Distributed Computer Systems = Making Computer Systems Scalable, Reliable, Performant, etc., Yet Able to Form an Efficient Ecosystem

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Co-sponsored by:

NWO Netherlands Organisation for Scientific Research

STW connecting innovators

@AIosup

Prof. dr. ir. Alexandru Iosup

Datacenters at Google

0.4s delay would lose 8 million searches/day

VU VRIJE UNIVERSITEIT AMSTERDAM

https://blog.kissmetrics.com/loading-time/?wide=1
http://www.fastcompany.com/1825005/how-one-second-could-cost-amazon-16-billion-sales

TU Delft

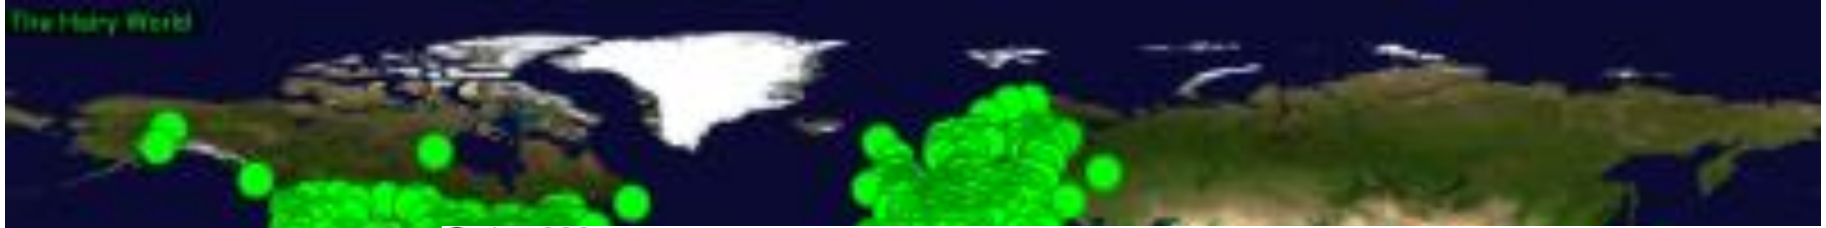

BitTorrent is

Is it robust?

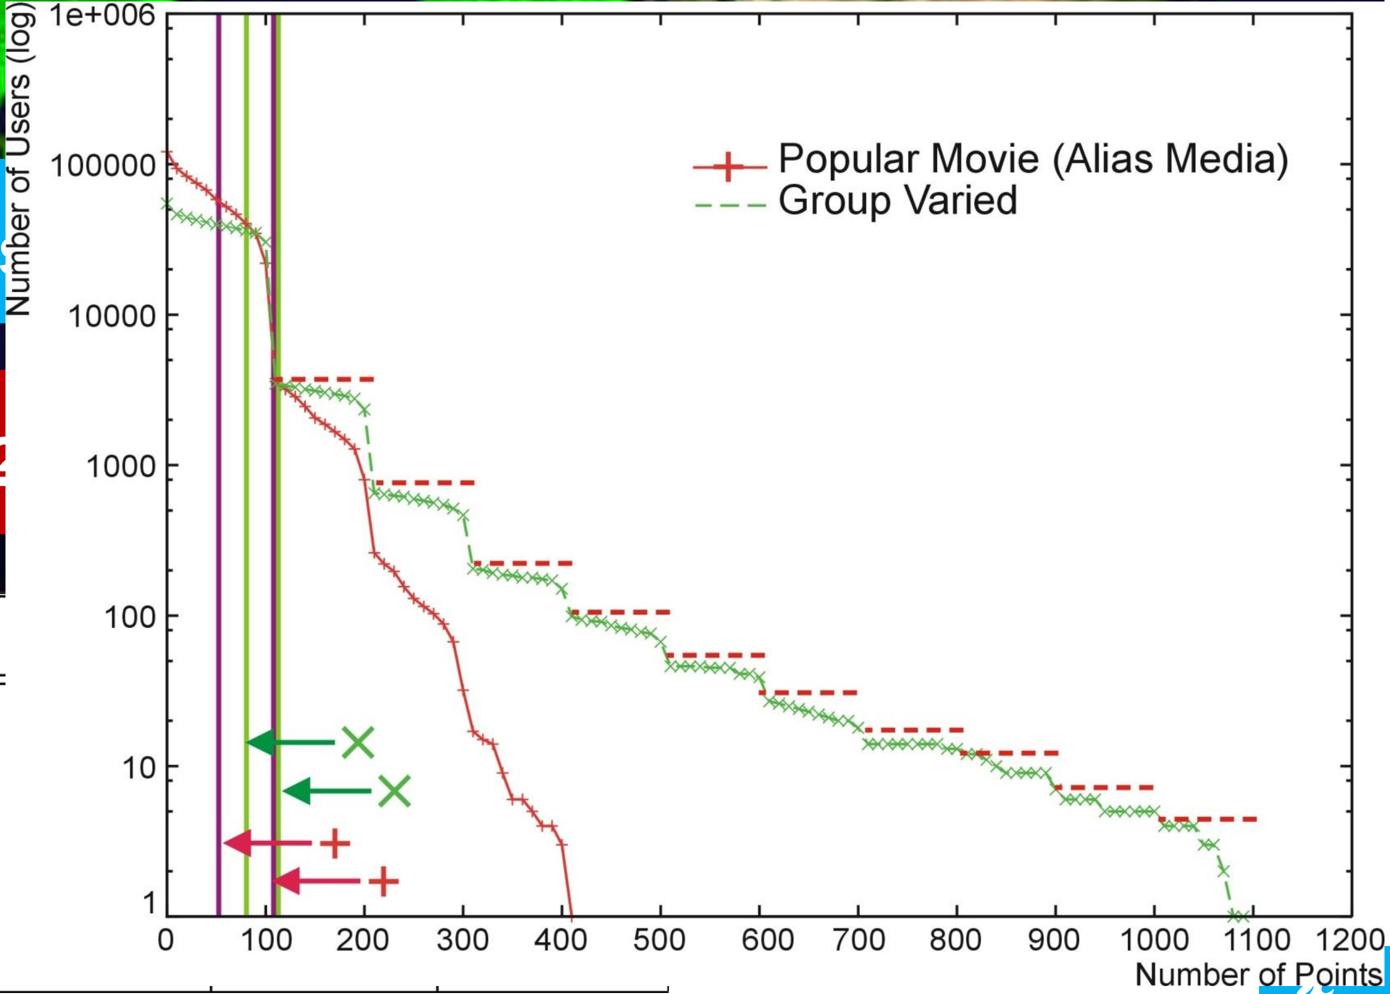

	All
Min	13 B
Q1 B	61.17 MB
Median	366.77 MB
Mean	1.50 GB
Q3 B	1.05 GB
Max	423.55 GB
StdDev	5.18 bn
Sum	3.43 PB

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Why Distributed Computing Systems?

- The Internet, email, cloud storage, …
- Media/Web Services and Search (and P2P File-Sharing)
- Banking systems
- Most IaaS clouds, Big Data systems
- High-Performance Computing: a cure for cancer, weather prediction, tsunami alerts, making a new car/plane/…,
- Plain Resources: Computation, Backup and Storage, etc.

> *“In my daily work,*
> *I work on very large, complex, **distributed systems** built out of many Python modules and packages.”*
> - Guido van Rossum

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

What Makes Distributed Computing Challenging?

- Failures
- Inconsistent State
- Scalability
- Multi-tenancy
- … Everything is difficult:
  distributed debugging, distributed resource management, …
- Active research area, industry needs advances and hires graduates

> *"You know you have a* ***distributed system*** *when the* *crash* *of a computer you've never heard of stops you from getting any work done."*
> - Leslie Lamport

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Course Goals

Lectures

1. Explain the basic concepts, objectives, and functions of distributed computing systems, e.g., communication, resource management and scheduling, data consistency, fault-tolerance, performance
2. Compare distributed computing with other computing paradigms
3. Identify the different flavors of modern distributed systems
4. Analyze the trade-offs inherent in the design of modern distributed systems
5. Design your portfolio distributed-system, with many basic and some complex operations of modern distributed systems
6. Implement your portfolio distributed-system
7. Analyze your portfolio distributed-system

Seminar

Design

Lab

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Course Topics (Lectures, Weeks 44-51)

Week	Topic
Week 44	1. Introduction to distributed computing systems
Week 45	2. Architecture: distributed workloads and resource management, scheduling mechanism and policy
Week 46	3. Communication: naming and messaging
Week 47	4. Fault-Tolerance: reliability, availability, and fault-tolerant protocols
Week 48	5. (Big) Data: consistency, replication, and distributed file systems
Week 49	6. Use Cases: scientific/engineering data processing
Week 50	7. Use Cases: massively multiplayer online gaming (consumer multimedia systems)
Week 51	8. Invited lecture: Databricks (whp)

VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Large Lab Exercise

Groups of max. 6

- Self-selection or matchmaking by TA

Create a **distributed system**

- Design from given requirements or from your own requirements (validate w TA)
- Implementation
- Validation/Testing/Evaluation

Zero tolerance for plagiarism

Create a **technical report** about the system (4—6 pages)

- **Report is worth as much as coding!**
- Explain design choices/trade-offs
- Analyse system characteristics
- **Exercise Technical writing**: quality, structure, presentation, graphing, etc.

Zero tolerance for plagiarism

Course Material

- Textbook
  - Maarten van Steen and Andrew S. Tanenbaum, Distributed Systems: Principles and Paradigms, Prentice Hall, 3nd Ed., 2017
    Free student edition: https://www.distributed-systems.net/index.php/books/distributed-systems-3rd-edition-2017/
- Lectures, Seminars, Design sessions
- Lab exercises
  - Plenty of material not in the book!
  - All available on Canvas, after the session

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Experimental Research Methodology
Our/Your Main Scientific Instrument: DAS-5

300+ scientists as users

Won IEEE Scale Challenge 2014

Course Schedule (Weeks 44-51, Exam Week 51)

- Lectures
  - Mon, 13:30-15:15, WN-Q105
- Design sessions (2x, groups of 40 students)
  - Wed, 09:00-10:45, WN-Q105 (WN-Q112 on Nov 15/Week 46)
- Seminar sessions (2x, groups of 40 students)
  - Thu, 15:30-17:15, WN-M655
- Lab exercise (schedule w TA, Demo Day)
- Exam (Dec 21, 15:15-18:00)

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Important Dates, Important Milestones

Guides: Additional recommended milestones, with recommended deadlines

- Nov 6 (Week 44)
  - Enroll for team exercise
  - No enrollment, No grade
- Dec 6 (Week 49)
  - Demonstrate system to TA
  - TA can run code, you can explain code and experiments
- Dec 15 (Week 50)
  - Lab results
  - Report ~ 50% Lab grade
  - Selected systems get to demo
  - Top-5 open-access booklet!
- Dec 20 (Week 51)
  - Demo Day
  - Databricks talk (whp)
- Dec 21 (Week 51)
  - Exam, written, open questions
  - Exam, design discussion

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Recommended Time Allocation

- Lectures
  - 2 hours per week in-class, Mon
- Design and Seminar sessions
  - 4 hours every 3 weeks in-class, Wed and Thu, respectively
- Large exercise (up to 6 persons)
  - ≤100 hours/person, including experiments and report writing
- Self-study
  - 4—6 hours/week. Check tech. blogs of Facebook, Google, etc.

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

The Grading System
(You can build your own path)

Course Points
10,000 for straight 10
4,000 for **large exercise**
6,000 for **Exam**
≤1,000 for **Design & Seminar**
≤1,000 for **Lab** bonuses
+50 activity in Lab/Design/Seminar/Lecture

+500 entry quiz

Competition
Demo Day (pitch to guest lecturer)
Top-5 Reports Published

Access Tokens
Discuss w Lecturer about MSc topics
Recommendation Letter

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Our Team


Laurens Versluis


Alexandru Uțã


Alexandru Iosup

- Course coordinator
  - Prof. dr. ir. Alexandru Iosup
- Co-teachers
  - Prof. dr. ir. Alexandru Iosup
  - Dr. Alexandru Uță
- Teaching Assistants
  - Ir. Laurens Versluis et al.

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

Protocol for Registering a Team for Lab/Seminar/Design

1. Choose when you would prefer to participate
   1. Systems: Weeks 1/2/3
   2. Ecosystems: Weeks 4/5/6
2. Send us an email by Mon, Nov 6, at 9am CET
   - To: DistributedSystems2017@atlarge-research.com
   - Subject: [vu] [Distributed Systems] Team <YourTeamName>
   - Content: Lab/Seminar/Design

     Team members: list_of <Name> (<ID>), …

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

The images used in this lecture courtesy of the Computer History Museum, Mountain View, California, USA, http://www.computerhistory.org/ ; the German Museum of Technology (Deutsches Technikmuseum Berlin, Germany, http://www.sdtb.de/Englisch.55.0.html ; the Science Museum, London, UK, http://www.sciencemuseum.org.uk/; and many anonymous contributors via Google Images. Many thanks!