The lecture "TMA4265: Stochastic Processes" is roughly divided into four (five) sections: discrete-time Markov chains, Poisson processes, continuous-time Markov chains, queuing theory, (Brownian motion). At the end of the first course segment the students should be able to

- explain what a discrete-time Markov chain is. You should be able to formulate a Markov chain and derive the corresponding transition graph linking the states of the state space.

- compute state probabilities and classify the states into equivalence classes and determine important class properties, such as transience, recurrence (positive and null) or periodicity.

- apply first-step analysis to answer questions such as in which state the Markov chain is ultimately trapped and how long, on average, it takes to reach this state.

- investigate long-term behavior of a Markov chain and distinguish between stationary and limiting probabilities.

- compute the expected time a Markov chain spends in transient states

- recognize and describe a branching process, determine if it converges to infinity or dies out, and compute certain quantities of interest.

- derive the transition probabilities of the reversed Markov chain, decide whether the process is time-reversible and exploit time-reversibility to compute limiting probabilities.

- simulate a Markov chain and check key-quantities derived theoretically using simulations. You should have a rough idea about Markov chain Monte Carlo procedures.

By the end of the second part the students will be able to

- define and illustrate a Poisson process, and compute key quantities such as arrival or inter-arrival times.

- work with the Poisson and exponential distribution, exploit their properties in the computation of conditional probabilities, and evaluate the superposition and intersection of several Poisson processes.