DynaSOAr: A Parallel Memory Allocator for Object-Oriented Programming on GPUs with Efficient Memory Access (Artifact)

Matthias Springer
Tokyo Institute of Technology, Japan
matthias.springer@acm.org

Hidehiko Masuhara
Tokyo Institute of Technology, Japan
masuhara@acm.org

Abstract

This artifact contains the source code of DynaSOAr, a CUDA framework for Single-Method Multiple-Objects (SMMO) applications. SMMO is a type of object-oriented programs in which parallelism is expressed by running the same method on all applications of a type.

DynaSOAr is a dynamic memory allocator, combined with a data layout DSL and a parallel do-all operation. This artifact provides a tutorial explaining the API of DynaSOAr, along with nine benchmark applications from different domains. All benchmarks can be configured to use a different memory allocator to allow for a comparison with other state-of-the-art memory allocators.

2012 ACM Subject Classification Software and its engineering → Allocation / deallocation strategies; Software and its engineering → Object oriented languages; Computer systems organization → Single instruction, multiple data

Keywords and phrases CUDA, Data Layout, Dynamic Memory Allocation, GPUs, Object-oriented Programming, SIMD, Single-Instruction Multiple-Objects, Structure of Arrays

Digital Object Identifier 10.4230/DARTS.5.2.2

Acknowledgements This work was supported by a JSPS Research Fellowship for Young Scientists (KAKENHI Grant Number 18J14726). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the TITAN Xp GPU used for this research.

Related Article Matthias Springer and Hidehiko Masuhara, "DynaSOAr: A Parallel Memory Allocator for Object-Oriented Programming on GPUs with Efficient Memory Access", in 33rd European Conference on Object-Oriented Programming (ECOOP 2019), LIPIcs, Vol. 134, pp. 17:1–17:37, 2019.
`https://dx.doi.org/10.4230/LIPIcs.ECOOP.2019.17`

Related Conference 33rd European Conference on Object-Oriented Programming (ECOOP 2019), July 15–19, 2019, London, United Kingdom

1 Scope

The main purpose of the artifact is (1) to support reproducibility of the experiments in the original paper and (2) to support programmers in using DynaSOAr in their own applications. In particular, this artifact provides experiments that support the following claims from the paper:

- SMMO applications that use DynaSOAr exhibit superior runtime performance and space efficiency compared to other GPU allocators (mallocMC, Halloc, CUDA allocator).
- DynaSOAr exhibits low fragmentation and is able to utilize almost the entire heap for allocations.
- As shown by our application examples, many programs can be expressed in SMMO.

2 Content

This artifact includes:

- The CUDA source code of DynaSOAr, the evaluation benchmarks and the infrastructure to run the benchmarks with other allocators.
- A description of all benchmark applications, highlighting their SMMO structure.
- A tutorial describing how to setup the artifact, how to use DynaSOAr, and how to interpret the program output in debug mode, which reflects the internal state of the allocator.
- Instructions for recreating a subset of the experiments.

3 Getting the Artifact

The most recent source code of DynaSOAr and the benchmark applications is available at: `https://github.com/prg-titech/dynasoar/`.

The remainder of the artifact is stored in the GitHub Wiki pages of the project: `https://github.com/prg-titech/dynasoar/wiki`.

We provide a zip file with the source code and a PDF export of the GitHub Wiki with this artifact.

4 Tested Platforms

The artifact was tested on Ubuntu 16.04.1 with three GPUs: NVIDIA GeForce 940MX, NVIDIA TITAN Xp and NVIDIA GeForce 1050 TI. A GPU with a minimum compute capability of 5.x and the CUDA Toolkit in version 9 or higher are required. Detailed software/hardware requirements can be found in the tutorial.

5 MD5 Checksum

The MD5 checksum of the zip file is e7b8bb33961ae66209581e2c58c6cad3.

6 Size of the Artifact

The size of the zip file is 3 MB.

7 License

DynaSOAr is licensed under the MIT License. DynaSOAr depends on the *CUB* library which is licensed under the BSD license. To allow for an easy setup, this artifact contains the source code of two other allocators, MallocMC and Halloc. These allocators are licensed under the BSD License. See the `LICENSE` file for details.