UNIVERSITY OF WATERLOO

What metrics do networked applications care about?

CS 456/656, Fall 2026

Mina Tahmasbi Arashloo

Outline

- **Where we were**
- Video streaming on a perfect network
- Video streaming on a real network
- What each application wants from the network
- Takeaways

Last time: we built a web browser

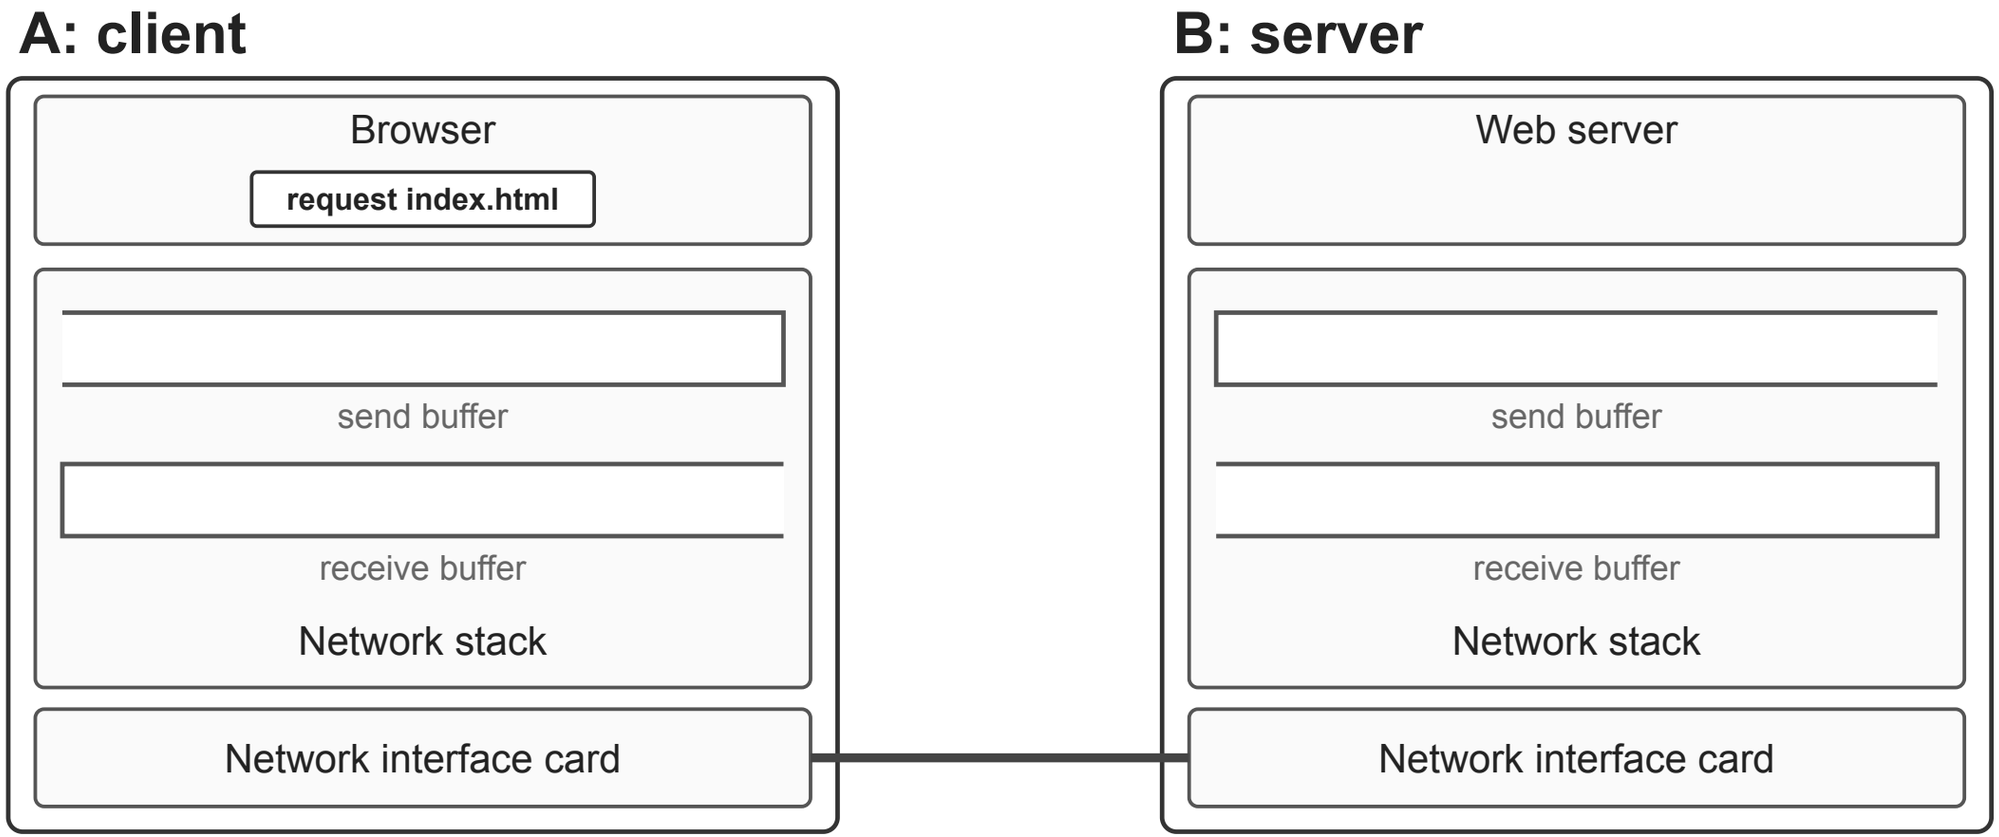

Last time: we built a web browser

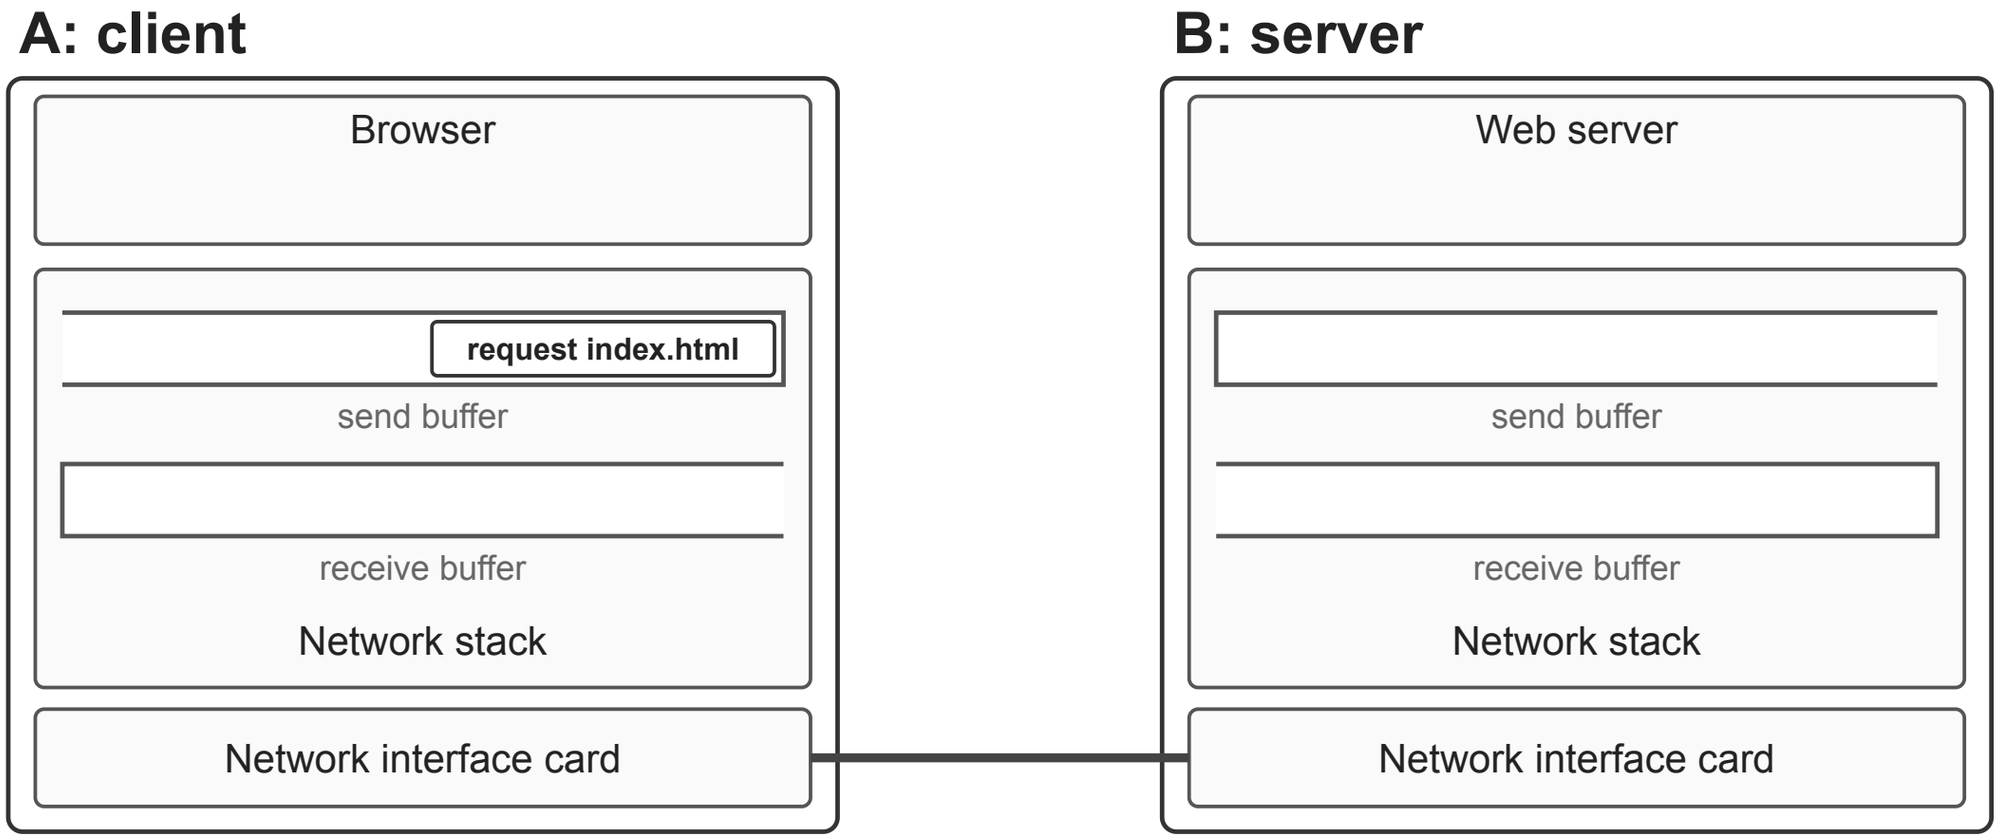

Last time: we built a web browser

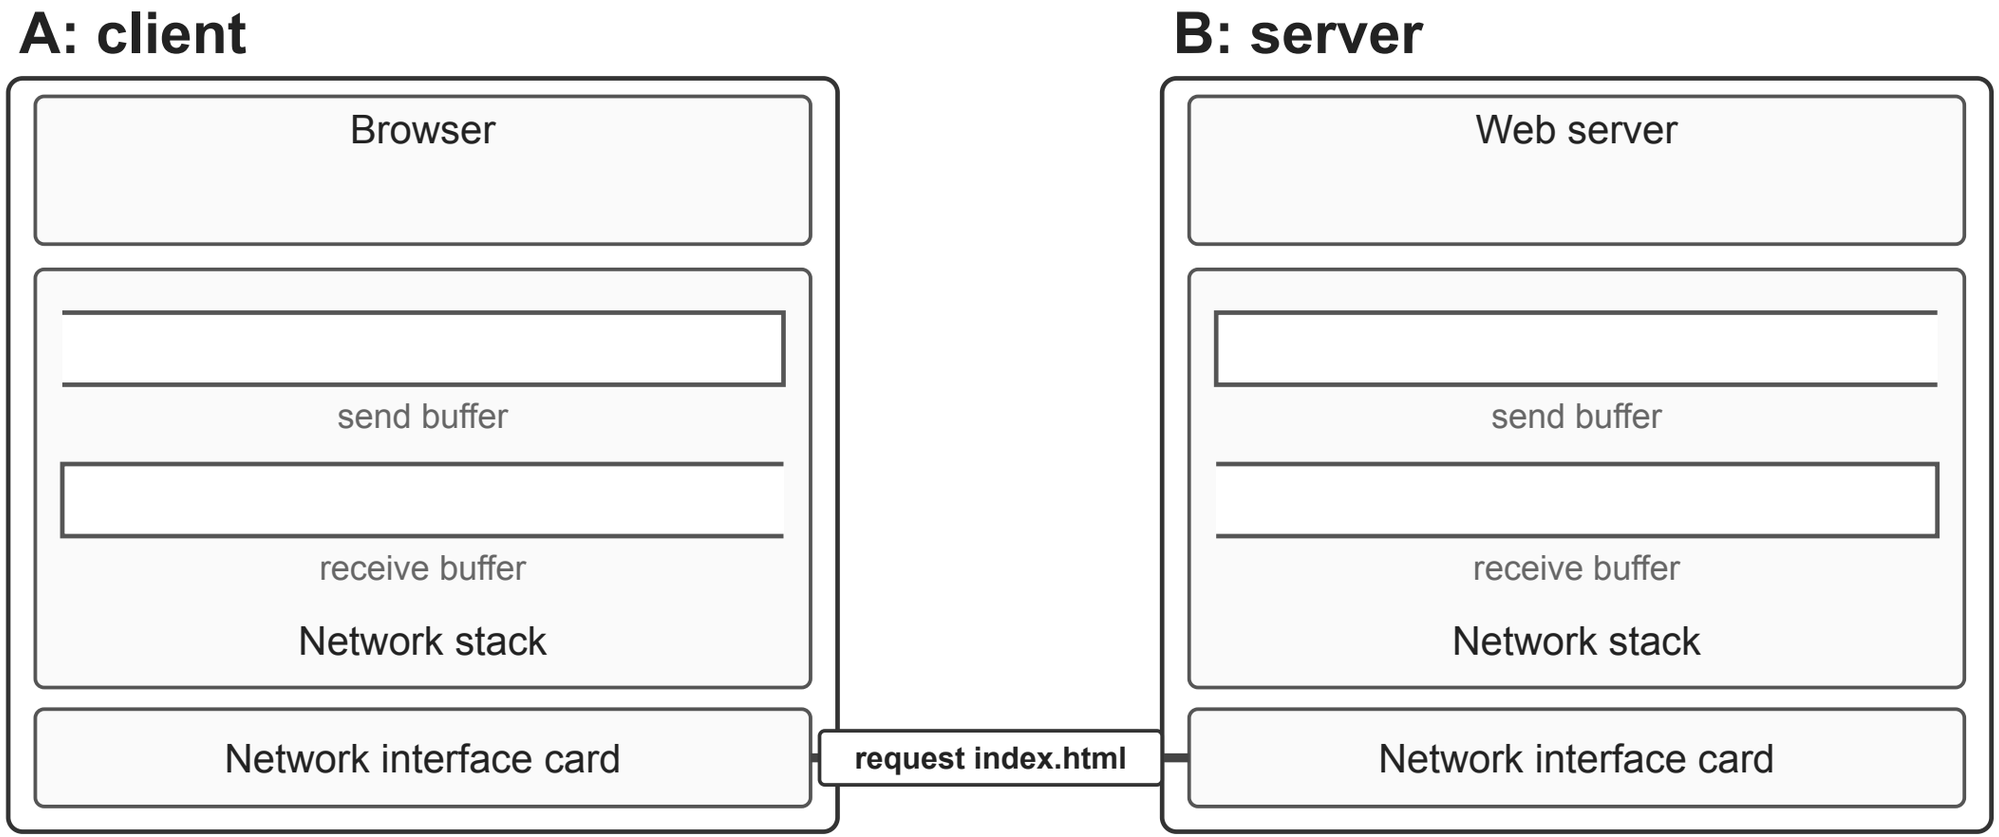

Last time: we built a web browser

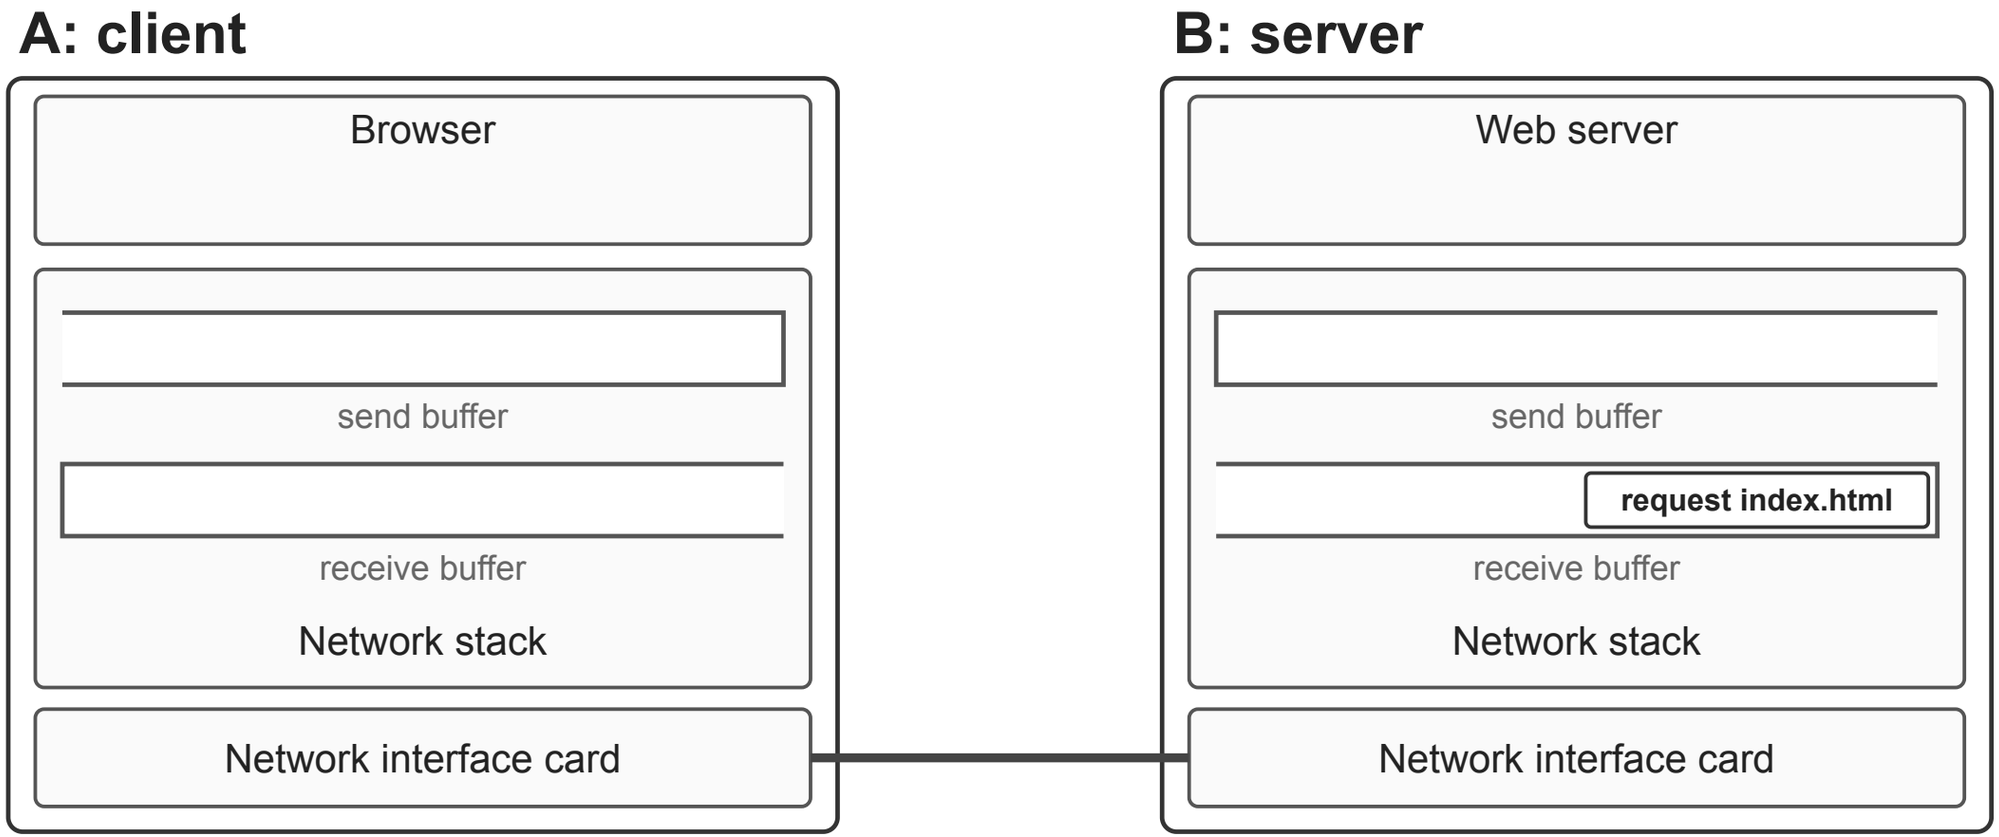

Last time: we built a web browser

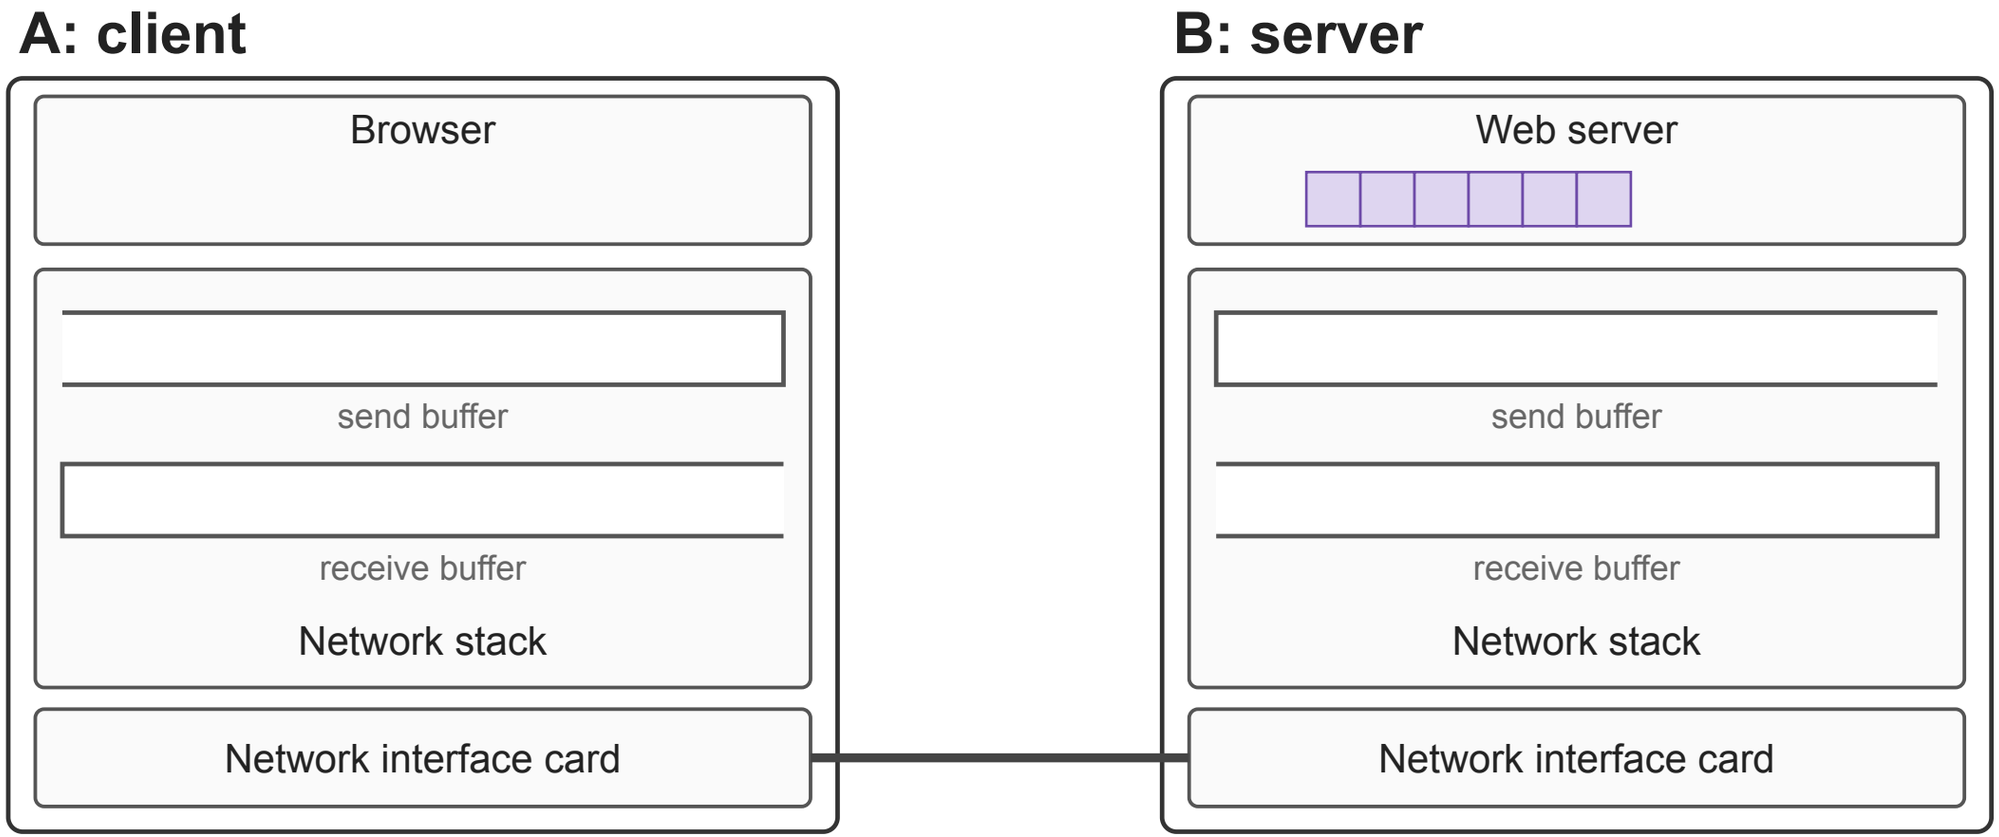

Last time: we built a web browser

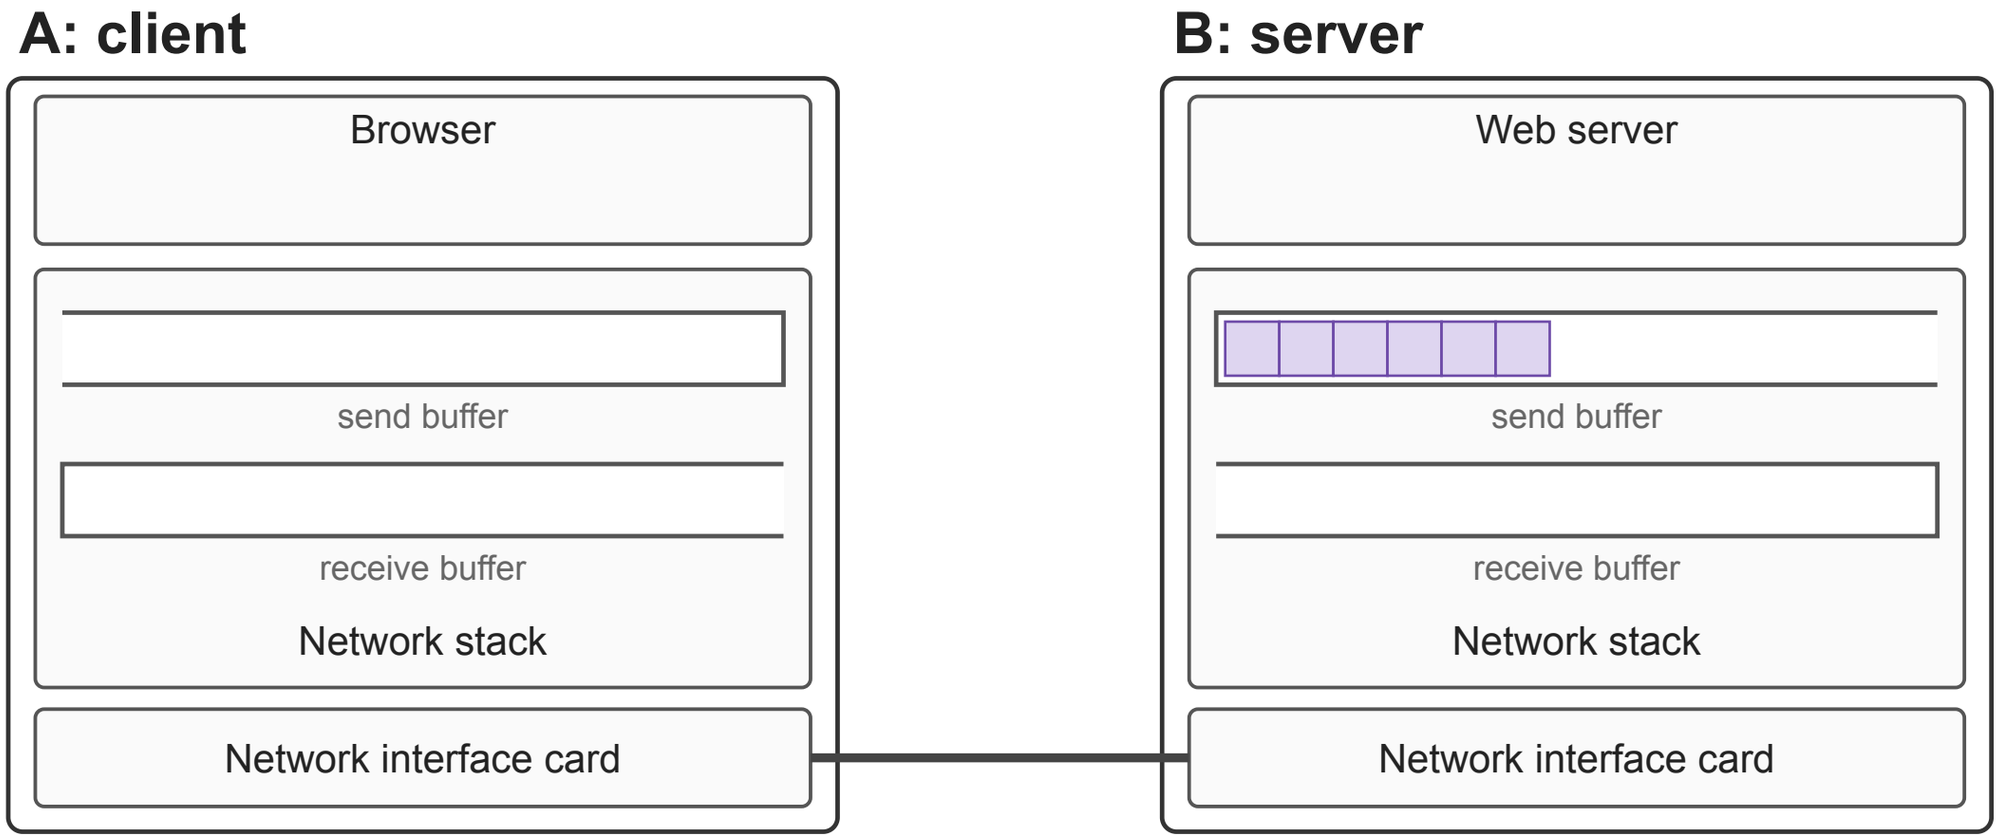

Last time: we built a web browser

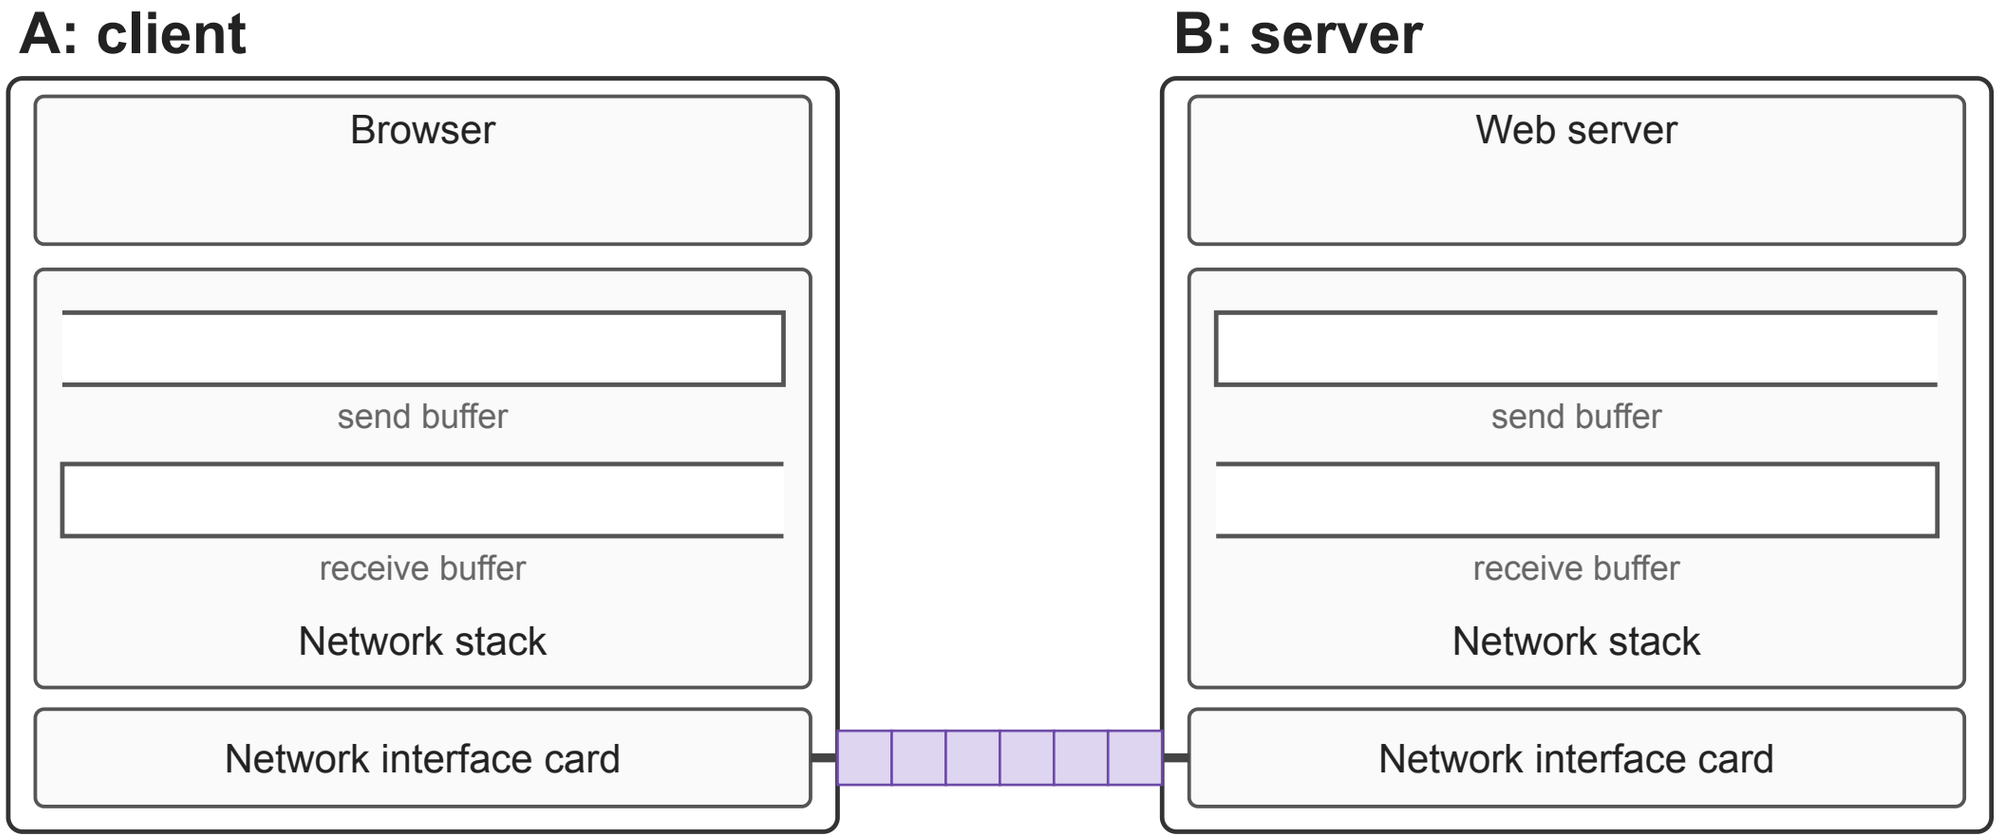

Last time: we built a web browser

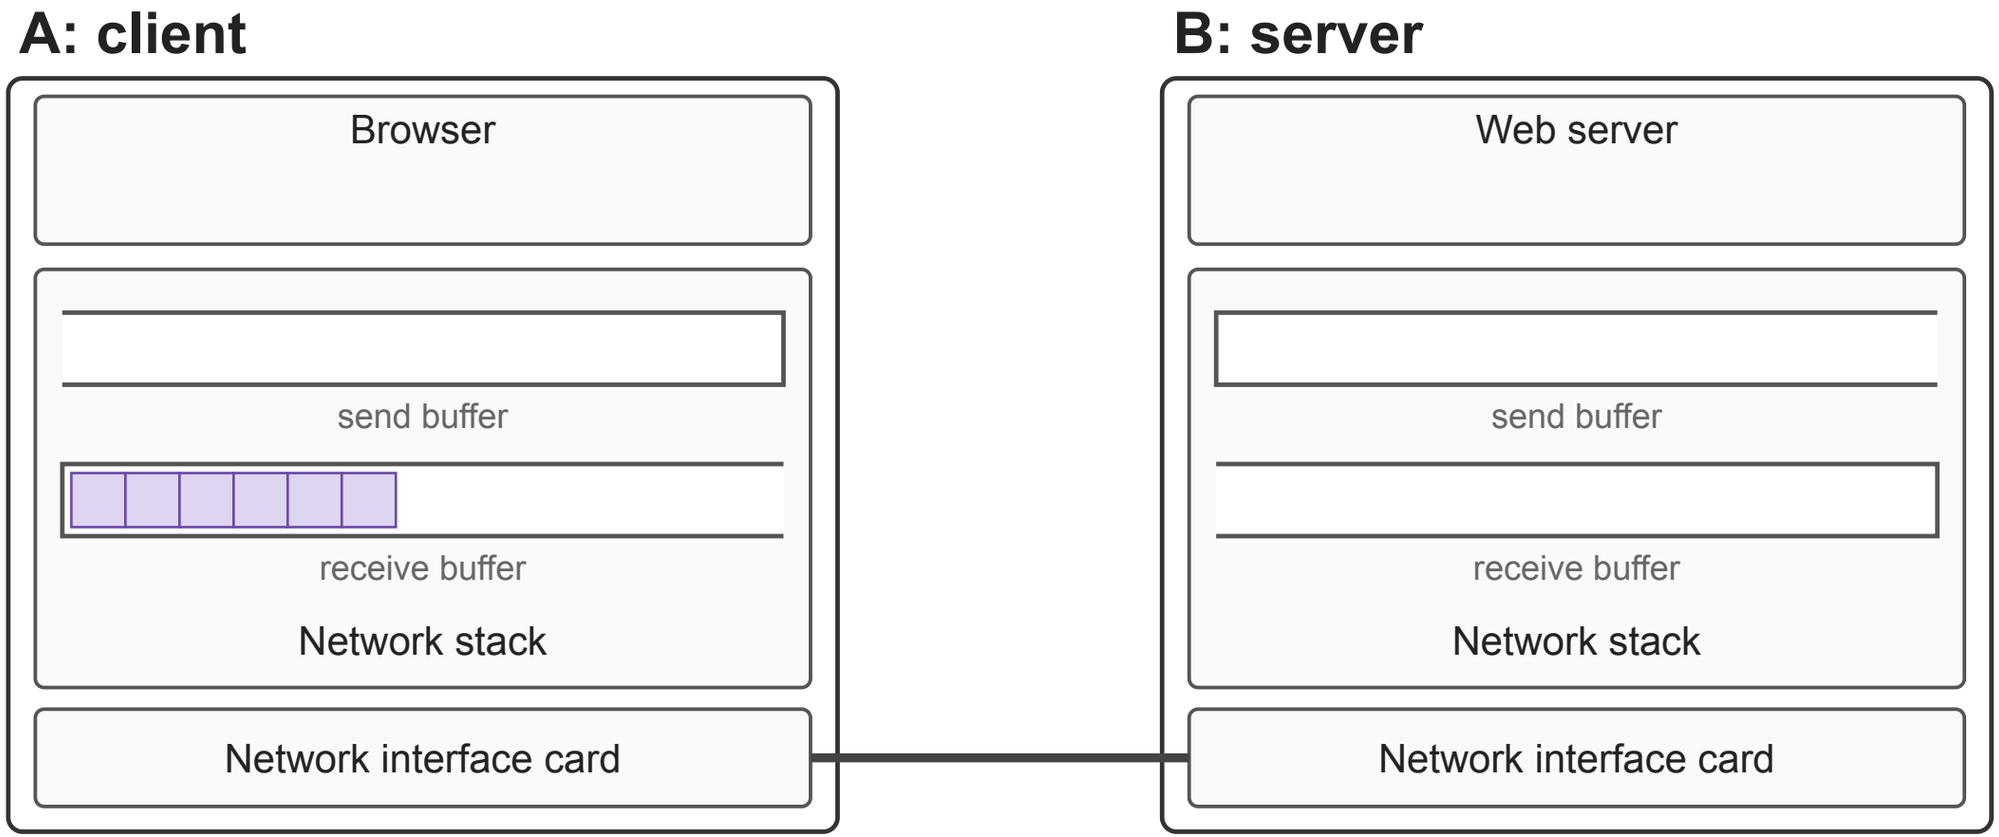

Last time: we built a web browser

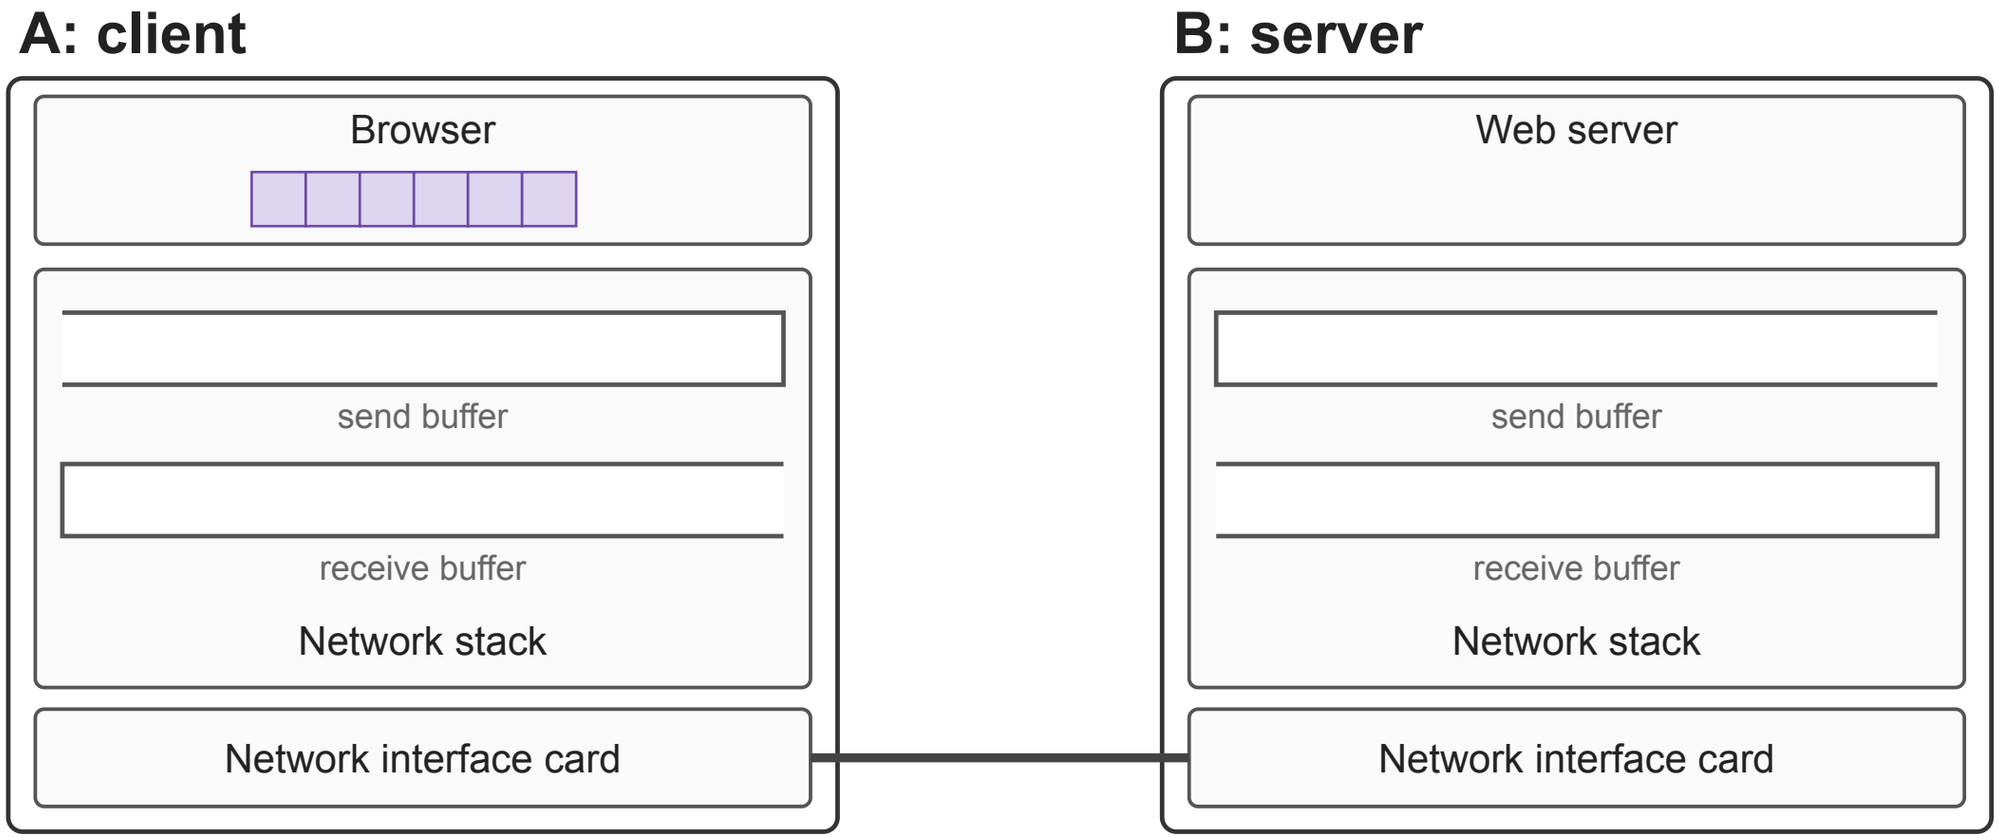

We also saw

- A page with more than one file, and head-of-line blocking
- Interleaving the files, in the application or in the network stack

Now let’s change the application

Outline

- Where we were
- **Video streaming on a perfect network**
- Video streaming on a real network
- What each application wants from the network
- Takeaways

Video streaming

- A video is also a file on a server that the client can ask for
- But nobody wants to wait for the whole file to download before playing the video
  - hence the term streaming
- Only need to show 30 frames every second
  - not the whole file
- The video is cut into chunks of a few seconds
  - the client asks for the next chunk when it needs it

What is one second of video?

30 frames, shown one every 1/30 s: that is 1 s of video

What is one second of video?

30 frames, shown one every 1/30 s: that is 1 s of video

the same 30 frames can arrive much faster, e.g., in 0.2 s

What is one second of video?

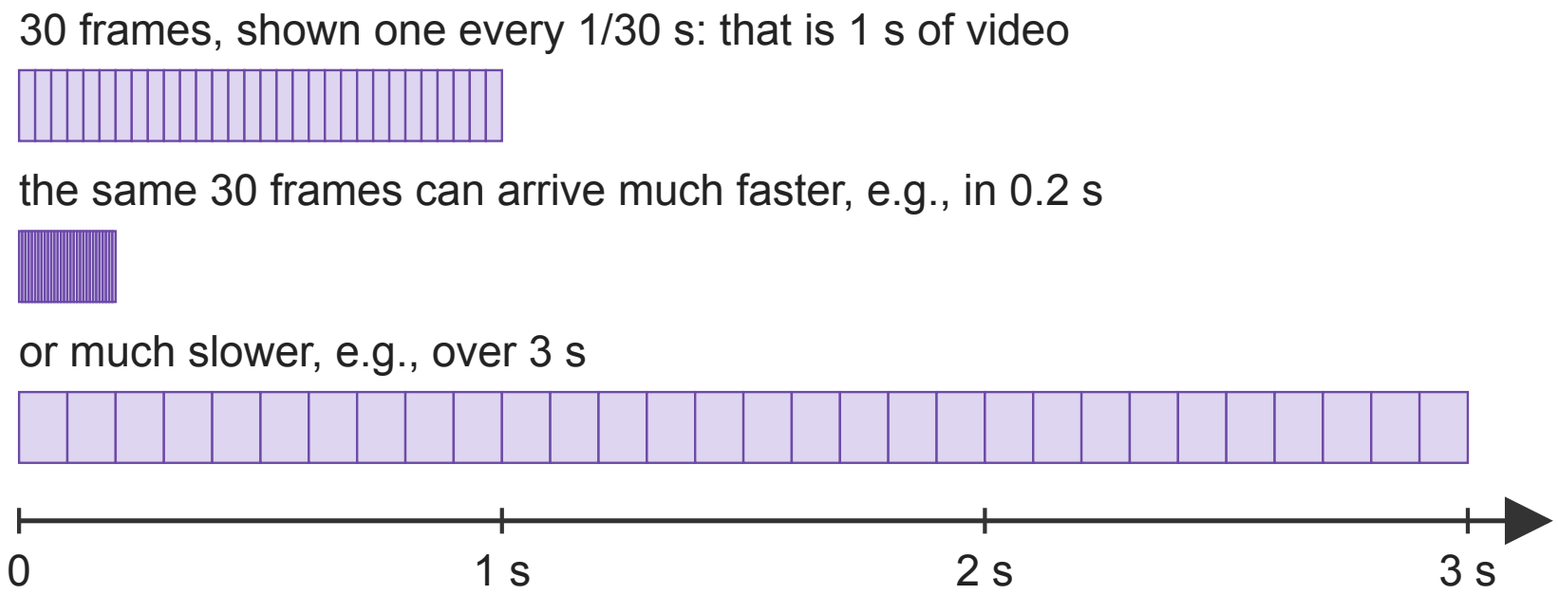

What is one second of video?

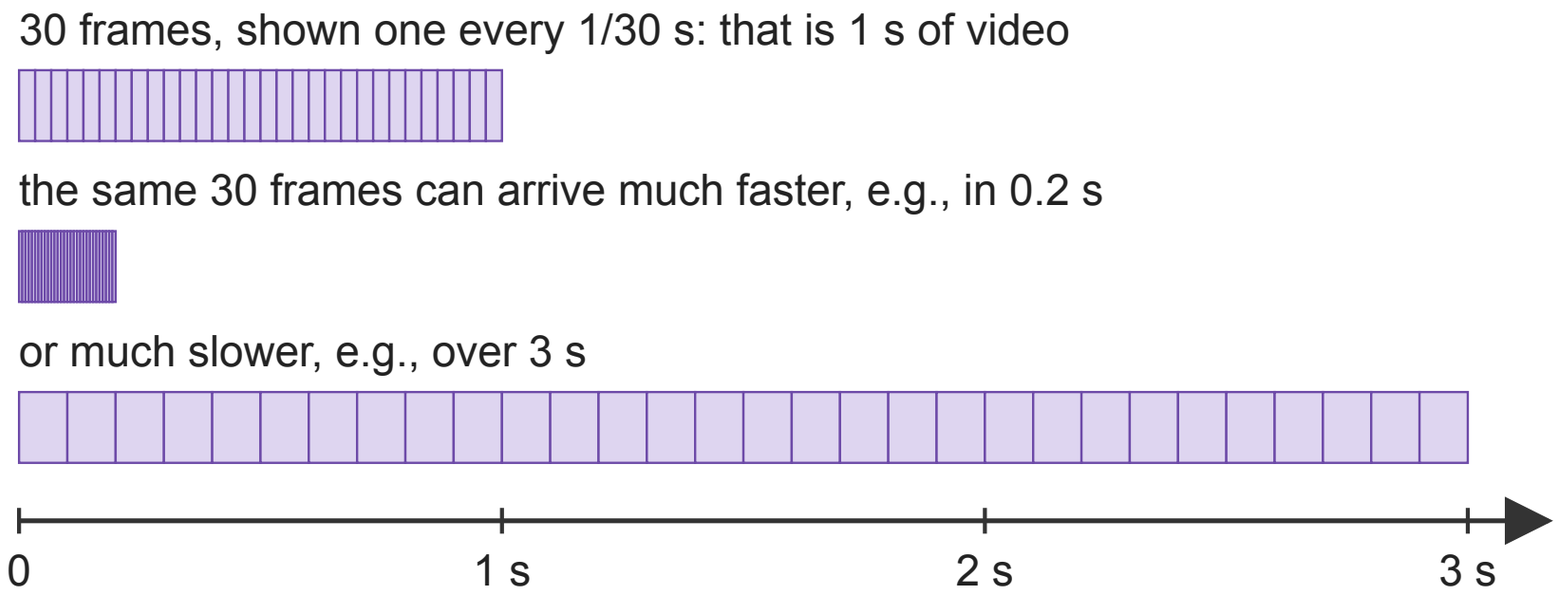

Seconds of video: how long it plays, not how long it takes to arrive

On a perfect network

these graphs simplify what really happens

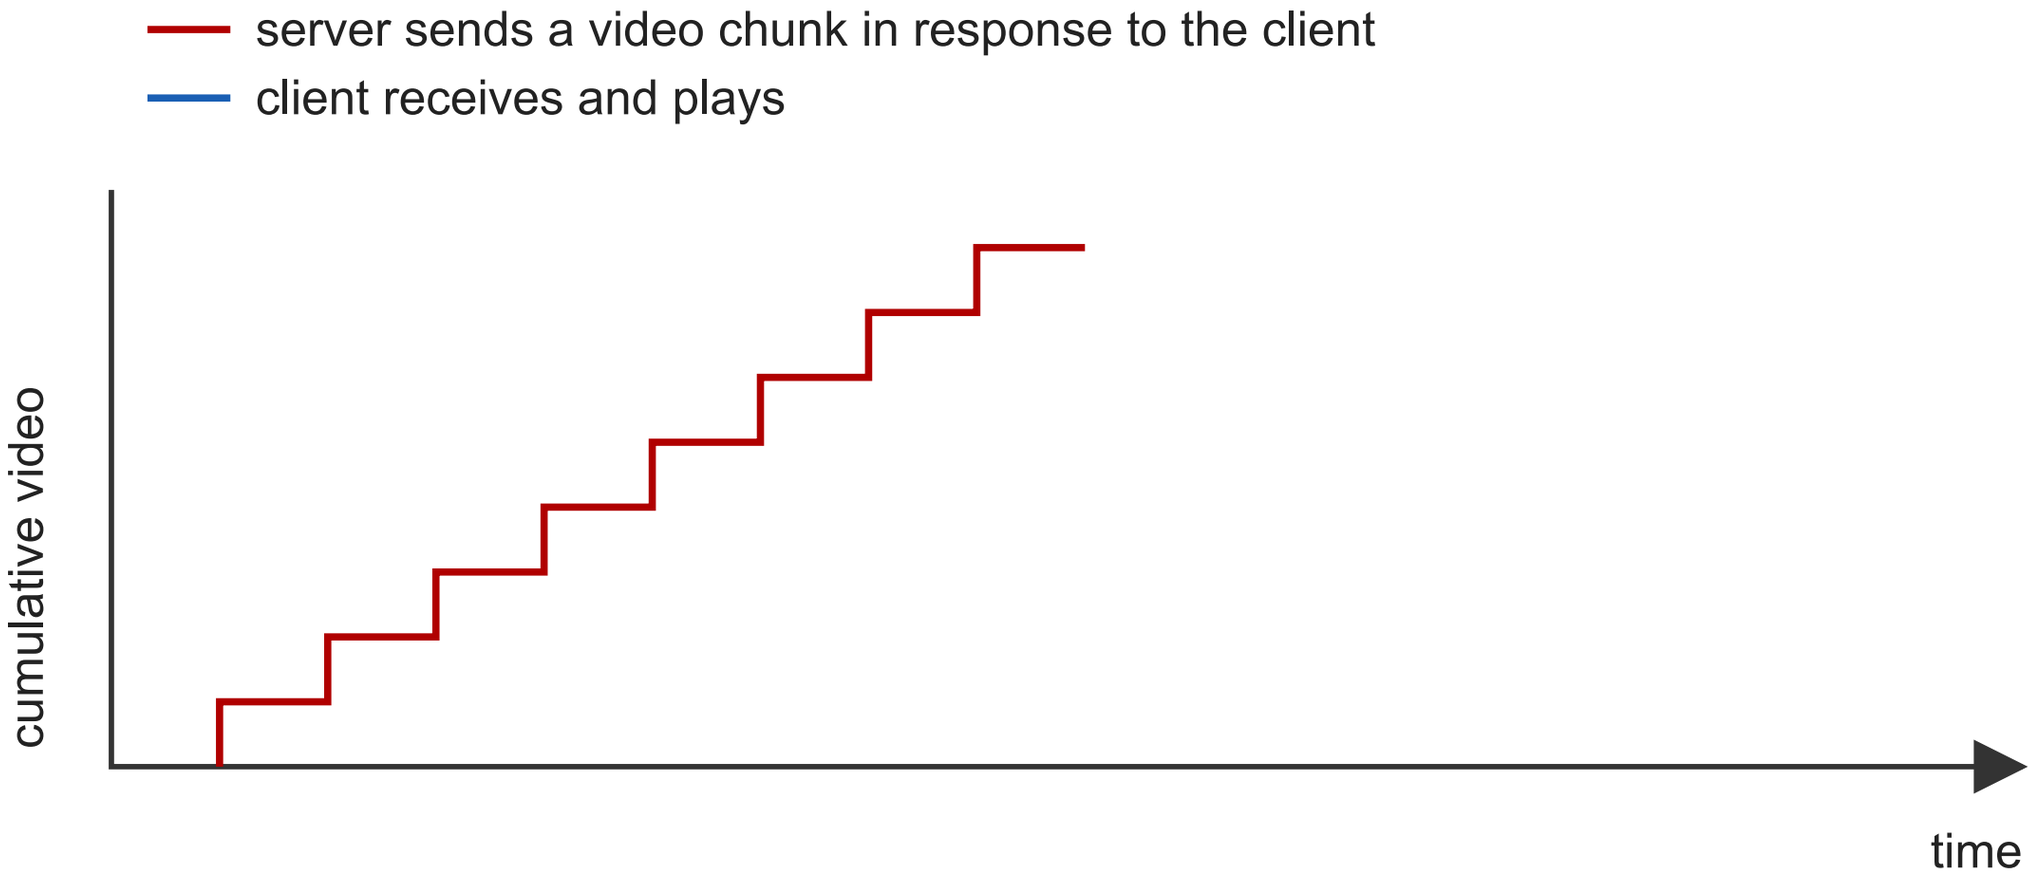

On a perfect network

these graphs simplify what really happens

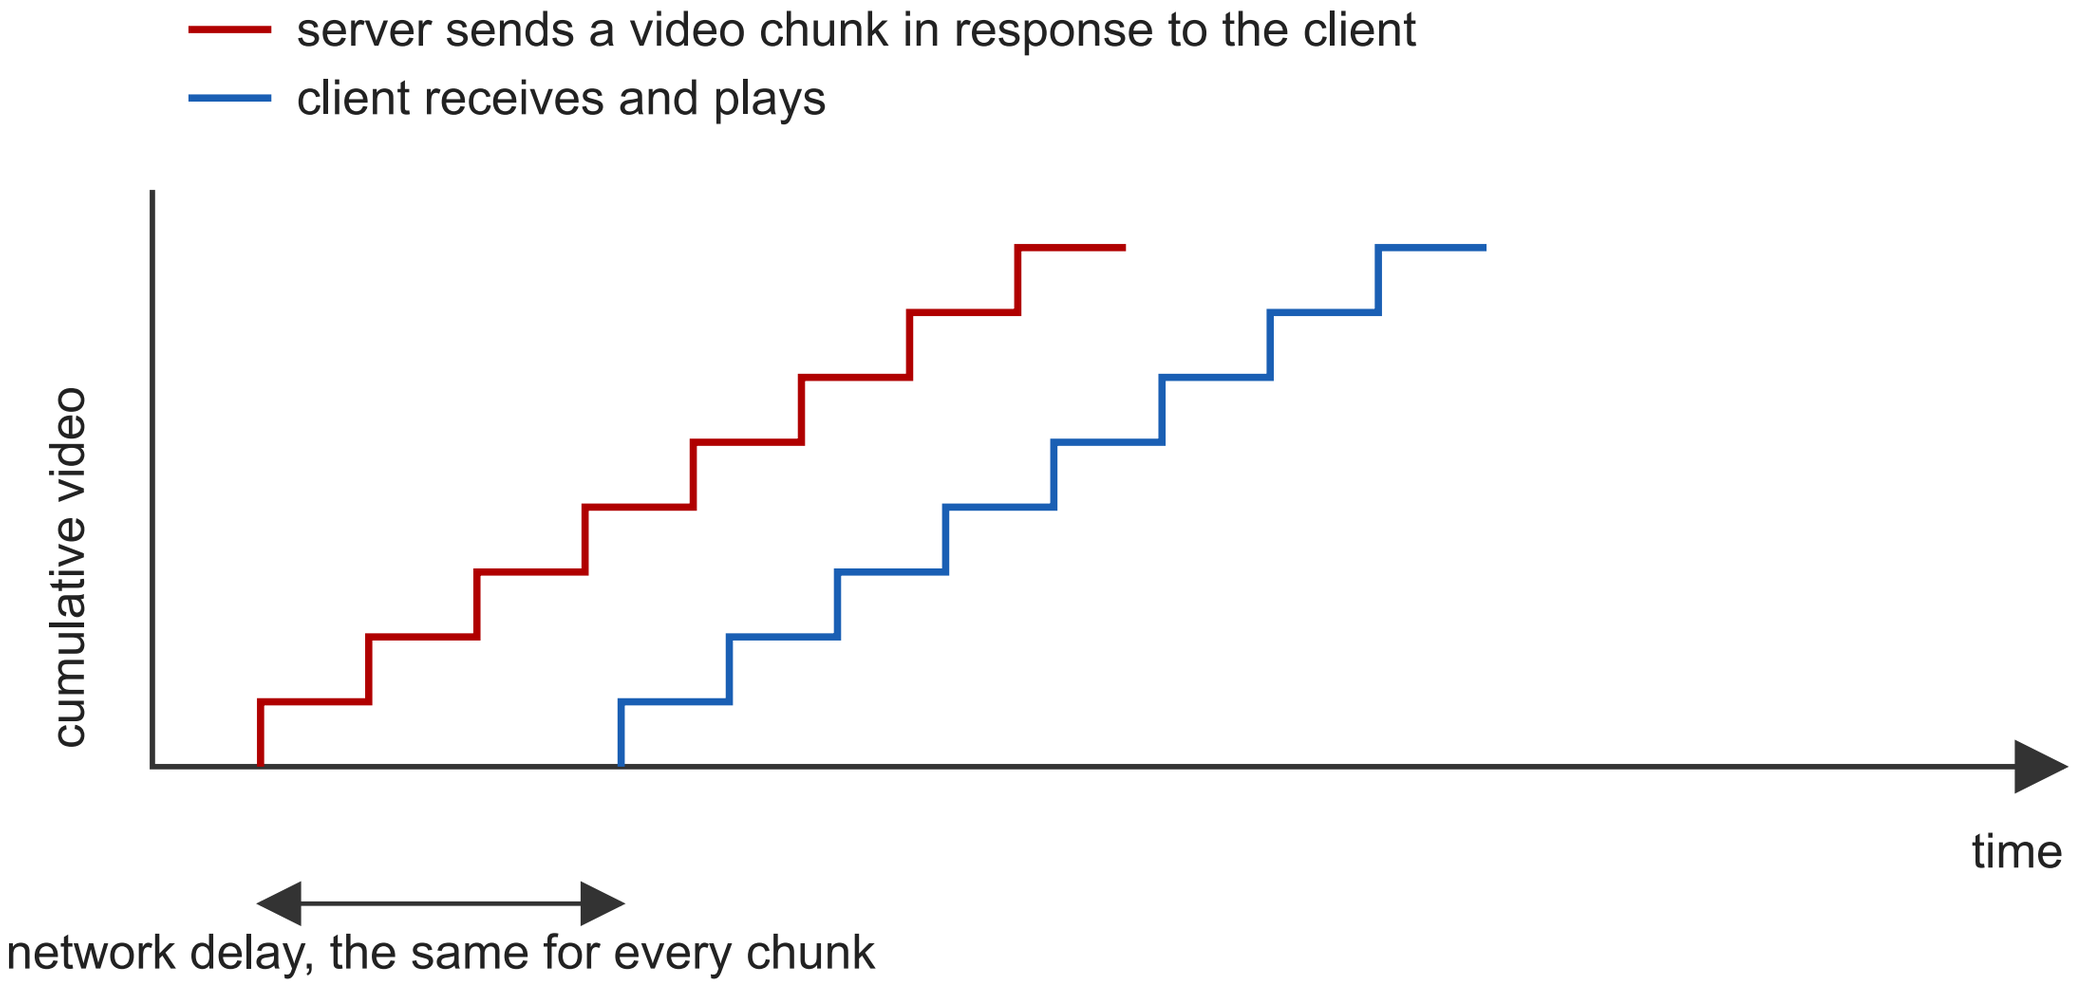

On a perfect network

these graphs simplify what really happens

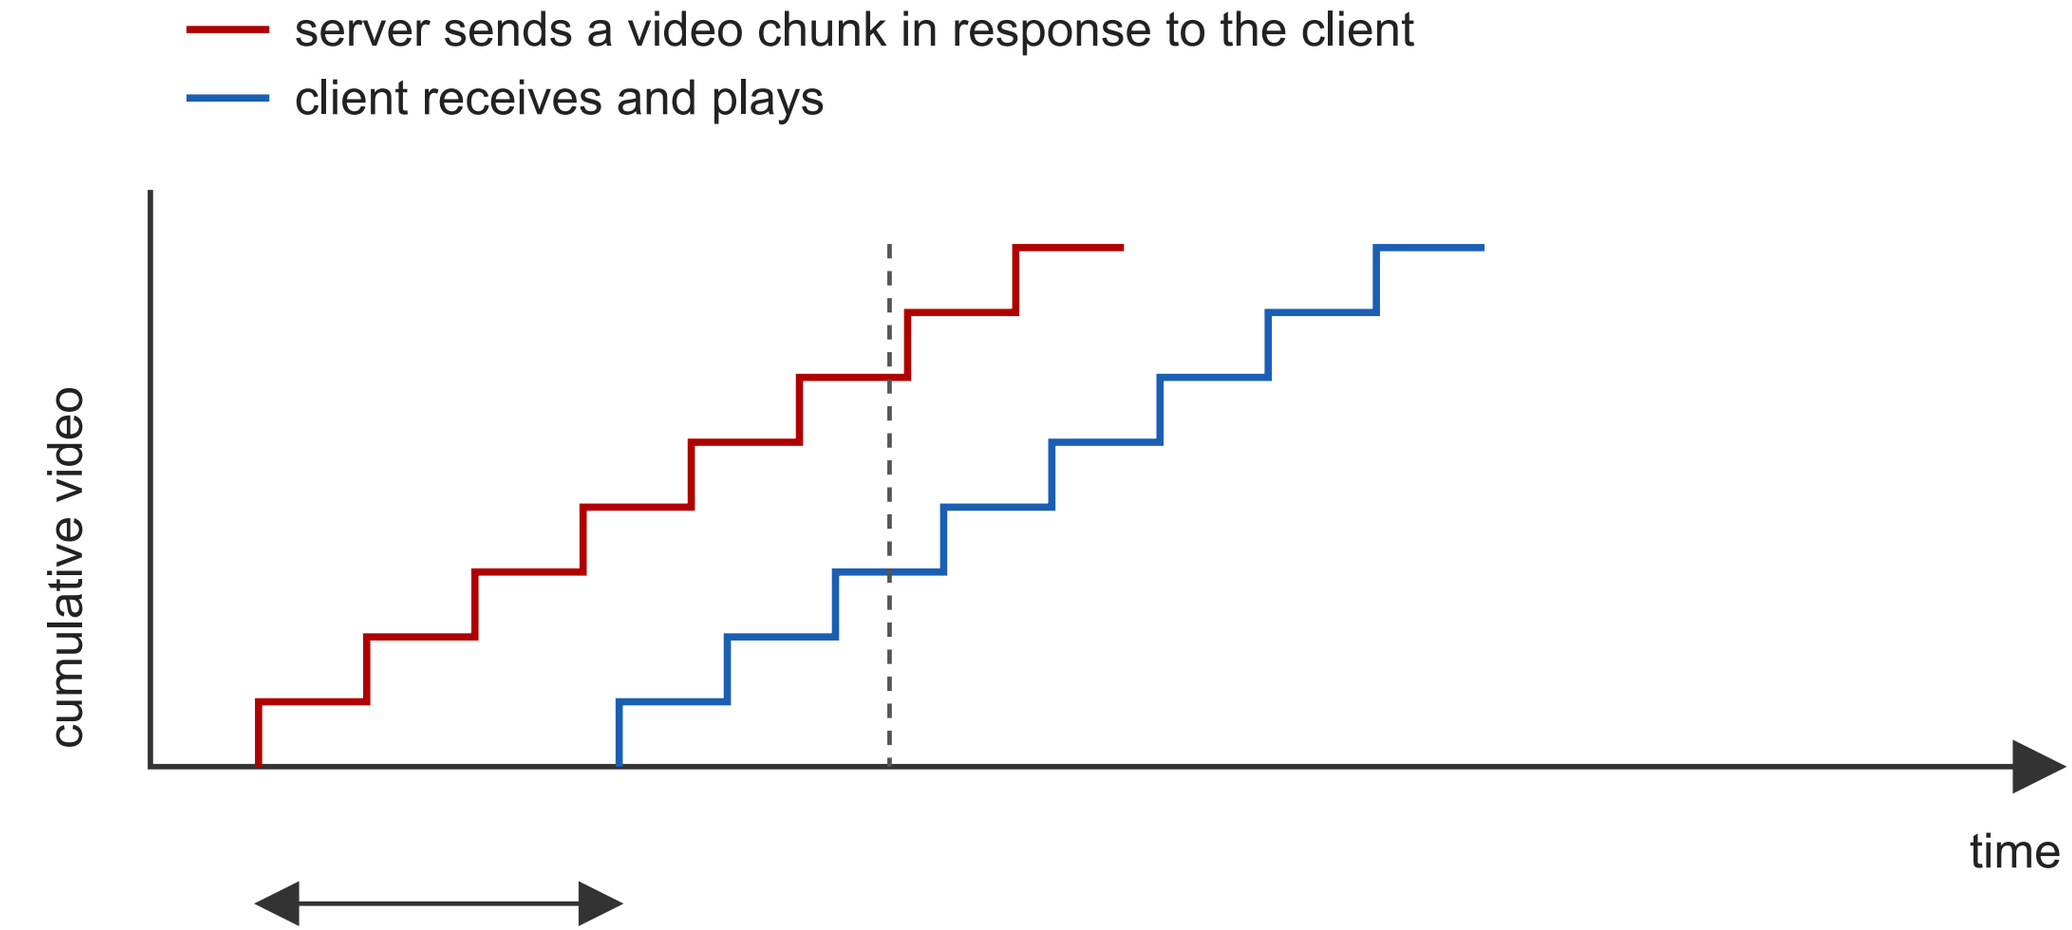

right now: the client plays early chunks
while it is still asking for later ones

How would you build the client?

Outline

- Where we were
- Video streaming on a perfect network
- **Video streaming on a real network**
- What each application wants from the network
- Takeaways

A real network is not a perfect pipe

- The delay may change from one video chunk to the next

A real network is not a perfect pipe

- The delay may change from one video chunk to the next

How much the delay changes is called jitter

Why does the delay change?

$$\text{delay} = 2 \times (d_{proc,A} + d_{prop} + d_{proc,B}) + L_{request} / R + L_{file} / R$$

Why does the delay change?

$$\text{delay} = 2 \times (d_{proc,A} + d_{prop} + d_{proc,B}) + L_{request} / R + L_{file} / R$$

delay = processing + propagation + transmission

Why does the delay change?

$$\text{delay} = 2 \times (d_{proc,A} + d_{prop} + d_{proc,B}) + L_{request} / R + L_{file} / R$$

delay = processing + propagation + transmission + waiting

Waiting delay: the time data waits in line because other data got there first and is still going into the pipe. We come back to this in Part 2.

Why does the delay change?

$$\text{delay} = 2 \times (d_{proc,A} + d_{prop} + d_{proc,B}) + L_{request} / R + L_{file} / R$$

delay = processing + propagation + transmission + waiting

Waiting delay: the time data waits in line because other data got there first and is still going into the pipe. We come back to this in Part 2.

Each of these can change, sometimes temporarily or for longer stretches of time throughout the streaming session

A chunk arrives late

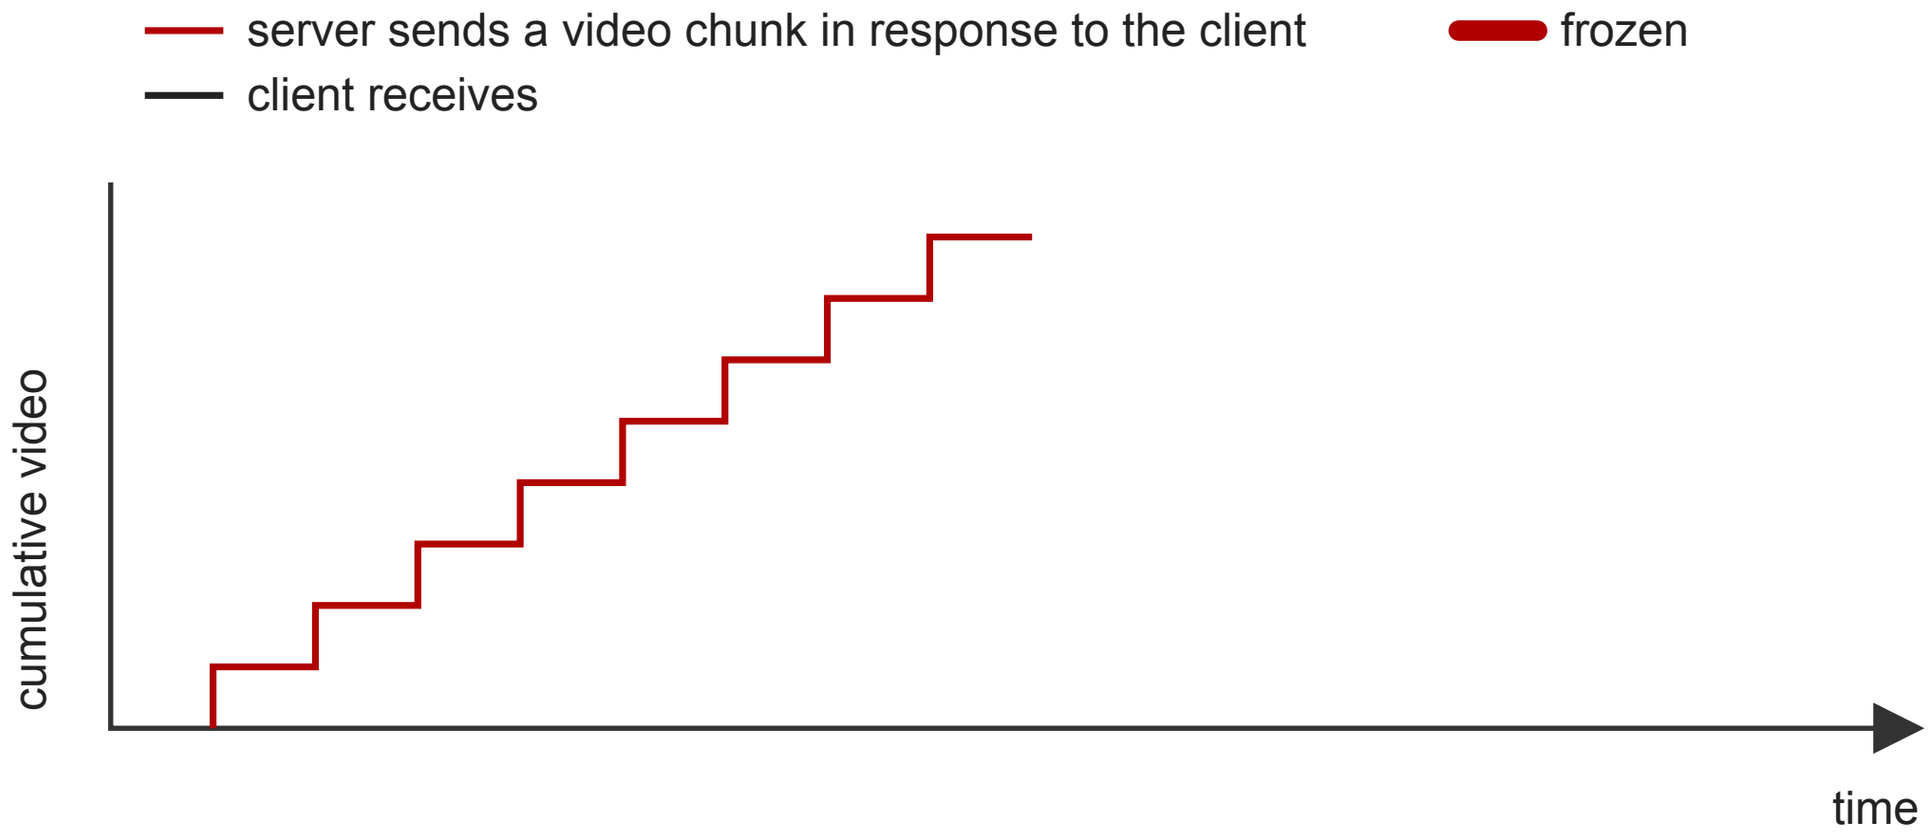

A chunk arrives late

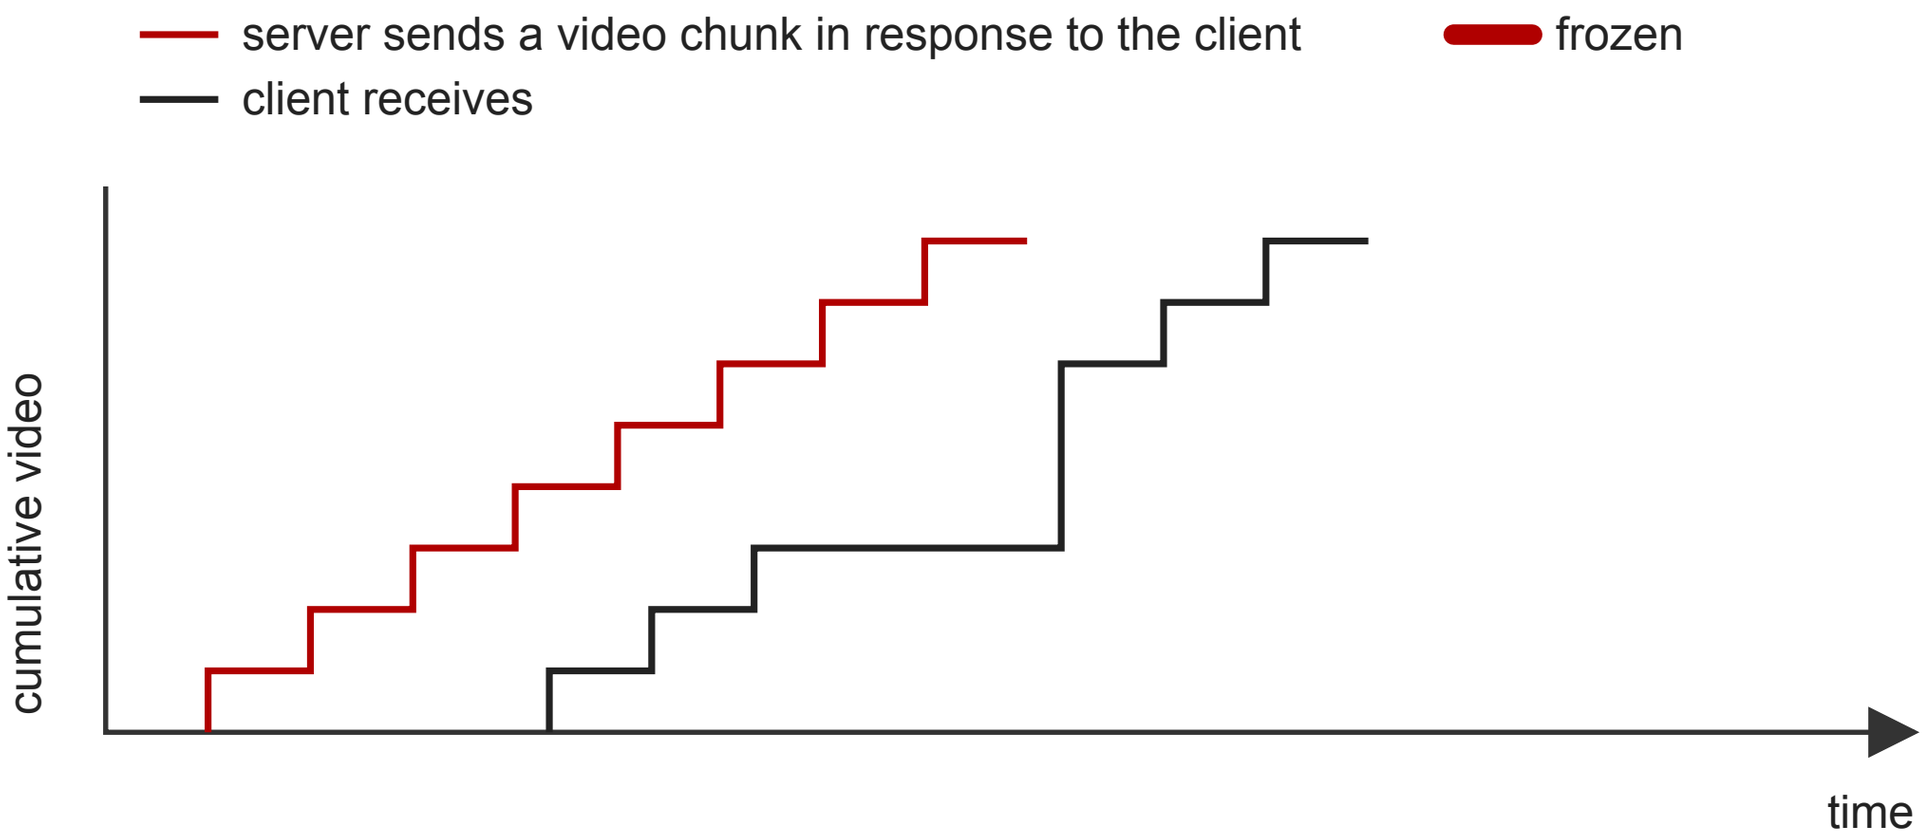

for a short while the delay goes up, then it goes back to normal

A chunk arrives late

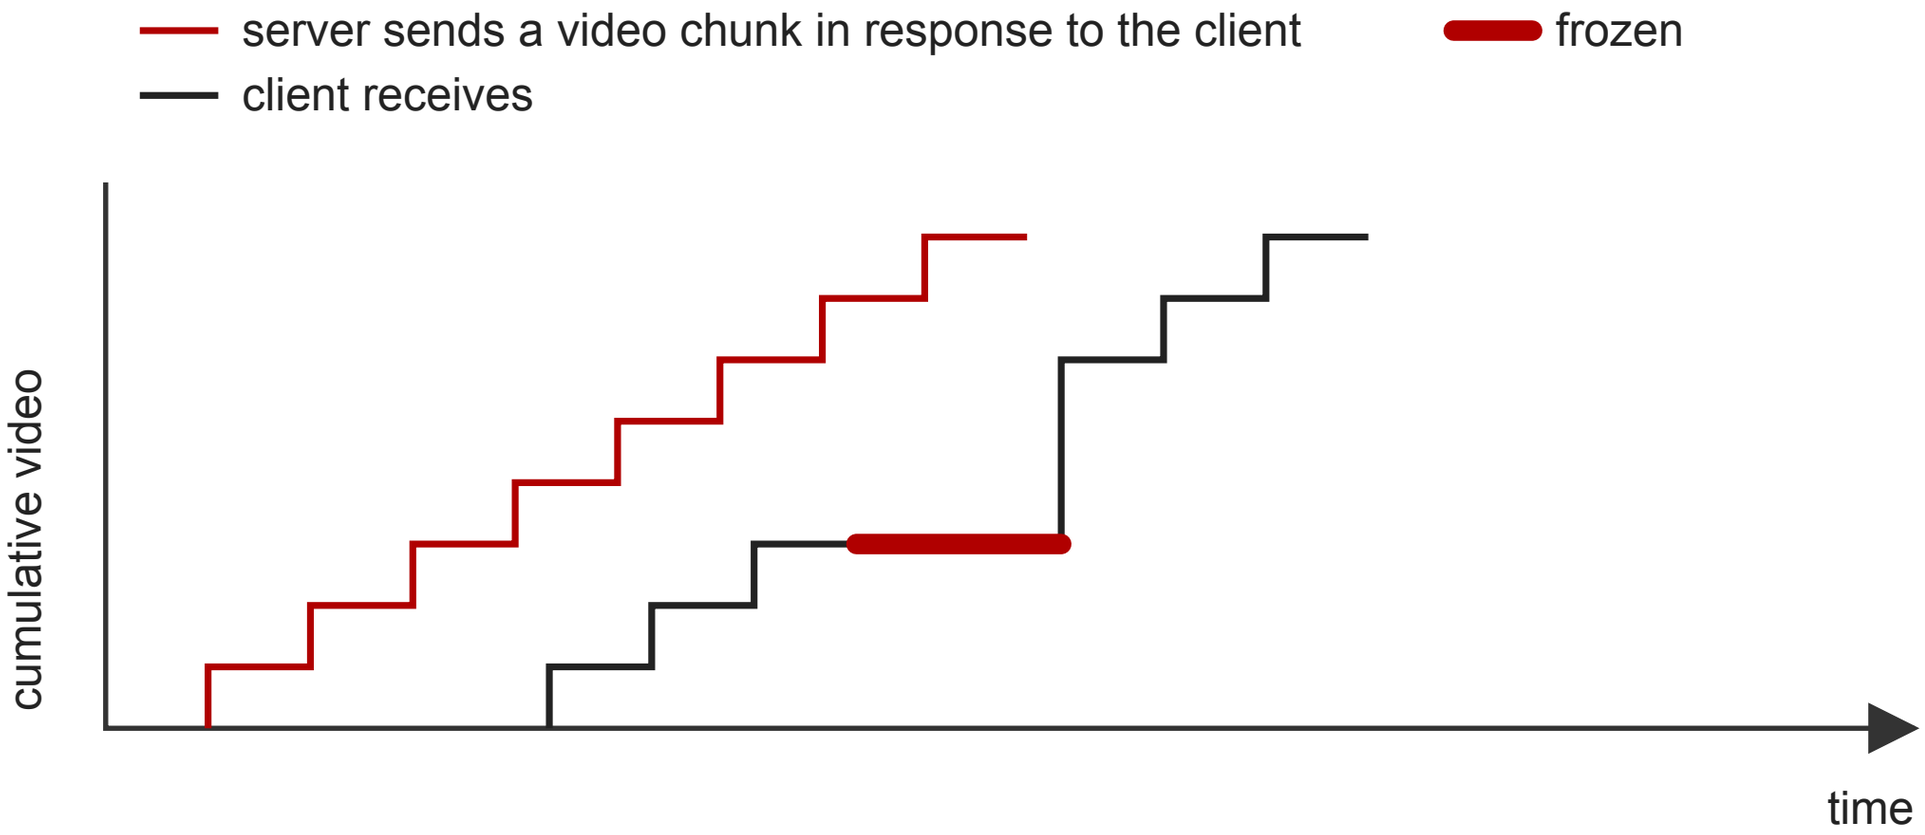

for a short while the delay goes up, then it goes back to normal

if the client starts playing the moment the chunks arrive,

when chunks are delayed, the video freezes until the late chunks show up

How would you change the client now?

What if the change lasts a while?

- So far, the delay went up and down only for a few chunks, so the playback buffer could smooth it out
- Sometimes it changes and stays that way for a while:
  - you walk from Wi-Fi to cellular
  - more people start using the network
  - your traffic starts taking a different path, longer or shorter, with a different rate

What if the change lasts a while?

- So far, the delay went up and down only for a few chunks, so the playback buffer could smooth it out
- Sometimes it changes and stays that way for a while:
  - you walk from Wi-Fi to cellular
  - more people start using the network
  - your traffic starts taking a different path, longer or shorter, with a different rate

Which factors that impact the delay can each of these change?

If the pipe gets longer for a while

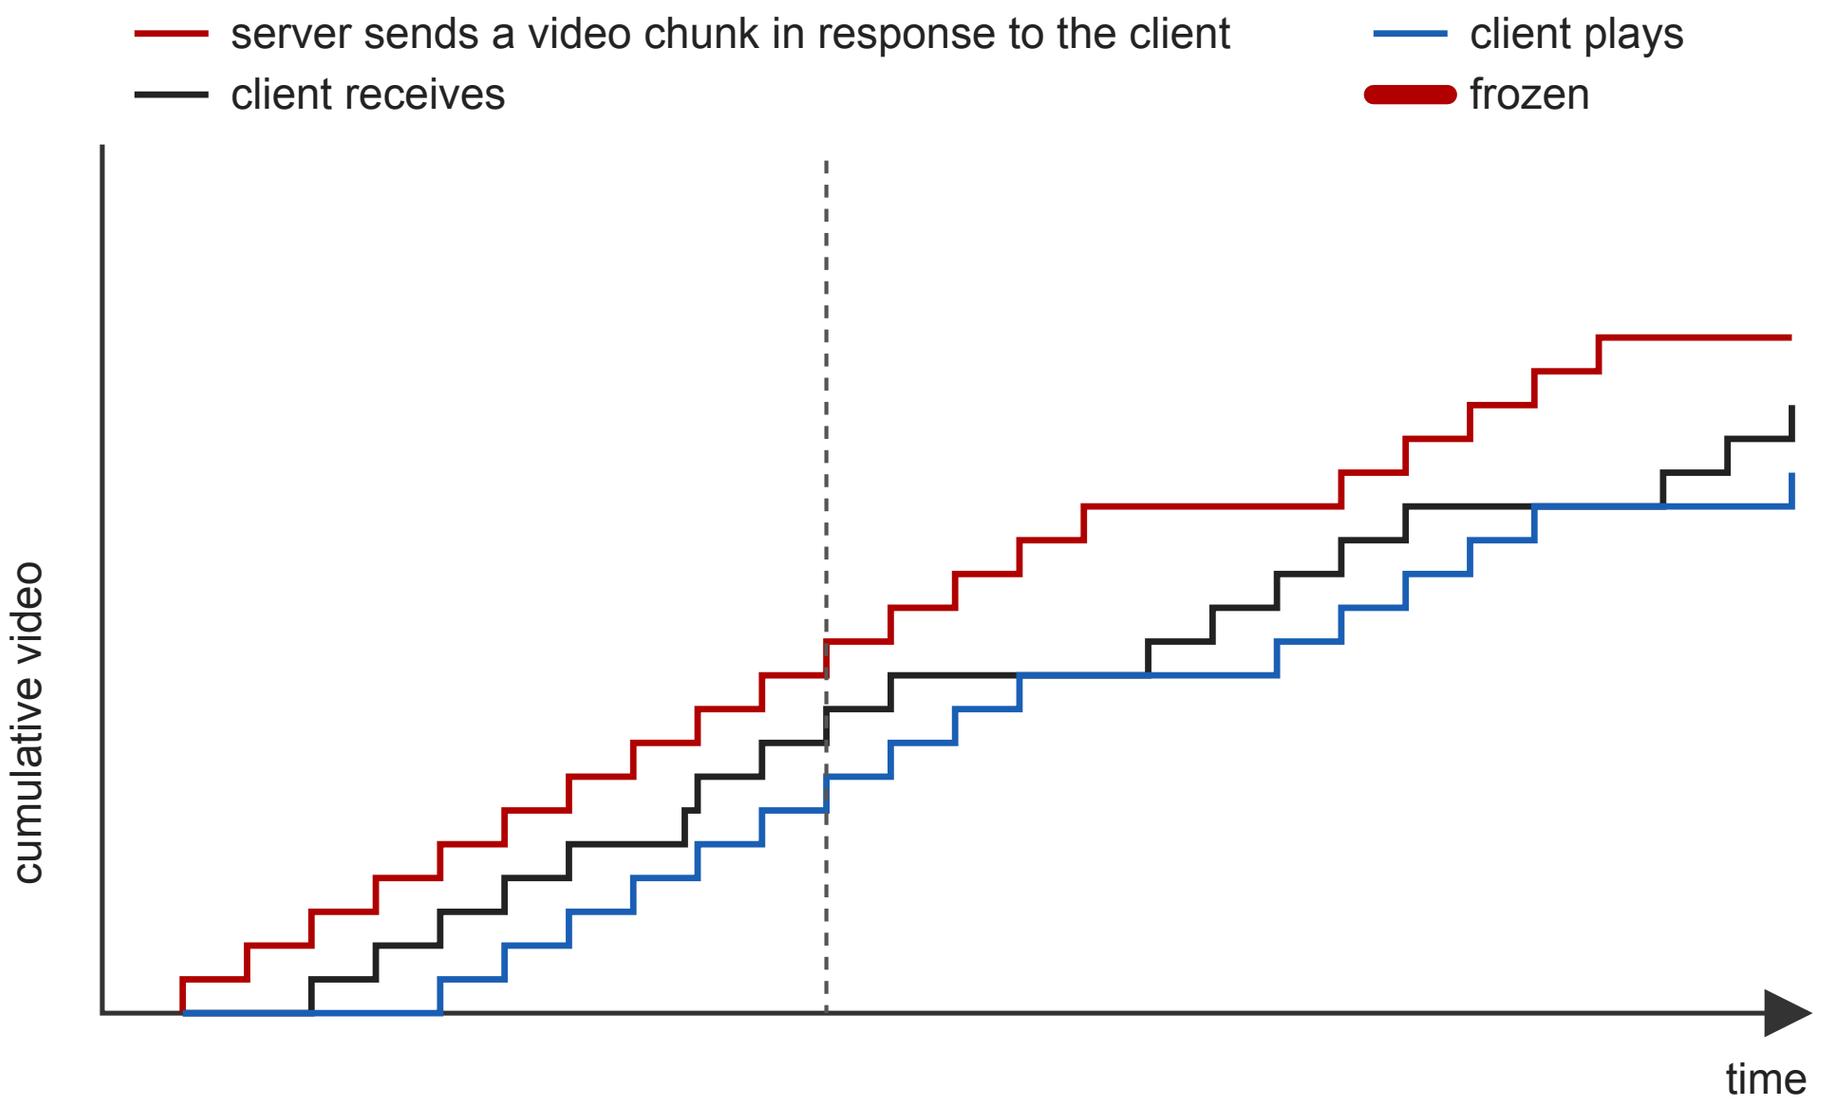

If the pipe gets narrower for a while

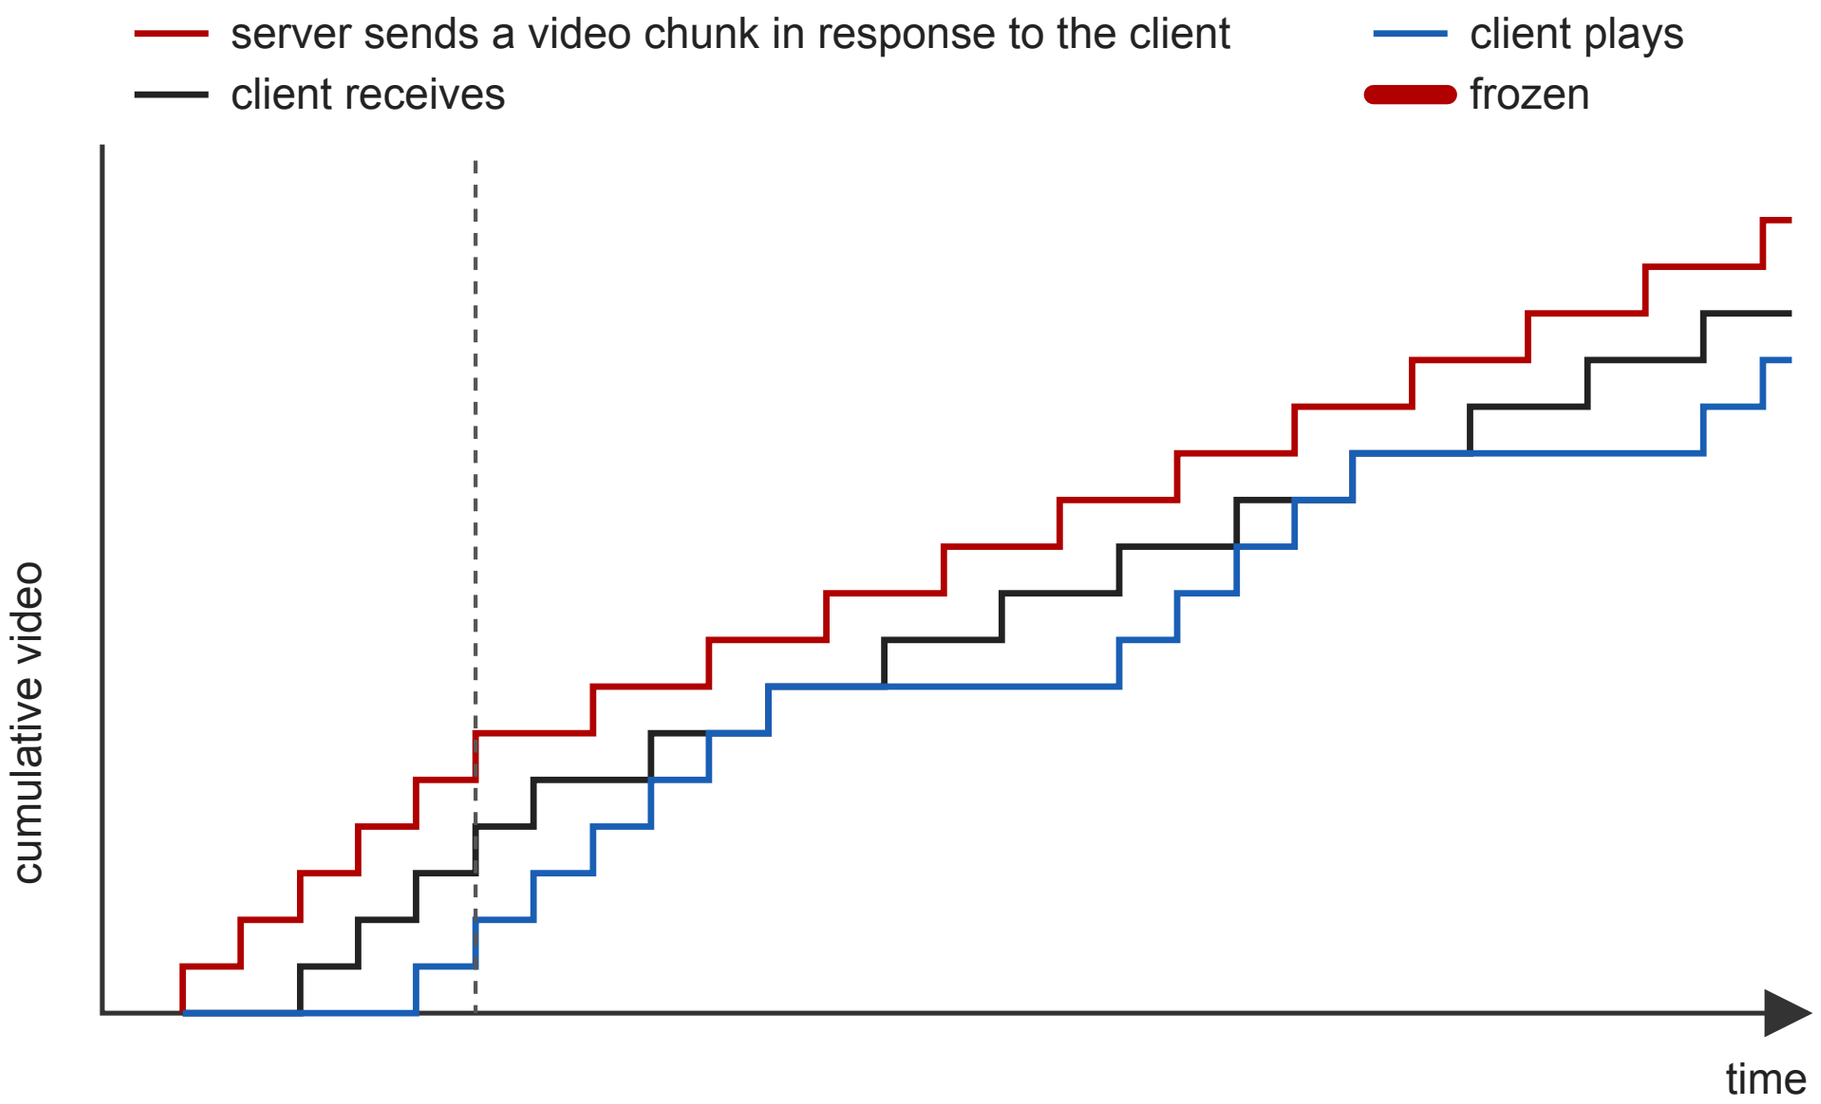

What if this was a video call instead of on-demand video streaming?

- Buffering is not feasible (no buffer or a very small one)
- If there is transient data loss or delay, we can just move on, no need to recover
- More reactive and conservative in reducing the quality to keep up with real-time
- Sacrifice video quality or just don't have the video keep up but have the audio keep up.

Reset

Outline

- Where we were
- Video streaming on a perfect network
- Video streaming on a real network
- **What each application wants from the network**
- Takeaways

What metrics and properties matter to an application?

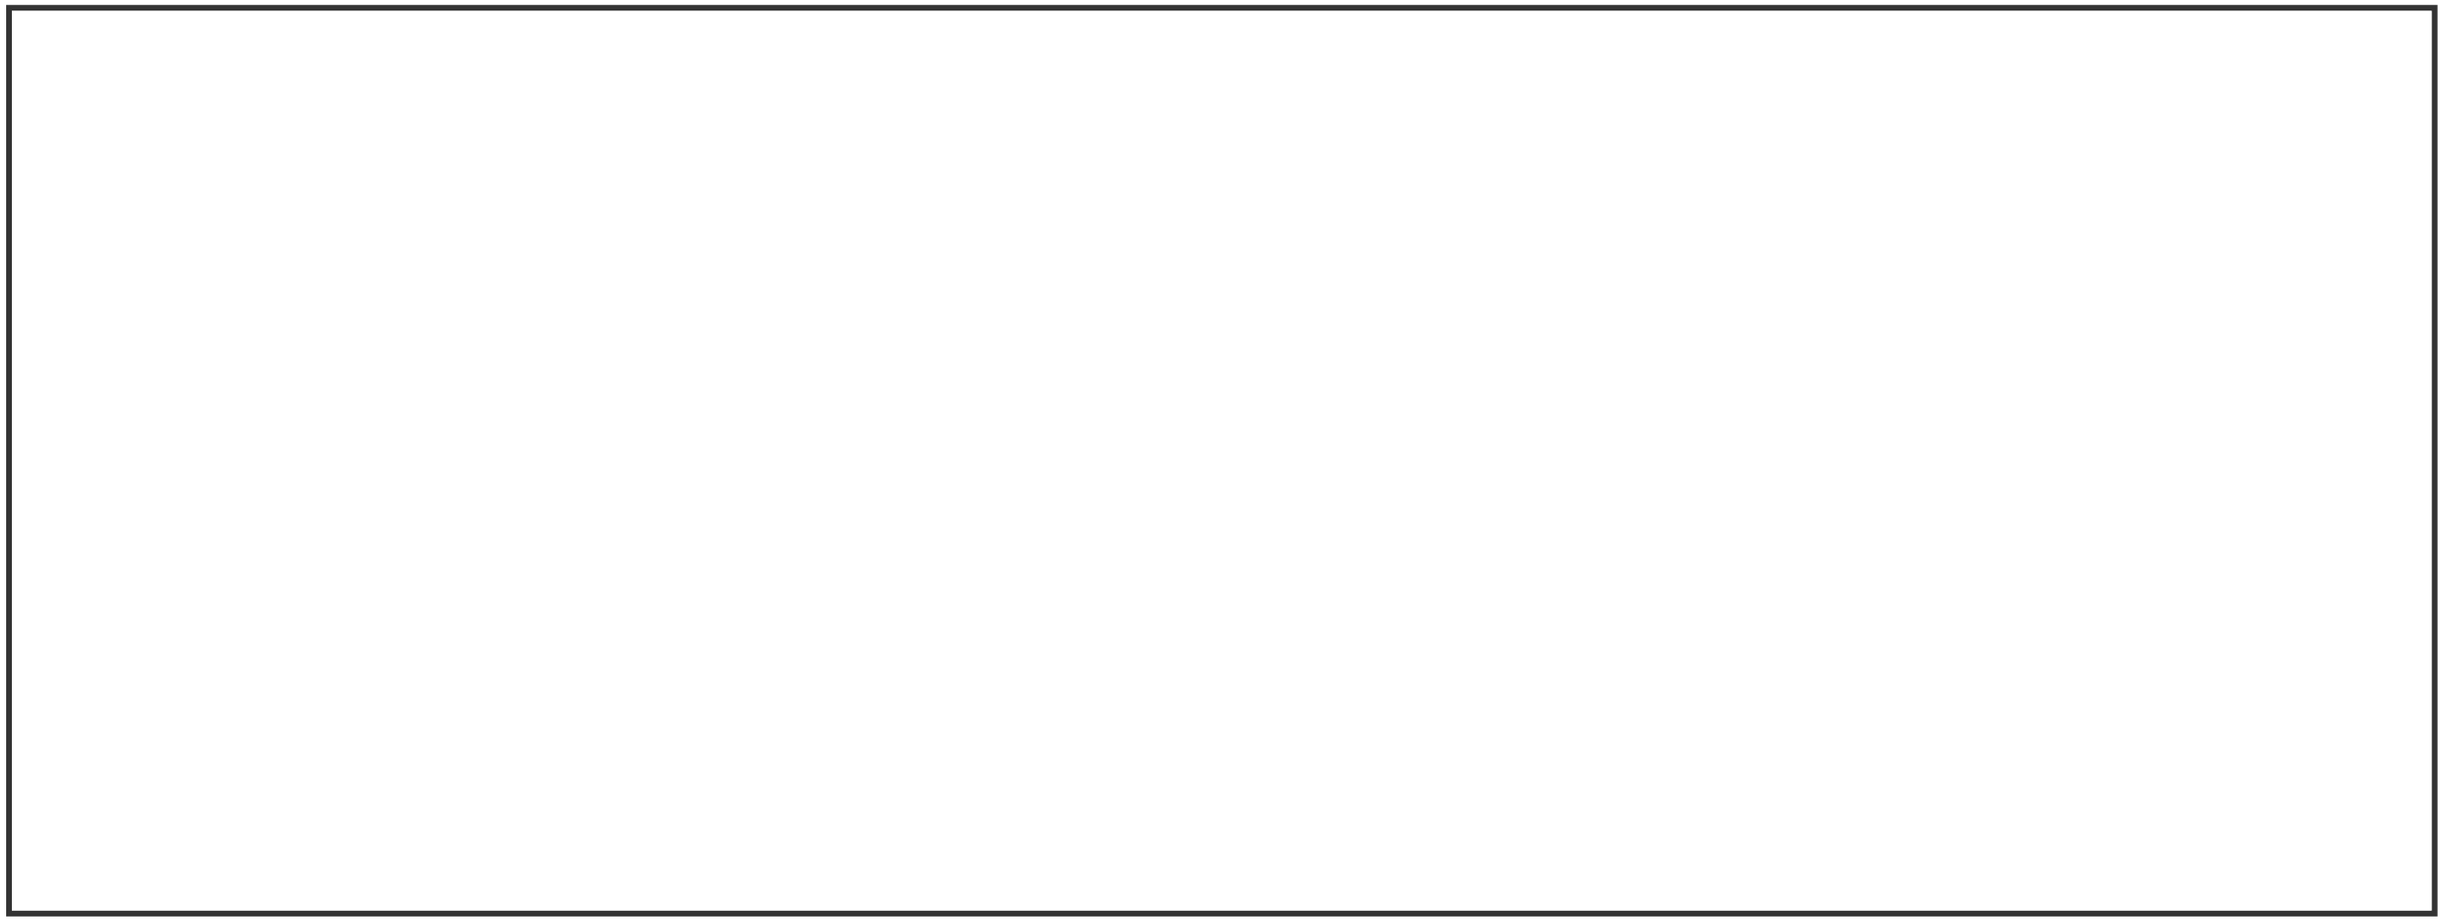

Reset

What do applications want from the network?

	Web browser	Video streaming	Video call
Throughput	Nice for light pages Needed for heavy page	Matters	Nice to have, especially for high quality video
Reliability (making sure all bits are	Ok to not show images? Matters for banking	Matters, not the end of the world if there is a bit	Matters the least
Jitter (variablity in latency)	Does not matter much	Matters (but we can deal with it with playback	Matters more
Latency	Nice to have (minimize time to load)	Not as much (buffer)	Matters most

Reset