Chair of Decentralized Systems Engineering
Department of Informatics

Persistent memory - GPU synergy and memory model analysis

General information	**Advisor**	Dimitrios Stavrakakis
	Email	dimitrios.stavrakakis@tum.de
	Date	17.01.2022
Type	Master / Bachelor / Guided Research	
Description	Persistent memory[1] is a recently added layer in the system stack of data centers. It resides on the memory bus and provides read and write access in byte granularity with latencies close to DRAM. Along with these features, it is also a non-volatile storage medium. This unique combination of features has opened up a way for a new programming model targeting to manipulate this medium and reap its benefits. Having that in mind, Intel developed PMDK[2], a collection of libraries and tools for system administrators and application developers to simplify managing and accessing persistent memory devices. PM files are mapped into the virtual address space of an application and can be directly accessed via conventional load/store instructions. This property opens up the way to discover the PM synergies with state-of-the-art technologies like GPUs. So far, GPUs are leveraged to offload demanding computations from the CPU and perform them in a highly-parallel, high-performance manner. CPU-GPU communication is achieved through memory mapped regions. With the introduction of PM	

	these regions can be stateful, but the GPU programming models currently do not incorporate persistency semantics which are crucial when interacting with PM. Towards this direction, a recent work [3] has proposed a programming model for GPU-PM systems. Another work explores memory persistency models for GPUs [4]. In this project, we are aiming to delve deeper into the synergies of GPU with PM, analyze their communication, examine further use-cases that can be benefitted by leveraging PM as a non-volatile medium, and strive to formally verify the memory model [5] to prove the correctness of the proposed approaches.
Keywords	Persistent memory, GPU, memory models, PMDK, formal verification
Goals	**Concrete outcomes** 1. Analysis of PM-GPU communication model 2. Formal verification of the memory model to prove the correctness of the PM-GPU programming model 3. Extension of the use-cases of hybrid PM-GPU systems
Prerequisites	**Compulsory** ● Basic operating system knowledge ● Good understanding of memory models & system components communication ● Previous experience with GPU programming **Preferred** ● Experience with PM programming based on PMDK ● Knowledge of memory models & formal verification methods
References	1. https://docs.pmem.io/persistent-memory/getting-started-guide/introduction 2. https://github.com/pmem/pmdk 3. https://github.com/csl-iisc/GPM-ASPLOS22 4. https://people.engr.ncsu.edu/hzhou/gpu_pm_pact19_final.pdf 5. https://dl.acm.org/doi/10.1145/2694344.2694391

Application process

Please send an email to the advisor including the following:

- Email subject: “Thesis application (DSE)”
- CV
- A copy of your transcript(s)
- A **motivation statement**, please include samples of your work that you are proud of (e.g., major projects, open-source contributions, Github page, etc.) and/or writing samples (e.g., your technical blog, project reports, etc.)