Modeling Round-off Error in the Fast Gradient Method for Predictive Control

Ian McInerney & Eric C. Kerrigan & George A. Constantinides *(Imperial College London)*

We present a method to determine the smallest data type for the Fast Gradient Method (FGM) to converge when solving a linear Model Predictive Control (MPC) problem in fixed-point arithmetic. We derive two models for the round-off error present in fixed-point arithmetic: a generic model and a structure-exploiting parametric model. We also propose a metric for measuring the amount of round-off error the FGM can tolerate before diverging. Using this metric and the round-off error models, we compute the minimum number of fractional bits needed for the fixed-point data type. We show that the structure-exploiting parametric model nearly halves the number of fractional bits needed to implement an example problem, significantly decreasing the resource usage for an implementation on a Field Programmable Gate Array (FPGA).

Recently, MPC has grown in popularity due to its ability to incorporate operating constraints in the computation of an optimal control action. At its core, MPC solves an optimization problem to determine the next control action, which for the Constrained Linear Quadratic Regulator (CLQR) is a Quadratic Program (QP). This popularity has led to MPC being implemented on smaller processors and FPGAs that do not contain hardware to accelerate floating-point computations.

The lack of floating-point acceleration necessitates the use of fixed-point representation, where the number is stored using a fixed number of bits seperated into two segments: integer bits and fraction bits. When implementing the QP solver in fixed-point, the number of fractional bits must be large enough to ensure that the round-off error does not cause the FGM iterations to diverge. Prior work has shown that the convergence of the FGM in fixed-point is dependent upon the eigenvalues of the fixed-point representation of the QP's Hessian lying in the range $(0, 1)$.

In this work we present a new metric, called the rounding stability margin, that uses the pseudospectrum of the Hessian to compute how much round-off error can be introduced before the FGM iterations diverge. Since the Hessian H of the QP is normal, the pseudospectrum can provide direct information about the eigenvalue spectrum. We compute the ϵ-pseudospectrum of H at the points $\lambda \in \{0, 1\}$ to determine how large a perturbation H must undergo to shift its spectrum to include those points. This is then used to bound the spectrum of an additive perturbation matrix on H that represents the round-off errors.

We present two models for the round-off error introduced by moving the Hessian into a fixed-point representation. The first is a generic model that finds the worst-case round-off error for any CLQR problem by using the worst-case perturbation for every element of H. The second is a parametric model that exploits the Toeplitz structure of the Hessian for Schur-stable systems to find the cutoff diagonal beyond which all elements of H round to zero. The round-off error is then modeled as the worst-case for all diagonals before the cutoff, and as the actual element value for all diagonals after the cutoff.

These round-off error models are then used to compute the spectrum of the additive perturbation matrix representing the transformation to fixed-point. Combining these models with the rounding stability margin allows for the computation of the number of fractional bits required for the FGM to be convergent. We demonstrate that exploiting the structure of the MPC problem can reduce the number of fractional bits needed by 30–45%, and allow for a reduction in hardware usage and solution time by up to 77% and 25%, respectively, for an implementation of FGM on a FPGA.