

Faces in the Clouds: Long-Duration, Multi-User, Cloud-Assisted Video Conferencing



[Imperial College](#): Richard G. Clegg, Peter Pietzuch

[UCL](#): Miguel Rio, David Griffin

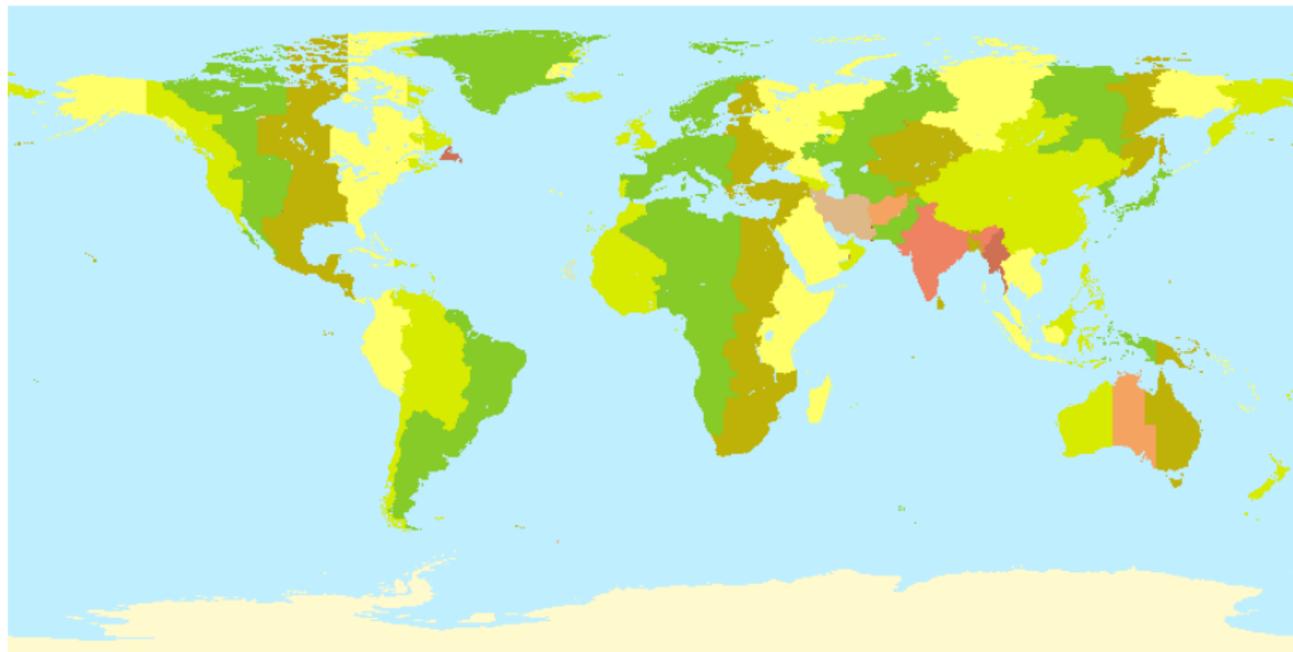
[Sky UK Network Services](#): Raul Landa

[British Telecom](#): Peter Hughes, Ian Kegel, Tim Stevens, Doug Williams

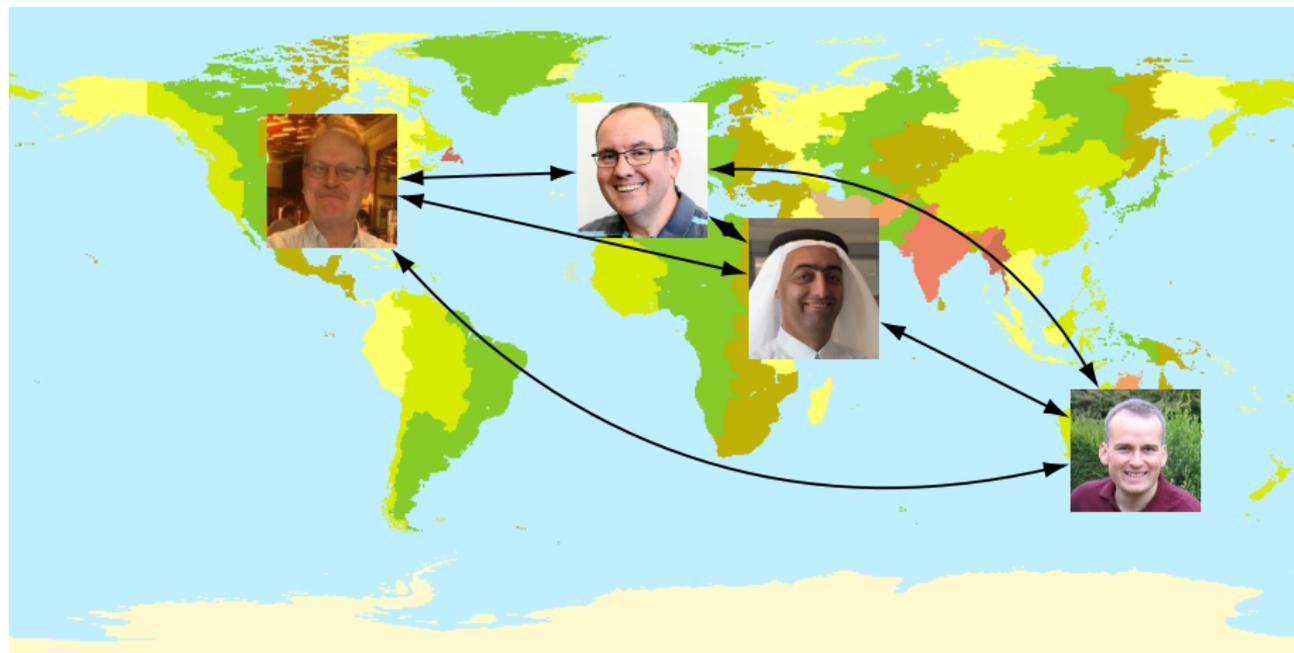
MoN15: Fifteenth Mathematics of Networks Meeting

- Website: <http://www.monmeetings.org/meeting15>
- Date: 23rd September.
- Location: Bath University.
- Topic: Dynamic processes on networks.
- Talks on any topic may be accepted.
- Cross disciplinary, not “hardcore” abstract mathematics.
- Abstracts: 19th August – title + abstract of 200-300 words.
- Free to attend.
- Email richard@richardclegg.org for more info.

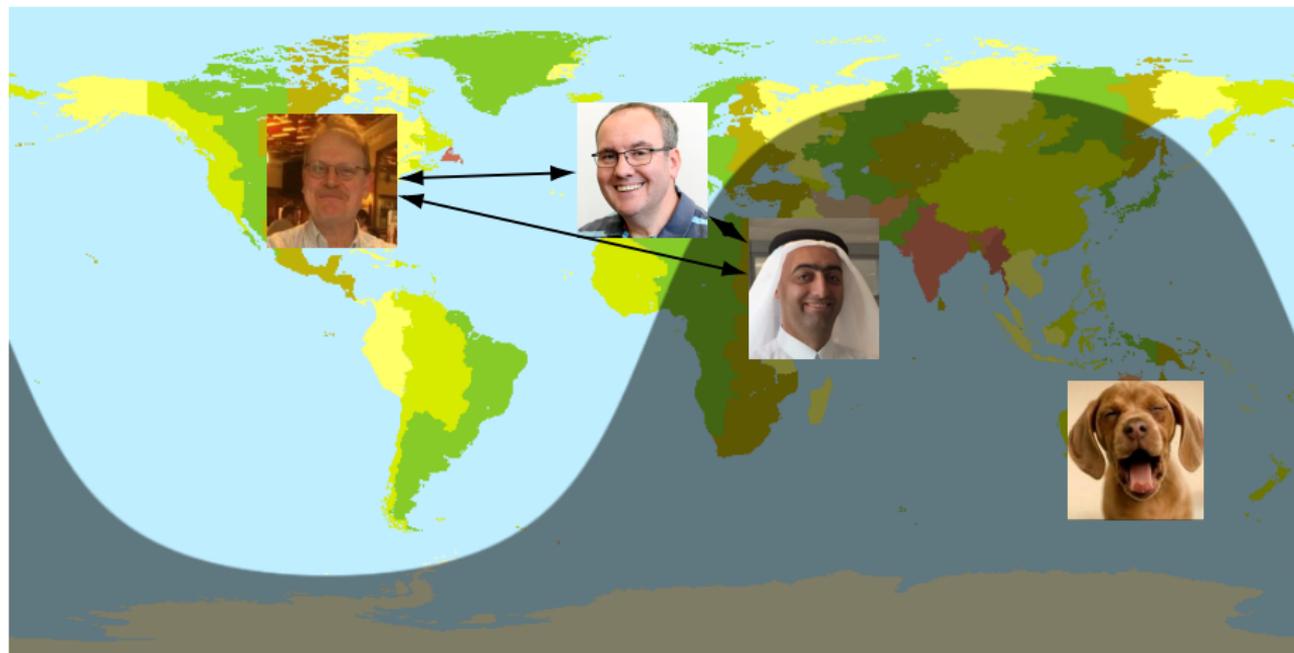
Long-duration (sun following) video chat



Long-duration (sun following) video chat



Long-duration (sun following) video chat



Long-duration (sun following) video chat



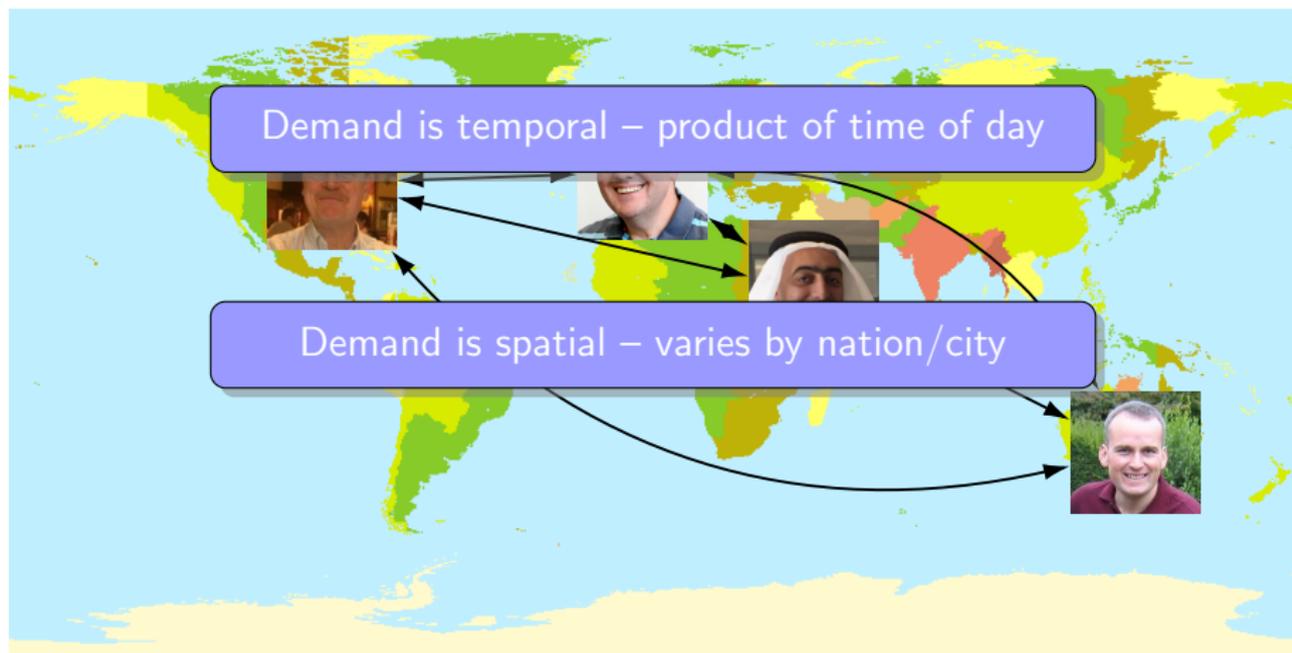
Long-duration (sun following) video chat



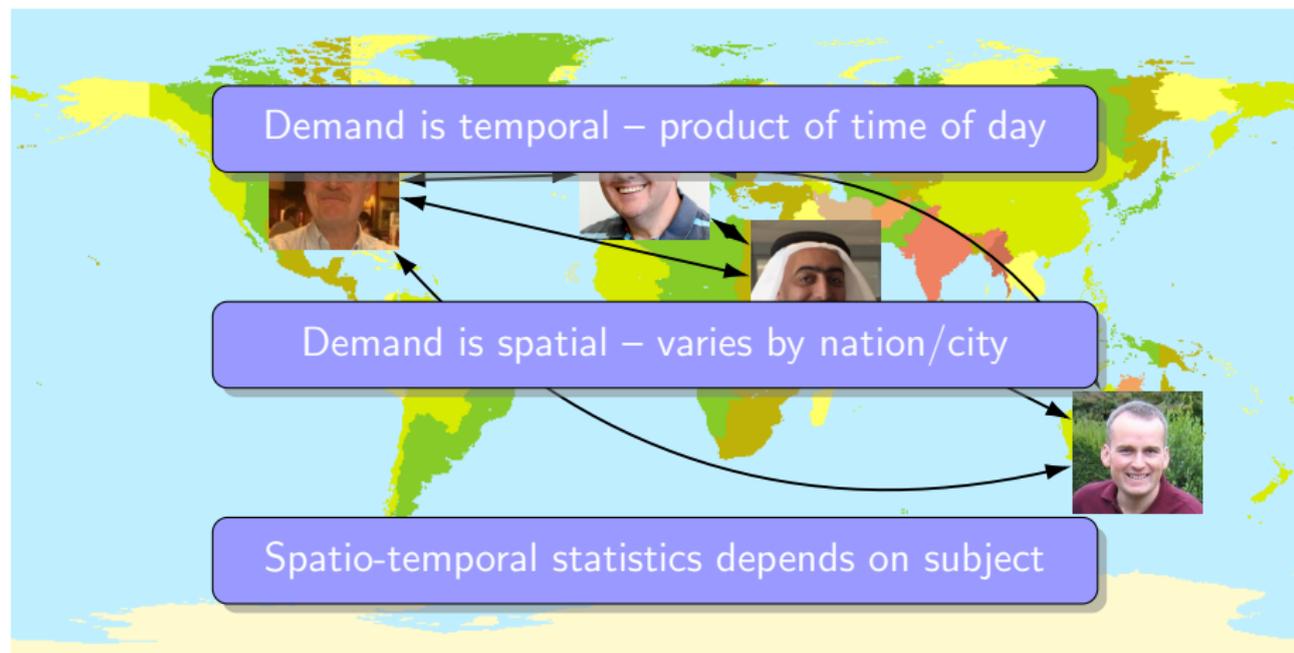
Long-duration (sun following) video chat



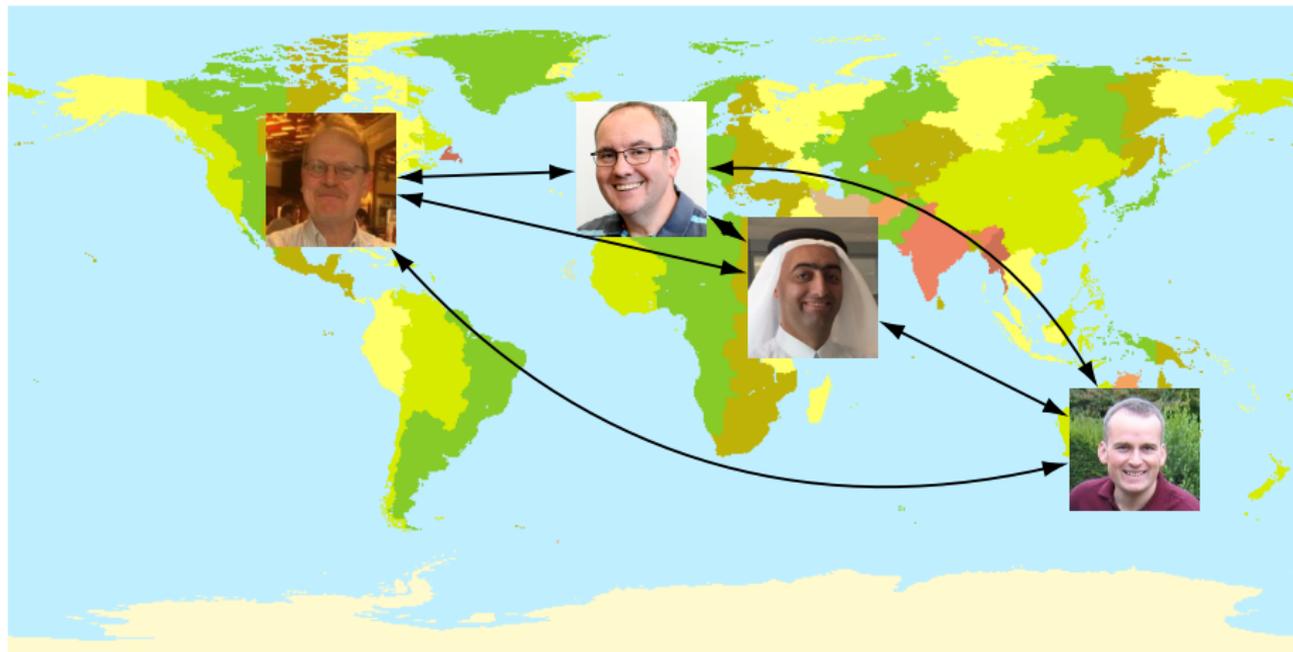
Long-duration (sun following) video chat



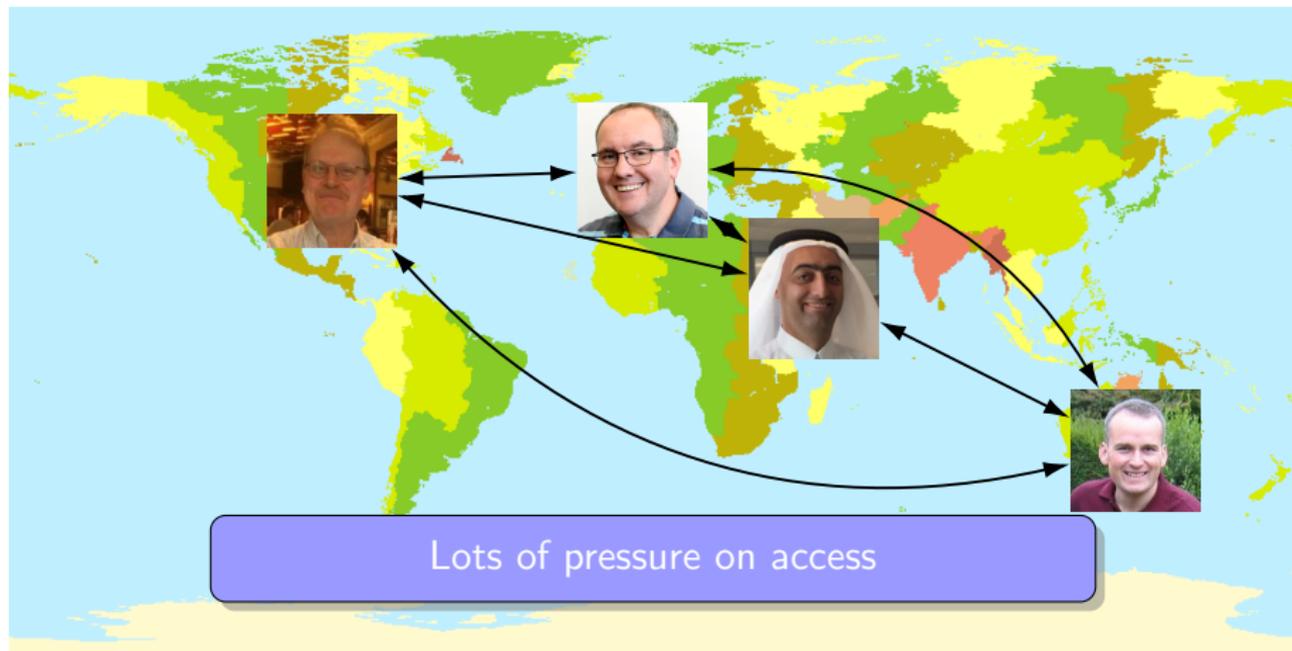
Long-duration (sun following) video chat



Cloud-assisted video chat



Cloud-assisted video chat



Cloud-assisted video chat



Hamed: Really enjoying Richard's talk at Coseners.

Typical video chat one person has "focus".

Cloud-assisted video chat



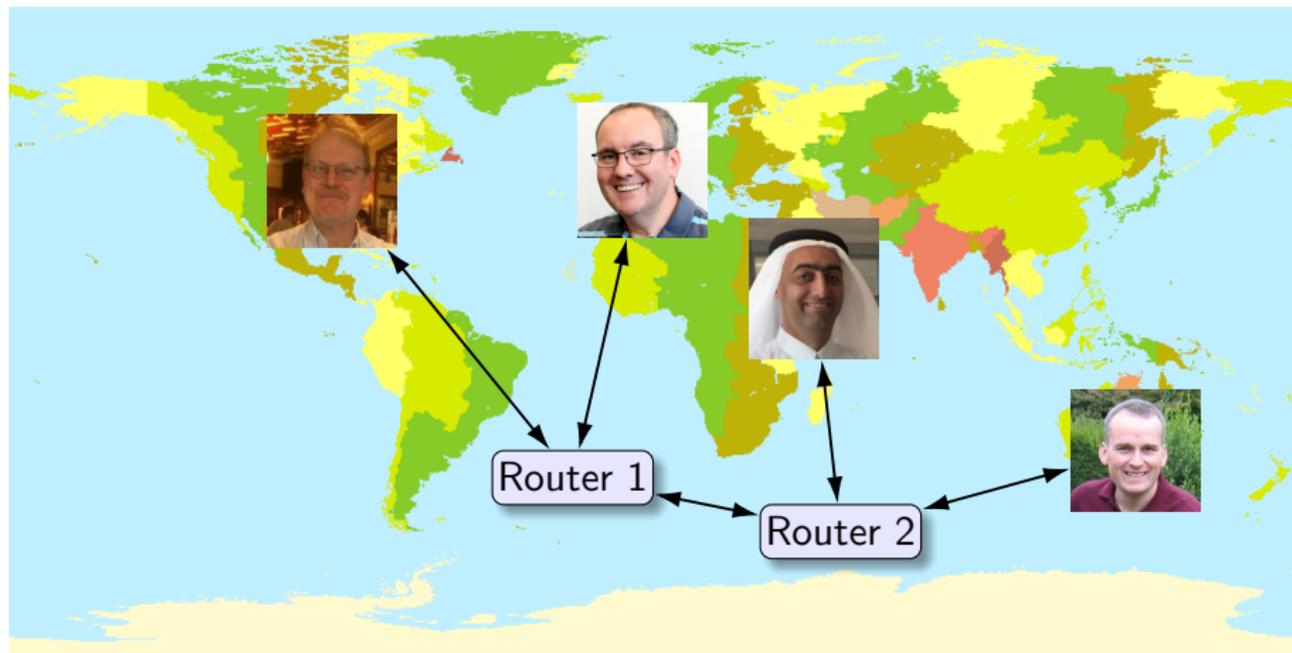
Andrew: Should be more about NetFPGA. Not mentioned cloud yet.

Only person with "focus" needs HD stream.

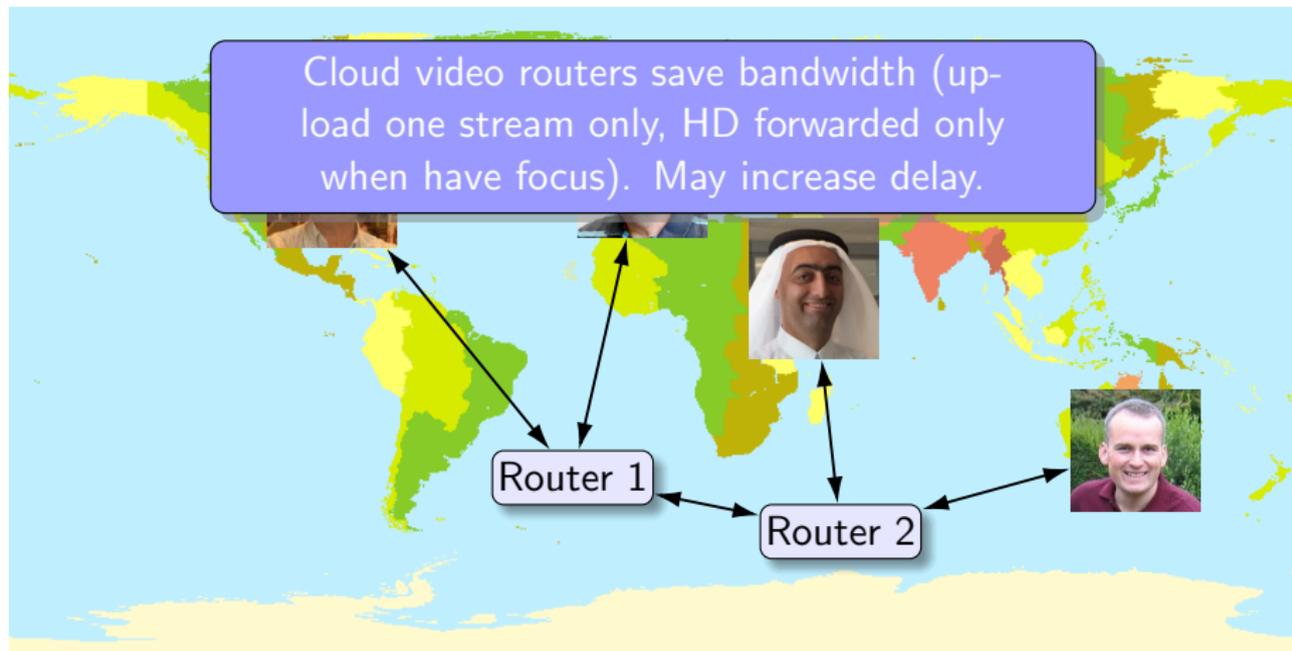
Cloud-assisted video chat



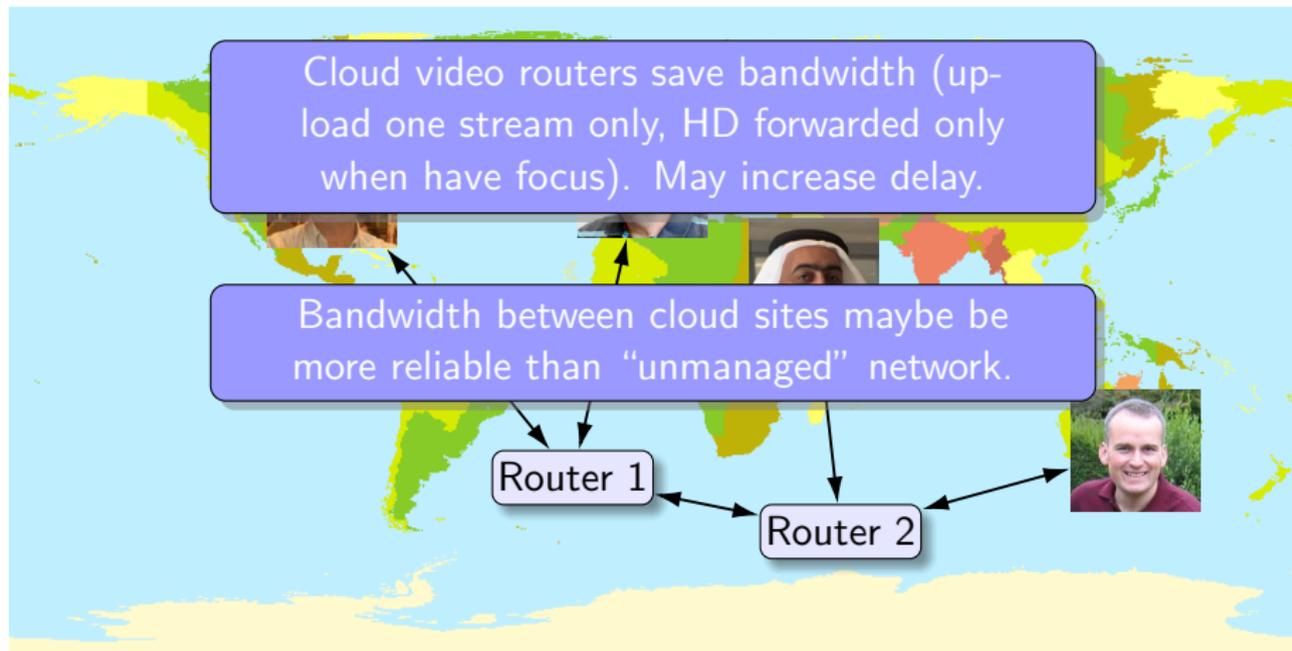
Cloud-assisted video chat



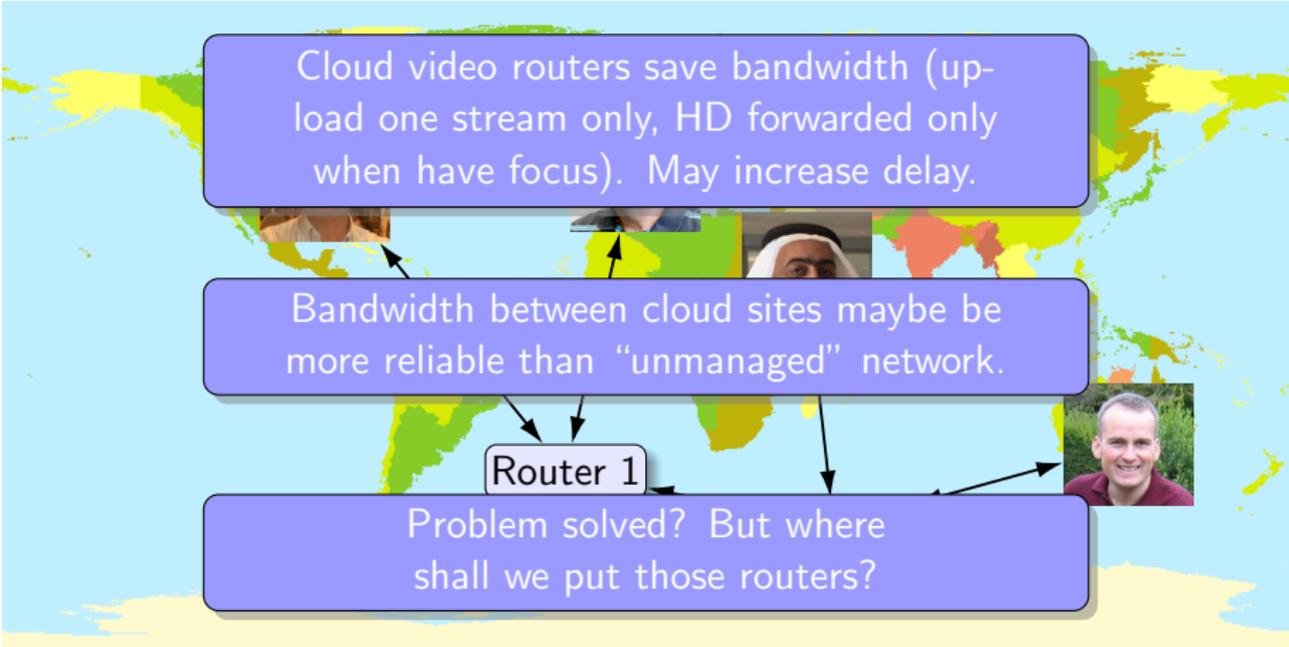
Cloud-assisted video chat



Cloud-assisted video chat



Cloud-assisted video chat



Cloud video routers save bandwidth (upload one stream only, HD forwarded only when have focus). May increase delay.

Bandwidth between cloud sites maybe be more reliable than “unmanaged” network.

Router 1

Problem solved? But where shall we put those routers?

Assumptions

Multi-user, long-duration, cloud-assisted video-chat.

Video routers situated in cloud can reduce bandwidth.

Router positioning can vary as chat population evolves.

Problem statement

Assumptions

Multi-user, long-duration, cloud-assisted video-chat.
Video routers situated in cloud can reduce bandwidth.
Router positioning can vary as chat population evolves.

Problem statement (static)

Choose n router locations to instantiate from m cloud locations.
Want to reduce end-to-end delay experienced by users.
Want to reduce time on “unmanaged” network.

Problem statement

Assumptions

Multi-user, long-duration, cloud-assisted video-chat.
Video routers situated in cloud can reduce bandwidth.
Router positioning can vary as chat population evolves.

Problem statement (static)

Choose n router locations to instantiate from m cloud locations.
Want to reduce end-to-end delay experienced by users.
Want to reduce time on “unmanaged” network.

Problem statement (dynamic)

Allow system to periodically update which n router locations are used.

Strategies

Random strategy

Choose n distinct router locations from m cloud locations completely at random.

Random strategy

Choose n distinct router locations from m cloud locations completely at random.

Advantage: Easy to code.

Advantage: Easy to come up with better strategy.

Strategies

Random strategy

Choose n distinct router locations from m cloud locations completely at random.

Advantage: Easy to code.

Advantage: Easy to come up with better strategy.

Dynamic strategy

Cluster current users of chat into n distinct clusters.

Pick n router locations from m cloud locations choosing cloud closest to each cluster centroid.

Update as users move.

Strategies

Random strategy

Choose n distinct router locations from m cloud locations completely at random.

Advantage: Easy to code.

Advantage: Easy to come up with better strategy.

Dynamic strategy

Cluster current users of chat into n distinct clusters.

Pick n router locations from m cloud locations choosing cloud closest to each cluster centroid.

Update as users move.

Static strategy

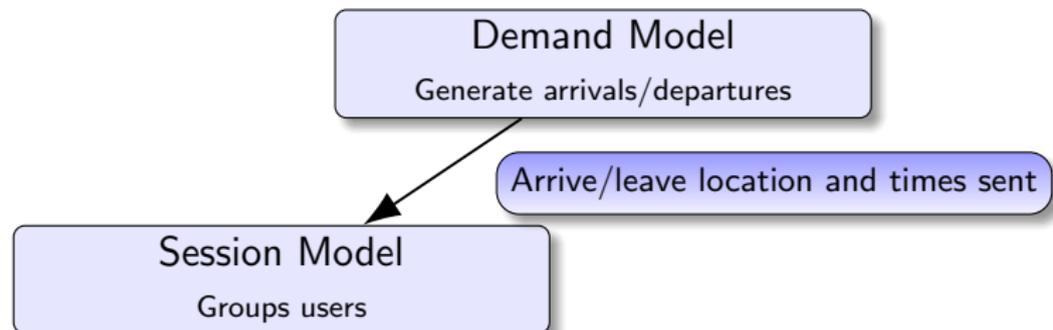
Cluster based on predicted not actual demand.

Pick n locations (and always use only those locations) from m cloud locations as before.

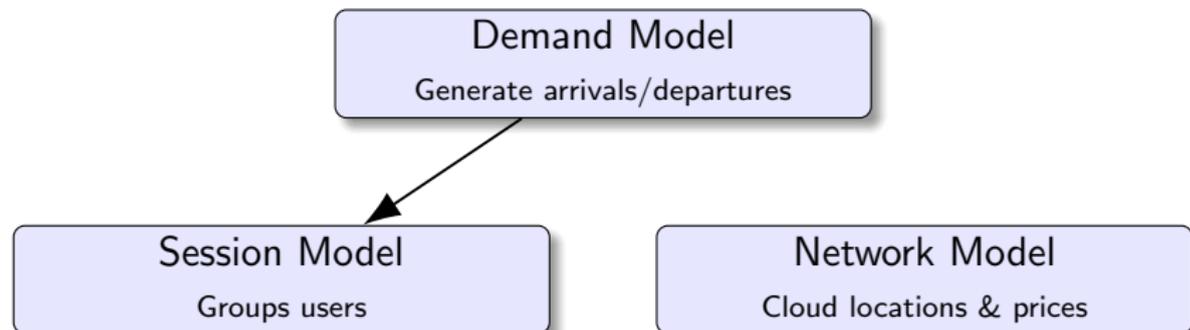
Demand Model

Generate arrivals/departures

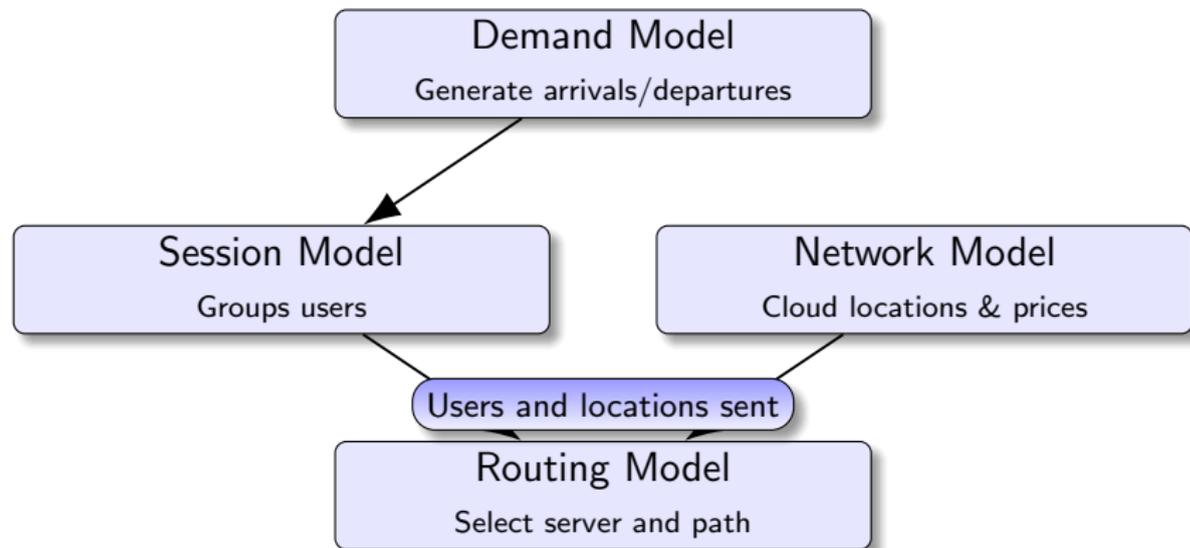
Data driven simulation



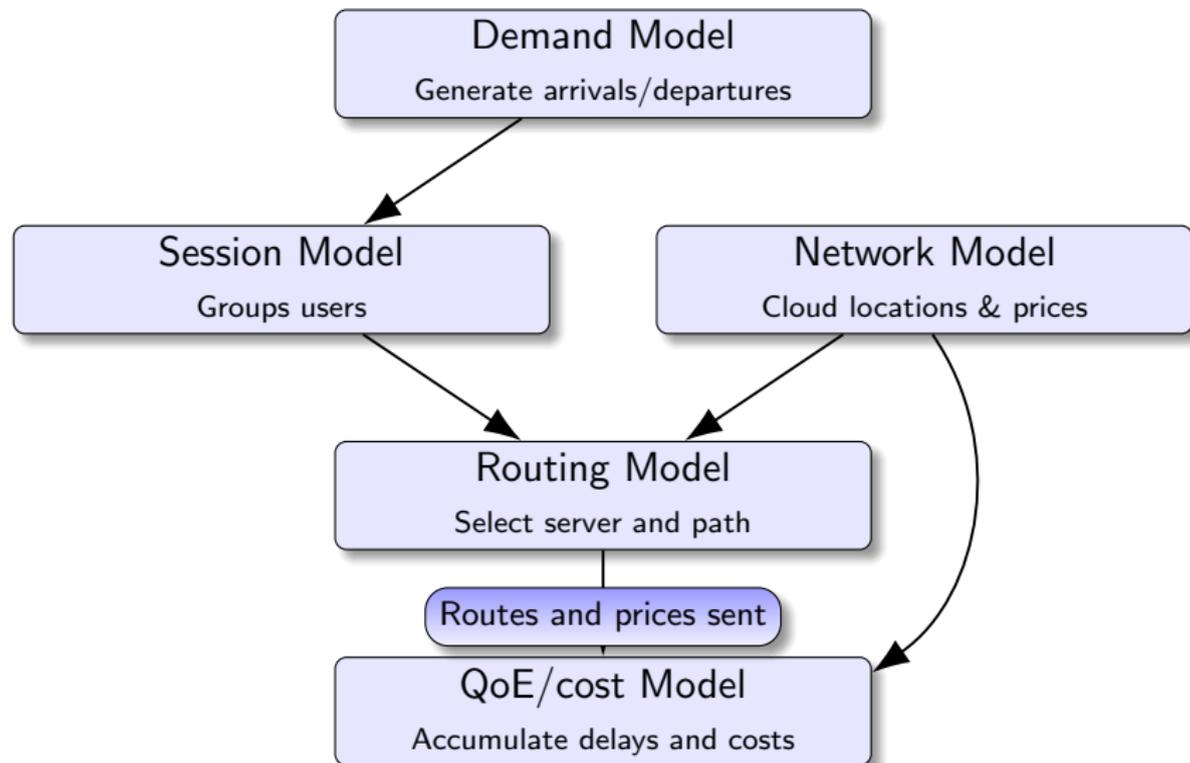
Data driven simulation



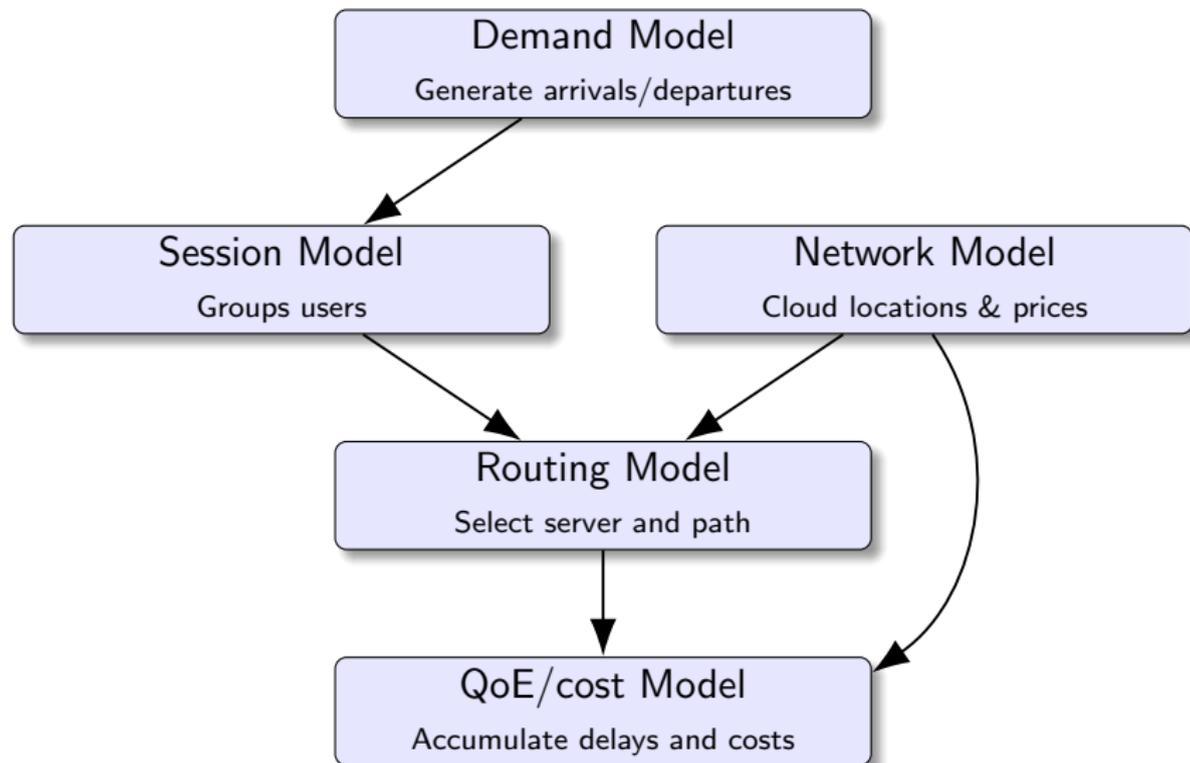
Data driven simulation



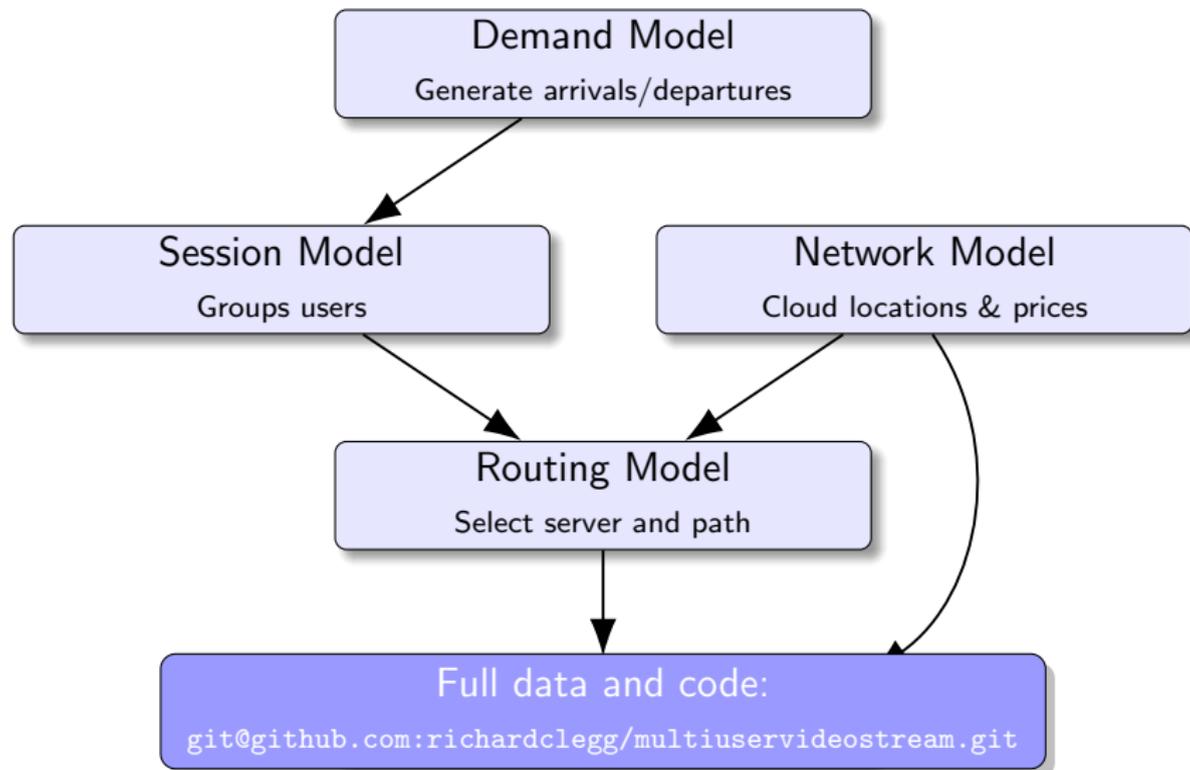
Data driven simulation



Data driven simulation



Data driven simulation



Scenarios

Video poker scenario



Video poker scenario



- Users grouped into “tables” .
- Data from online poker sites about countries of origin.
- Data about length of play sessions (lognormal dist).

Scenarios

Video poker scenario



- Users grouped into “tables”.
- Data from online poker sites about countries of origin.
- Data about length of play sessions (lognormal dist).

MOOC chat scenario



Scenarios

Video poker scenario



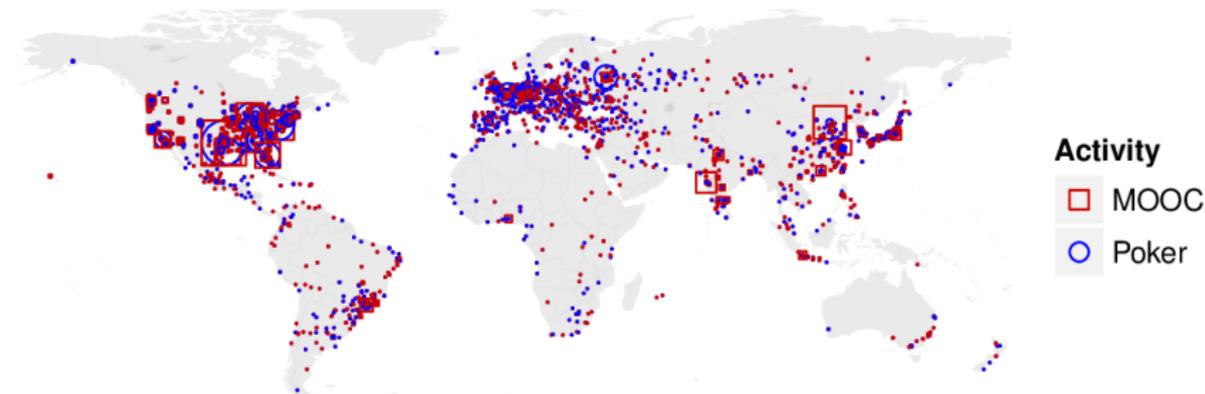
- Users grouped into “tables”.
- Data from online poker sites about countries of origin.
- Data about length of play sessions (lognormal dist).

MOOC chat scenario



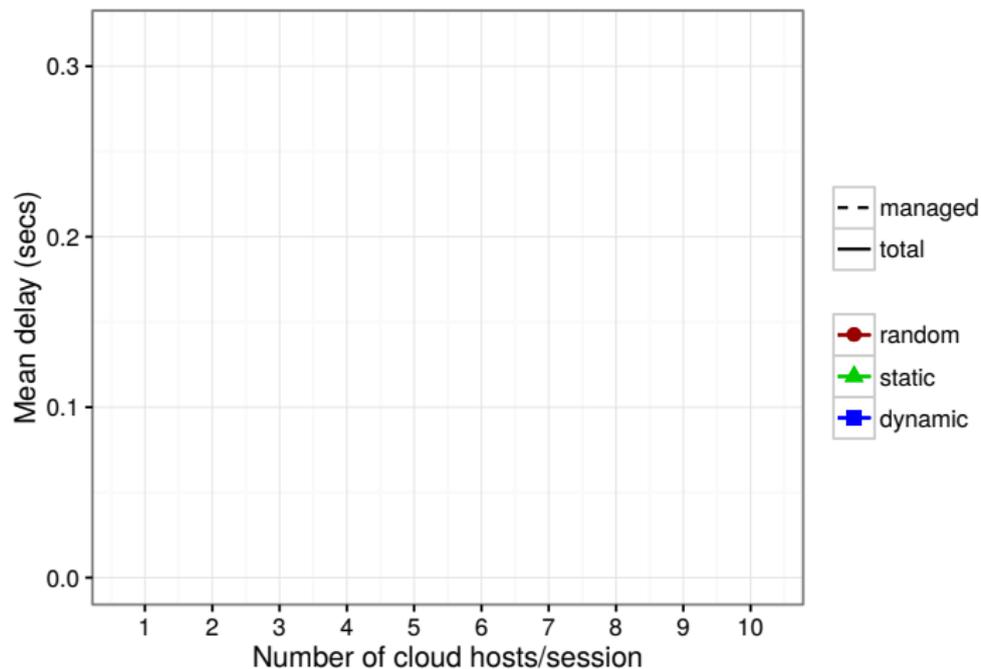
- Single large chat room.
- Data from several MOOCs about country of origin.
- Data from telcos about call durations (lognormal dist).

Scenario demands

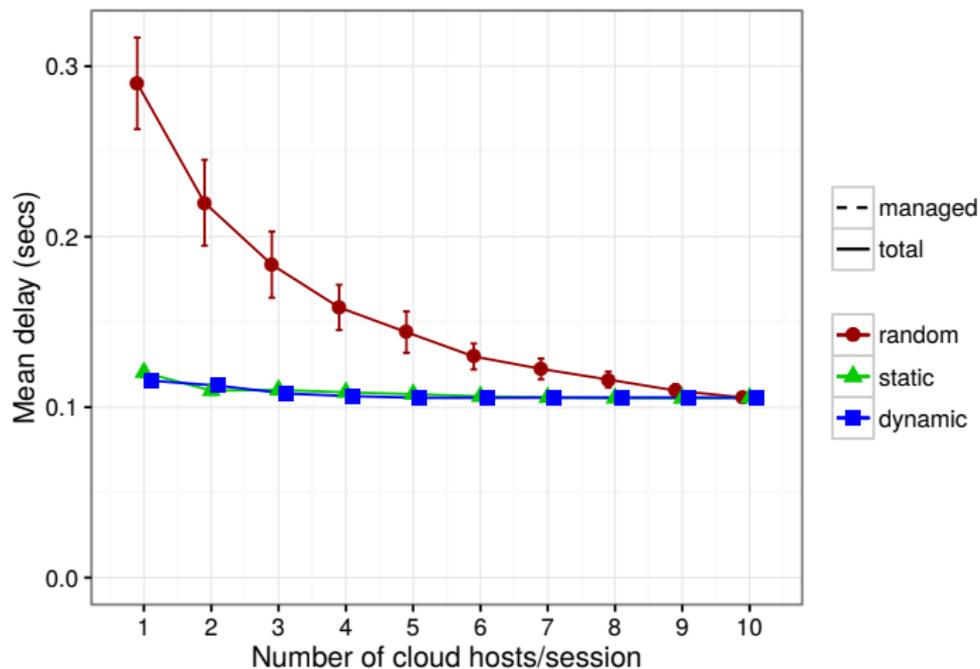


- Combine data about national demand with data about server locations.
- Add location/pricing for 10 Amazon EC2 data centres.

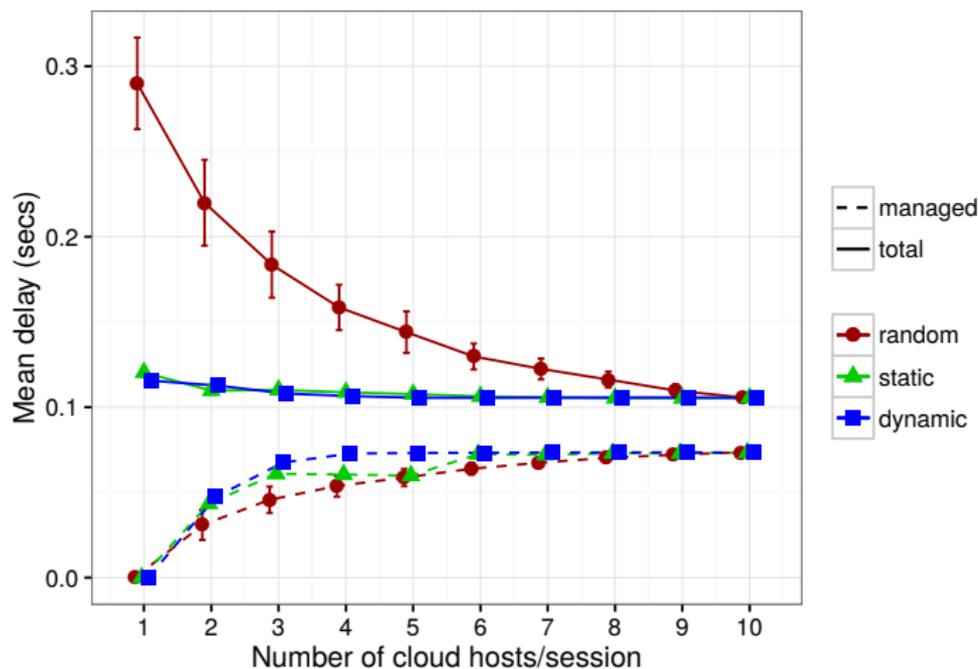
QoE vs no cloud hosts/session – poker scenario



QoE vs no cloud hosts/session – poker scenario

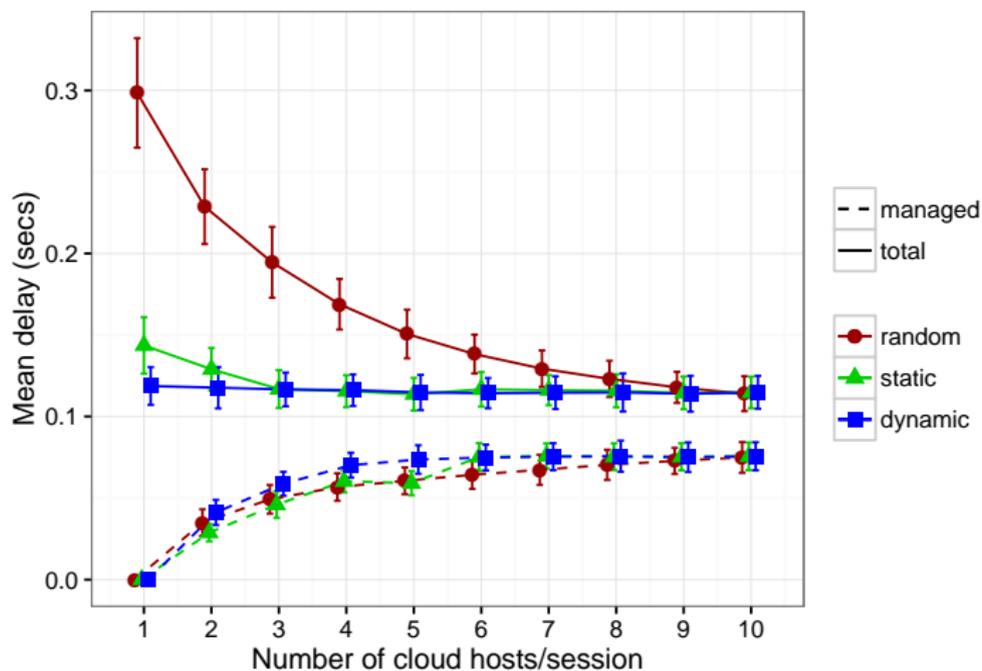


QoE vs no cloud hosts/session – poker scenario



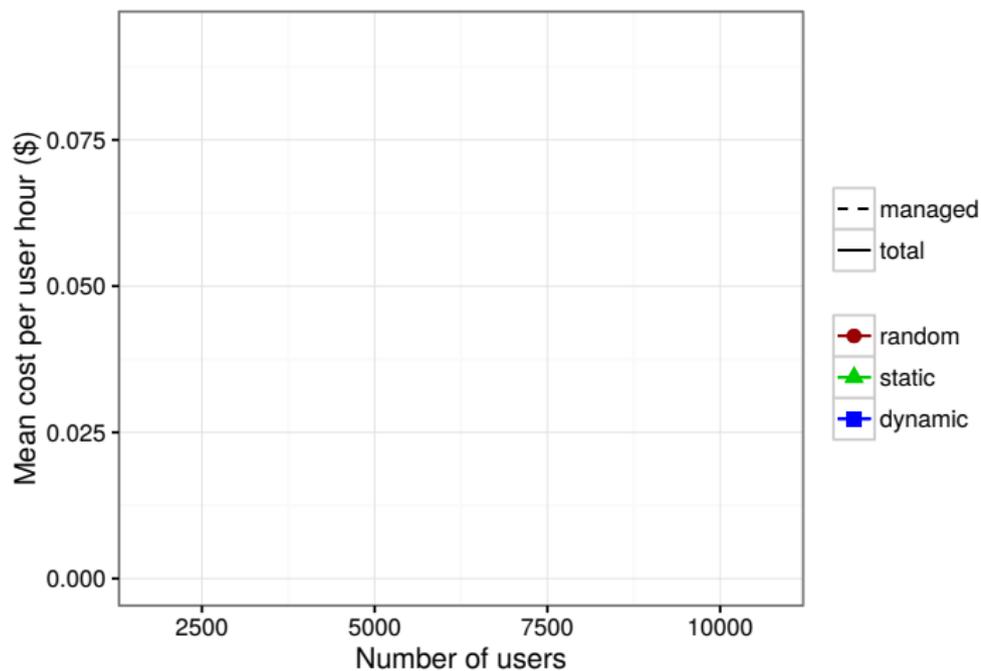
- Static and dynamic perform similarly.
- More than 3 servers gets most of benefit.
- Gap shows time on “unmanaged” network.

QoE vs no cloud hosts/session – MOOC scenario

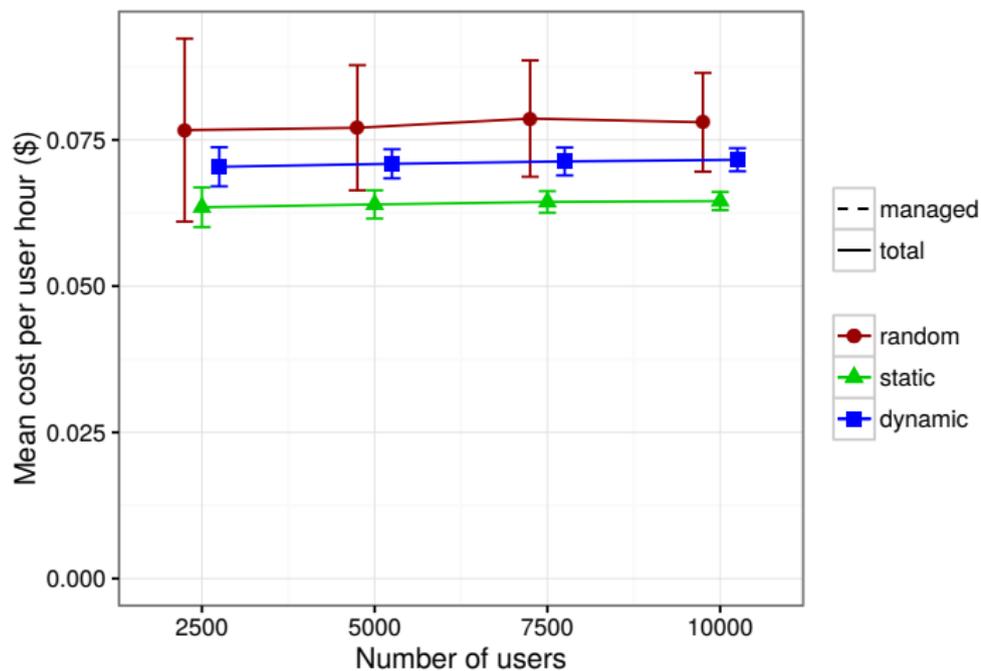


■ Similar story (more variability).

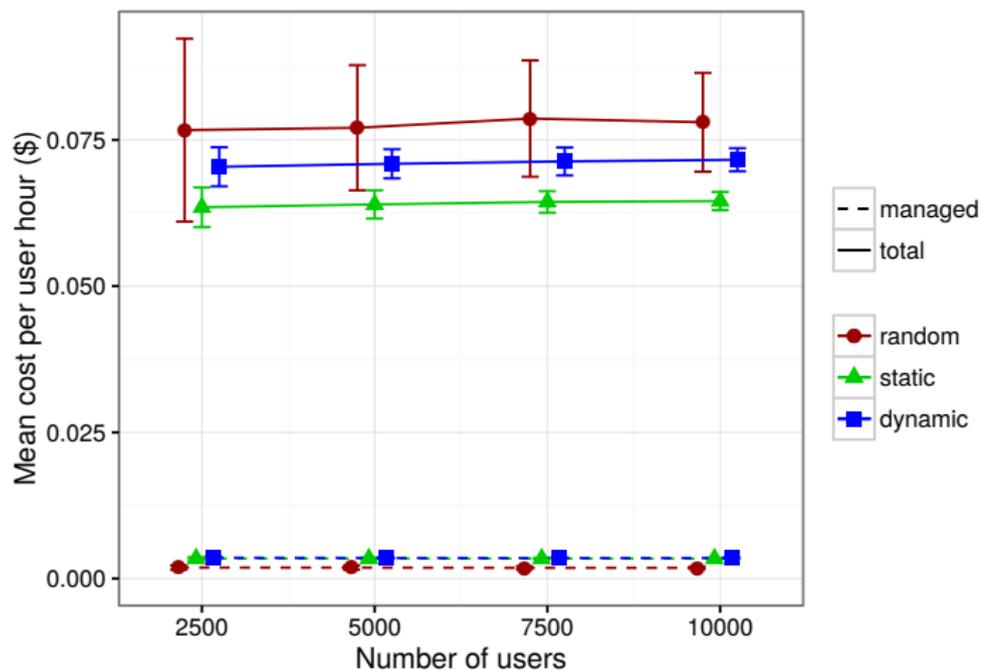
Cost vs no users (3 cloud hosts) – poker scenario



Cost vs no users (3 cloud hosts) – poker scenario

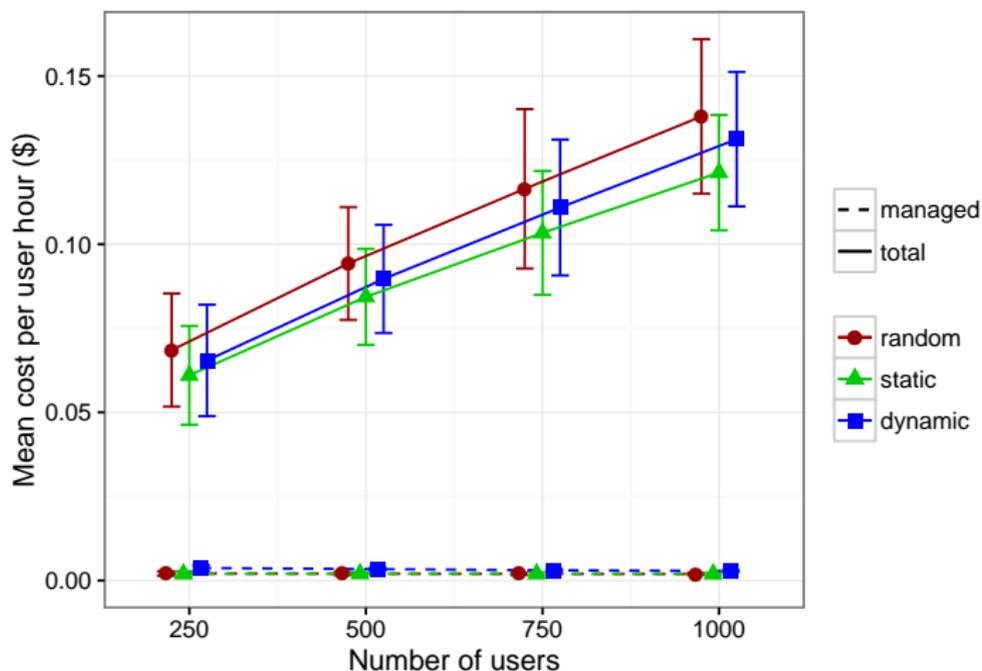


Cost vs no users (3 cloud hosts) – poker scenario



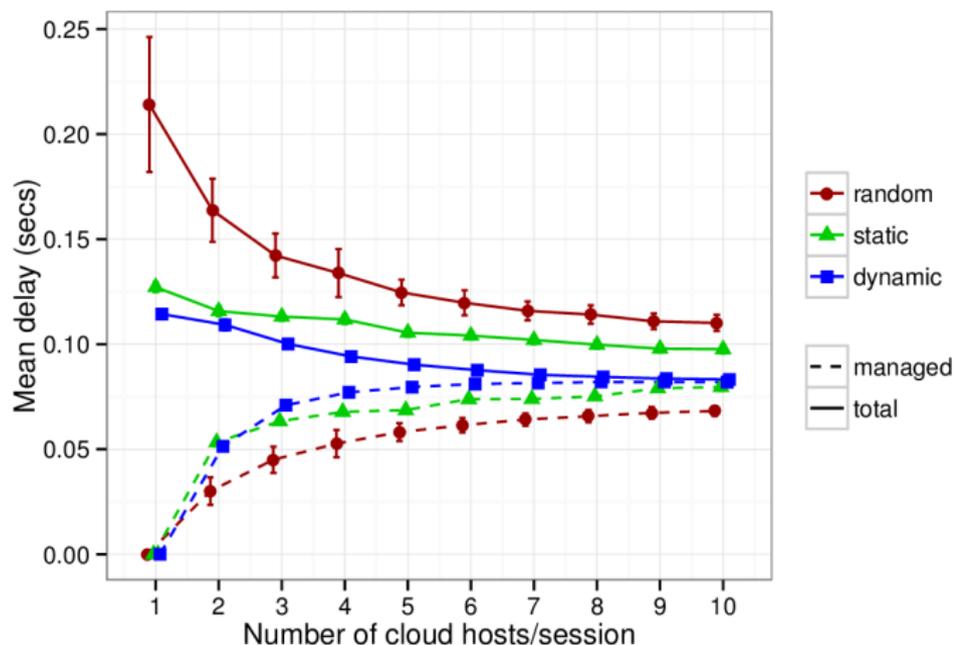
- Cost mainly data into network (no CPU cost).
- Cost per user scales well with number of users.

Cost vs no users (3 cloud hosts) – MOOC scenario



■ Cost per user increases as size of chat room increases.

QoE vs no cloud hosts/session (fog computing) – poker



- Choice from larger group 2,507 real-world server locations.
- Less overall delay. Less time on unmanaged network.
- Dynamic has advantage over static server choice.

Conclusions

Scaling

Mean delay not affected by number of users.

Cost per user increases linearly with number of users in same session.

Conclusions

Scaling

Mean delay not affected by number of users.

Cost per user increases linearly with number of users in same session.

Node choice strategy

Random is not very good.

With predictable demand and little choice of nodes static performs similarly to dynamic.

Conclusions

Scaling

Mean delay not affected by number of users.

Cost per user increases linearly with number of users in same session.

Node choice strategy

Random is not very good.

With predictable demand and little choice of nodes static performs similarly to dynamic.

Fog computing scenario

Choosing data centres from larger set gives better QoE less time on unmanaged.

Dynamic gets advantage over static with more choice.

Conclusions

Scaling

Mean delay not affected by number of users.

Cost per user increases linearly with number of users in same session.

Node choice strategy

Random is not very good.

With predictable demand and little choice of nodes static performs similarly to dynamic.

Fog computing scenario

Choosing data centres from larger set gives better QoE less time on unmanaged.

Dynamic gets advantage over static with more choice.

Thank you! Questions?

richard@richardclegg.org