

Memphis

Many-core Modeling Platform for Heterogenous SoCs

Memphis Manycore

Apresentação Angelo Elias Dalzotto

Coordenador: Prof. Dr. Fernando G. Moraes

Adaptado da apresentação de Dr. Marcelo Ruaro



GAPH-PUCRS

Overview

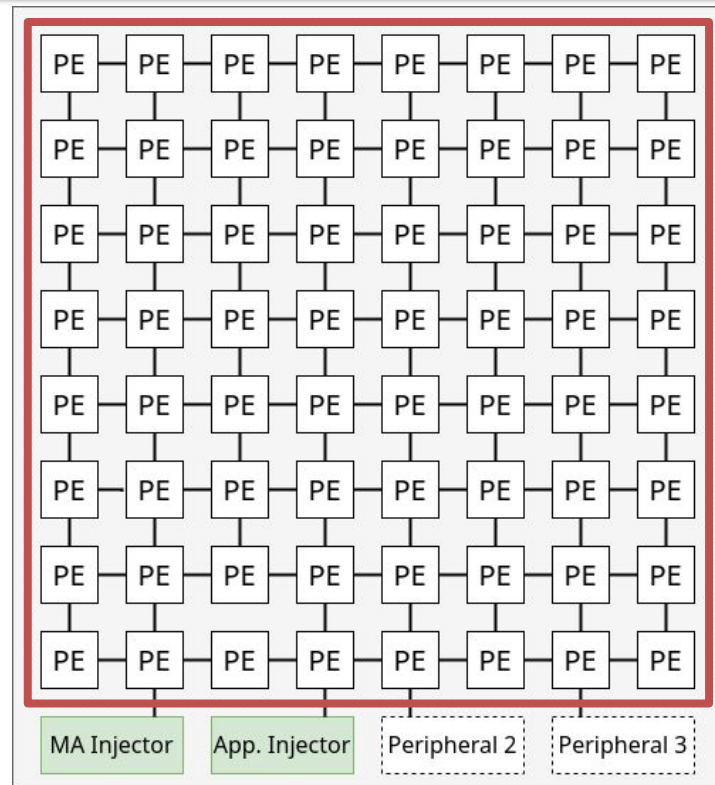
Memphis

Many-core Modeling Platform for Heterogeneous SoCs

3

GPPC – **Homogeneous** region

General Purpose Processing Core



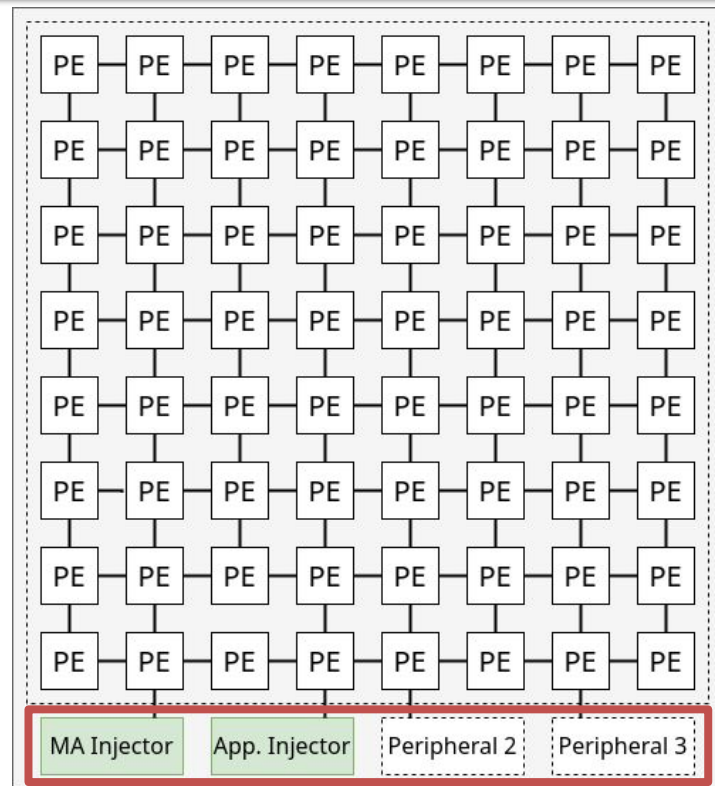
Memphis

Many-core Modeling Platform for Heterogeneous SoCs

4

Peripherals – **Heterogeneous** region

- Connected to **PE** borders
- Provides:
 - I/O Interface
 - Default: **Application Injector**
 - Default: **MA Injector**
 - Hardware acceleration



Hardware

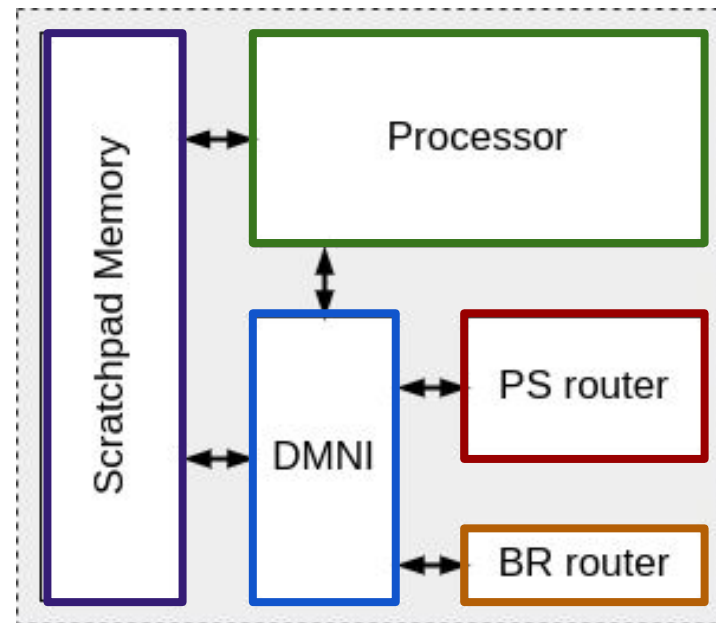
Memphis¹

Many-core Modeling Platform for Heterogeneous SoCs

6

PE – Homogeneous region
 Processing Element

- *RISC-V* **CPU**
- *DMNI* (DMA + NI)
- *Scratchpad* **memory**
- *Hermes* **router**
- *BrNoC* **router**



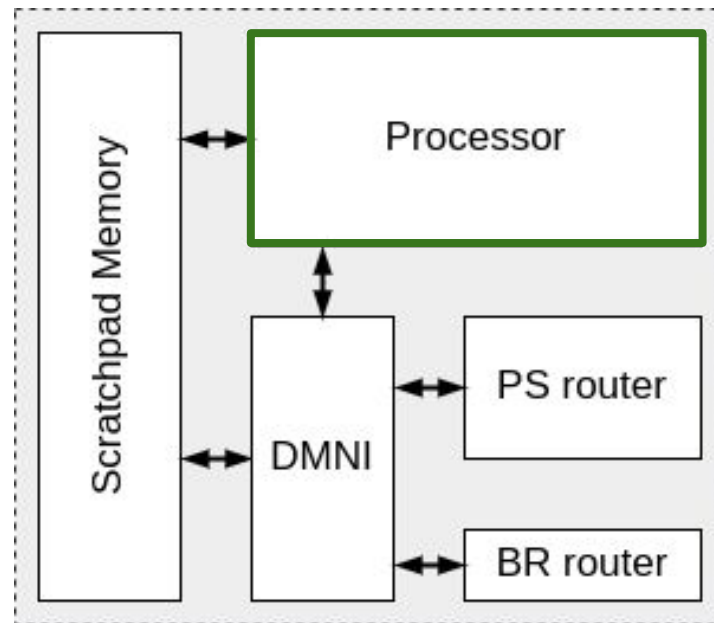
1. Ruaro, M., Caimi, L. L., Fochi, V., and Moraes, F. G. (2019). Memphis: a framework for heterogeneous many-core SoCs generation and validation. Design Automation for Embedded Systems, 23(3-4):103-122.

Processor

7

RISC-V CPU

- *RV32IM Instruction Set Architecture*
 - 32-bit
 - 32 registers
 - **No floating-point**
- Currently ISS-only
- Simple paging support




Local Memory

8

Scratchpad memory

- Dual-port
 - CPU access
 - DMNI access

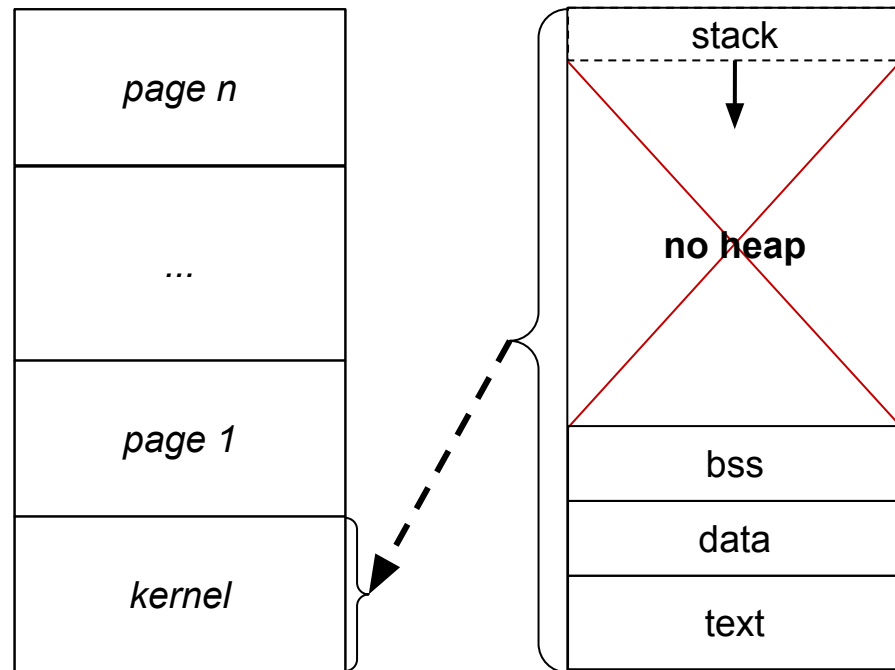


Local Memory

9

Scratchpad memory

- Dual-port
 - CPU access
 - DMNI access
- Parameterizable pages
 - Equal sizes
 - Currently no dynamic memory

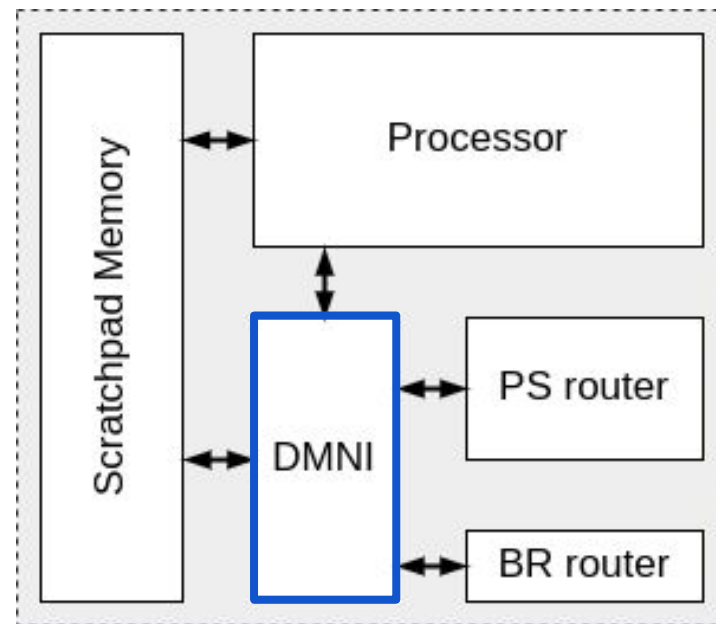


DMNI

10

Direct Memory Network Interface²

- Interface between local memory and NoC
- Specialized for NoC-based manycores



2. Ruaro, M., Lazzarotto, F. B., Marcon, C. A., and Moraes, F. G. (2016). DMNI: A specialized network interface for NoC-based MPSoCs. In Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS), pages 1202-1205.

Packet-Switching Router

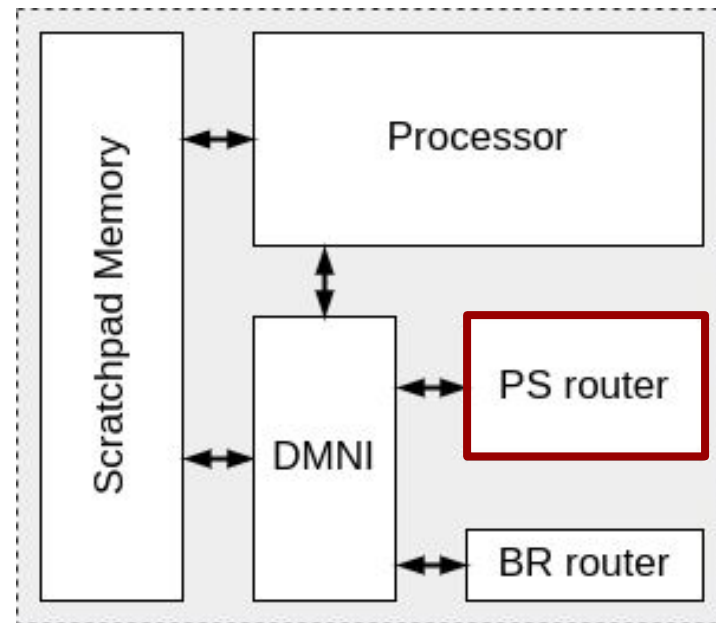
11

Hermes Router³

- XY addressing and routing
- Packet switching
- Wormhole credit-based control flow

Hermes has other derivations:

- Asynchronous
- Virtual-channel
- Circuit-switching
- ...



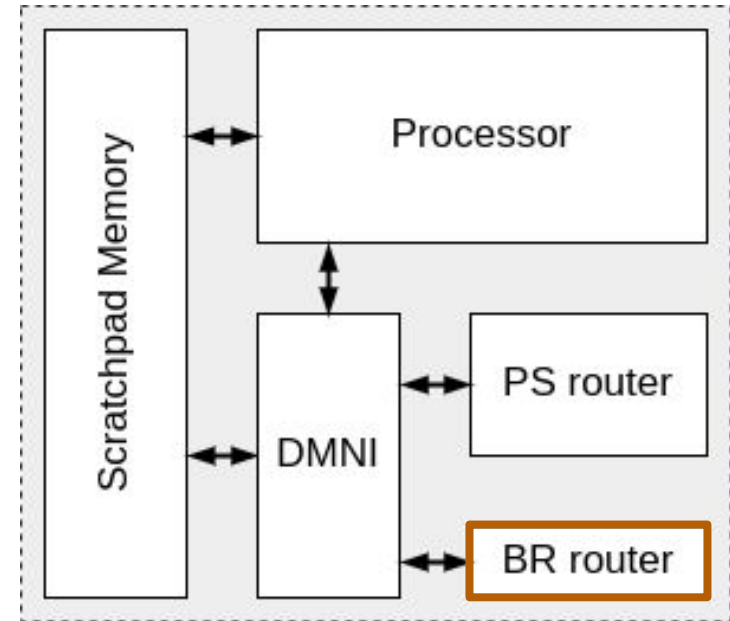
3. Moraes, F., Calazans, N., Mello, A., Möller, L., and Ost, L. (2004). HERMES: an infrastructure for low area overhead packet-switching networks on chip. *Integration*, 38(1):69-93.

Broadcast Router

12

BrNoC Router⁴

- Broadcast transmission only
 - Interrupts ALL or a TARGET PE
- Single-flit messages



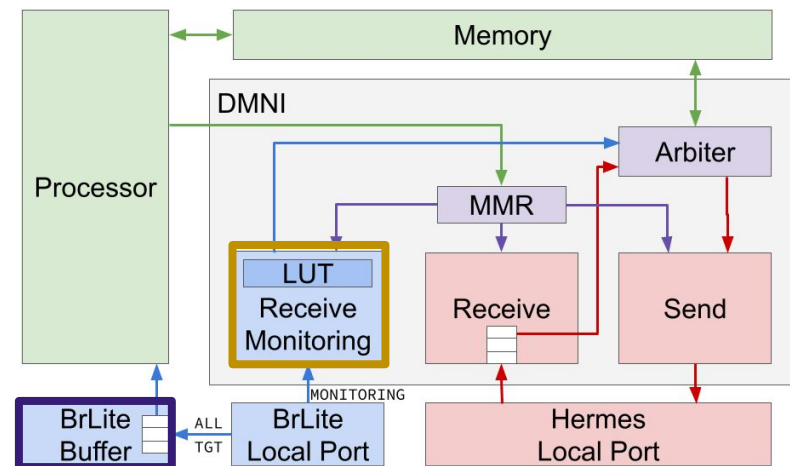
4. Wachter, E., Caimi, L., Fochi, V., Munhoz, D., Moraes, F. (2017). BrNoC: A broadcast NoC for control messages in many-core systems. *Microelectronics Journal*, 68:69-77.

Broadcast Router

13

BrNoC Router⁴

- Broadcast transmission only
 - Interrupts ALL or a TARGET PE
- Single-flit messages
- Buffered output
 - Transmission is faster than processing
- Monitoring messages through DMNI




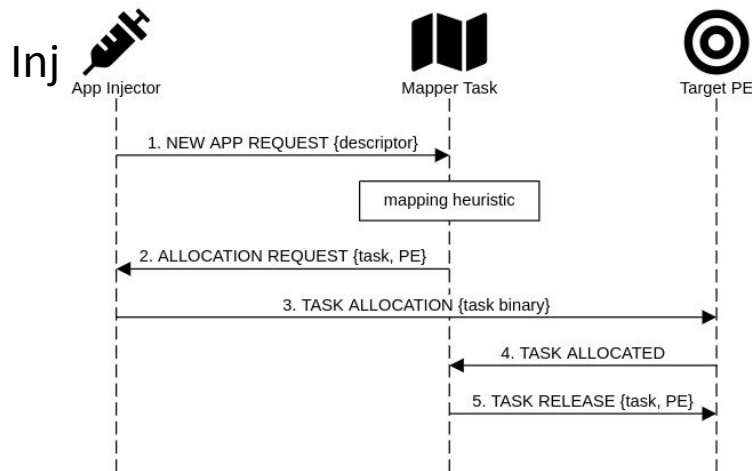
4. Wachter, E., Caimi, L., Fochi, V., Munhoz, D., Moraes, F. (2017). BrNoC: A broadcast NoC for control messages in many-core systems. *Microelectronics Journal*, 68:69-77.

Application Injector

14

Injectors

- Interacts with external off-chip memory
- Loads applications into the manycore
- Separated injection for secure content (MA




Software

Management Application⁵

16

- Division of management tasks
- **Observer tasks**
- **Decider tasks**
- **Actuator tasks**
- **Mapper task⁶**



5. Ruaro, M., Santana, A., Jantsch, A., and Moraes, F. G. (2021). Modular and Distributed Management of Manycore SoCs. ACM Transactions on Computer Systems (TOCS), 38(1-2):1-16.

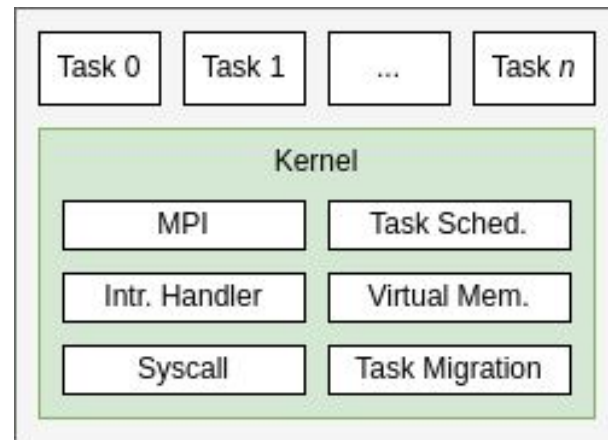
6. Dalzotto, A. E., Ruaro, M., Erthal, L. V., and Moraes, F. G. (2021). Dynamic Mapping for Many-cores using Management Application Organization. In Proceedings of the International Conference on Electronics, Circuits, and Systems (ICECS), pages 1-6.

Operating System

17

Kernel

- User tasks
 - Virtual Memory
 - Scheduler
- System Calls
- Message Passing Interface
- Actuation Mechanisms
 - Task migration support



Scheduler

18

Best effort

- Round-robin


Real-time⁷

- LST

Least Slack Time

NoC packet and message structure

19



Packet Header - seen by NoC:

- Target PE
- Payload Size

Packet Payload - seen by kernel:

- **Service**
- 10 configurable flits
- Optional payload with variable size

BrNoC message structure

20



Header - NoC view only:

- ID - retransmission control
- SVC - ALL/TARGET/CLEAR
- SRC and TGT
- PROD - producer task ID

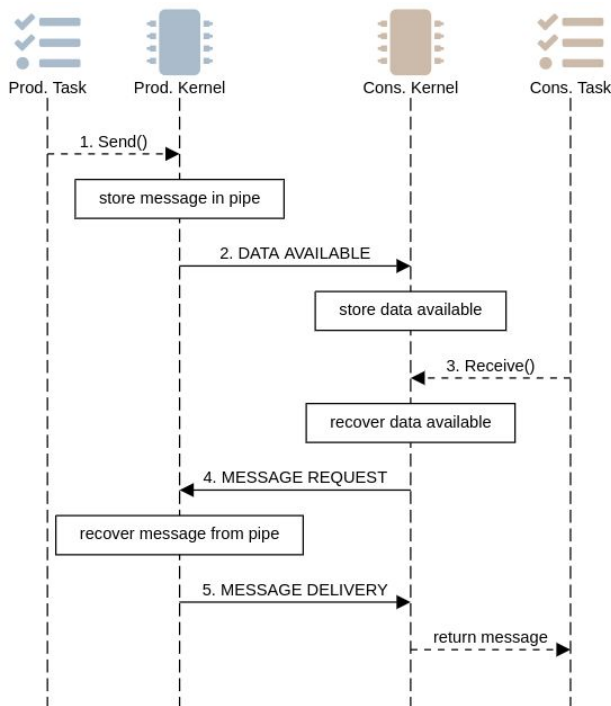
Payload - seen by kernel:

- **Additional 40-bit payload**
 - 8-bit system service
 - 32-bit actual payload

Inter-task Communication

21

Tasks communicate using **Send** and **Receive** primitives

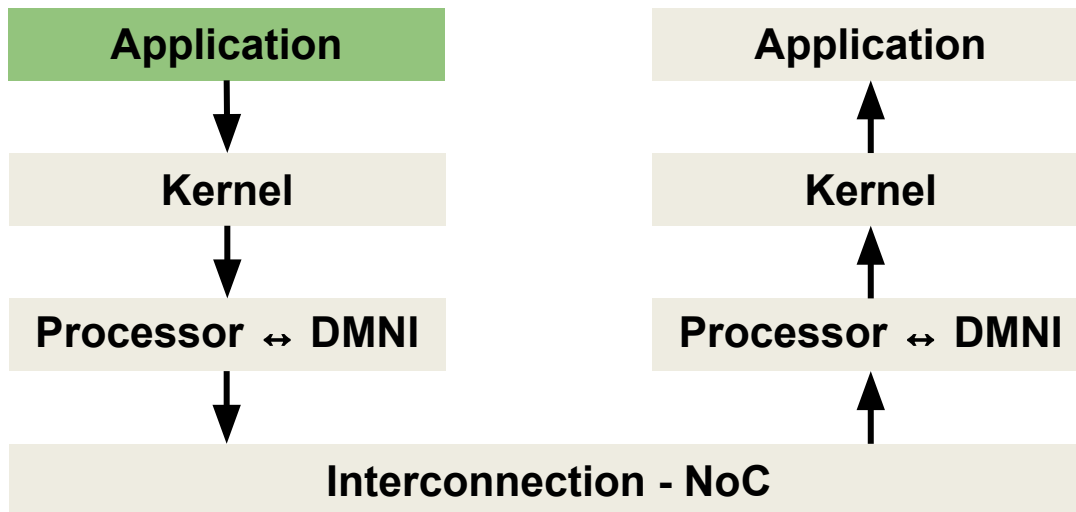


Communication Layers

22

Application

- Sends the message by calling the **Send** API primitive

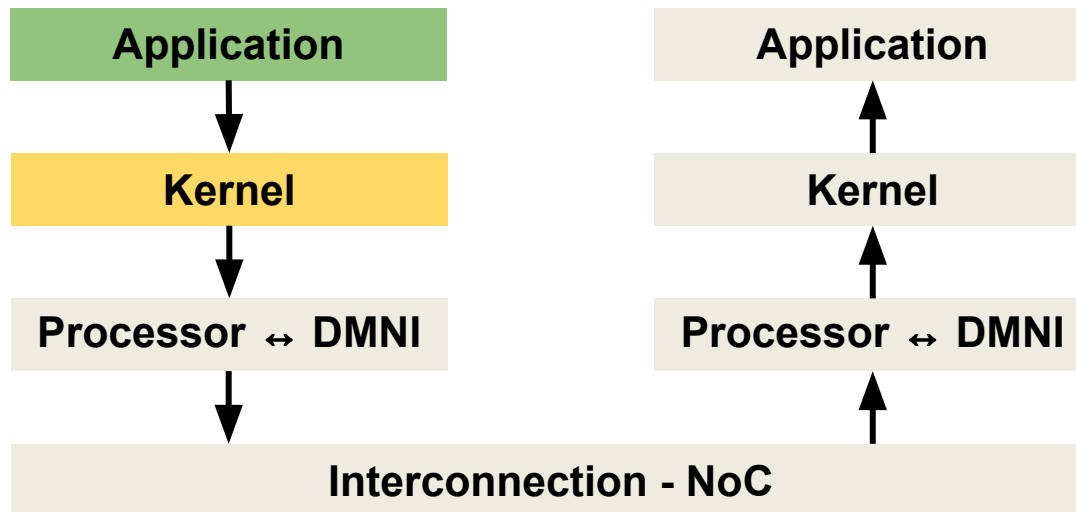


Communication Layers

23

Kernel

- Stores the message in the **pipe** and creates a **packet**
- Configures the DMNI to send the message

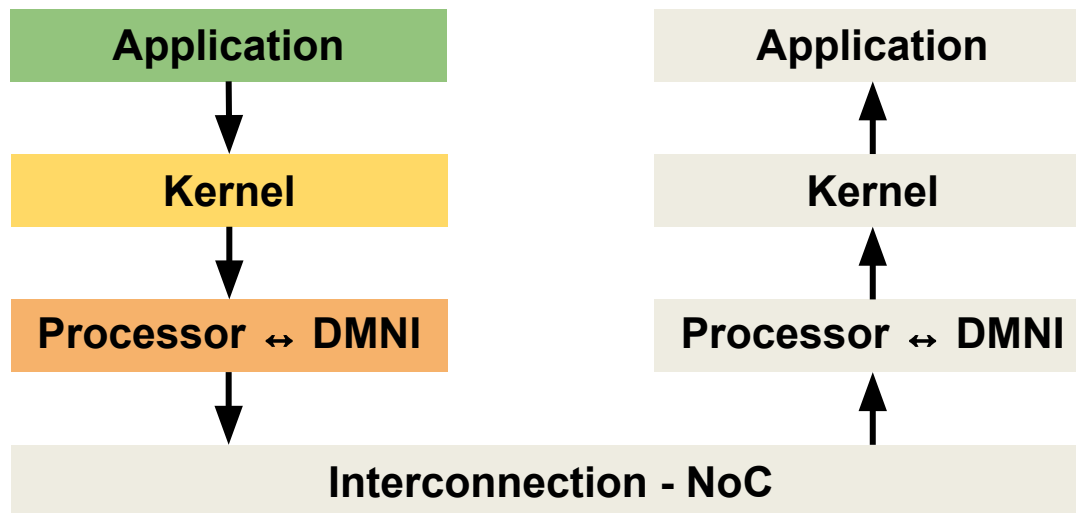


Communication Layers

24

DMNI

- Copies **packet** and **pipe** to NoC
- Serialization and flow control

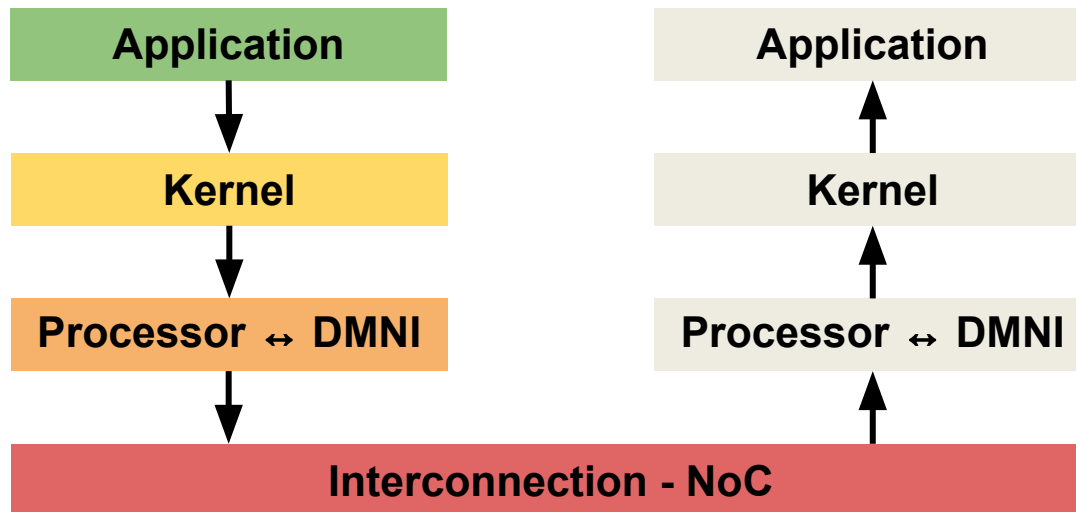


Communication Layers

25

NoC

- Sends the message to the target PE
- Divided in 32-bit flits

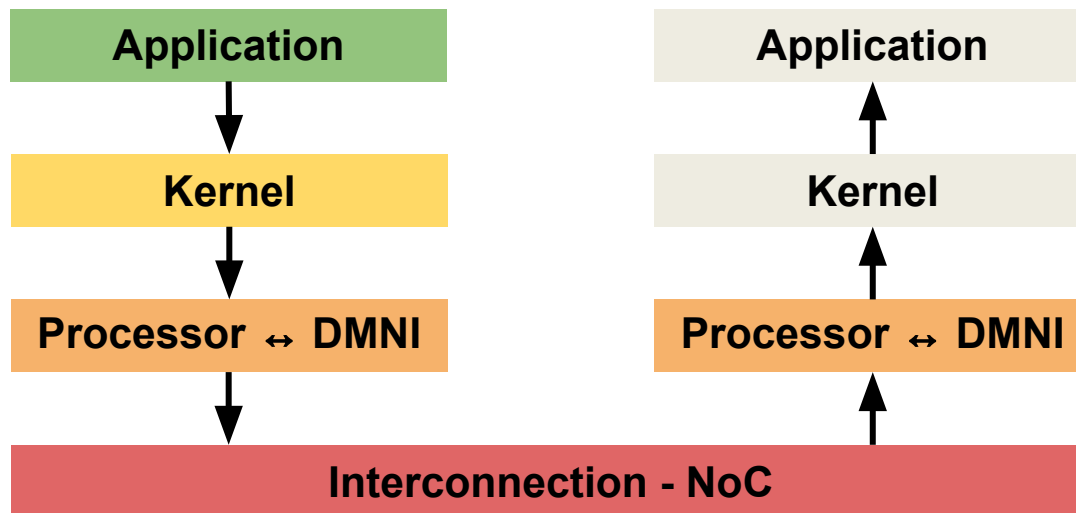


Communication Layers

26

DMNI

- Receives packet containing **MESSAGE_DELIVERY**
- Deserialization and flow control

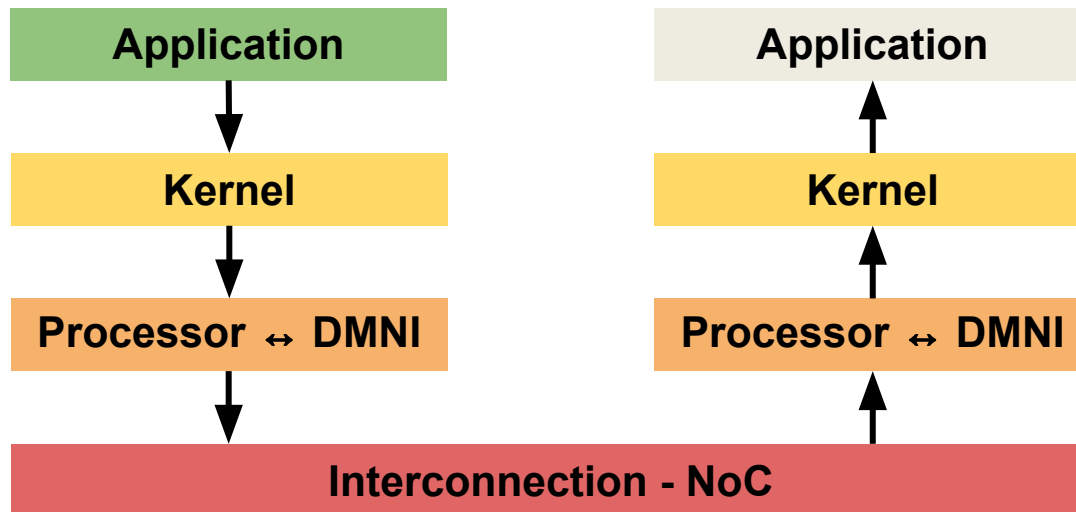


Communication Layers

27

Kernel

- Handles the packet
- Programs the DMNI to receive the message

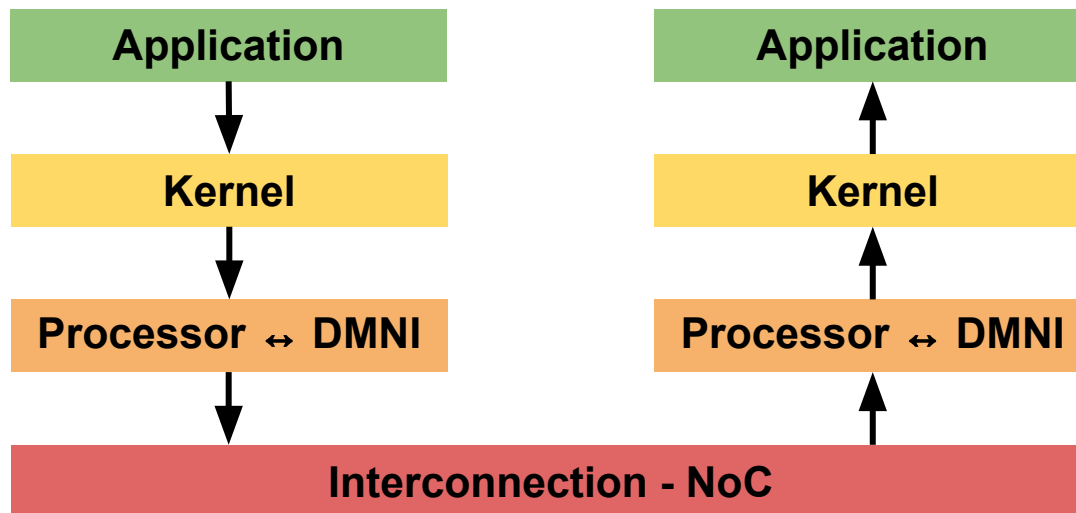


Communication Layers

28

Application

- The execution is returned to the application
- Message is stored in the pointer set by the **Receive** syscall



39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
RELEASE_PERIPHERAL								—					—			—										—																
CLEAR_LUT								—					—			—										TASK_ID																
ANNOUNCE_MONITOR								—					MON_TYPE			—										TASK_ID																
DATA_AV								—					—			PROD_ADDR										CONS_ID																
MESSAGE_REQUEST								—					—			CONS_ADDR										PROD_ID																
TASK_MIGRATION								—					—			ADDR										TASK_ID																
MONITORING								MON_PAYLOAD																																		

Hermes Services

DATA_AV (to/from periph.)	Avoid broadcast DoS
MESSAGE_REQUEST (to/from periph.)	
MESSAGE_DELIVERY	Deliver message with variable-size payload
MIGRATION_*	

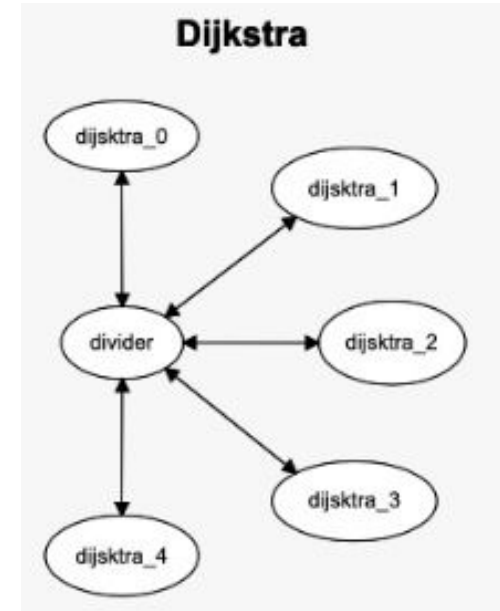

Applications

Application

32

Application

- Set of **tasks**
- Modeled as a CTG
(Communicating Task Graph)



Task

33

Task

- Performs computation and communication
- Each .c file is a task

idct.c

iquant.c

ivlc.c


mpeg.cfg

mpeg_std.h

print.c

start.c

MPEG



```

int main()
{
    printf("MPEG frames:\n");

    for(int i = 0; i < MPEG_FRAMES; i++)
    {
        memphis_receive(&msg1, idct);
        printf("%d\n", memphis_get_tick());
    }

    printf("End Task PRINT - MPEG\n");

    return 0;
}
  
```

#include <memphis.h>

34

```
typedef struct {  
    int length;  
    int msg[MSG_SIZE];  
} message_t;  
void memphis_send(message_t *msg, unsigned cons_id);  
void memphis_receive(message_t *msg, unsigned prod_id);  
void memphis_send_any(message_t *msg, unsigned cons_id);  
void memphis_receive_any(message_t *msg);  
void memphis_br_send_all(unsigned payload, char service);  
void memphis_br_send_tgt(unsigned payload, unsigned cons_id, char service);  
unsigned memphis_get_tick();  
unsigned memphis_get_id();  
unsigned memphis_get_addr();  
void memphis_real_time(unsigned period, unsigned deadline, unsigned exec_time);
```

Generation⁸

8. Castilhos, G., Wachter, E., Madalozzo, G., Erichsen, A., Monteiro, T., and Moraes, F. (2014). A framework for mpsoe generation and distributed applications evaluation. In Proceedings of the International Symposium on Quality Electronic Design (ISQED), pages 408-411.

Testcase

36

Features of the platform

File: example_testcase.yaml

sw:

max_tasks_app: 10

hw:

page_size_KB: 32

stack_size: 1024

tasks_per_PE: 4

mpsoc_dimension: [3,3]

Peripherals:

- name: APP_INJECTOR

pe: 2,2

port: N

- name: MAINJECTOR

pe: 0,0

port: S

\$ memphis testcase example_testcase.yaml

- Creates folders and copies files
- Builds hardware and kernel

Scenario

37

Applications to evaluate

File: example_scenario.yaml

management:

- **task:** mapper_task
static_mapping: [0,0]
- **task:** migration_task
static_mapping: [0,1]
- **task:** rt_monitor
static_mapping: [0,2]

apps:

- **name:** synthetic
- **name:** prod_cons
start_time_ms: 5
static_mapping:
prod: [1,1]

\$ memphis scenario example_testcase/ example_scenario.yaml

- Copies applications
- Builds applications
- Generates repository (MA/Application Injector memory)

Debugging

Debugging Framework⁹

39

Data Extraction (back-end)

- Extracts simulated data from platform
- Inserts into logs

Graphical Debugging (front-end)

- Read extracted data from logs
- Enable easy debuggability by the graphical features

9. Ruaro, M., Carara, E. A., and Moraes, F. G. (2014). Tool-set for NoC-based MPSoC debugging—A protocol view perspective. In Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS), pages 2531-2534.

Main View

40

Debug possibilities:

- Communication flows
- Routing
- Link utilization
- Management Protocols
- Parallel communications



Mapping View

41

Debug possibilities:

- Task mapping algorithm
- PEs occupation
- Task execution status


<input type="checkbox"/> All tasks status <input checked="" type="checkbox"/> Only running <input type="checkbox"/> Only terminated <button>Updating</button> <input type="checkbox"/> Without Task ID			
Slave 0x3 idct 256 RUN iquant 257 RUN	Slave 1x3 start 260 RUN	Slave 2x3 dijkstra_0 512 RUN dijkstra_1 513 RUN	Slave 3x3 dijkstra_4 516 RUN divider 517 RUN
Cluster M 0x2	Slave 1x2 ivlc 258 RUN print 259 RUN	Cluster M 2x2	Slave 3x2 dijkstra_2 514 RUN dijkstra_3 515 RUN
Slave 0x1 bank 768 RUN p1 769 RUN	Slave 1x1 p4 772 RUN recognizer 773 RUN	Slave 2x1 cons 0 RUN prod 1 RUN	Slave 3x1
Global M 0x0	Slave 1x0 p2 770 RUN p3 771 RUN	Cluster M 2x0	Slave 3x0

CPU Utilization View

42

Debug possibilities:

- Scheduling algorithms
- OS and task bugs
- Other software malfunctions



Resources

Resources

44

Code, installation and usage instructions:

<https://github.com/gaph-pucrs/MA-Memphis>

Debugger tutorial (pt):

<https://youtu.be/nvgtvFcCc60>