

# Optimus: A Parallel Optimization Framework With Topology Aware PSO and Applications

Sarat Sreepathi  
North Carolina State University  
sarat\_s@ncsu.edu

Advisor: G.Kumar Mahinthakumar  
North Carolina State University  
gmkumar@ncsu.edu

**Abstract**—This research presents a parallel metaheuristic optimization framework, Optimus (Optimization Methods for Universal Simulators) for integration of a desired population-based search method with a target scientific application. Optimus includes a parallel middleware component, PRIME (Parallel Reconfigurable Iterative Middleware Engine) for scalable deployment on emergent supercomputing architectures. Additionally, we designed TAPSO (Topology Aware Particle Swarm Optimization) for network based optimization problems and applied it to achieve better convergence for water distribution system (WDS) applications. The framework supports concurrent optimization instances, for instance multiple swarms in the case of PSO. PRIME provides a lightweight communication layer to facilitate periodic inter-optimizer data exchanges. We performed scalability analysis of Optimus on Cray XK6(Jaguar) at Oak Ridge Leadership Computing Facility for the leak detection problem in WDS. For a weak scaling scenario, we achieved 84.82% of baseline at 200,000 cores relative to performance at 1000 cores and 72.84% relative to one core scenario.

## I. OPTIMUS

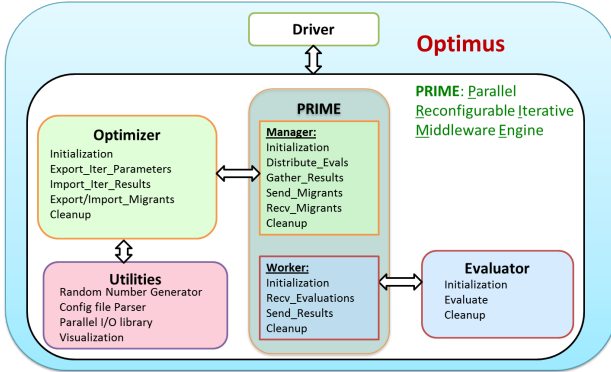


Fig. 1: Architecture of the Optimus framework

The architecture of Optimus is shown in Figure 1. Optimus facilitates seamless integration of optimizer and simulation components on high performance computing (HPC) platforms. It is designed to support multiple optimizer instances (islands or swarms) as required by the optimization algorithm.

Figure 2 illustrates the algorithm-system mapping of Optimus when using multiple cooperative swarms. To the left, the diagram shows the system architecture of a Cray XK6 supercomputer and the corresponding mapping of optimization algorithm components is shown on the right. This is a decentralized scheme where all swarms optimizers are peers of equal

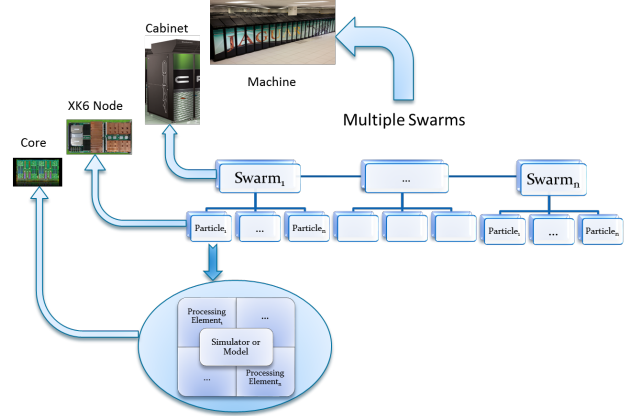


Fig. 2: Cooperative Multi-Swarm Optimization Framework with algorithm to system architecture mapping

status. Periodically, swarms communicate search information (best solution found etc.) with each other.

## II. WATER DISTRIBUTION SYSTEMS APPLICATIONS

Water Distribution Systems (WDSs) are the primary supplier of safe drinking water to the public. They form a critical infrastructure that is vulnerable to accidental and intentional contamination incidents that could result in adverse human health and safety impacts. When a contamination event is detected via the first line of defense, e.g., a water quality surveillance sensor, the municipal authorities need to rapidly identify contamination source characteristics (such as location, time of release etc.) and undertake mitigation actions. Furthermore, leaks due to breakage of pipes etc. can be potential intrusion points for contamination. Leaks also lead to significant loss of water during distribution before reaching the customers resulting in significant economic losses for the utilities. Therefore contaminant source characterization and leak detection are important issues for water supply utilities.

## III. TAPSO

We designed a new Particle Swarm Optimization (PSO) technique, TAPSO (Topology Aware PSO) for network-based optimization problems. Standard PSO is designed to operate in real valued space where the optimization parameters are real numbers. These ideas can be extended to discrete variables where the search space is suitably represented through

integers. But this may not result in good performance for discrete parameters when the search space is governed by a topology as adjacent numbers may still be further apart in search space. The TAPSO algorithm is designed to incorporate such topology/connectivity information and combine it with canonical PSO paradigm to design a technique that is more amenable for graph based optimization problems. We conducted a comparative analysis of standard PSO and TAPSO on the leak detection problem for a large WDS network with 12,457 nodes (Figure 3) and found TAPSO to be more effective.




Fig. 3: TAPSO: Visualization results for leak detection problem in WDS. The color coding is as follows: sensors (orange), true leak (red), best estimate so far (green) and candidate solution (cyan).

#### IV. PARALLEL PERFORMANCE RESULTS

The computational platform used for parallel performance analysis study is Jaguar, the Cray XK6 supercomputer at Oak Ridge Leadership Computing Facility.

##### A. Weak Scaling: Problem Setup

We performed a weak scaling study of the Optimus framework using the WDS simulator as the target application. The WDS network used in this study is Micropolis, a medium-sized water network with 1,834 nodes. We used a simulated scenario with four concurrent leaks resulting in eight optimization parameters, where four leak locations are discrete parameters and corresponding leak magnitudes are continuous parameters.

Figure 4 describes the weak scaling performance from 1000 cores (evaluators) to 200,000 cores. In each scenario, we divided the total number of processes into groups, where every 1,000 evaluators constitute a group and communicate with corresponding Swarm optimizer. For instance, there are 10 swarms (each 1000 cores) at 10,000 cores and 200 such

swarms at 200,000 cores. Swarm optimizers share search information between each other using the inter-optimizer communicator that connected them in a ring topology every five iterations. In this case, they share information every five iterations. The total number of iterations in these test cases is 100. Each swarm in turn contains 2,000 particles which translates to two WDS simulations per core per iteration. We achieved 84.82% of baseline at 200,000 cores relative to performance at 1000 cores.




Fig. 4: Optimus: Weak Scaling Performance - X-axis represents the number of cores and number of swarms in each case are indicated in parenthesis.

Figure 5 depicts the number of simulations completed per minute in each case for the weak scaling problem setup. Evidently, the number of simulations increases as the problem size increases.




Fig. 5: Weak Scaling: Number of simulations per minute

#### V. CONCLUSIONS

The Optimus framework can serve as a platform for algorithmic exploration, specifically multi-population techniques such as multi-island genetic algorithms and multi-swarm methods. Furthermore, we have developed a decentralized manager-worker model using asynchronous communication primitives and adapted it for optimization methods.

In this work, we tried to demonstrate that such techniques could scale computationally to a very large number of cores. Future work includes further investigation of collaborative parallel search algorithms for solving previously intractable optimization problems. The TAPSO algorithm could potentially be adapted for other network-based optimization problems, for instance in Epidemiology, Cybersecurity, Logistics and Geospatial domains.