ABSTRACT. Shanks' baby step-giant step algorithm is known to be a generic method for computing element orders and discrete logarithms in any finite abelian group. We show how this method can be optimized in various applications where more is known about the underlying group than just the group operation itself. Additional information includes situations when the distribution of the solution in a given search interval is known, when symmetries in the search interval or less expensive inversions exist, and when baby steps are less expensive than giant steps. In each case, this additional information can be readily used to obtain a speed-up of the algorithm. We apply the optimized methods to the computation of invariants of hyperelliptic function fields. Hereby, we are able to verify the speed-up in an important application.