"Applying the cavity method to design efficient load balancing schemes"

Marc Lelarge INRIA, Paris

Motivated by load balancing and caching problems in large scale distributed systems, such as peer-to-peer video-on-demand systems, we introduce a simple model of random bipartite graphs and study spanning subgraphs with degree constraints. Through a rigorous formulation of the cavity method inspired from the objective method of Aldous and Steele, we analyze the optimal size of such subgraphs and relate it to the efficiency of the system. Our analysis allows a characterization of content replication policies that minimize the inefficiency. These optimal policies, which differ markedly from proportional placement, have a simple structure which makes them implementable in practice.

Joint work with Charles Bordenave, Justin Salez, Mathieu Leconte and Laurent Massoulie