先端科学技術研究科　修士論文要旨

所属研究室 (主指導教員)	コンピューティング・アーキテクチャ (中島　康彦　(教授))		
学籍番号	2411010	提出日	令和 8年 1月 16日
学生氏名	飯田　昌直		
論文題目	HIBIKI: An FPGA Accelerator for Large-Scale TSP with Batched Parallel Local Search and Hierarchical Divide-and-Conquer HIBIKI: バッチ並列局所探索と階層的分割統治を用いた大規模TSP向けFPGAアクセラレータ		

要旨

Edge robots and drones demand real-time, high-quality solutions to large-scale Traveling Salesman Problems (TSP) under tight power and on-chip memory constraints. I present HIBIKI, a hierarchical parallel TSP solver realized through hardware-algorithm co-design and implemented as an FPGA accelerator, and I validate it end-to-end on an AMD Versal VMK180 board. HIBIKI recursively decomposes an instance into fixed-size open-tour (path) subproblems capped at 512 nodes to match BRAM/URAM capacity, enabling predictable memory access and deep pipelining. Each subproblem is solved using a crossover-free (1+λ)-EA-style iterated local search (ILS): λ = 16 offspring are generated from a single parent, improved via batched parallel 2-opt, and evaluated with a k-nearest-neighbor candidate set (k = 16) to reduce search cost. To sustain exploration without stalling the hardware pipeline, HIBIKI triggers a constant-time double-bridge perturbation upon stagnation. I implement 32 parallel solver cores on the VMK180 operating at 300 MHz. Across TSPLIB instances with 1,002-7,397 nodes, HIBIKI delivers 9-22× higher throughput than a clustering-based divide-and-conquer GA baseline. On the largest instance pla7397, the prototype reaches a sub-10% gap to the TSPLIB optimum in about 0.1 seconds. I further report time-to-quality curves based on optimality gap over time, as well as post-route timing and resource utilization. Using total board power measurements, I show that HIBIKI improves energy-delay product (EDP) and time-to-quality compared to CPU-based baselines.