STEVE KIRKLAND, University of Regina, Regina, Saskatchewan
Conditioning of the Entries in Google's PageRank Vector

The internet search engine Google approaches the problem of ranking web pages by computing an estimate of the Perron vector of a certain non-negative matrix associated with the world wide web. That Perron vector, known as the PageRank vector, is then used to rank the importance of the corresponding web pages. In this talk, we discuss the conditioning of the entries in the PageRank vector, and our results in turn lead to error bounds for Google's computed estimate of that vector. Our techniques involve non-negative matrices, generalized inverses, and combinatorial considerations.