

## CLLOUD COMPUTING

Syllabus for the course [CS 4365/CS 5354](#), Spring 2015

Instructor: Vladik Kreinovich, office CCSB 3.0404, email [vladik@utep.edu](mailto:vladik@utep.edu), office phone (915) 747-6951

Class time: TR 1:30-2:50 pm, CCSB 1.0204

Office hours: Tuesdays 12:30-1:30 pm, Tuesdays and Thursdays 3-4:30 pm, or by appointment.

### MAIN OBJECTIVES:

- to teach the foundations of cloud computing, and
- to enable students to design and use the cloud.

**CONTENTS** In traditional computing, the user decides where to store the data and which processor (s) should process the data. The user's selections are not always optimal -- since the users may not have access to appropriate data storage and data processing resources, and even if the user has this access, the user may have the skills to optimally allocate these resources. The main idea behind cloud computing is to let sophisticated algorithms decide where the data is stored and where it is processed, the user just sends his or her data to the "cloud".

This make life easier for the user, but for us computer scientists who have to develop the corresponding algorithms, this setting leads to important challenges.

**Planning and scheduling.** First, we need to decide where to place the servers on which the data will be stored, on which server to place the data provided by each user, which processor(s) to use for processing the data, when to relocate this data, etc. Algorithms for doing this planning and scheduling are already very sophisticated, many of the details are so important that the companies that run commercial clouds do not disclose these details. In the class, we will overview the corresponding problems and the main ideas behind the planning and scheduling algorithms used in cloud computing.

**Growth.** Servers are a big investment. When locating a server for storing and/or processing data, we need to take into account not only the current needs, but also the expected future needs for such services. It is therefore important to be able to accurately predict the growth of the internet, the growth of the tasks that can be handled by the cloud. Such prediction techniques will be overviewed in the class.

**Privacy and security.** When the user stores the data on his or her own computer, security is largely the responsibility of the user. When the data is stored on the cloud, maintaining security and privacy of this data is the cloud's problem. The cloud must have built-in mechanisms for guaranteeing security and privacy -- otherwise, no one will use the cloud's services. In the class, we will discuss methods for guaranteeing such security and privacy.

**Cost.** The cloud does not make computers run faster or perform the computations more accurately, its main advantage is that it can allow users to save money. Thus, when designing and maintaining a cloud, we must pay special attention to the issues that are often overlooked in computer science -- the issues of cost. The cloud must be beneficial for the user, and it must be profitable for the company that provides the cloud services. In the class, we will discuss how such a win-win situation can be achieved.

**Self-healing.** Even when we run a simple computer network -- like a departmental network -- we need full-time system administrators to make sure that everything works. For a cloud, that supports several orders of magnitude more computations, the problem of reliability becomes even more important. It is therefore desirable to make a system that would automatically detect if one of the computers failed and would then re-route all the computations so that these failures do not affect the users. In the class, we will discuss machine learning and other techniques that are used to detect the faults, and techniques for re-routing.

**Parallelization.** A cloud has a large number of processors, so it is desirable to take advantage of this number by parallelizing the computing tasks as much as possible. Parallelization is difficult even when we have a fixed number of computers with fixed connections -- e.g., most modern computers have 4 or more processors with a potential of working in parallel, but most compilers under-utilize this parallelization potential. The problem becomes even more complex if we have thousands of processors at different locations. In the class, we will discuss basic algorithms for cloud-related parallelization.

**Green computing.** The last but not the least are the issues of energy consumption and environmental impact. Already up to 15-20% of the overall energy consumption goes into computer activities. This may not be very noticeable when we work on a single laptop, but this becomes crucial when we consider locations where large amounts of cloud information are stored and processed. In the class, we overview the main techniques for estimating and minimizing the energy impact of cloud computing.

**Practice.** Just like with any other type of programming, it is difficult to learn cloud computing without an experience of storing data in the actual cloud. We plan to have such an experience as a part of the course.

### Topics:

- what is cloud computing, the problem of optimal server location
- main terms of cloud computing
- growth models, least squares techniques for predicting growth
- security, RSA algorithm
- possibility of parallelization, parallel algorithms, the main ideas of MapReduce
- Bayesian approach to fault detection
- clustering, "jungle" computing, green computing
- practical cloud computing

**PROJECTS.** An important part of the class is a project. There are three possible types of projects:

- The most common type of a project is that you can select a technical paper related to cloud computing, a paper that interests you, and report about what you learned from this paper -- both orally and in written form. Feel free to form teams, this way you can it will be easier to handle papers. Make sure that the paper has technical contents -- that it has either formulas or code.
- If, as part of your current research project, you are interested in doing some research related to cloud computing, a research that will be useful for your future thesis, project, or dissertation, this will be an excellent project for the class.
- Finally, if you have not yet selected any topic for your thesis or dissertation, and you think that maybe something related to cloud computing will be a good idea for such topic, then we may start your research -- and this will be a very good project for the class.

**TESTS AND GRADES:** There will be two tests and one final exam. Each topic means home assignments -- both theoretical and programming. Maximum number of points:

- first test: 10
- second test: 25
- home assignments: 10
- final exam: 35
- project: 20

(smart projects with ideas that can turn into a serious scientific publication get up to 40 points).

A good project can help but it cannot completely cover possible deficiencies of knowledge as shown on the test and on the homeworks. In general, up to 80 points come from tests and home assignments. So:

- to get an A, you must gain, on all the tests and home assignments, at least 90% of the possible amount of points (i.e., at least 72), and also at least 90 points overall;
- to get a B, you must gain, on all the tests and home assignments, at least 80% of the possible amount of points (i.e., at least 64), and also at least 80 points overall;
- to get a C, you must gain, on all the tests and home assignments, at least 70% of the possible amount of points (i.e., at least 56), and also at least 70 points overall.

**STANDARDS OF CONDUCT:** Students are expected to conduct themselves in a professional and courteous manner, as prescribed by the Standards of Conduct. Students may discuss programming exercises in a general way with other students, but the solutions must be done independently. Similarly, groups may discuss project assignments with other groups, but the solutions must be done by the group itself. Graded work should be unmistakably your own. You may not transcribe or copy a solution taken from another person, book, or other source, e.g., a web page). Professors are required to - and will - report academic dishonesty and any other violation of the Standards of Conduct to the Dean of Students.

**DISABILITIES:** If you feel you may have a disability that requires accommodation, contact the The Center for Accommodations and Support Services (CASS) at 747-5148, go to Room 106 E. Union, or [e-mail to cass@utep.edu](mailto:cass@utep.edu). For additional information, please visit the [CASS website](#).