

FPGA Accelerators for Robust Visual SLAM on Humanoid Robots

Maria Rafaela Gkeka, Alexandros Patras, ECE

Department, University of Thessaly

Nikolaos Tavoularis, Stylianos Piperakis, Emmanouil

Hourdakis, Panos Trahanias, *Foundation for Research and Technology - Hellas (FORTH)*

Christos D. Antonopoulos, Spyros Lalis, Nikolaos

Bellas, *ECE Department, University of Thessaly*

Contact: margkeka@uth.gr

Visual Simultaneous Localization and Mapping (vSLAM) is the process of mapping the robot's observed environment using an optical sensor, while concurrently determining the robot's pose with respect to that map. For humanoid robots, the implementation of vSLAM is particularly challenging, due to the intricate motions of the robot. In this work, we present a pose graph optimization module based on RGB features, as an extension on the KinectFusion pipeline (a well-known vSLAM algorithm), to help recover the robot's pose during unstable gait patterns where the KinectFusion tracking system fails. We implement and evaluate a plethora of embedded MPSoC FPGA designs and we explore several architectural optimizations, both precise and approximate, highlighting their effect on performance and accuracy. Properly designed approximations, which exploit domain knowledge and efficient management of CPU and FPGA fabric resources, enable real-time vSLAM (at more than 30 fps) in humanoid robots without compromising robot tracking and map construction. We show that a combination of precise and approximate optimizations and tuning of algorithmic parameters provide a speedup of up to 15.7X and 22.5X compared with the precise FPGA and ARM-only implementations, respectively, without violating the tight accuracy constraints.

Keywords: SLAM; Humanoid Robots; FPGA; Approximate Computing

DOI: <http://dx.doi.org/10.1145/3490422.3502331>