

EE 599: Error-Correcting Codes for Distributed Storage

Fall 2016

Lec: MW 8:00-11:20 am

3 units

Course Description

Given the current interest in gathering and processing an ever-increasing amount of data and the move towards cloud storage, a single data center where much of this data is actually stored, may be required to store several billion gigabytes of information. Within the data center, information pertaining to a single file is stored in a distributed manner, across multiple storage units. Each storage unit or node in the network, is prone to failure, could be down for maintenance or simply be busy, responding to other concurrent demands for data. To handle node failure or the non-availability of a node, the data center stores information in redundant fashion, using techniques ranging from simple replication to the more sophisticated use of erasure-correcting codes. There is considerable recent interest in keeping to a minimum, the amount of storage overhead needed to ensure reliability against data loss. While traditional codes such as Reed-Solomon codes do well in this respect, they are not the most efficient, when it comes to dealing with node repair and non-availability of a node.

This shortcoming has led to the recent introduction of two new classes of error-correcting codes, known respectively as regenerating codes and locally-recoverable codes. This class will begin by providing the requisite background on error-correcting codes before going on to present these recent exciting developments in coding theory. The application of these new classes of codes to the distributed storage of data will also be discussed.

Website

Blackboard: <https://blackboard.usc.edu/>

Course Material

- Class notes and journal articles

Reference Text Book

- Ron M. Roth, *Introduction to Coding Theory*, Cambridge University Press, 2007, ISBN: 978-0-521-84504-5.

Instructor:**Professor P. Vijay Kumar**

Email: pvk1729@gmail.com

Phone: (213) 740 4668

Class Location: EEB 110

Office: EEB 534

Office Hours: 2:30-3:30 pm MW

Prerequisite EE 441 and EE 503 (or equivalent).**Teaching Assistants and Graders TBA****Grading Policy (will be announced in class, but will likely be)**

Midterm Exam,	50%
Final Exam	50%

List of Topics:

<u>Week 1:</u>	<u>Lec 1</u>	Binary Block Codes, Groups, Rings and Fields
	<u>Lec 2</u>	Finite Fields
<u>Week 2:</u>	<u>Lec 3</u>	Linear binary codes, generator and parity-check matrices; the dual code, the Hamming code
	<u>Lec 4</u>	Systematic codes, minimum distance; bounds on code size
<u>Week 3:</u>	<u>Lec 5</u>	Standard-array decoder; decoding in the presence of erasures
	<u>Lec 6</u>	Reed-Solomon and RAID codes
<u>Week 4:</u>	<u>Lec 7</u>	Midterm Exam (90 minutes); Regenerating Codes: motivation,
	<u>Lec 8</u>	Bound on file size, storage-repair bandwidth tradeoff
<u>Week 5:</u>	<u>Lec 9</u>	MBR code example, product-matrix construction
	<u>Lec 10</u>	High-Rate MSR Codes, other recent developments
<u>Week 6:</u>	<u>Lec 11</u>	Locally recoverable codes, bound on minimum distance, the pyramid code
	<u>Lec 12</u>	The Windows Azure code, a general code construction
<u>Week 7:</u>	<u>Lec 13</u>	Recent work including codes for multiple erasures
	<u>Lec 14</u>	Course Review and Final Exam (90 minutes)

*(Note: first meeting Wed, Aug 2,4 in EEB 539 at 8 am,**make-up class (for Mon Aug 22) on Fri, Aug 26, 8-11:20, in EEB 110)*

Statement on Academic Conduct and Support Systems

Academic Conduct

Plagiarism – presenting someone else’s ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in *SCampus* in Section 11, *Behavior Violating University Standards* <https://scampus.usc.edu/1100-behavior-violating-university-standards-and-appropriate-sanctions/>. Other forms of academic dishonesty are equally unacceptable. See additional information in *SCampus* and university policies on scientific misconduct, <http://policy.usc.edu/scientific-misconduct/>.

Discrimination, sexual assault, and harassment are not tolerated by the university. You are encouraged to report any incidents to the *Office of Equity and Diversity* <http://equity.usc.edu/> or to the *Department of Public Safety* <http://capsnet.usc.edu/departments/public-safety/online-forms/contact-us>. This is important for the safety whole USC community. Another member of the university community – such as a friend, classmate, advisor, or faculty member – can help initiate the report, or can initiate the report on behalf of another person. *The Center for Women and Men* <http://www.usc.edu/student-affairs/cwm/> provides 24/7 confidential support, and the sexual assault resource center webpage sarc@usc.edu describes reporting options and other resources.

Support Systems

A number of USC’s schools provide support for students who need help with scholarly writing. Check with your advisor or program staff to find out more. Students whose primary language is not English should check with the *American Language Institute* <http://dornsife.usc.edu/ali>, which sponsors courses and workshops specifically for international graduate students. *The Office of Disability Services and Programs* http://sait.usc.edu/academicsupport/centerprograms/dsp/home_index.html provides certification for students with disabilities and helps arrange the relevant accommodations. If an officially declared emergency makes travel to campus infeasible, *USC Emergency Information* <http://emergency.usc.edu/> will provide safety and other updates, including ways in which instruction will be continued by means of blackboard, teleconferencing, and other technology.