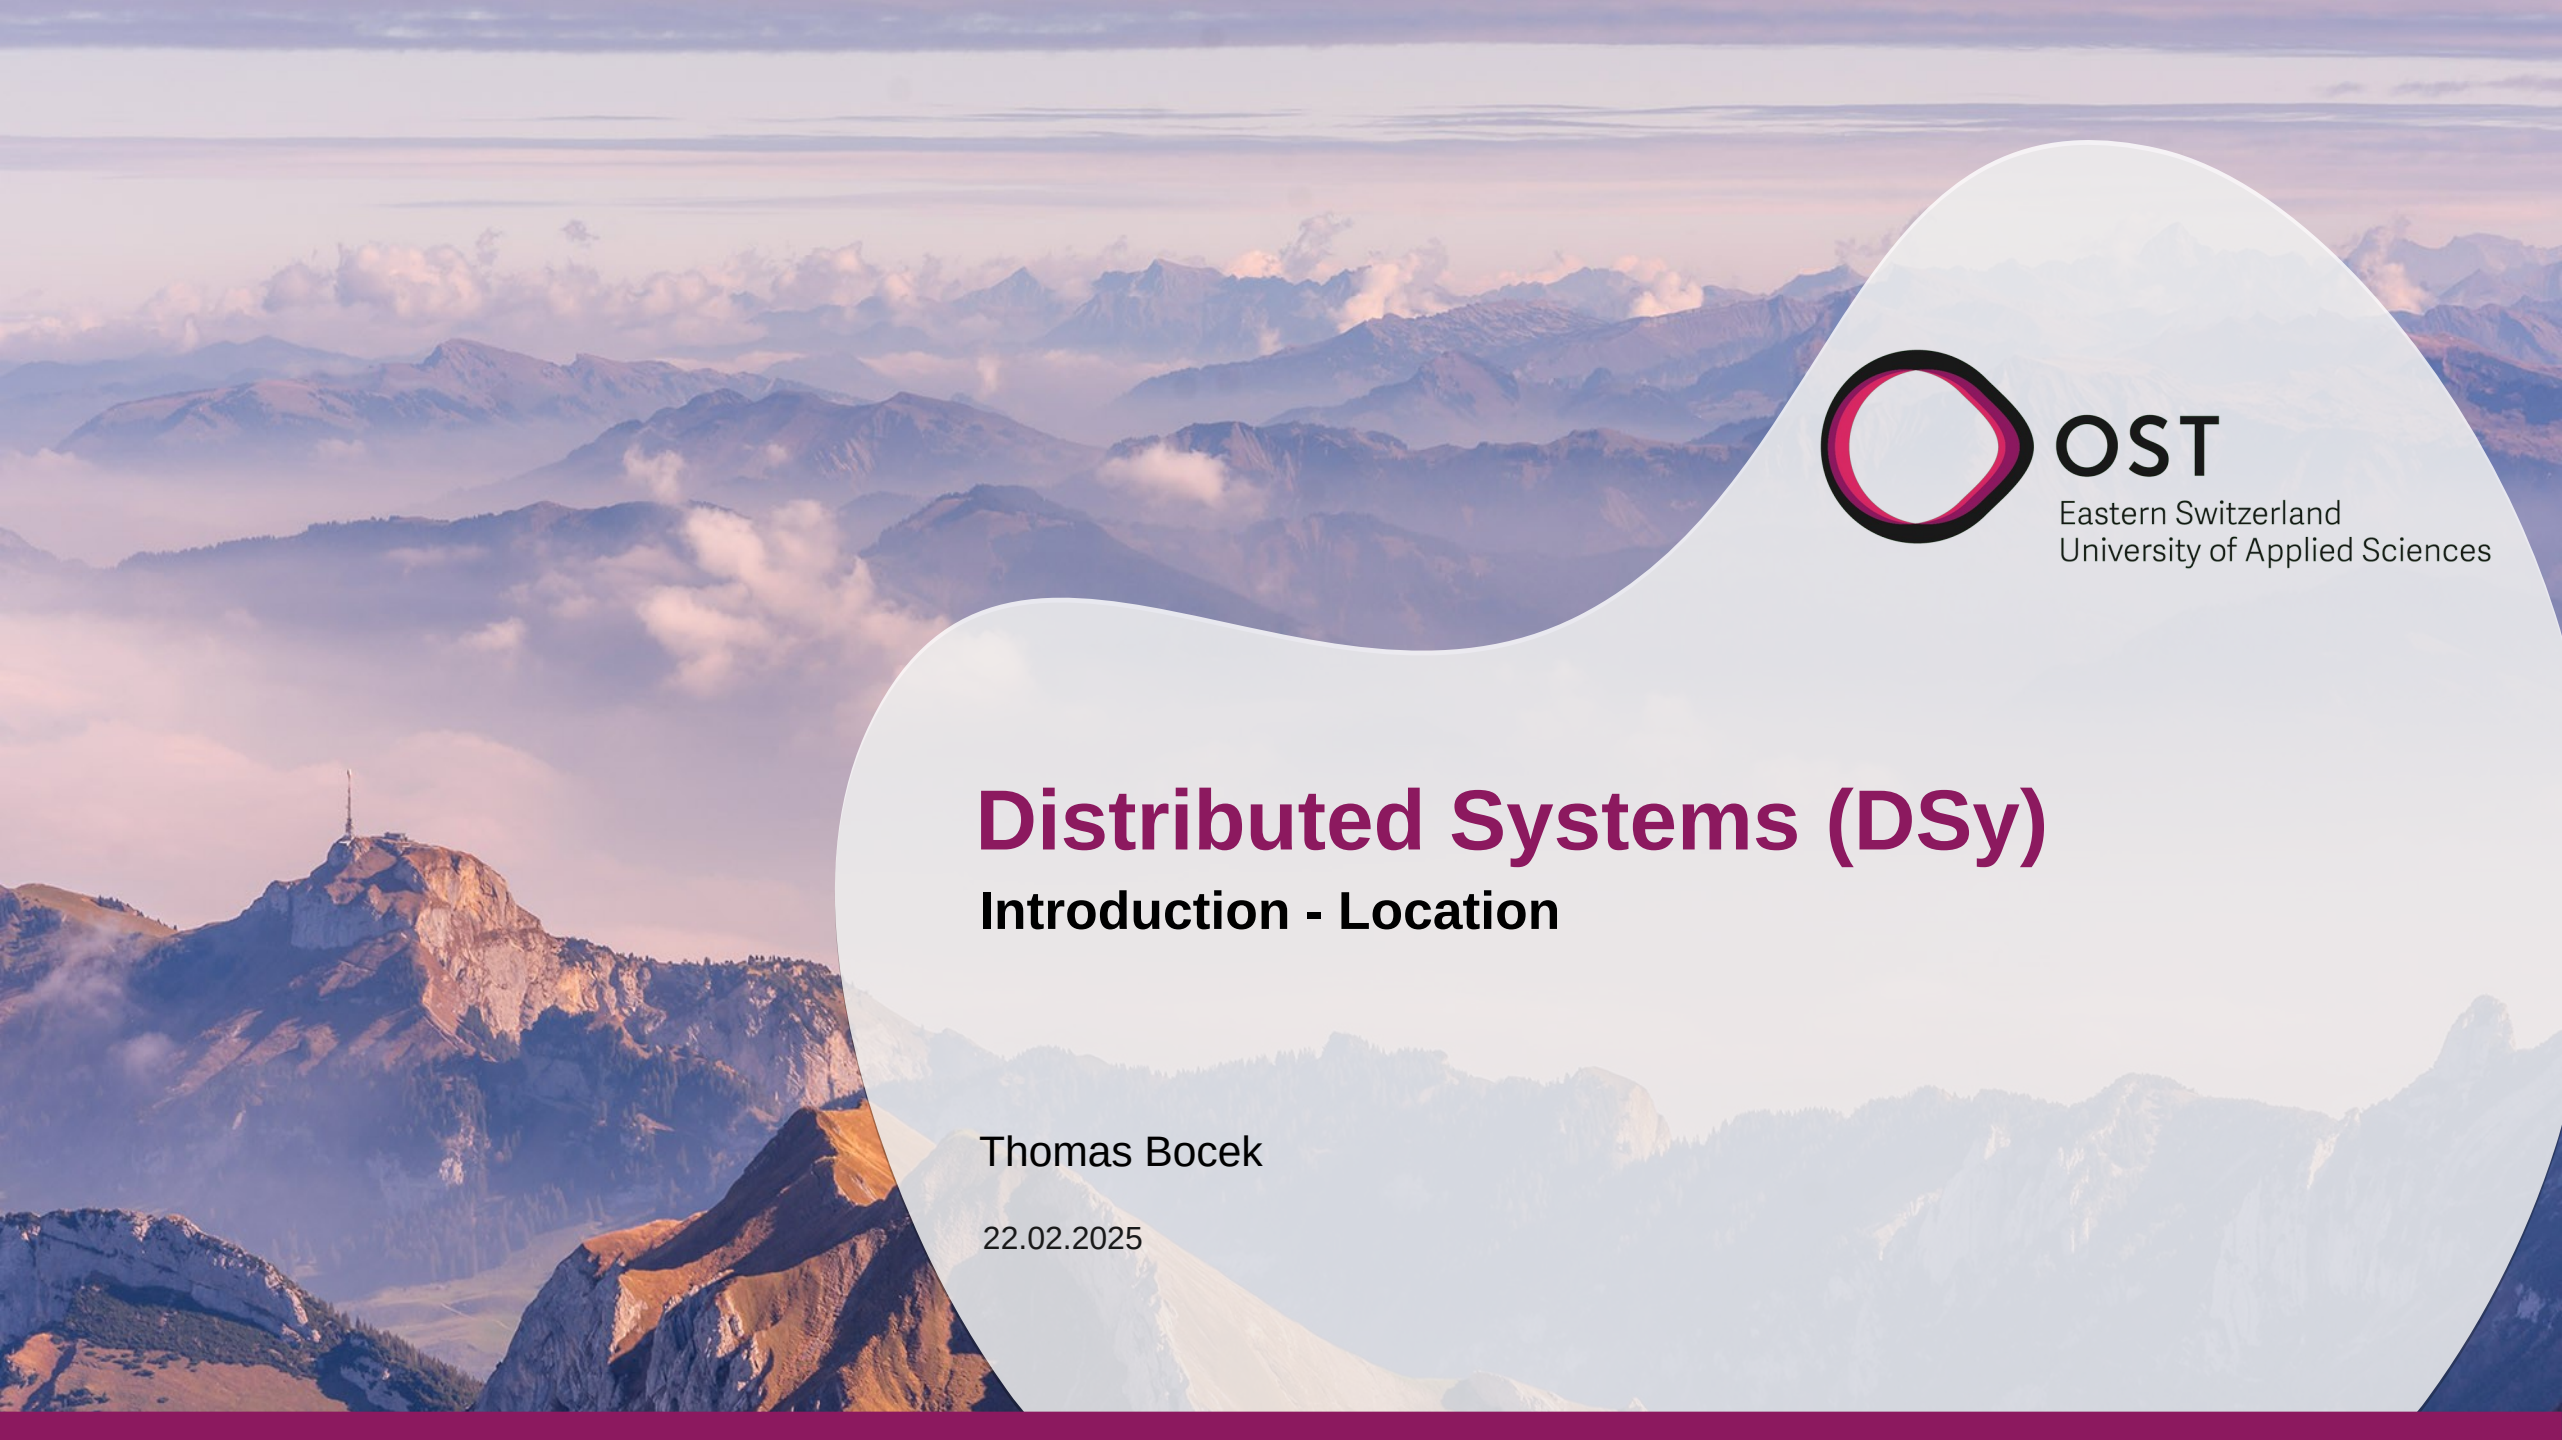

OST
Eastern Switzerland
University of Applied Sciences

Distributed Systems (DSy)

Introduction - Location

Thomas Bocek

22.02.2025

Learning Goals

- Distributed systems add complexity. Avoid complexity!
- Why do we need distributed systems?

  1) Scaling (if one machine is not enough)
  2) Location (to move closer to the user)
  3) Fault-tolerance (HW will fail eventually)

Distributed Systems Motivation

- Why Distributed Systems
  - Location
    - Everything gets faster (CPU, bandwidth, SSD), but latency stays
    - Einstein: nothing in nature is faster than the speed of light → you will always have latency

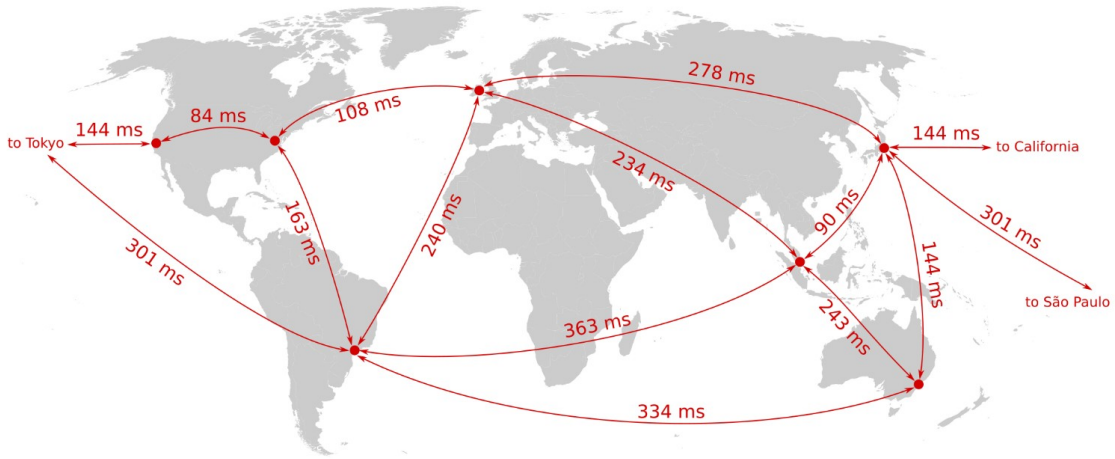

- Speed of light (c) in vacuum is ~300’000 km/s
  - Physical limit
- Latency: time for signal to travel from source to destination and back (round-trip time)
  - Perfect vacuum light tube to Sydney: RTT
    - (16540÷300000)×1000×2 = ~110ms
  - In practice: ~298ms [link] (ping au-ln.metercdn.net)
  - Space? Starlink satellite altitude: LEO ~550km [link]
    - Perfect condition, optimal location, no processing delay, no handoffs between satellites: theoretical latency: 7.3ms,
    - In practice: latency 20-60ms [link]

16,540 km

Distance from Rapperswil-Jona to Sydney

https://www.inkandswitch.com/local-first.html

Speed of Light

- Practice vs. theoretical limit - 298ms vs 110ms / 20-60ms vs 7.3ms
  - No direct path (fiber)
    - Land route in Europe (Switzerland to Mediterranean coast): ~1’000 km
    - Maybe SeaMeWe-5? ~16’000km
    - Singapore to Sydney (undersea cable): ~7,000 km
    - Total estimate: ~24,000 km
  - (24000÷300000)×1000×2 = 160ms
    - Still not 298ms
- Signal travels only speed of light in vacuum
  - Fiber = signal travels in glass [link] ~200’000 km/s

- Single mode fibers provide lower latency than multimode fibers, refractive index, wavelength of the light
- Hollow core fiber e.g. [link] with less latency
- Other materials [link]

Media	% of c	Description
Thick coaxial cable	77%	Originally used for ethernet, referred to as “thicknet”
Thin coaxial cable	65%	Referred to as ethernet “thinnet” or “cheapernet”
Unshielded twisted pair	59%	Multipaired copper cabling used for LAN and telecom applications
Microstrip	57%	PCB trace on FR4 dielectric, μr = 3.046
Stripline	47%	PCB trace in FR4 dielectric, μr = 4.6
Optical fiber	67%	Silica waveguide used to transport optical energy
Vacuum	100%	Vacuum or free space

OST

Speed of Light

- (24000÷200000)×1000×2 = 240ms
  - Non-optimal routing, queuing delays, routing delays and traffic inspection, signal repeating, protocol overhead
    - ~50-60ms - plausible
- Satellites have direct connection, light/radio travels through air/space almost at ~300’000km/s
  - Wifi with lowest latency? No
    - CSMA/CA, wait times before transmission, acknowledgment packets, retransmissions, signal processing at transmitter, processing at receiver, MAC layer processing, protocol stack traversal, DCF (Distributed Coordination Function) backoff, channel busy waiting
    - Typical case: +5ms latency

- Starlink in theory with lower latency than fiber?
  - Yes, latency to cover distance may be smaller using satellites [link]

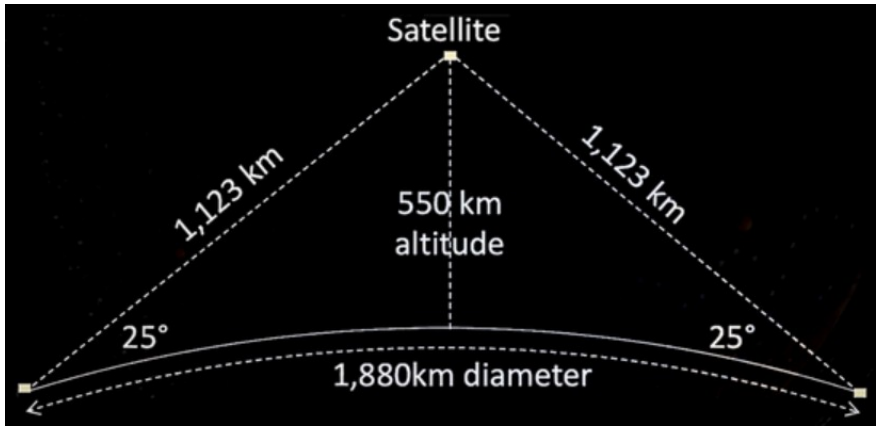

  - Latency satellite (vacuum):
    - ((2×1123)÷300000)×1000×2=15ms
  - Latency fiber (glass)
    - (1880÷150000)×1000×2=18.8ms

OST

Speed of Light

- Bandwidth much higher with fiber ~23Pb/s
  - Laser: NASA ~200Gb/s
    Starlink Inter-satellite ~100Gb/s, can be multiroute
  - Weather conditions affecting signal strength (ground – satellite), geomagnetic storms
  - Protocol overhead, network processing, signal encoding/decoding, queuing
- Geostationary satellite: 477ms latency
- Inter-satellite communication [youtube]

OST

Distributed Systems Motivation

- Copper vs Fiber
  - Copper propagates faster [link], but not much
  - Depending on the fiber material, latency can change

Minimum velocity factors allowed for network cable standards

VF (%)	Cable type	Ethernet physical layer
74~79%	Cat-7 twisted pair	
77%	RG-8/U	Minimum for 10BASE5[4]
67%	Optical fiber (silica glass)	Minimum for 10BASE-FL,[5] 100BASE-FX, ...
67%	Plastic optical fiber	1000BASE-RH*x* PMMA
63%	Plastic optical fiber	polystyrene
65%	RG-58A/U	Minimum for 10BASE2[6]
65%	Cat-6A twisted pair	10GBASE-T
64%	Cat-5e twisted pair	100BASE-TX, 1000BASE-T
58.5%	Cat-3 twisted pair	Minimum for 10BASE-T[7]

- Importance of latency
  - Amazon: +100ms latency → 1% sales loss [link]
  - Google: +500ms latency → 20% drop in traffic [link]
  - Bing: +500ms latency → revenue down 1.2% [link]
- Gaming

Sensitivity to latency in online gaming
>300 ms – game unplayable
>150 ms – player performance degraded
>100 ms – player performance affected
50 ms – target performance
13 ms – lower limit of detectibility

Source: PubNub

http://web.archive.org/web/20160729190529/https://www.pubnub.com/blog/2015-02-09-how-fast-is-realtime-human-perception-and-technology/

OST

Distributed Systems Motivation

- Gaming / e.g., Esports LoL, price ~$2.25m:
  - Human reaction time 200ms
  - Total from keypress to display:
    - Thinkpad 13 ChromeOS: 70ms
    - Lenovo X1 carbon 2016: 150ms
  - TV output lag ~8ms (random TV)
  - Keyboard 15-60ms
    - Key travel time!
    - PS/2 vs USB keyboard
    - USB polling ~8ms, PS/2 interrupt based, direct path to CPU, USB gaming keyboard with 1ms polling
- 60hz display frame rate: 8ms delay

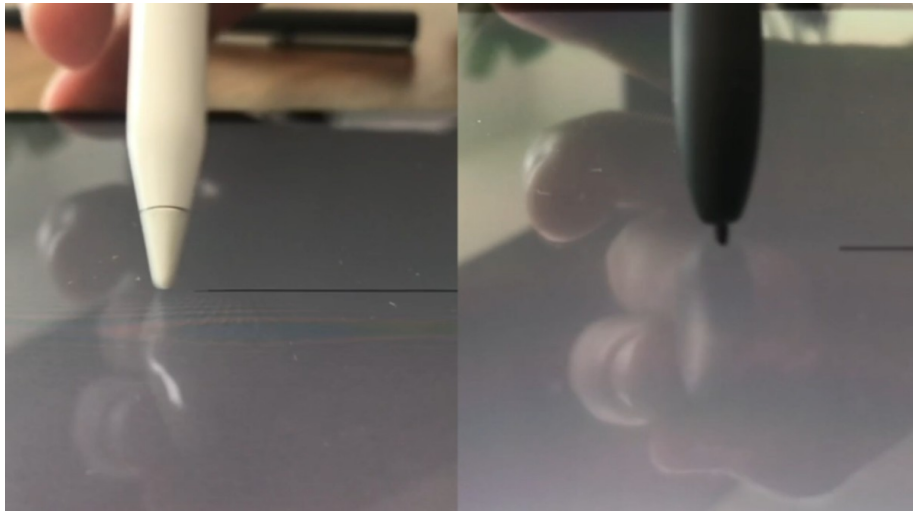

- Tablet pen, latency 20-80ms [link]
- Competitive gaming: use special hardware!
  - 120 or 240hz, low latency mouse/keyboard

OST

Distributed Systems Motivation

- Reducing latency
  - Assumption: perfect repeater, switch, router with no latency
  - Perfect mouse, keyboard, display
- RTT to Sydney still 110ms with perfect (unrealistic) conditions
  - nothing in nature is faster than the speed of light → you will always have latency
- Place services closer to user → distributed system
  - Reduced latency
  - Can increased bandwidth and throughput
  - Can improved reliability and availability
  - Drawback: coordination of data replication and caching
- e.g., CDN: Content delivery network
  - Place your images, sites, scripts close to your users
- New protocols can decrease nr. of RTT
  - Upcoming lecture