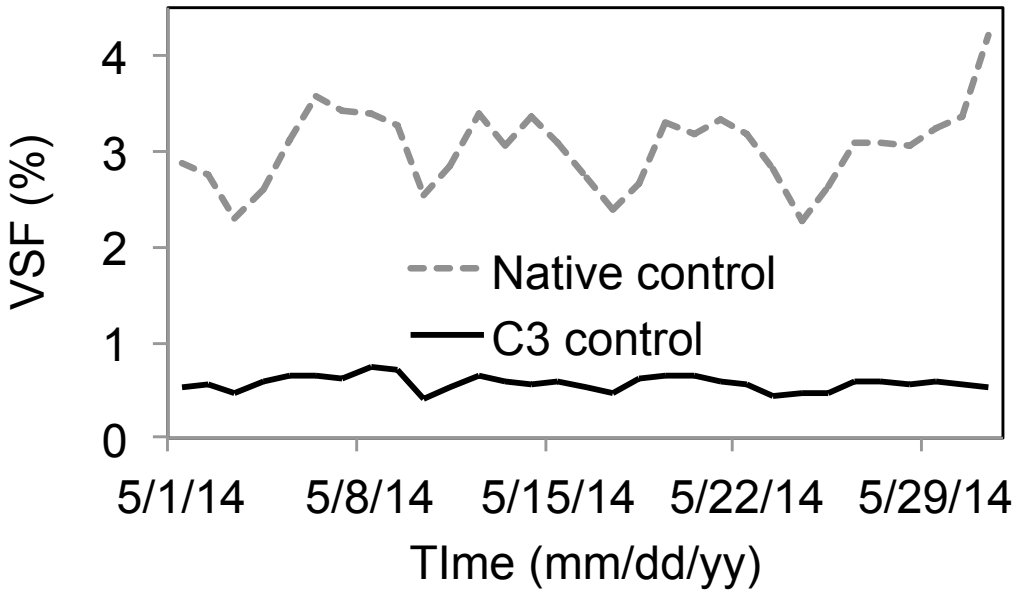

QUALITY IMPROVEMENT ACROSS PROVIDERS

Rebuffering Ratio improvement consistent results across 5 different content providers

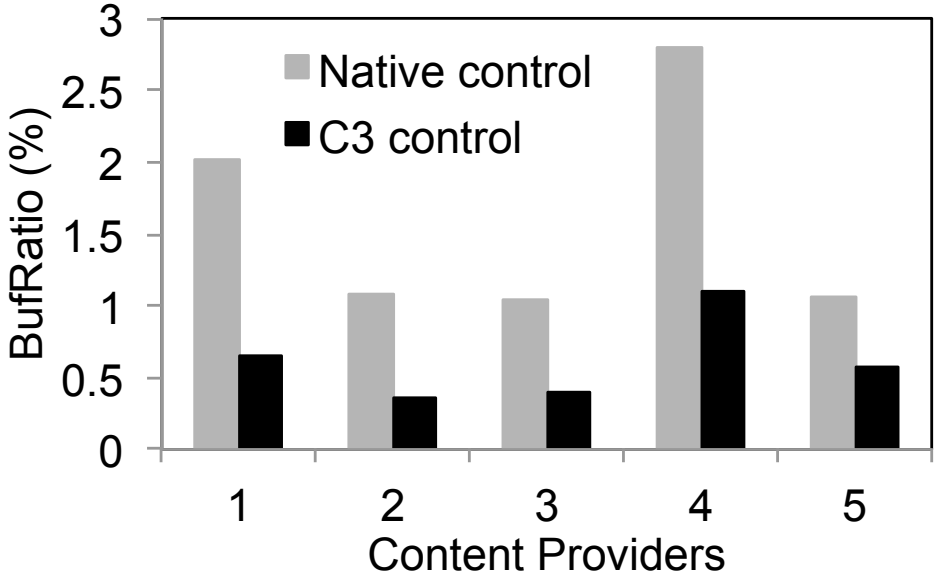

LESSONS

- Validates centralized control at an unprecedented scale
- Reinforces the case for centralized control
  - Drivers: client heterogeneity, global policies, real-time monitoring
  - Enablers: big-data, cloud, application-level resilience
- Unique challenges driving new ideas
  - Split control balances global view and real-time
  - Minimal client critical to handle heterogeneity
- Broader applicability of C3 architecture
  - Other applications (e.g., gaming, voice, conferencing)
  - Network layer control (e.g., SDN, CDN)

CONCLUSIONS

- Television is coming to the Internet, high expectations on quality
- Large optimization space for video quality
  - ➢ Drives the need for a proactive approach
- C3 is a centralized predictive control approach
- Solves key challenges: (a) client heterogeneity, (b) Internet-scale
  - Minimal sensing/actuation layer & move all computation to the controller
  - Split control-plane scale-out solution exploiting application-level resilience
- C3 used by many content providers with quality improvement results