CLOUD COMPUTING CONCEPTS

with Indranil Gupta (Indy)

Part 1 CONCLUSION

And we Take a Pause

- What you've sampled so far has (hopefully) given you a good taste for what's underneath cloud computing systems

What you've Learnt so far

- Introduction: Clouds, Mapreduce, Key-value stores
- Classical Precursors: Peer-to-peer systems, Grids
- Widely-used algorithms: Gossip, Membership, Paxos
- Classical algorithms: Time and Ordering, Snapshots, Multicast
- Fun: Interviews with leading managers and researchers, from both industry and academia

After the Pause: Coming Up in C3 - Part 2

- Classical algorithms: Leader Election, Mutual Exclusion, Scheduling
- Scalability: Concurrency control, Replication Control
- Trending Areas: Stream processing, Graph processing, Structure of Networks, Sensor Networks
- Miscellaneous: Distributed File systems, Distributed shared memory, Security, Datacenter outage studies
- Fun: Interviews with leading managers and researchers, from both industry and academia

The Goal of C3 - Part 2 Remains the Same

- Course about the internals of cloud computing
- Distributed systems and algorithms that underlie today's cloud computing technologies
- We'll discuss
  - Concepts
  - Techniques
  - Industry systems, including open-source (from the inside)

Format of C3 – Part 2 is Familiar to you

- 2 Homeworks
- (Optional) 1 Programming Assignment
  - Implement a key-value store inside an emulator
  - Builds on the Programming assignment from C3 – Part 1
- 1 Exam

Onward!

- C3 – Part 1 is a prerequisite for C3 – Part 2, so you're perfectly placed to move forward to Part 2
- Cloud computing is an exciting area to be studying, very dynamic and continuously changing
- Come, let's continue our journey through the landscape.
- I'm looking forward to seeing you in Cloud Computing Concepts – Part 2!