



**OST**

Eastern Switzerland  
University of Applied Sciences

# Distributed Systems (DSy)

**In the news**

Thomas Bocek

24.02.2024

# Distributed Systems & Blockchain in the News

- 20.02.2024: The Billion Row Challenge (1BRC) - Step-by-step from 71s to 1.7s [\[link\]](#)
  - How fast can YOU aggregate 1B rows using modern #Java? tl;dr – 1.5s
    - Write a Java program for retrieving temperature measurement values from a text file and calculating the min, mean, and max temperature per weather station
  - Its Java, but with black magic, aka Unsafe
    - Winner [\[link\]](#) vs. regular [\[link\]](#) 71s
  - 0: Use good VM 71s → 66s
    - 1: Parallelize I/O 66s → 17s
    - 2: Directly parse temperature as int 17s → 11s
    - 3: Custom hashtable 11s → 6.8s
    - 4: sun.misc.Unsafe, SWAR 6.8s → 2.6s
    - 5: Analyze branch prediction 2.6s → 2s
    - 6: Eliminate startup/cleanup costs → eliminate 200ms 2s → 1.8s
    - 7: Use smaller chunks and work stealing → micro optimizations 1.8s → 1.7s
    - “... keep in mind that a lot of the gains come from dispensing with all the best practices that apply to production code...”

# Distributed Systems & Blockchain in the News

- 22.02.2024: Bitcoin Inventor Satoshi Nakamoto Emails Revealed in Court [\[link\]](#)
  - Revealed emails between Satoshi Nakamoto and Adam Back
- 21.02.2024: Lots of new LLM
  - Gemma: Introducing new state-of-the-art open models [\[link\]](#) - probably censored, biased? [\[link\]](#), 7B model
  - Mistral-next [\[link\]](#) - uncensored, no information yet, ?B model
  - Phind [\[link\]](#), a 70B model
- A 24GB 7900XTX can run a quantized 30B model. This is currently to top consumer product with 24GB, best HW maxed out at 192GB [\[link\]](#) (MI300X), H200 141GB [\[link\]](#) → multi GPU in one machine → then second machine
- 11.01.2024: Intel Processor Instability Causing Oodle Decompression Failures [\[link\]](#)
  - Distributed system: fault tolerance