

Memento

Energy-Aware Memory Placement in Operating Systems

Sven Köhler^{1*}, Benedict Herzog^{2*}, Henriette Hofmeier², Manuel Vögele², Lukas Wenzel¹
*both authors contributed equally to this work

¹ Hasso Plattner Institute, Potsdam


² Ruhr University Bochum

Motivation

- New memory technologies (NVM, HBM, ...), cell-types (e.g. PCRAM), and coherent interconnects (e.g. CXL) challenge existing system and programming abstractions of a homogenous memory space
- Like functional and non-functional properties, all memory technologies highly differ in their energy demand
- The OS community needs to react and enable optimisation strategies, while leveraging energy saving potentials

Memory Energy Model

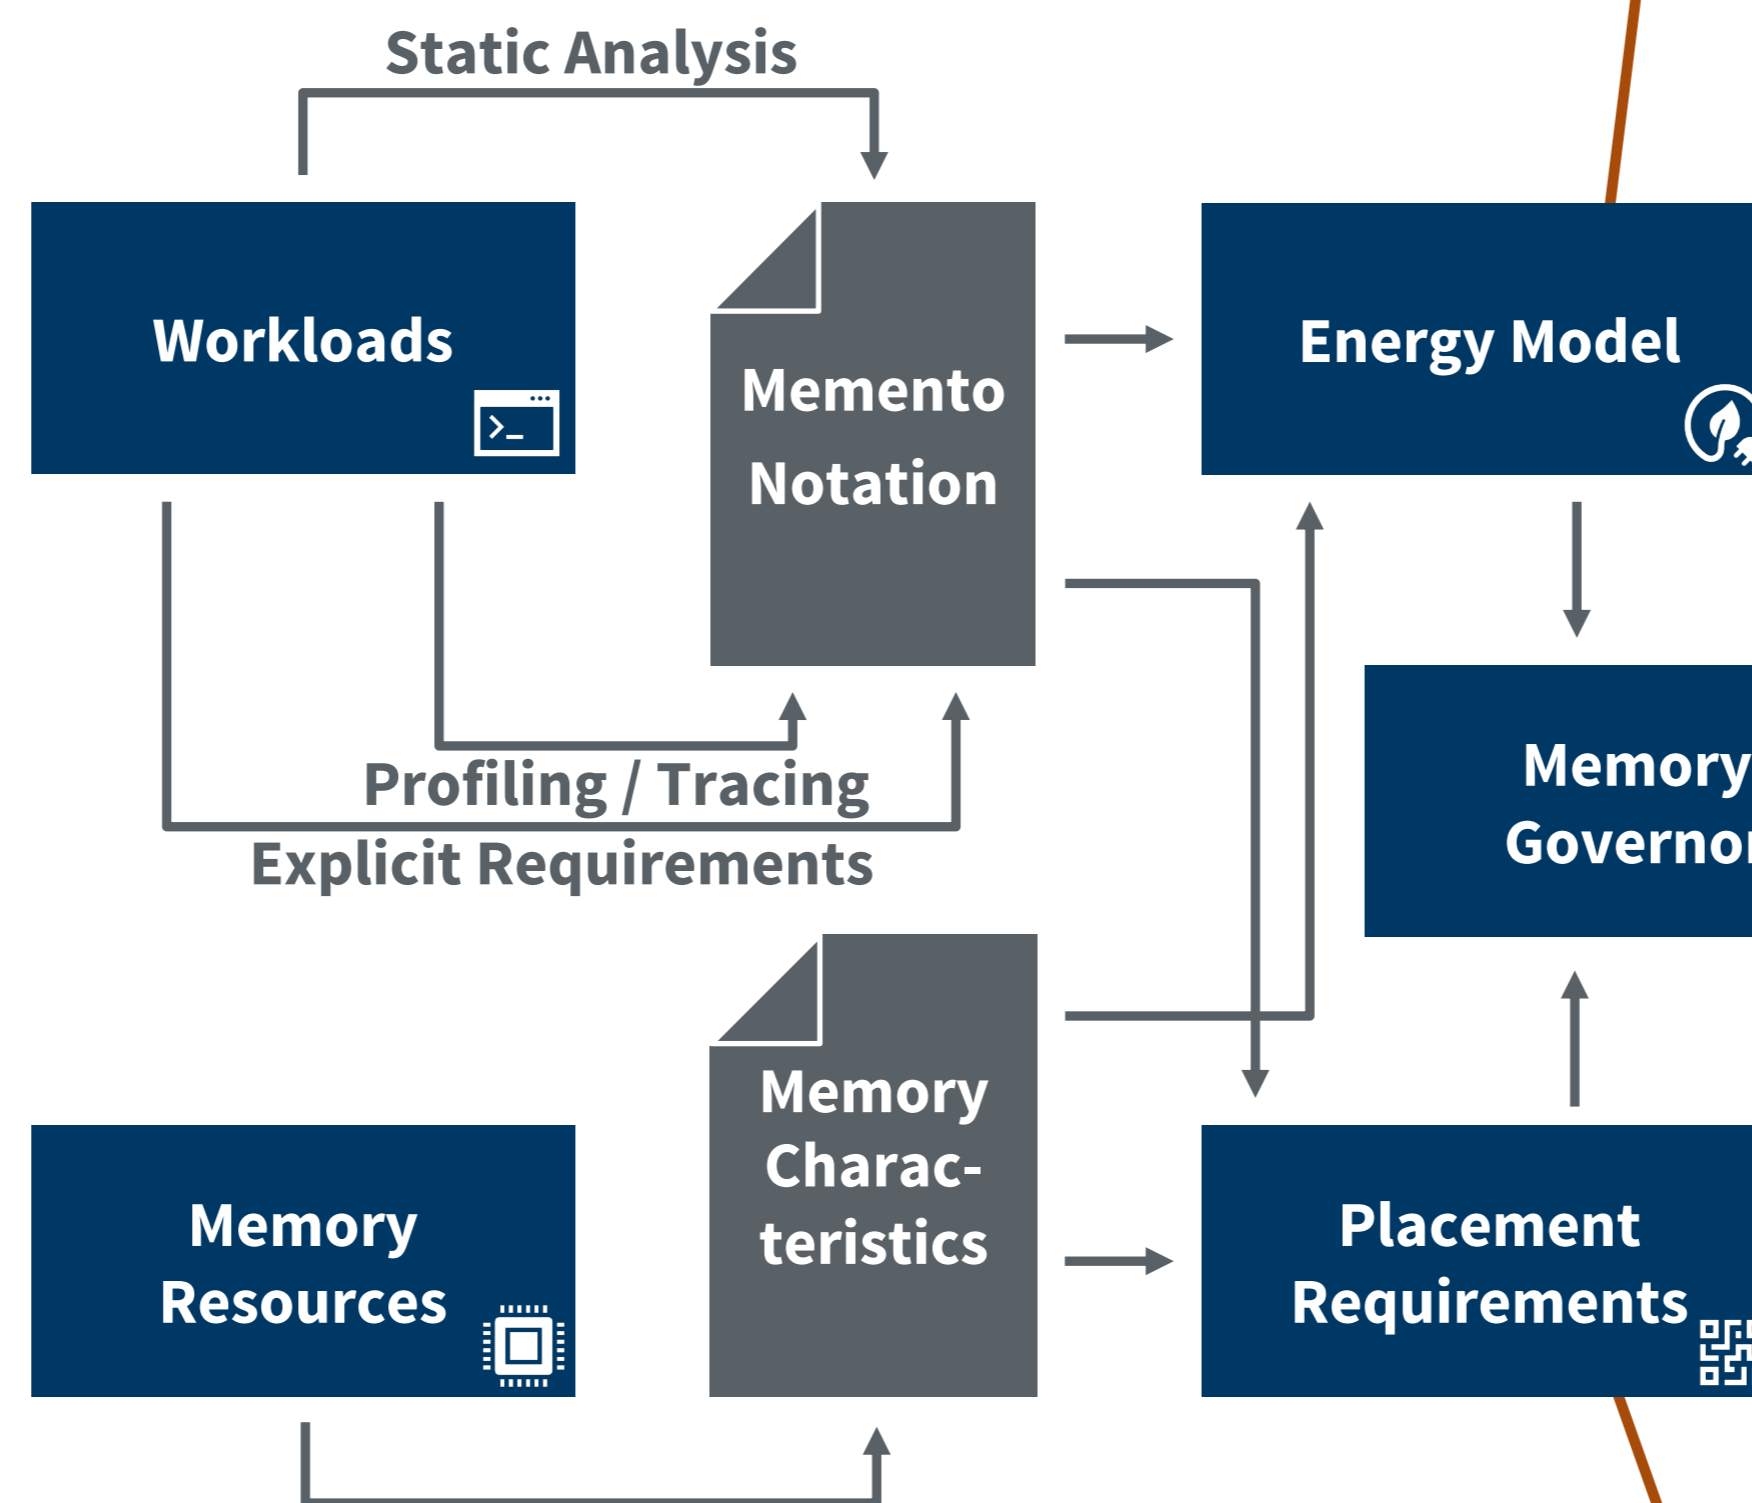
- Model of potential placement impact, based on:



- Approaches with ML techniques, i.e. linear, ensemble, neural networks
- We will evaluate expressiveness and accuracy for each, and additional costs for training and inference at runtime

The Memento Approach

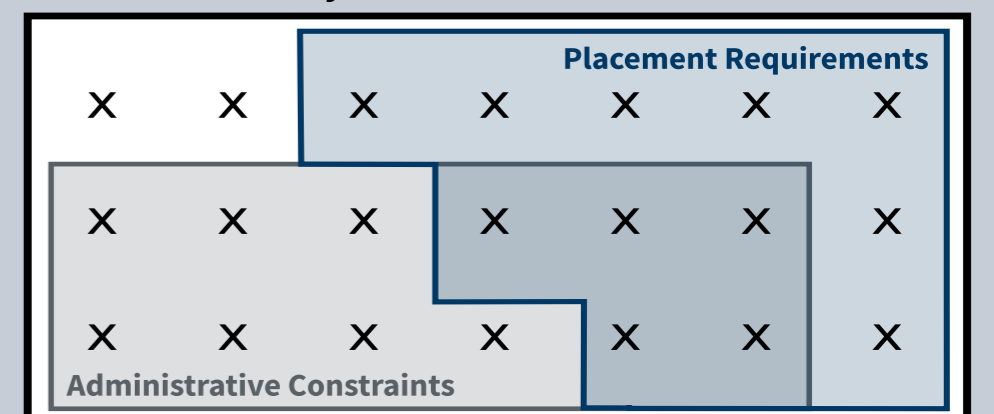
- The central **Memory Governor** matches allocation requests with the best available memory resource
- Using an energy model we try to minimize the entire system's energy demand at runtime
- We determine a workload's **sensitivity**, i.e. how it reacts to different **memory types**' bandwidth, latencies, or volatility
- Workload classification from code analysis, tailored micro benchmarks, and profiling at run-time [2]
- Deduced requirements are denoted in an exchangeable format, the **Memento Notation** (e.g. persisted across runs in dedicated ELF sections)



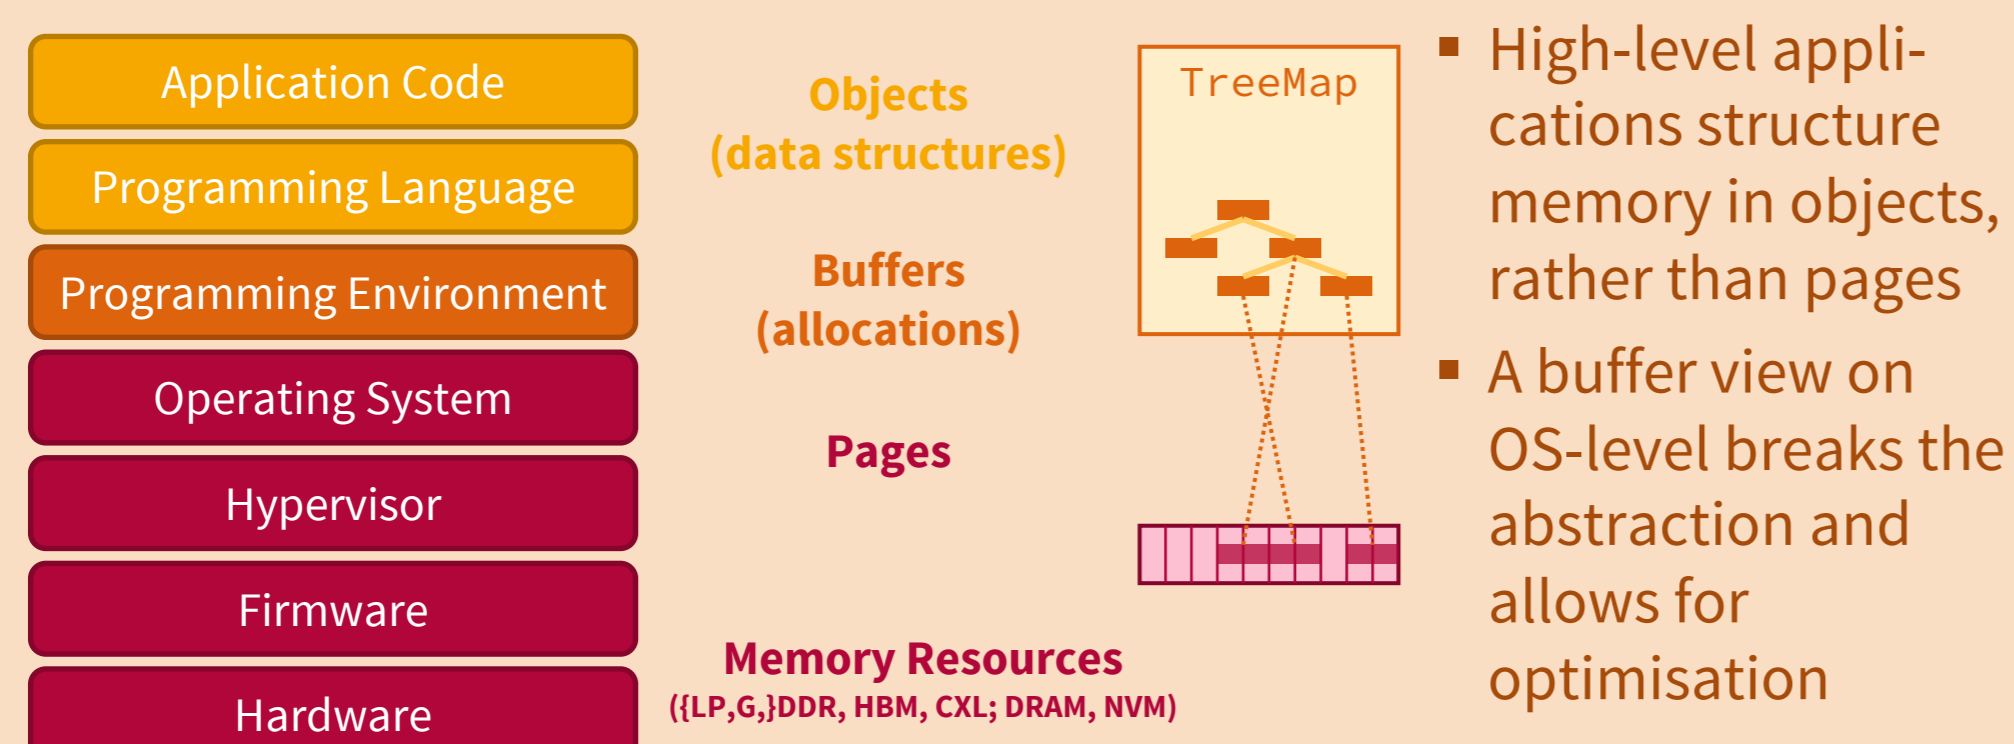
Memory Governor

- Part of the OS (service + module)
- Keeps track of available / free memory hardware resources
- Chooses most energy-efficient buffer location, considering availability, placement requirements, and administrative constraints
- Inspired by current frequency governors for compute devices
- Works on allocation granularity (not threads or processes)

Available Memory Resources



Choosing The Right Granularity



Placement Requirements

- Developers are experts on the dynamics of their application, not on available system resources
- Legacy applications do not necessarily make (good) placement decisions for new memory resource types

- First, offer **explicit API** with manually annotated sensitivity as starting point, e.g. with scoped allocators
- For legacy workloads, use the **implicit** knowledge persisted in the Memento Notation

```
AllocCharacteristic LatencyNVSpec {
    // uninitialised fields default to 0
    weight_latency = 3.f,
    weight_randomAccess = 1.f,
    non_volatile = true
};
{
    // all allocations in this scope are
    // associated with LatencyNVSpec
    MemGuard guard(LatencyNVSpec);
    Index * jumpListA = new Index[4096];
}
```

Carbon-Efficient Placement

- Memento tools, notation and benchmarks also applicable in other fields, e.g. a carbon-efficient placement following the SCI specification [1]:

$$\text{Software Carbon Intensity} = \frac{\text{Operational Carbon} + \text{Embodied Carbon}}{\text{Unit of Work}}$$

- case study (gCO₂e/alloc) best highlighted

	10 μs allocation 1 access volatile			50 ns allocation 100,000 accesses persistent			10 s allocation 5 accesses volatile		
	DRAM	NVDIMM	Optane	DRAM	NVDIMM	Optane	DRAM	NVDIMM	Optane
Wind	1.3e-16	1.5e-16	1.6e-11	7.3e-13	1.6e-6	1.3e-10	1.5e-10	8.0e-11	
Mix	1.6e-15	1.6e-15	5.4e-10	8.2e-12	5.4e-5	1.6e-9	1.6e-9	2.7e-9	
Gas	2.0e-15	2.1e-15	6.9e-10	1.0e-11	6.9e-5	2.0e-9	2.1e-9	3.4e-9	
Coal	4.2e-15	4.2e-15	1.4e-9	2.1e-11	1.4e-4	4.2e-9	4.2e-9	7.1e-9	

References

1. Green Software Foundation. 2021. Software Carbon Intensity Standard. https://github.com/Green-Software-Foundation/sci/blob/main/Software_Carbon_Intensity/Software_Carbon_Intensity_Specification.md
2. Julia Lawall and Gilles Muller. 2018. Coccinelle: 10 Years of Automated Evolution in the Linux Kernel. In *Annual Technical Conference (ATC'18)*. USENIX, 601–614

