CORR 2000-02

Baby-Step Giant-Step Algorithms for Non-uniform Distributions

Simon R. Blackburn* & Edlyn Teske

Abstract The baby-step giant-step algorithm, due to Shanks, may be used to solve the discrete logarithm problem in arbitrary groups. the paper explores a generalisation of this algorithm, where extra baby steps may be computed after carrying out giant steps (thus increasing the giant step size). the paper explores the problem of deciding how many, and when, extra baby steps should be computed so that the expected cost of the generalised algorithm is minimised. When the logarithms are uniformly distributed over an interval of length n, the expected cost of the generalised algorithm is 6the expense of a slightly larger worst case cost). In some situations where logarithms are far from uniformly distributed, any baby-step giant-step algorithm that computes all its baby steps before taking a giant step must have infinite expected cost, but the generalised algorithm has finite expected cost. the results are heuristic, but are supported by evidence from simulations.