

Otter: Towards Data-Driven Taskification of Parallel Programs

A. Tuft¹, T. Weinzierl¹, A. Chalk², L. Morgenstern¹, R. Ford²

¹ Dept. of Computer Science, Durham University, ² STFC Hartree Centre, Daresbury Laboratory
December 7, 2023

- ▶ Will tasks pay off in my code?
- ▶ Can I do better than trial-and-error when adding tasks?
- ▶ Do task runtimes make reasonable task scheduling decisions?
- ▶ Can user knowledge about the task graph be used to improve schedulers?

PROJECT GOALS

1. Determine task graph at runtime (without major performance costs).
2. Runtime-agnostic task-tracing API.
3. Scheduling simulator to explore possible task schedules.
4. Estimate potential benefits of tasking.
5. Feed scheduling insights back into runtime.



www.exahype.org



<https://swift.strw.leidenuniv.nl/>

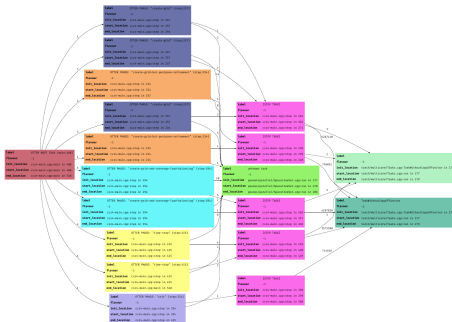


<https://github.com/stfc/PSyclone>



1. Tracing API:
 - direct annotation
 - OMPT plugin
 - Fortran wrapper
2. Post-processing & visualisation
3. Prototype simulator
4. Tracing demonstrated with realistic science workload (ExaHyPE SBH benchmark, 31M events).

```
EnclaveTask::EnclaveTask(int taskTypeID)
: Task(taskTypeID) {
    OTTER_DEFINE_TASK(enclave_task,
    OTTER_NULL_TASK,
    otter_add_to_pool,
    "[enclave task (type=%d)]",
    taskTypeID);
}

bool EnclaveTask::run() {
    OTTER_REMOVE_FROM_POOL(enclave_task,
    "[enclave task (type=%d)]",
    getTaskId());
    OTTER_TASK_START(enclave_task);
    computeTask();
    OTTER_TASK_END(enclave_task);
    return false;
}
```



- ▶ ExaHyPE: Exascale hyperbolic PDE engine (OpenMP+MPI).
- ▶ Once per timestep, spawn low-priority, ready "enclave" tasks.
- ▶ Intended to allow overlap with MPI communication.
- ▶ Observation: task runtime simply consumes ready tasks, so threads later left spinning.



(a) OpenMP tasks consumed eagerly, so few left during communication phase and threads left to spin (in red).

Figure: From H. Schulz, G. Gadeschi, O. Rudy, T. Weinzierl: Task Inefficiency Patterns for a Wave Equation Solver. IWOMP 2021

Next

- ▶ Tracing interface
 - unify direct & OMPT event sources
 - add Fortran wrapper
- ▶ Prototype scheduling simulator
 - predict execution time & experiment with scheduling algorithms
 - validate on simple benchmarks [→] realistic workloads

Later

- ▶ Simulator refinements
 - NUMA effects, user-specified task scheduling algorithm
- ▶ Simulator-driven analysis of where tasks pay off
 - Recommend where to add tasks in realistic application & predict speedup
- ▶ Improve runtime performance
 - Explore whether task-graph information can be used to improve runtime performance.
- ▶ Accelerator support
 - When do *offloaded* tasks pay off?

- ▶ The authors acknowledge the support through the embedded CSE programme of the ARCHER2 UK National Supercomputing Service (<http://www.archer2.ac.uk>) under grant no ARCHER2-eCSE04-2, Durham's oneAPI Academic Centre of Excellence made by Intel, ExCALIBUR's Phase 1a grant ExaClaw (EP/V00154X/1) and ExCALIBUR's cross-cutting project EX20-9 Exposing Parallelism: Task Parallelism (grant ESA 10 CDEL).
- ▶ <https://excalibur.ac.uk/projects/exposing-parallelism-task-parallelism/>
- ▶ Joint work by Durham's Department of Computer Science, Department of Physics, Durham's Advanced Research Computing and the Hartree Centre

