

# Design-Time Memory Subsystem Optimization for Low-Power Multi-Core Embedded Systems

Manuel  
Strobel

# CONTEXT AND MOTIVATION

- Complex and compute-intensive applications motivate the use of multi-core devices
- Limitations in the embedded field call for design concepts that allow for example to:
  - Respect predictability and real-time constraints
  - Minimize energy and power consumption
- Focus on the optimization potential of the memory subsystem in MPSoC platforms:
  - Utilization of scratchpads
  - Code/data placement
  - Utilization of memory low-power modes
- Objective: Automated optimization workflow

# PROBLEM STATEMENT

- Assuming a given MPSoC platform that defines the memory architecture
- Application-specific optimization method:
  - Applied at system design-time
  - Based on memory access information (e.g. from Instruction Set Simulation)

Three-fold problem formulation can be given as follows:

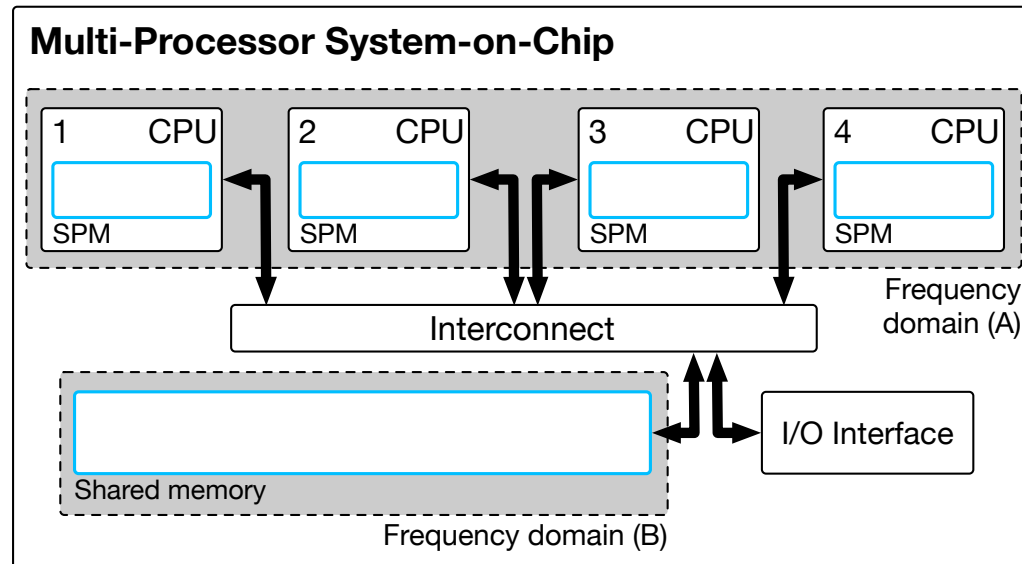
1. Find an **application to memory binding  $\beta$**  to reduce the **dynamic energy** consumption from memory accesses.
2. Determine a **memory operation mode schedule  $\gamma$**  based on low-power modes for a **static energy** consumption reduction.
3. Ensure the **implementation** of these results on the **software level**.

# CONTENT

1. MPSoC architecture and system model
2. Optimization workflow
  - Overview
  - Application to memory binding  $\beta$
  - Memory operation mode scheduling  $\gamma$
  - Integration of optimization results
3. Experimental results
4. Summary and Outlook

# MPSOC ARCHITECTURE AND SYSTEM MODEL

- All processing units are equipped with core-local scratchpad (SPM)
- Shared memory on the global level
- All cores have access to all memories (core-local, remote, and global)
- Different frequency domains

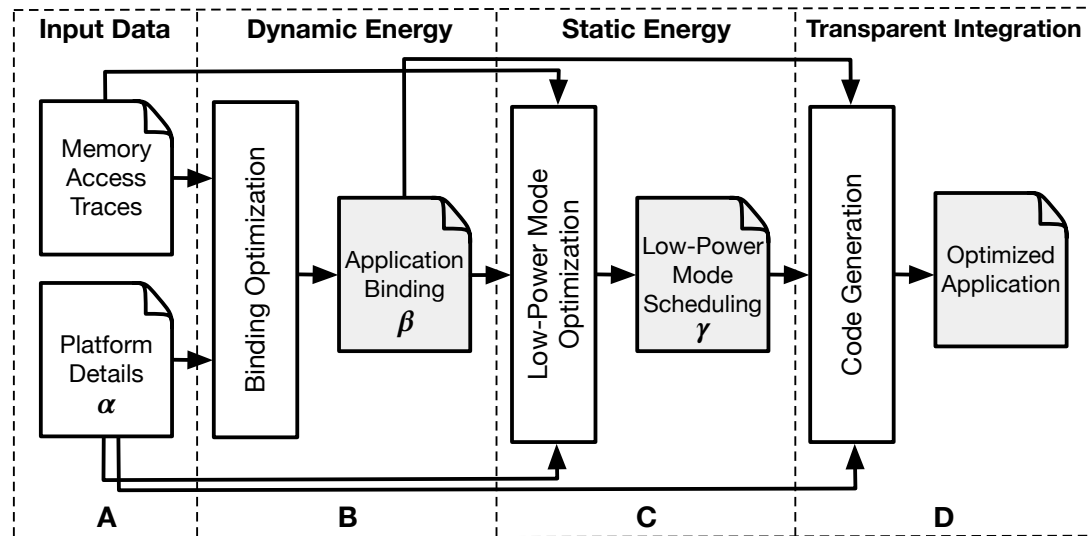


In line with:  
[1], [2], and [3]

# OPTIMIZATION WORKFLOW

## Overview

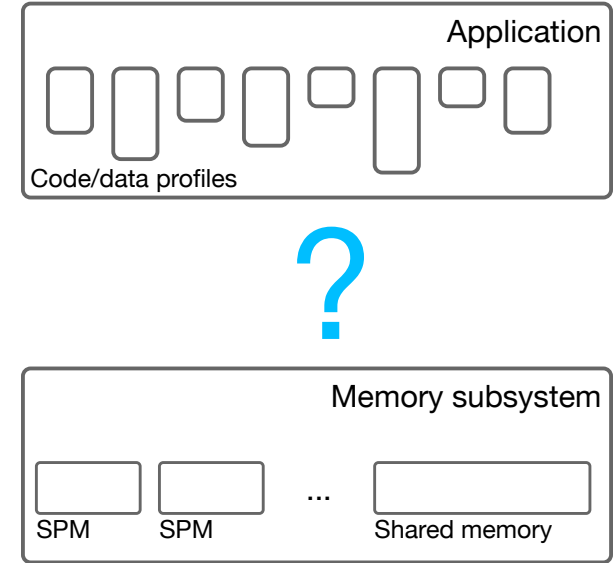
- A) Input data includes platform and simulation details
- B) First optimization step for binding optimization targets dynamic energy consumption
- C) Second, subsequent step for static energy optimization using low-power modes
- D) Integration into software using a compiler backend tool on the basis of LLVM



# OPTIMIZATION WORKFLOW

## Application to memory binding $\beta$



- Realized as Integer Linear Program (ILP)
- Map application profiles (code/data) to memory blocks
- Respect the following constraints:
  - Ensure memory blocks are large enough for content
  - Bind every code profile at least to one memory
  - Bind every data profile to one and only one memory
- Other rules specify energy and time consumption
- Optimization goal: Minimize energy consumption due to memory accesses
- Detailed mathematical notation of the ILP is given in the paper



# OPTIMIZATION WORKFLOW

## Memory operation mode scheduling $\gamma$

- Realized as Mixed-Integer Quadratic Program (MIQP)
- Find one operation mode vector per code profile
- Respect the following constraints:
  - Assign one and only one configuration per profile
  - All required memories (per profile) must be active
- Other rules compute static energy consumption:
  - Depending on the operation mode vector
  - Plus energy penalty from mode switching
- Optimization goal: Minimize consumed static energy
- Detailed model is again given in the paper




Operation mode vector example:  
(ACT, LS, ... , DS)



# OPTIMIZATION WORKFLOW

## Integration of optimization results

- Integrated with compiler → LLVM backend extension (more details available in [4])
- Basic concept: Software-programmable (memory) configuration registers




- Compiler backend is used for insertion of code snippets for:
  - Code and data placements (using linker directive `.section`)
  - Operation mode changes

```
push {r5,r6}      @ secure utilized registers
ldr r5, =MCR      @ load core-local MCR address
ldr r6, =MODE     @ load power mode of profile
str r6, [r5,#0]  @ update MCR accordingly
pop {r5,r6}      @ restore registers
```

# EXPERIMENTAL RESULTS

- Experiments for ARM quad-core platform and benchmarks from MiBench suite [5]
- Two possible operation scenarios: run-to-idle (top) and run-to-sleep (bottom)...




# SUMMARY AND OUTLOOK

- Two-stage optimization workflow for memory subsystems in MPSoC design
  - Includes dynamic and static energy consumption minimization
  - Provides optimal code/data placement
  - Impact of memory low-power modes depends on application and operation scenario
- Automated solution for efficient utilization of scratchpads in multi-core embedded systems
- Identified future work:
    - Investigation of applications with higher degree of inter-core communication
    - Co-operation of scratchpads and traditional caches
    - Integration of peak power limitations

# REFERENCES

- [1] Infineon Technologies AG: 32-bit aurix microcontroller based on tricore (2019).  
URL <https://www.infineon.com/cms/en/product/microcontroller/32-bit-tricore-microcontroller/>
- [2] Casini, D., Biondi, A., Nelissen, G., Buttazzo, G.: Memory feasibility analysis of parallel tasks running on scratchpad-based architectures. In: 2018 IEEE Real-Time Systems Symposium (RTSS) (2018). DOI 10.1109/RTSS.2018.00047
- [3] Gu, S., Zhuge, Q., Yi, J., Hu, J., Sha, E.H.: Optimizing task and data assignment on multi-core systems with multi-port spms. IEEE Transactions on Parallel and Distributed Systems 26(9) (2015). DOI 10.1109/TPDS.2014.2356194
- [4] Strobel, M., Radetzki, M.: A backend tool for the integration of memory optimizations into embedded software. In: Proc. of the 2019 Forum on Specification & Design Languages (FDL) (2019)
- [5] Guthaus, M.R., Ringenberg, J.S., Ernst, D., Austin, T.M., Mudge, T., Brown, R.B.: Mibench: A free, commercially representative embedded benchmark suite. In: Proc. of the 2001 IEEE International Workshop on Workload Characterization. IEEE (2001). DOI 10.1109/WWC.2001.990739



**University of Stuttgart**  
Chair of Embedded Systems

# Thank you!

**Manuel Strobel**

e-mail [manuel.strobel@informatik.uni-stuttgart.de](mailto:manuel.strobel@informatik.uni-stuttgart.de)

phone +49 (0) 711 685-88215

University of Stuttgart  
Chair of Embedded Systems  
Pfaffenwaldring 5b  
70567 Stuttgart  
Germany