

# DDR Memory Interference Optimization on Heterogeneous Multicore Platforms

---

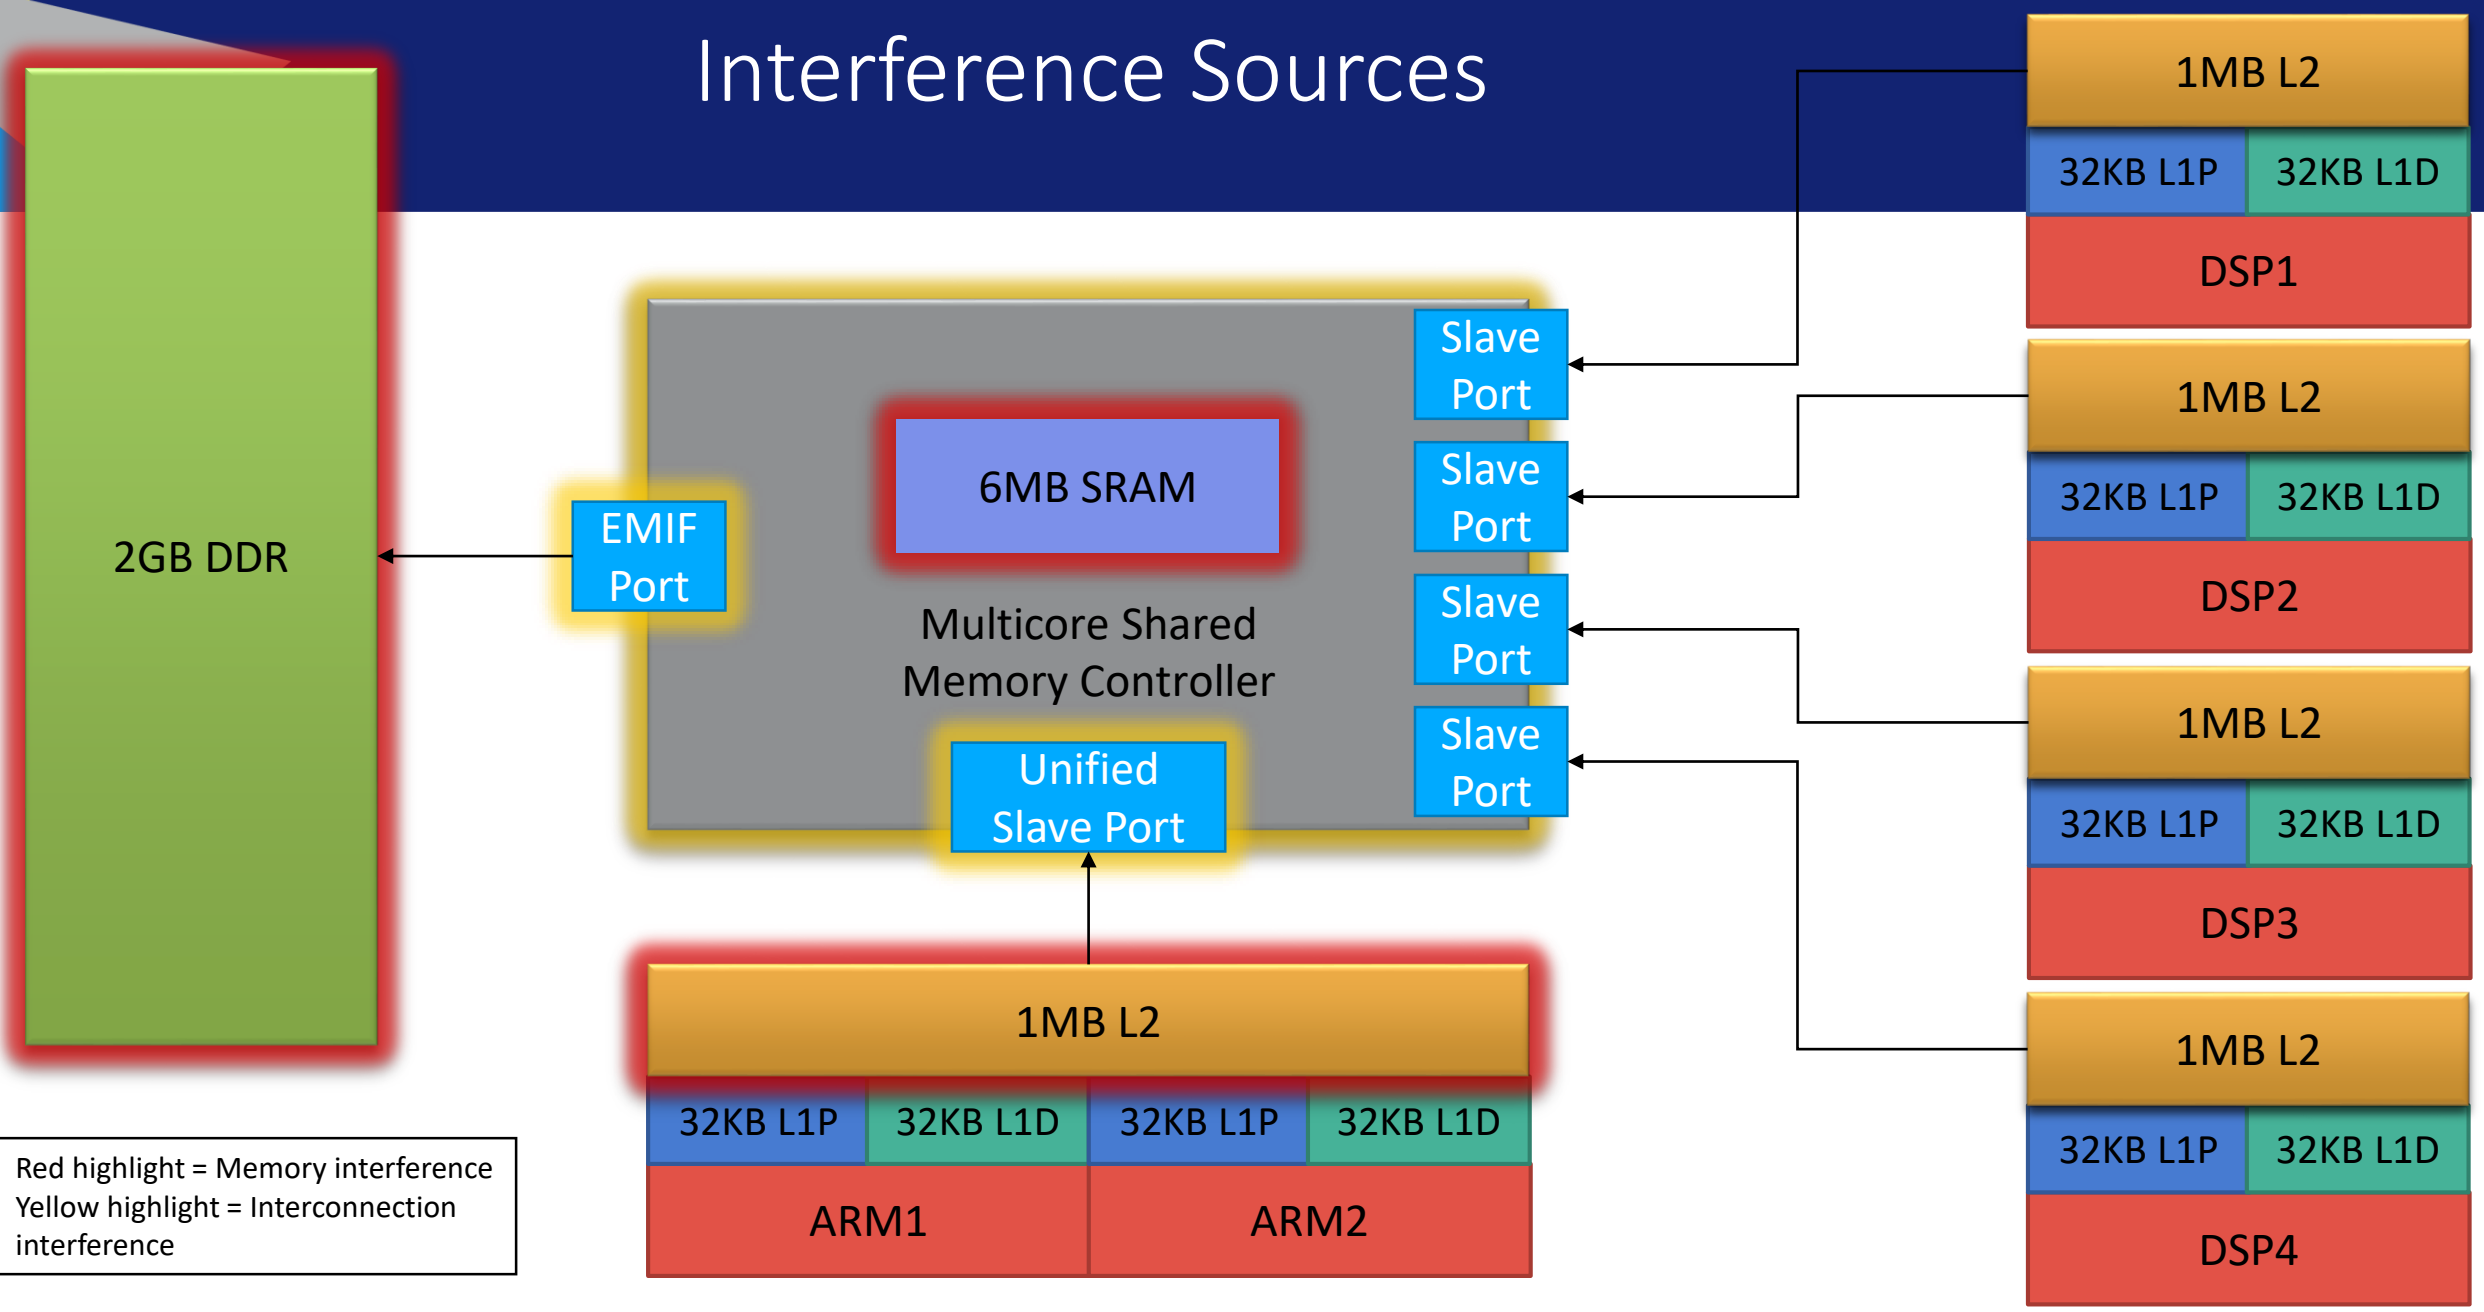
*CONCORDE PROJECT*

*Alfonso Mascareñas González*

*17/02/2022*



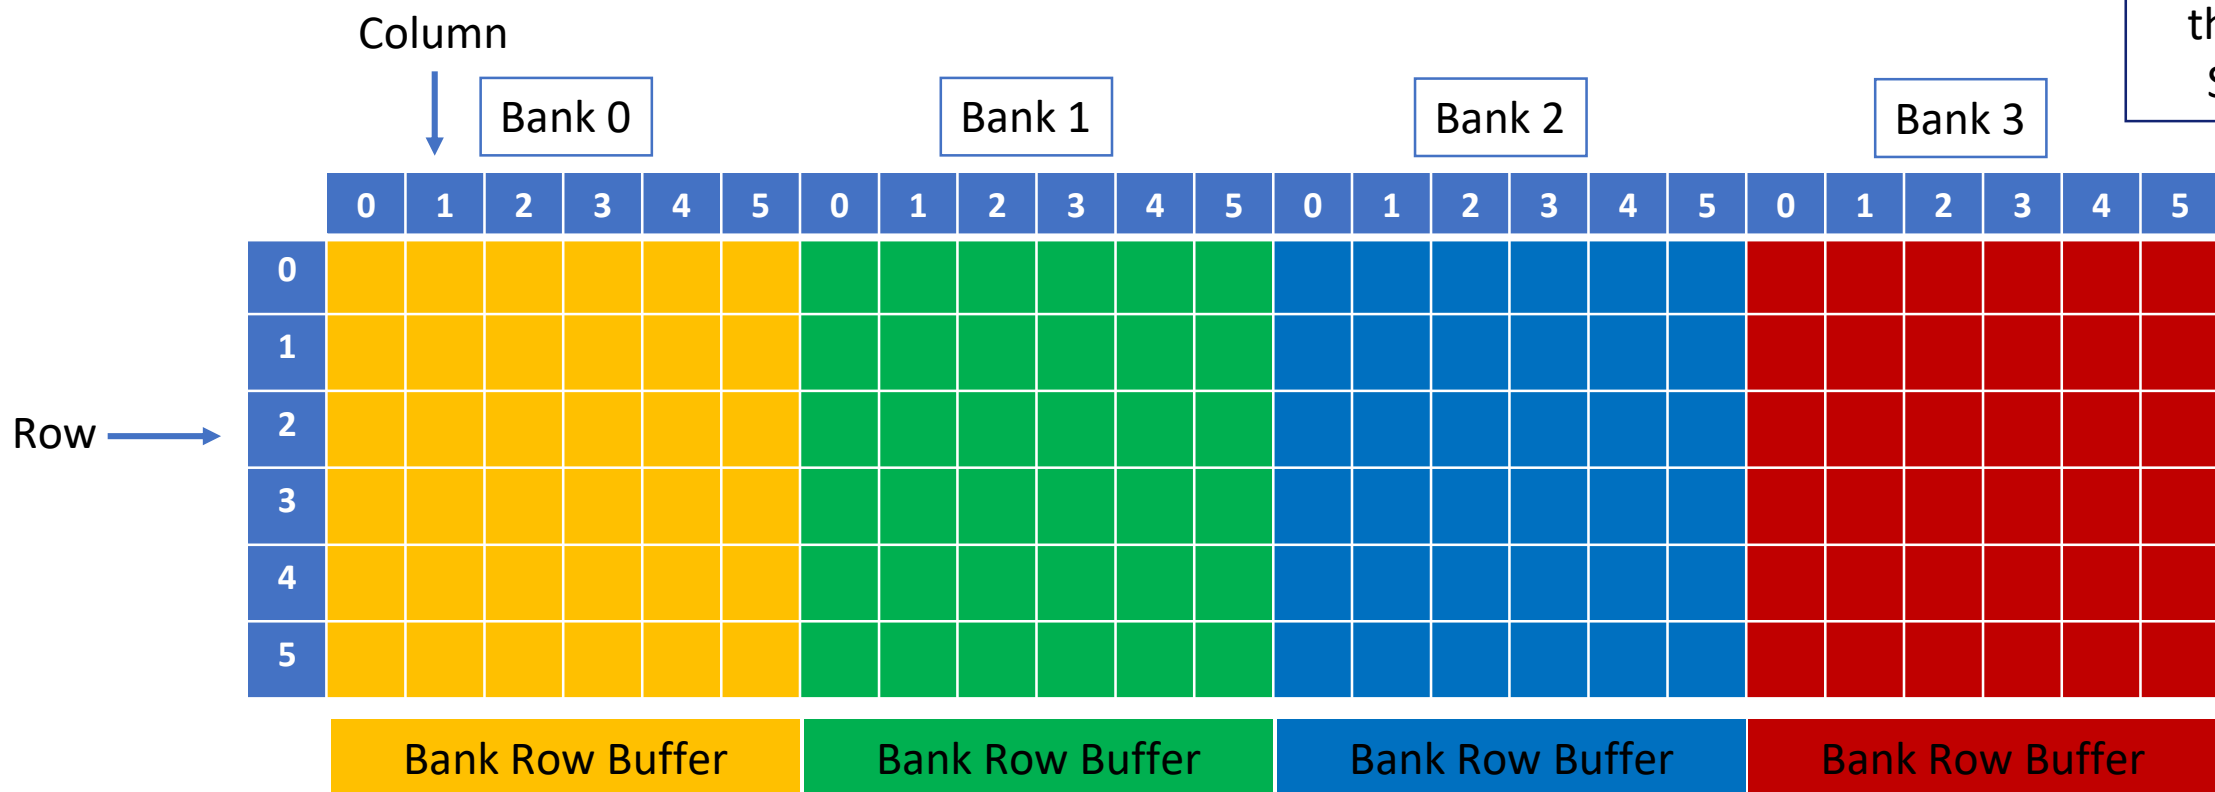
# Interference Sources





Red highlight = Memory interference  
Yellow highlight = Interconnection interference

# DDR3 SDRAM Device: DDR3 Addressing


- Minimum division: Columns, rows, banks and ranks.
- Banks can treat commands independently from each other



# Task and Memory Mapping



Core-Bank Map



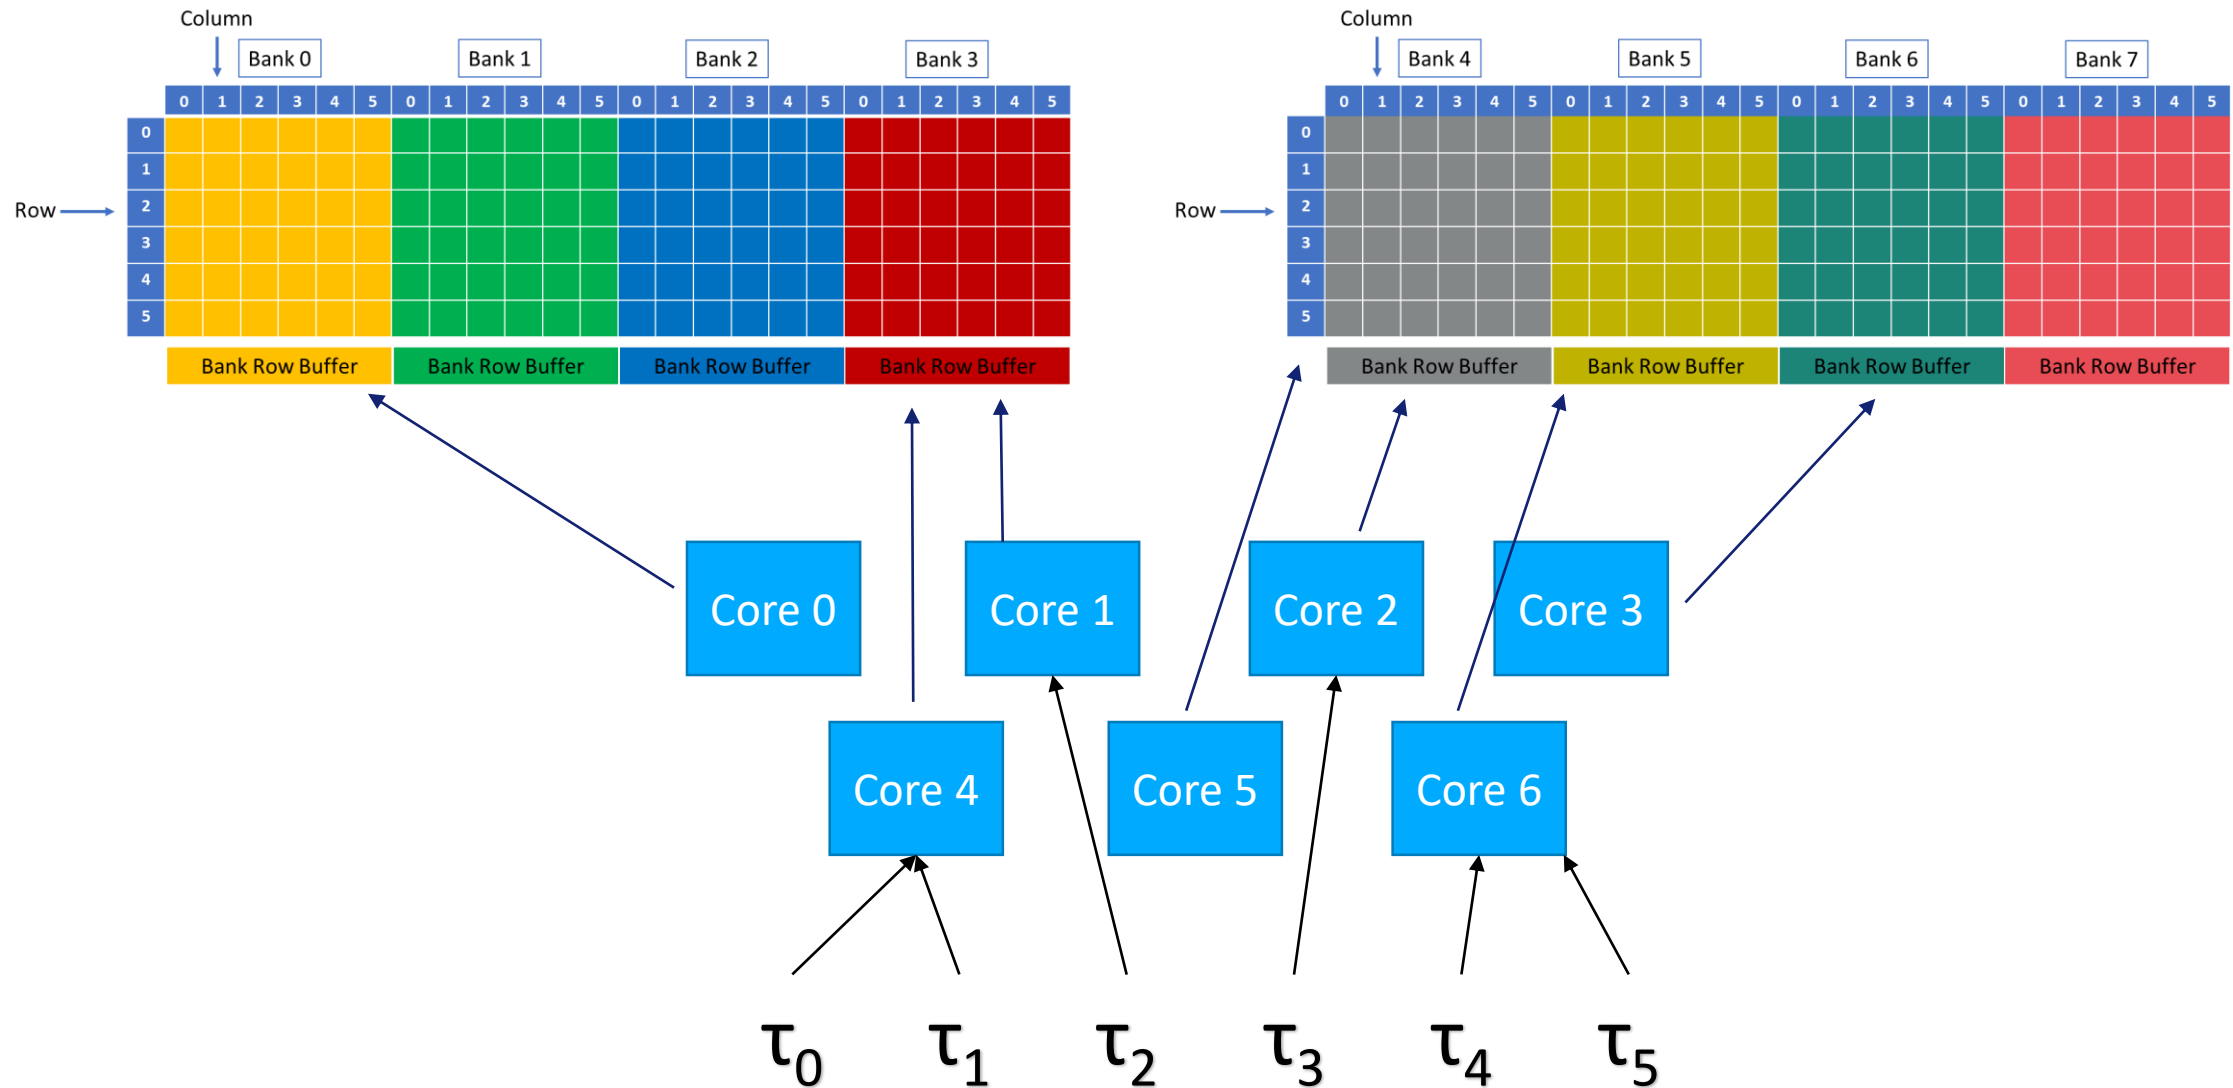
Task-Core Map

$C = \text{core}$   
 $\tau = \text{task}$

$\tau_0$   $\tau_1$   $\tau_2$   $\tau_3$   $\tau_4$   $\tau_5$

Task Set

# Task and Memory Mapping

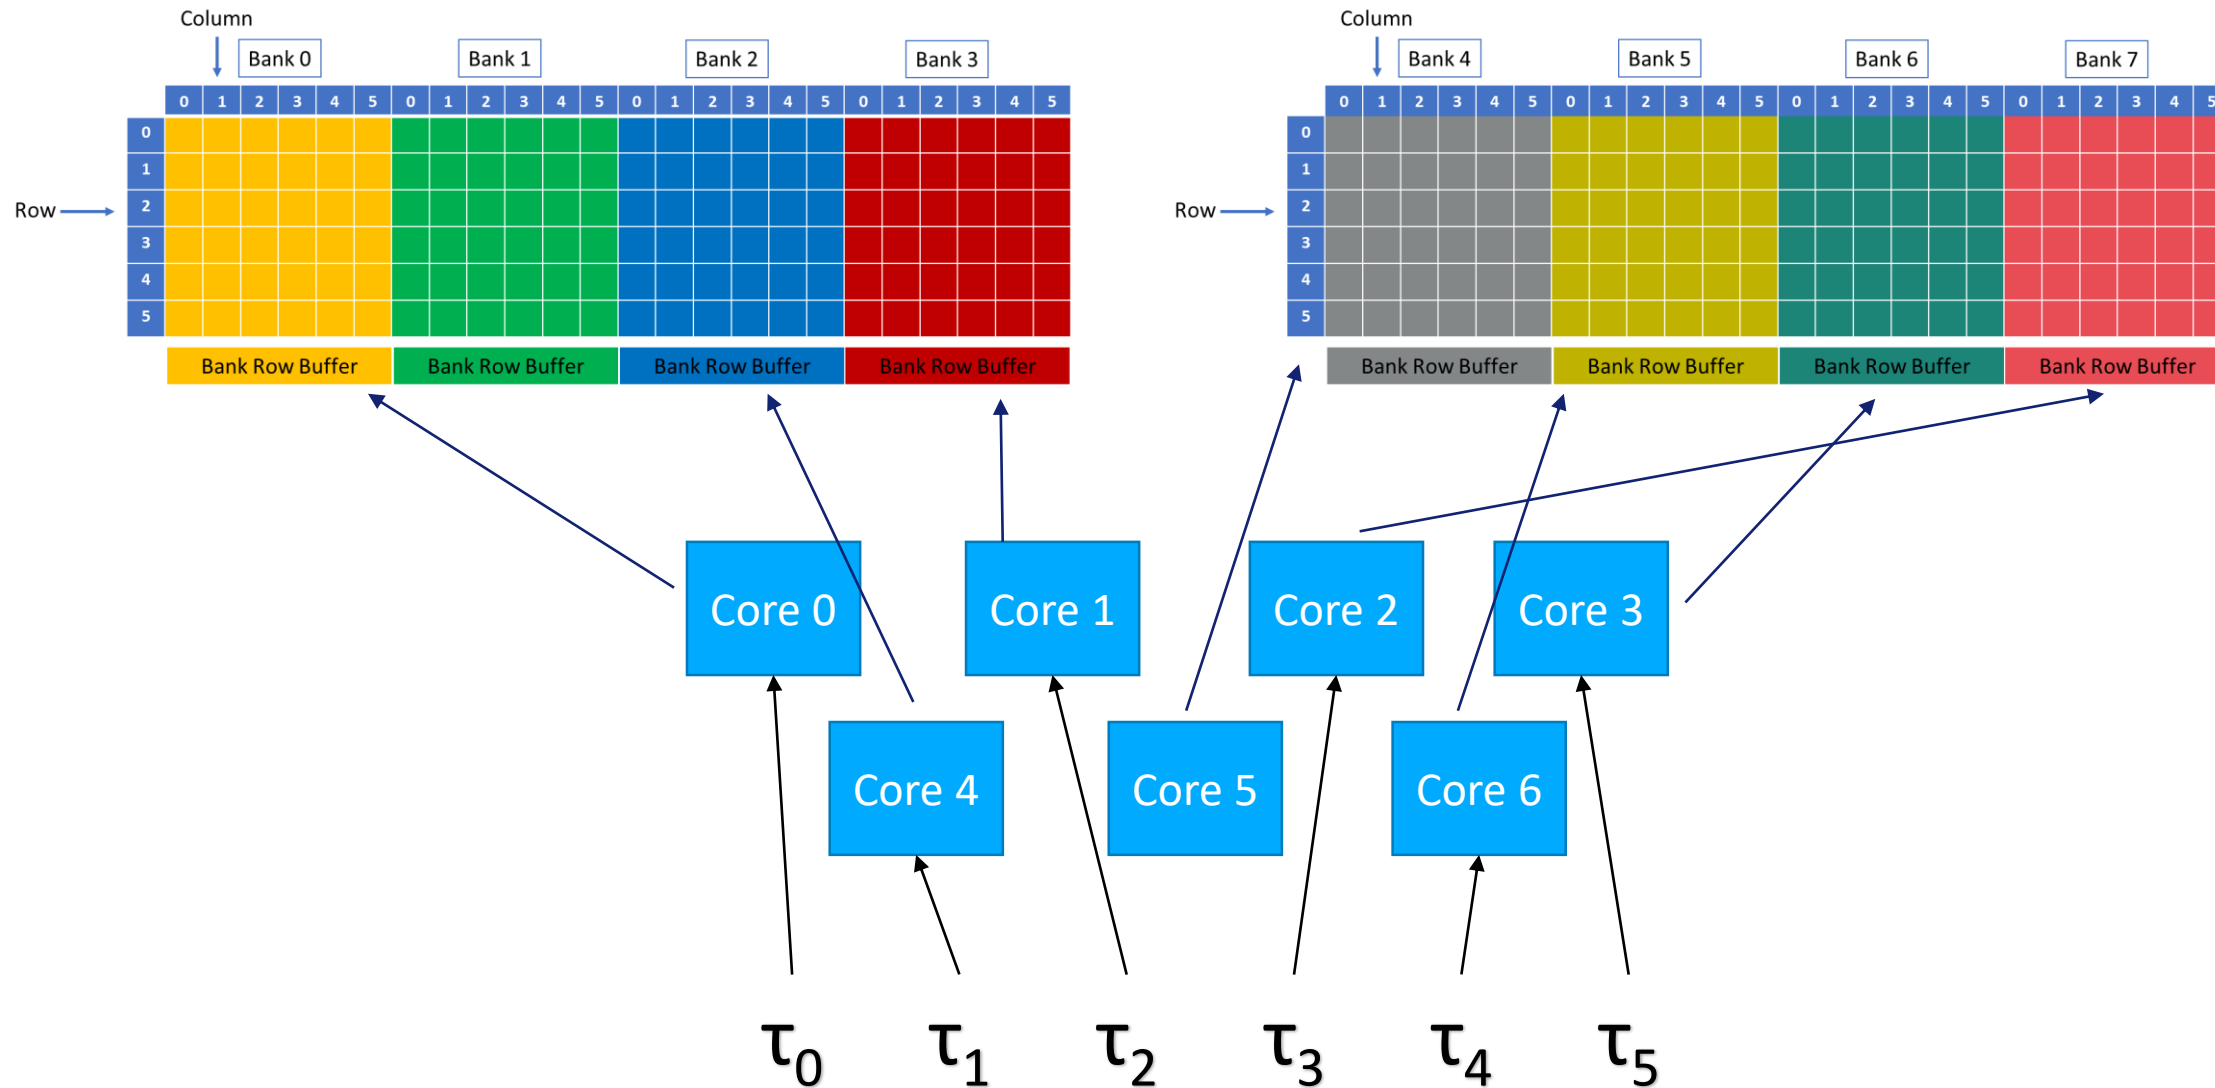


Core-Bank Map

Task-Core Map

Task Set

# Task and Memory Mapping

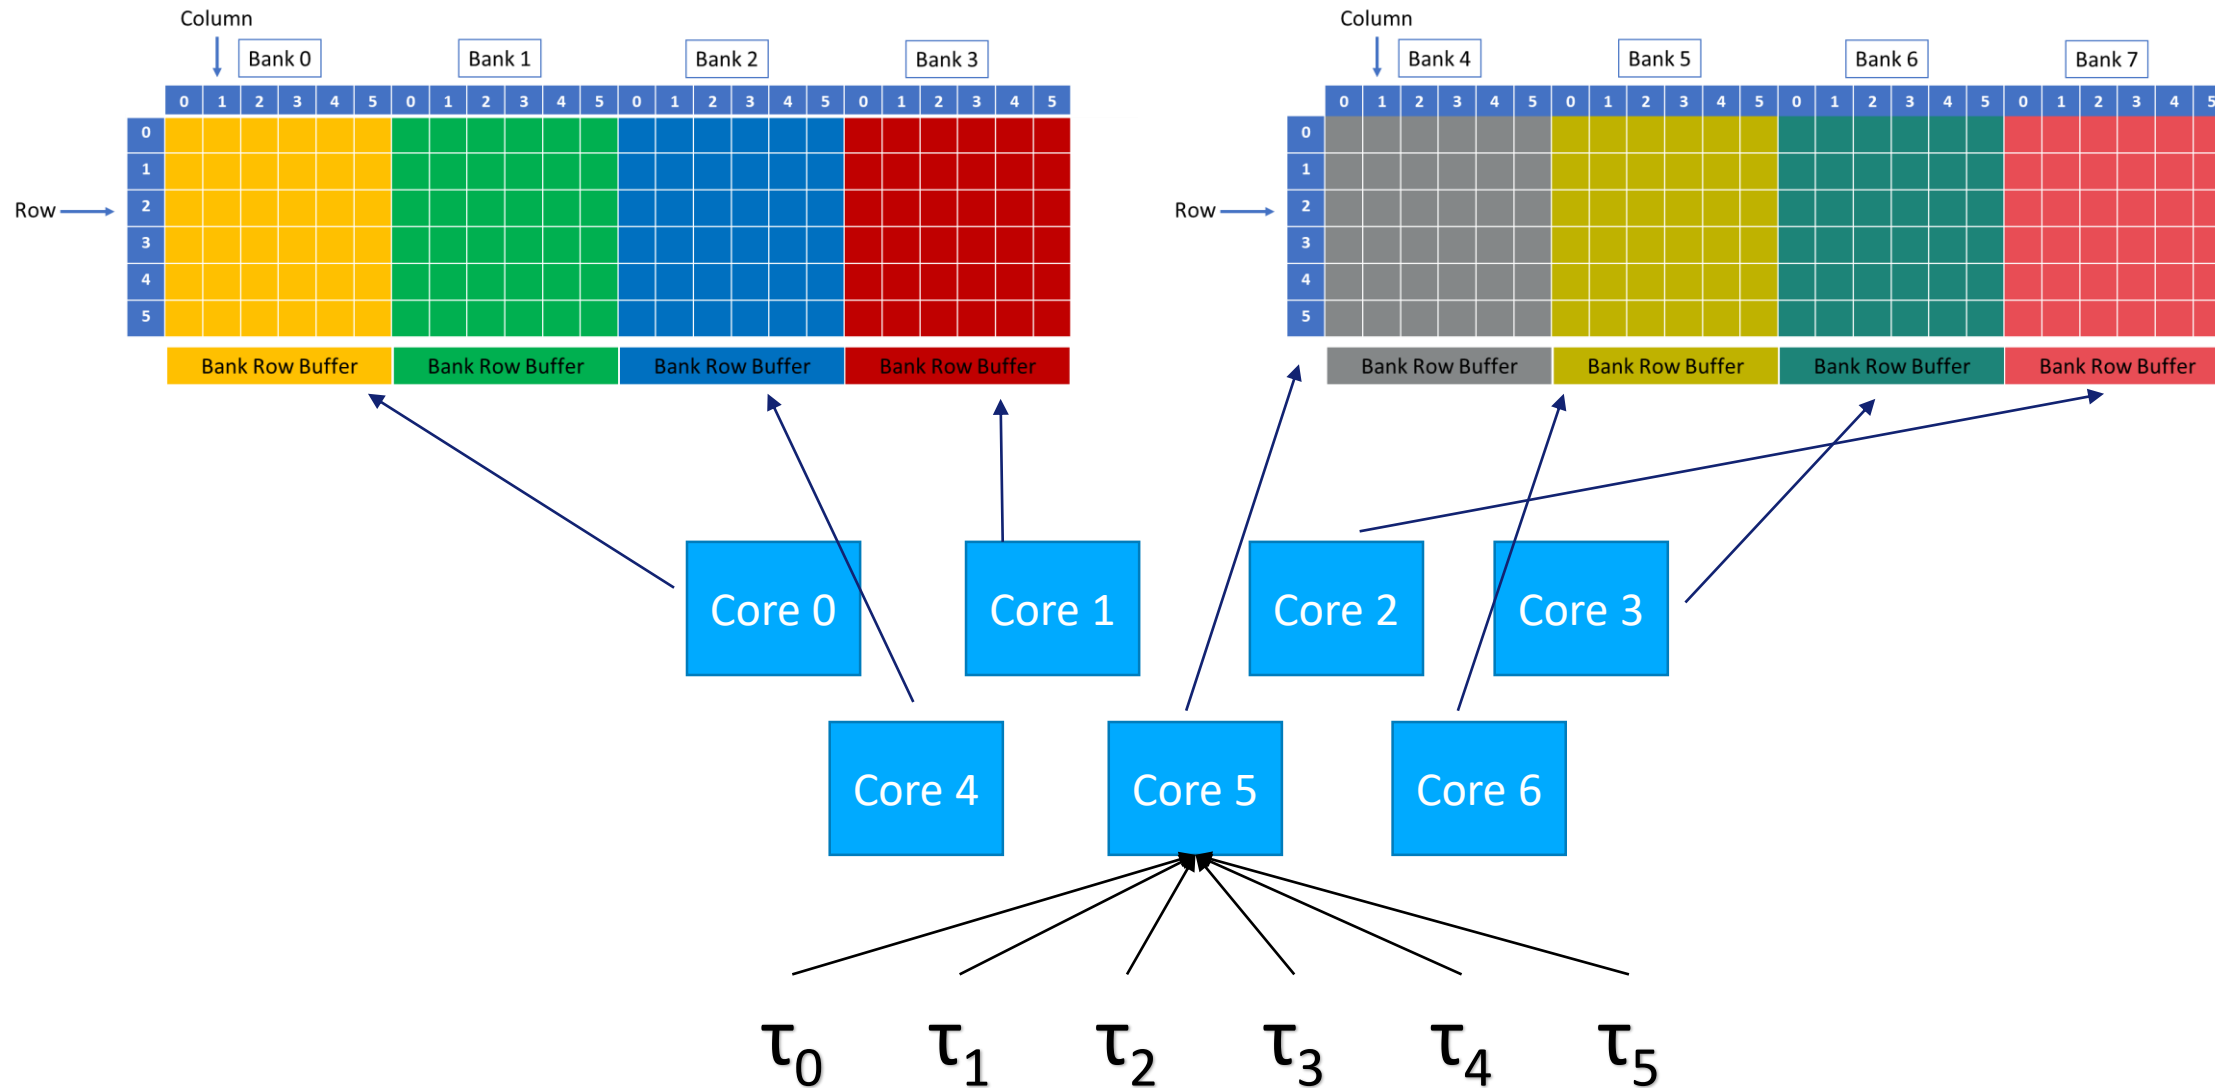


Core-Bank Map

Task-Core Map

Task Set

# Task and Memory Mapping



Core-Bank Map

Task-Core Map

Task Set

# Task and Memory Mapping Representation

The previous kind of maps can be represented as follows:

<b>Bank</b>	7	0	5	4	4	1	2	1	3	6	6	7
<b>Core</b>	0	7	4	5	5	1	3	1	2	6	6	0
<b>Task</b>	0	1	2	3	4	5	6	7	8	9	10	11

Alternatively, in one single row:

<b>Map</b>	0	7	4	5	5	1	7	0	5	4	4	1
	$\tau_0$	$\tau_1$	$\tau_2$	$\tau_3$	$\tau_4$	$\tau_5$	$\tau_0$	$\tau_1$	$\tau_2$	$\tau_3$	$\tau_4$	$\tau_5$

Core:  $\tau_0$  to  $\tau_5$  (tasks 0-5)  
Bank:  $\tau_0$  to  $\tau_5$  (tasks 6-11)

Only 6 tasks are shown







# Task and Memory Mapping Representation

$C = core$   
 $\tau = task$

Or even:



In heterogeneous  
platforms



Blue = ARM core  
Green = DSP

# Multi-objective Optimization

How do we map tasks to cores and cores to banks? Is there a logic?

Objectives:

- Increase Task Execution Parallelism → *Minimize(Workload Variance)*
- Decrease Maximum DDR Interference → *Minimize(Maximum Interference Cost)*



*Minimize(Workload Variance)* ⇒ **Maximize** *(Maximum Interference Cost)*  
*Minimize(Maximum Interference Cost)* ⇒ **Maximize** *(Workload Variance)*



# Multi-objective Optimization

$WCET_i$  = Worst Case Execution Time of  $\tau_i$  in isolation  
 $IC(\tau_i)$  = DDR memory interference cost for  $\tau_i$   
PE = Processing element = Core

- Cost functions:

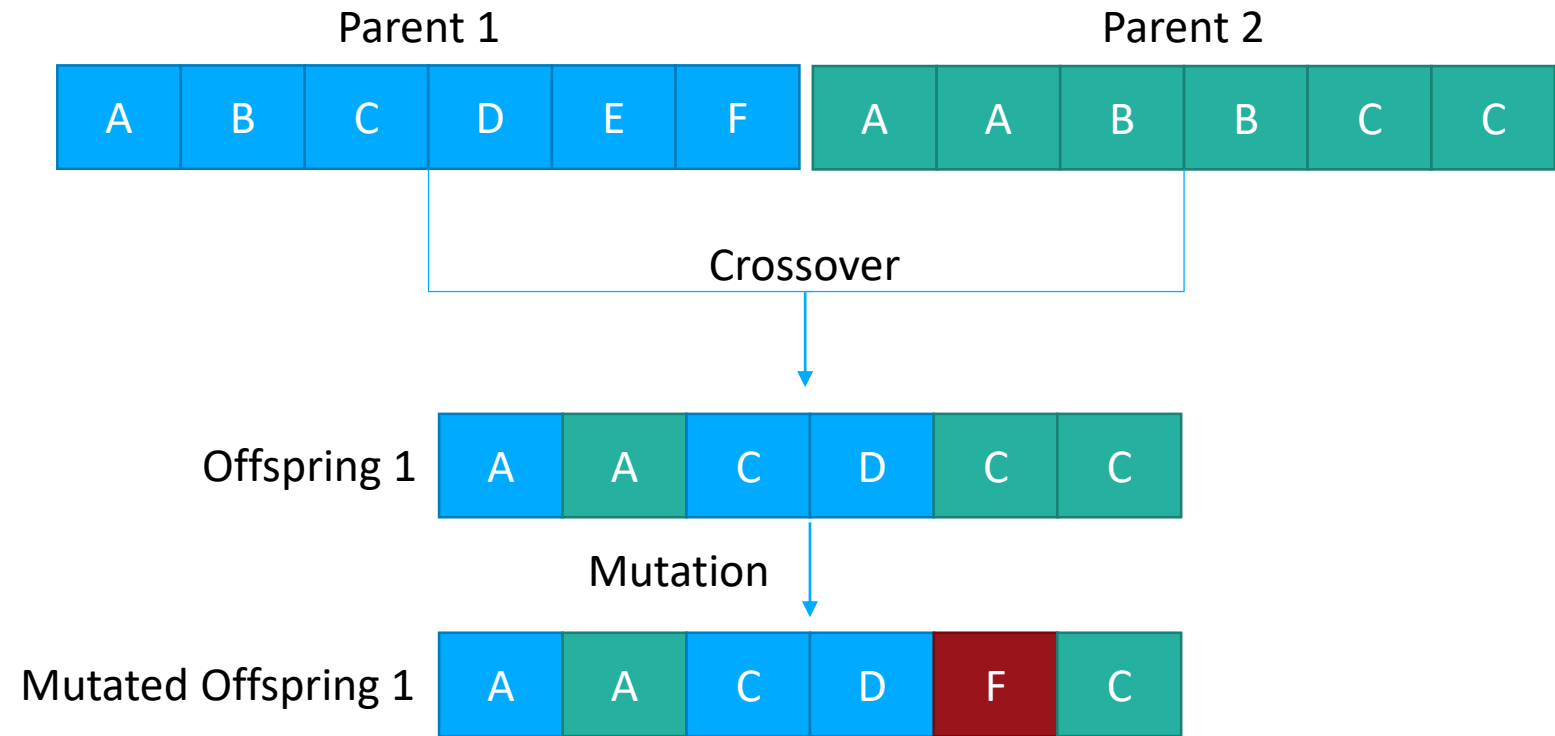
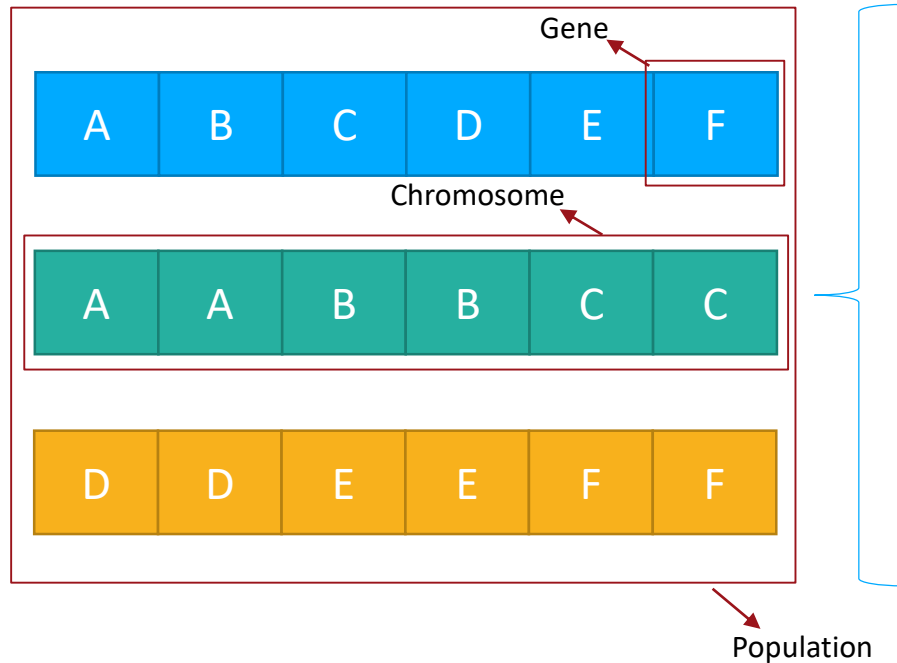
$$\textit{Workload Variance} = \sqrt{\frac{\sum_{i=0}^{n-1} (WCET_i - \textit{avg}(WCET))^2}{n}} \quad \Downarrow \Downarrow$$

$$\textit{Maximum Interference Cost} = \max([IC(\tau_0), \dots, IC(\tau_{n-1})]) \quad \Downarrow \Downarrow$$

- Constraints:

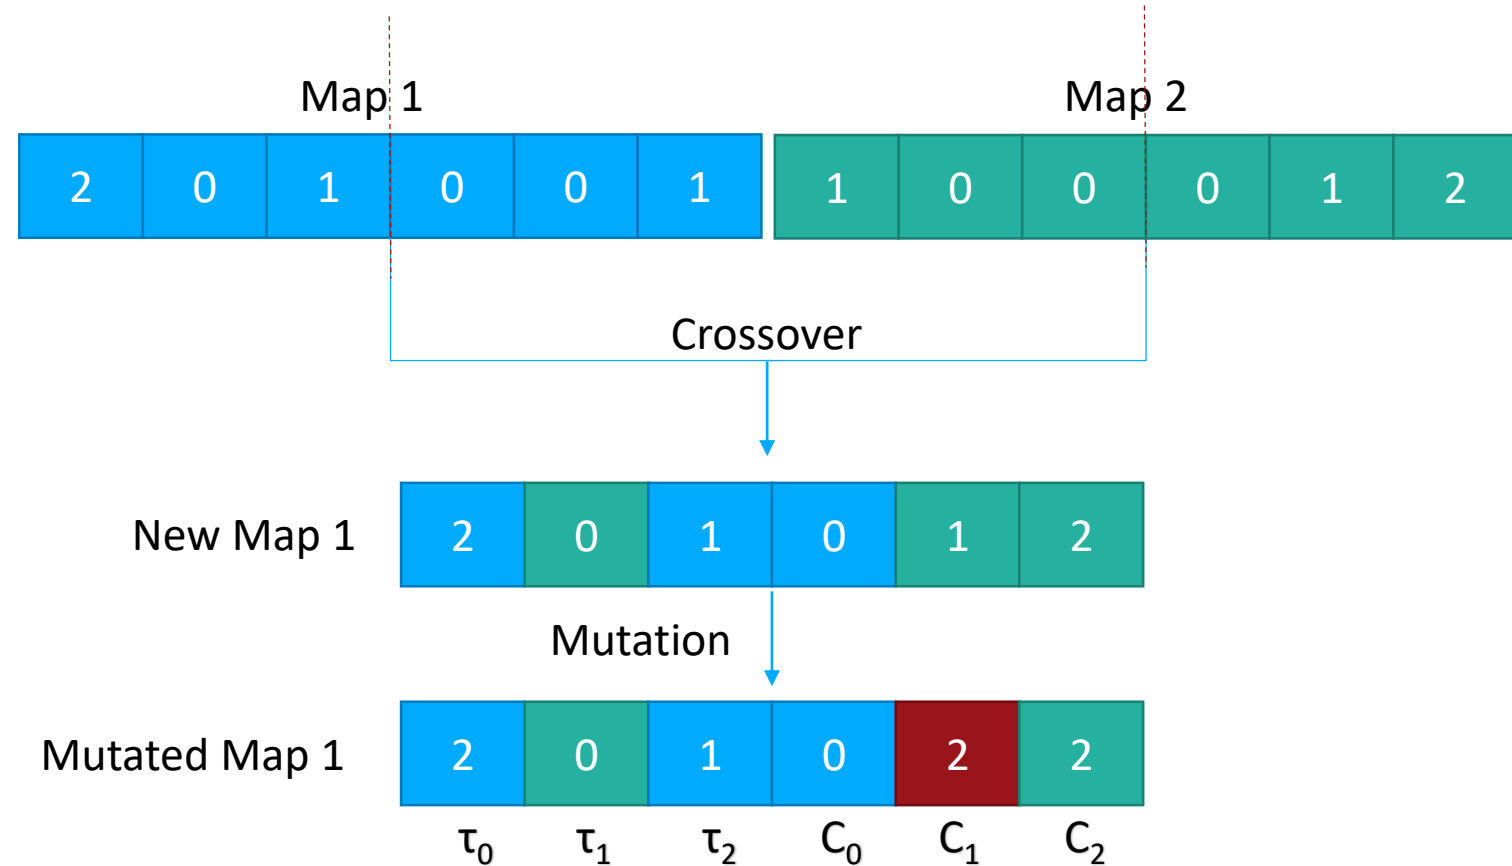
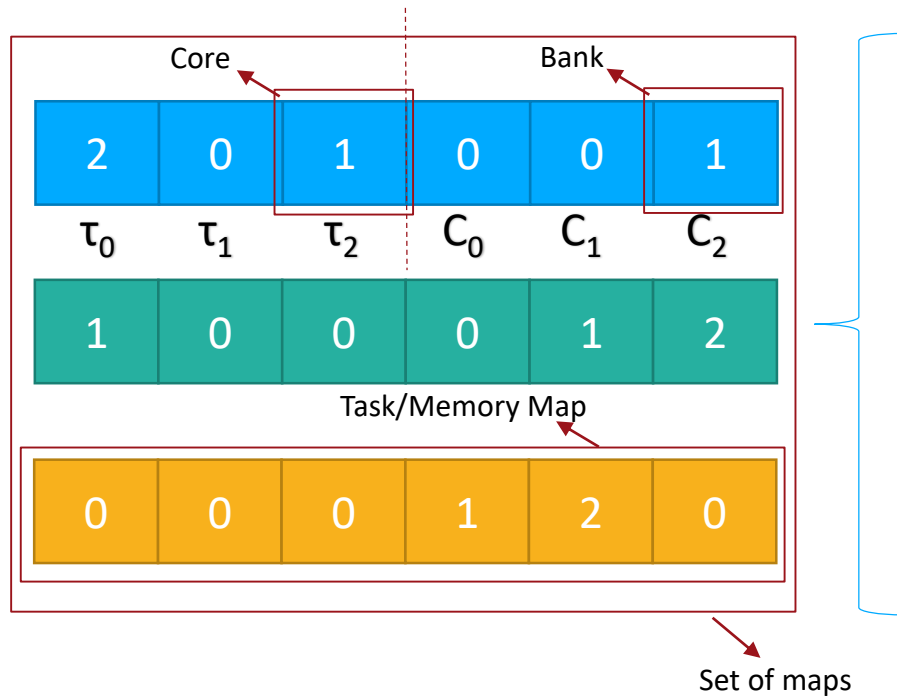
$$\textit{Deadline}_i > WCET_i + IC(\tau_i) + (WCET_j | PE_j = PE_i)$$

# Meta-heuristics Multi-objective Optimization: Genetic algorithm



- Of all candidate offspring, the worst are discarded

# Meta-heuristics Multi-objective Optimization: Genetic algorithm




- Of all candidate offspring, the worst are discarded

# Task/Memory Pareto Front – Keystone II


At the end of the algorithms execution, we obtain different Pareto front

Task mapping on platform cores



The dominated solutions from all the algorithms are removed. Another Pareto front is obtained

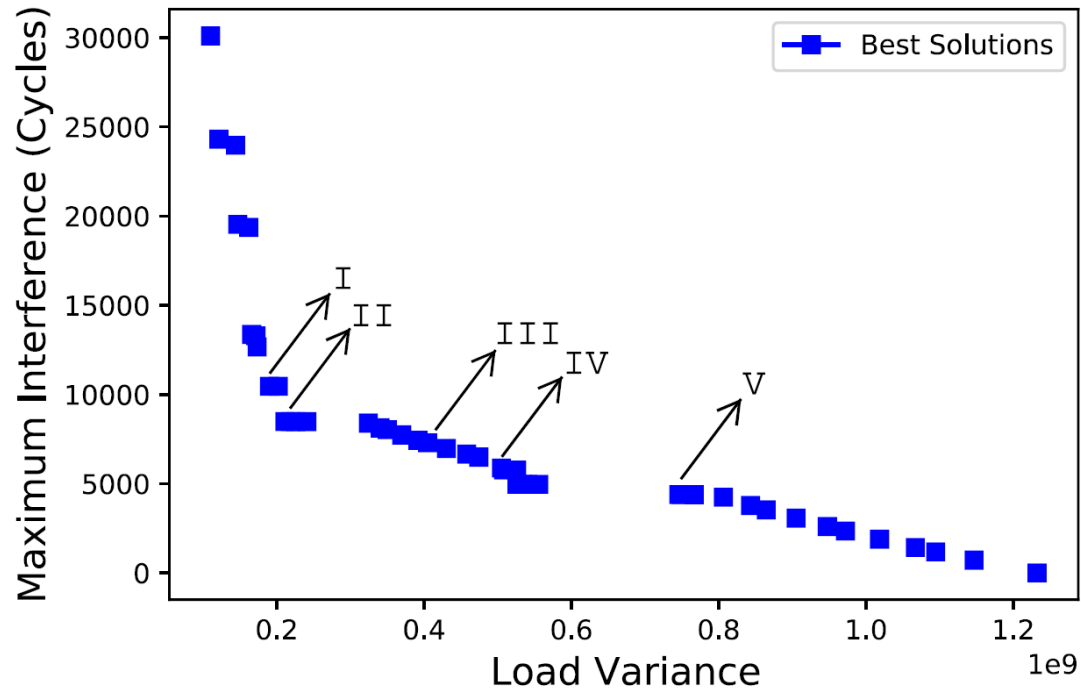
Task mapping on platform cores



# Task/Memory Map Evaluation – Keystone II


Some solutions are tagged for their evaluation

Task mapping on platform cores



The maximum interference is measured for each solution


Keystone II Task Mapping Validation



# Task/Memory Pareto Front – Sitara AM5728

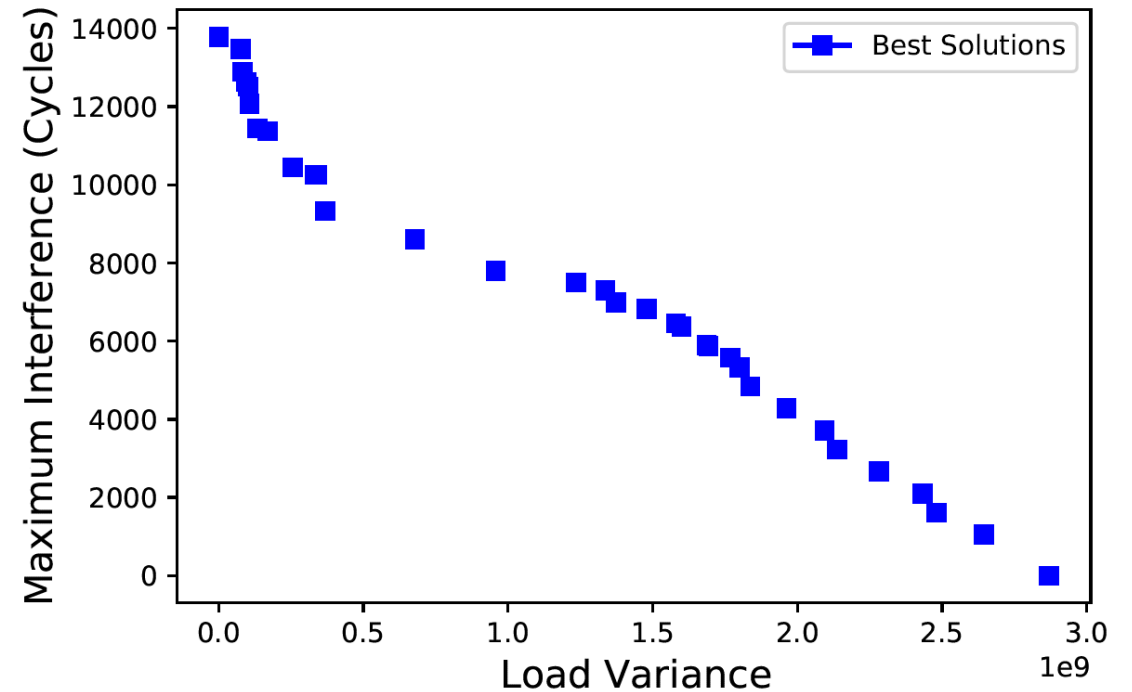
At the end of the algorithms execution, we obtain different Pareto front

Task mapping on platform cores



The dominated solutions from all the algorithms are removed. Another Pareto front is obtained

Task mapping on platform cores






# Task/Memory Map Evaluation – Sitara AM5728

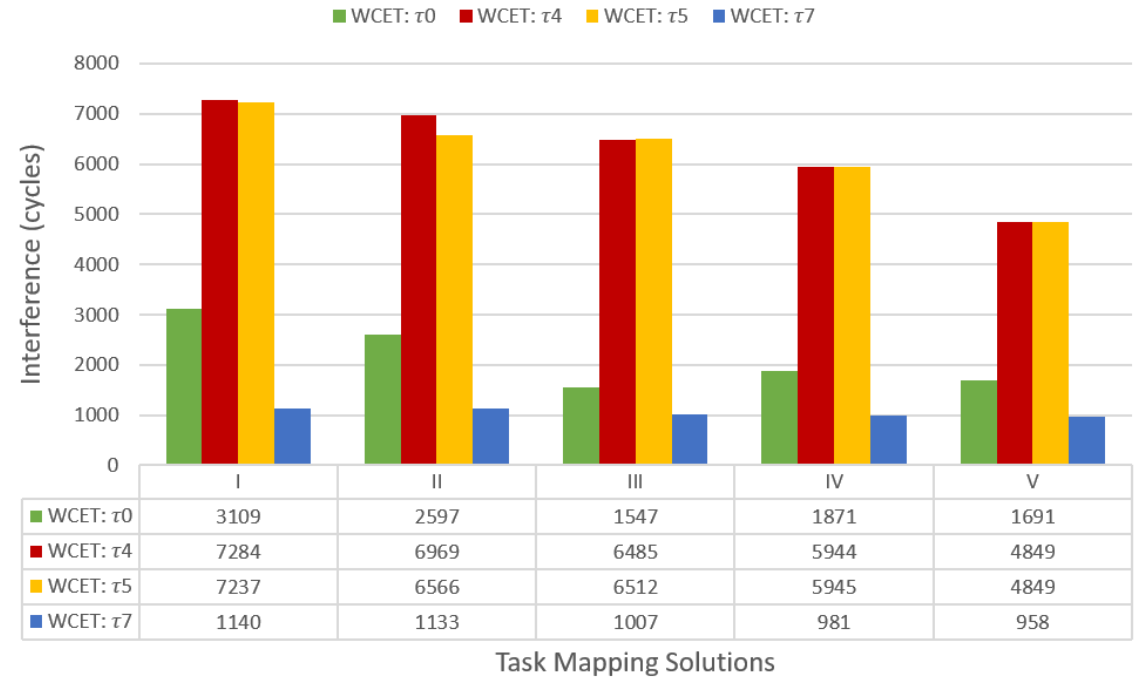
Some solutions are tagged for their evaluation

The maximum interference is measured for each solution

Task mapping on platform cores



Sitara AM5728 Task Mapping Validation




# Task/Memory Mapping Logic

- Task-Core mapping:
  - Stack most intensive tasks on the same core (sequential execution, i.e., no interference) → Drawback = Less parallelism
  - Equally spread the tasks among the cores → Drawback = memory Interference
  - Increase the number of active cores (more parallelism) → Drawback = Generally memory Interference
  - Select the correct core type for the task (ARM,DSP): the execution time and interference vary → Drawback = None
- Core-bank mapping: Always try with a private bank and, when all of these are occupied ( $N^{\circ}C > N^{\circ}B$ ), share bank


# Why should we go for the private banks?

- For instance, let's compare the worst-case read transmission time cost, for the intra-bank and inter-bank case:

$$RT^{\text{intra}} = CL + \frac{BL}{2} + t_{RTW}$$



$$RT^{\text{inter}} = CL + \frac{BL}{2} + t_{RTW} - WL$$



# Why should we go for the private banks?

- Another example is the DDR memory PREcharge command worst-case interference cost (PRE):
  - $PRE^{intra} = t_{RP}$
  - $PRE^{inter} = 1$where  $t_{RP} \gg 1$
- And yet another example related to the DDR memory ACTive command worst-case interference cost (ACT):
  - $ACT^{intra} = t_{RCD}$
  - $ACT^{inter} = t_{RRD}$where  $t_{RCD} > t_{RRD}$

Thank you for your attention