

Optimal Score Aggregation Algorithms

Ronald Fagin
IBM Research – Almaden
fagin@us.ibm.com

ABSTRACT

Assume that there is a set of “voters” and a set of “candidates”, where each voter assigns a numerical score to each candidate. There is a scoring function (such as the mean or the median), and a consensus ranking is obtained by applying the scoring function to each candidate’s scores. The problem is to find the top k candidates, while minimizing the number of database accesses. The speaker will present an algorithm that is optimal in an extremely strong sense: not just in the worst case or the average case, but in every case! Even though the algorithm is only 10 lines long (!), the paper containing the algorithm won the 2014 Gödel Prize, the top prize for a paper in theoretical computer science.

CCS Concepts

• Theory and algorithms for application domains → Database theory → Data structures and algorithms for data management


Keywords

Aggregation algorithms; score aggregation; instance optimal

Short Bio

Ronald Fagin is an IBM Fellow at IBM Research – Almaden. IBM Fellow is IBM’s highest technical honor. There are currently around 90 active IBM Fellows (out of around 400,000 IBM employees worldwide), and there have been only around 250 IBM Fellows in the over 50-year history of the program. Ron received his B.A. in mathematics from Dartmouth College and his Ph.D. in mathematics from the University of California at Berkeley.

He is a Fellow of IEEE, ACM, and AAAS (American Association for the Advancement of Science).



He has co-authored four papers that won Best Paper Awards and three papers that won Test-of-time Awards, all in major conferences. He was named Docteur Honoris Causa by the University of Paris. He won the IEEE Technical Achievement Award, IEEE W. Wallace McDowell Award, and ACM SIGMOD Edgar F. Codd Innovations Award (a lifetime achievement award in databases). He is a member of the US National Academy of Engineering and the American Academy of Arts and Sciences.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author(s). Copyright is held by the owner/author(s).

PODS'16, June 26–July 1, 2016, San Francisco, California, USA.

ACM 978-1-4503-4191-2/16/06.

DOI: <http://dx.doi.org/10.1145/2902251.2902308>