

CS 656: Networked Systems for Cloud Computing

Ramesh Govindan

Logistics

Class Information

- **Location:** KAP 147
- **Time:** MW 4-5:50pm

Instructor: Ramesh Govindan

- **Office Location:** SAL 212
- **Office Hours:** TBD
- **Contact mode:** TBD

Teaching Assistant: TBD

- **Office Location:** TBD
- **Office Hours:** TBD
- **Contact mode:** TBD

Course Description

Today, most online services are hosted on large public clouds. Beneath the hood, cloud computing is powered by large, complex networked systems. Most cloud compute and storage is hosted in massive data centers connected by extremely high capacity networks. In turn, these data centers are interconnected by global private WANs. On this infrastructure, cloud providers run innovative networked systems that enable them to offer high availability, high throughput and low latency to their customers at low cost. These networked systems form the subject matter for the course. The course will be based on readings from representative papers published in top systems and networking conferences.

Note: This is the MS-equivalent of the PhD-level CS 656 (Networked Systems for Cloud Computing). The two courses are substantially the same, and are intended to be co-offered to MS and PhD students. PhD students will be evaluated slightly differently: they will be required to prepare and present 1-2 papers each in class.

Catalogue Description

Systems and network designs and protocols for cloud computing and data center networks.

Learning Objectives

In this course, we will visit the critical technology trends and new challenges in cloud and data center designs for different trade-offs of performance, scalability, manageability, and cost across the application, infrastructure, and

network layers. The course will include student reviews, discussions, lectures, and in-class quizzes. The readings will be selected from recent papers published in top systems and networking conferences. Critical evaluation of the readings will be conducted through independent reviews and student led discussions.

Prerequisite(s) CSCI 551 or CSCI 651 preferred, but an undergraduate networks course (CS 353 equivalent) might be acceptable with instructor permission.

Co-Requisite (s) None

Concurrent Enrollment None

Required Readings

An initial list of readings follows. This list will likely change in future years to include significant advances in the field.

Data center design

[JupiterRising] Arjun Singh et al. (2016). “Jupiter rising: a decade of clos topologies and centralized control in Google’s datacenter network”. In: *Commun. ACM* 59.9, pp. 88–97. DOI: 10.1145/2975159. URL: <https://doi.org/10.1145/2975159>

[JupiterEvolving] Leon Poutievski et al. (2022). “Jupiter evolving: transforming google’s datacenter network via optical circuit switches and software-defined networking”. In: *SIGCOMM ’22: ACM SIGCOMM 2022 Conference, Amsterdam, The Netherlands, August 22 - 26, 2022*, pp. 66–85. DOI: 10.1145/3544216.3544265. URL: <https://doi.org/10.1145/3544216.3544265>

Routing

[FacebookBGP] Anubhavnidhi Abhashkumar et al. (2021). “Running BGP in Data Centers at Scale”. In: *18th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2021, April 12-14, 2021*, pp. 65–81. URL: <https://www.usenix.org/conference/nsdi21/presentation/abhashkumar>

[Orion] Andrew D. Ferguson et al. (Apr. 2021). “Orion: Google’s Software-Defined Networking Control Plane”. In: *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*. USENIX Association. ISBN: 978-1-939133-21-2. URL: <https://www.usenix.org/conference/nsdi21/presentation/ferguson>

[Fastpass] Jonathan Perry et al. (2014). “Fastpass: a centralized "zero-queue" datacenter network”. In: *ACM SIGCOMM 2014 Conference, SIGCOMM’14, Chicago, IL, USA, August 17-22, 2014*, pp. 307–318. DOI: 10.1145/2619239.2626309. URL: <https://doi.org/10.1145/2619239.2626309>

WAN design

[B4] Chi-Yao Hong, Subhasree Mandal, et al. (2018). “B4 and after: managing hierarchy, partitioning, and asymmetry for availability and scale in google’s software-defined WAN”. in: *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM 2018, Budapest, Hungary, August 20-25, 2018*, pp. 74–87. DOI: 10.1145/3230543.3230545. URL: <https://doi.org/10.1145/3230543.3230545>

[SWAN] Chi-Yao Hong, Srikanth Kandula, et al. (2013). “Achieving high utilization with software-driven WAN”. in: *ACM SIGCOMM 2013 Conference, SIGCOMM 2013, Hong Kong, August 12-16, 2013*, pp. 15–26. DOI: 10.1145/2486001.2486012. URL: <https://doi.org/10.1145/2486001.2486012>

Transport protocols

[DCTCP] Mohammad Alizadeh et al. (2010). “Data center TCP (DCTCP)”. in: *Proceedings of the ACM SIGCOMM 2010 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, New Delhi, India, August 30 -September 3, 2010*, pp. 63–74. DOI: 10.1145/1851182.1851192. URL: <https://doi.org/10.1145/1851182.1851192>

[Swift] Gautam Kumar et al. (2020). “Swift: Delay is Simple and Effective for Congestion Control in the Datacenter”. In: *SIGCOMM '20: Proceedings of the 2020 Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication, Virtual Event, USA, August 10-14, 2020*, pp. 514–528. DOI: 10.1145/3387514.3406591. URL: <https://doi.org/10.1145/3387514.3406591>

Virtualization

[Andromeda] Michael Dalton et al. (2018). “Andromeda: Performance, Isolation, and Velocity at Scale in Cloud Network Virtualization”. In: *15th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2018, Renton, WA, USA, April 9-11, 2018*, pp. 373–387. URL: <https://www.usenix.org/conference/nsdi18/presentation/dalton>

[VMware] Teemu Koponen et al. (2014). “Network Virtualization in Multi-tenant Datacenters”. In: *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2014, Seattle, WA, USA, April 2-4, 2014*, pp. 203–216. URL: <https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/koponen>

Fast communication

[Firestone] Daniel Firestone et al. (2018). “Azure Accelerated Networking: SmartNICs in the Public Cloud”. In: *15th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2018, Renton, WA, USA, April 9-11, 2018*, pp. 51–66. URL: <https://www.usenix.org/conference/nsdi18/presentation/firestone>

[Snap] Michael Marty et al. (2019). “Snap: a microkernel approach to host networking”. In: *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019, Huntsville, ON, Canada, October 27-30, 2019*, pp. 399–413. DOI: 10.1145/3341301.3359657. URL: <https://doi.org/10.1145/3341301.3359657>

[IRMA] Arjun Singhvi et al. (2020). “IRMA: Re-envisioning Remote Memory Access for Multi-tenant Datacenters”. In: *SIGCOMM '20: Proceedings of the 2020 Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication, Virtual Event, USA, August 10-14, 2020*, pp. 708–721. DOI: 10.1145/3387514.3405897. URL: <https://doi.org/10.1145/3387514.3405897>

Load-balancers

[Maglev] Daniel E. Eisenbud et al. (2016). “Maglev: A Fast and Reliable Software Network Load Balancer”. In: *13th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2016, Santa Clara, CA, USA, March 16-18, 2016*, pp. 523–535. URL: <https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/eisenbud>

[Ananta] Parveen Patel et al. (2013). “Ananta: cloud scale load balancing”. In: *ACM SIGCOMM 2013 Conference, SIGCOMM 2013, Hong Kong, August 12-16, 2013*, pp. 207–218. DOI: 10.1145/2486001.2486026. URL: <https://doi.org/10.1145/2486001.2486026>

Global traffic control

[Mapping] Fangfei Chen, Ramesh K. Sitaraman, and Marcelo Torres (2015). “End-User Mapping: Next Generation Request Routing for Content Delivery”. In: *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM 2015, London, United Kingdom, August 17-21, 2015*, pp. 167–181. DOI: 10.1145/2785956.2787500. URL: <https://doi.org/10.1145/2785956.2787500>

[AnycastCDN] Matt Calder et al. (2015). “Analyzing the Performance of an Anycast CDN”. in: *Proceedings of the 2015 ACM Internet Measurement Conference, IMC 2015, Tokyo, Japan, October 28-30, 2015*, pp. 531–537. DOI: 10.1145/2815675.2815717. URL: <https://doi.org/10.1145/2815675.2815717>

[Espresso] Kok-Kiong Yap et al. (2017). “Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering”. In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM 2017, Los Angeles, CA, USA, August 21-25, 2017*, pp. 432–445. DOI: 10.1145/3098822.3098854. URL: <https://doi.org/10.1145/3098822.3098854>

[EdgeFabric] Brandon Schlinker et al. (2017). “Engineering Egress with Edge Fabric: Steering Oceans of Content to the World”. In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM 2017, Los Angeles, CA, USA, August 21-25, 2017*, pp. 418–431. DOI: 10.1145/3098822.3098853. URL: <https://doi.org/10.1145/3098822.3098853>

Reliability

[TailatScale] Jeffrey Dean and Luiz André Barroso (2013). “The tail at scale”. In: *Commun. ACM* 56.2, pp. 74–80. DOI: 10.1145/2408776.2408794. URL: <https://doi.org/10.1145/2408776.2408794>

[EvolveorDie] Ramesh Govindan et al. (2016). “Evolve or Die: High-Availability Design Principles Drawn from Googles Network Infrastructure”. In: *Proceedings of the ACM SIGCOMM 2016 Conference, Florianopolis, Brazil, August 22-26, 2016*, pp. 58–72. DOI: 10.1145/2934872.2934891. URL: <https://doi.org/10.1145/2934872.2934891>

Programmable switches

[P4] Pat Bosshart et al. (2013). “Programming Protocol-Independent Packet Processors”. In: *CoRR* abs/1312.1719. arXiv: 1312.1719. URL: <http://arxiv.org/abs/1312.1719>

[FBOSS]: Sean Choi et al. (2018). “FBOSS: building switch software at scale”. In: *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM 2018, Budapest, Hungary, August 20-25, 2018*, pp. 342–356. DOI: 10.1145/3230543.3230546. URL: <https://doi.org/10.1145/3230543.3230546>

Technological Proficiency and Hardware/Software Required

Proficiency in C/C++ and Python.

Grading Breakdown

Reviews (20%) Students are expected to write reviews for each class. Your review should look similar to reviews of papers submitted to conferences, which usually have four parts:

1. Paper summary: What problem does the paper address (1-2 sentences or bullets)? What are the core novel ideas or technical contributions of the work (1-2 sentences or bullets)? What is the paper’s approach, what specific techniques/mechanisms does it use, and what are its main findings? (3-5 sentences)
2. Strengths: 2-4 bulleted points (Explain in more details in the detailed comments)

3. Weaknesses: 2-4 bulleted points (Explain in more details in the detailed comments)
4. Detailed comments on the paper: Elaborate in a few sentences each of the strengths and weaknesses.

Quizzes (30%) There will be four short (30 min) quizzes, evenly spaced through the semester, to test your understanding of the papers in the class.

Research project (50%) The semester-long project is an open-ended systems research project. Project topics are of your choice but should be related to networked systems for cloud computing. Projects should be done in groups of two or three and should include a systems building component.

Schedule

Week	Topic	Other Milestones
1	Data center designs	Identify team
2	Routing	
3	Routing	Project proposal
4	WAN Designs	Quiz 1
5	Transport protocols	
6	Transport protocols	
7	Virtualization	
8	Fast communication	Quiz 2
9	Fast communication	Mid-term report
10	Load balancers	
11	Global traffic control	
12	Global traffic control	Quiz 3
13	Reliability	
14	Programmable switches	
15	Student project presentations	
16	Student project presentations	Quiz 4, Final Report

Research project

The purpose of the class project is for students to have some research experience on problems related to cloud computing. Students are encouraged to identify new research projects based on their own reading of papers from recent conferences. The instructor will also provide sample projects for students at the beginning of the semester. Students are encouraged to work in groups of 2 or 3 students.

Project Timeline

- Week 1-2: Identify team members and project topics
- Week 3: Proposal due (team member, topics and milestone)
- Week 9: Mid-term report due (data description, preliminary results)
- Week 15: Project presentation (open to all students)
- Week 16: Final report due (similar to a 6-page workshop research paper)

Grade breakdown of the course project

- Proposal: 5%
- Mid-term report: 10%
- Final report: 20%
- Presentation: 15%

Statement on Academic Conduct and Support Systems

Academic Integrity:

The University of Southern California is a learning community committed to developing successful scholars and researchers dedicated to the pursuit of knowledge and the dissemination of ideas. Academic misconduct, which includes any act of dishonesty in the production or submission of academic work, compromises the integrity of the person who commits the act and can impugn the perceived integrity of the entire university community. It stands in opposition to the university's mission to research, educate, and contribute productively to our community and the world.

All students are expected to submit assignments that represent their own original work, and that have been prepared specifically for the course or section for which they have been submitted. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s).

Other violations of academic integrity include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), collusion, knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

The impact of academic dishonesty is far-reaching and is considered a serious offense against the university. All incidences of academic misconduct will be reported to the Office of Academic Integrity and could result in outcomes such as failure on the assignment, failure in the course, suspension, or even expulsion from the university.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

Students and Disability Accommodations:

USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Support Systems:

Counseling and Mental Health - (213) 740-9355 – 24/7 on call

Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

988 Suicide and Crisis Lifeline - 988 for both calls and text messages – 24/7 on call

The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline is comprised of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL) – 24/7 on call

Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

Office for Equity, Equal Opportunity, and Title IX (EEO-TIX) - (213) 740-5086

Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-5086 or (213) 821-8298

Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

The Office of Student Accessibility Services (OSAS) - (213) 740-0776

OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

USC Campus Support and Intervention - (213) 740-0411

Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity, Equity and Inclusion - (213) 740-2101

Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call

Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-1200 – 24/7 on call

Non-emergency assistance or information.

Office of the Ombuds - (213) 821-9556 (UPC) / (323-442-0382 (HSC)

A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

Occupational Therapy Faculty Practice - (323) 442-2850 or otfp@med.usc.edu

Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.