

System Design and Methodology / Embedded Systems Design

I. Modeling and Design of Embedded Systems

**TDTS07/TDDI08
VT 2026**

Ahmed Rezine

(Based on material by Petru Eles and Soheil Samii)

**Institutionen för datavetenskap (IDA)
Linköpings universitet**

Course Information

Web page: <http://www.ida.liu.se/~TDT07>
<http://www.ida.liu.se/~TDDI08>

Examination: March 24th digital written exam
(see instructions linked on course page)

Labs (see course page and lesson notes)

Lecture notes: made available on the web page

Course Information

Recommended literature:

- *Embedded System Design*, by Peter Marwedel. Springer, 2nd edition 2011, 3d edition 2018, 4th edition 2021. The 4th edition is open access and available online via Springer.com
- *Introduction to Embedded Systems – A Cyber-Physical Systems Approach*, Edward Lee and Sanjit Seshia, 1st edition 2011, 2nd edition 2017 ' (available online: LeeSeshia.org)

Course Information

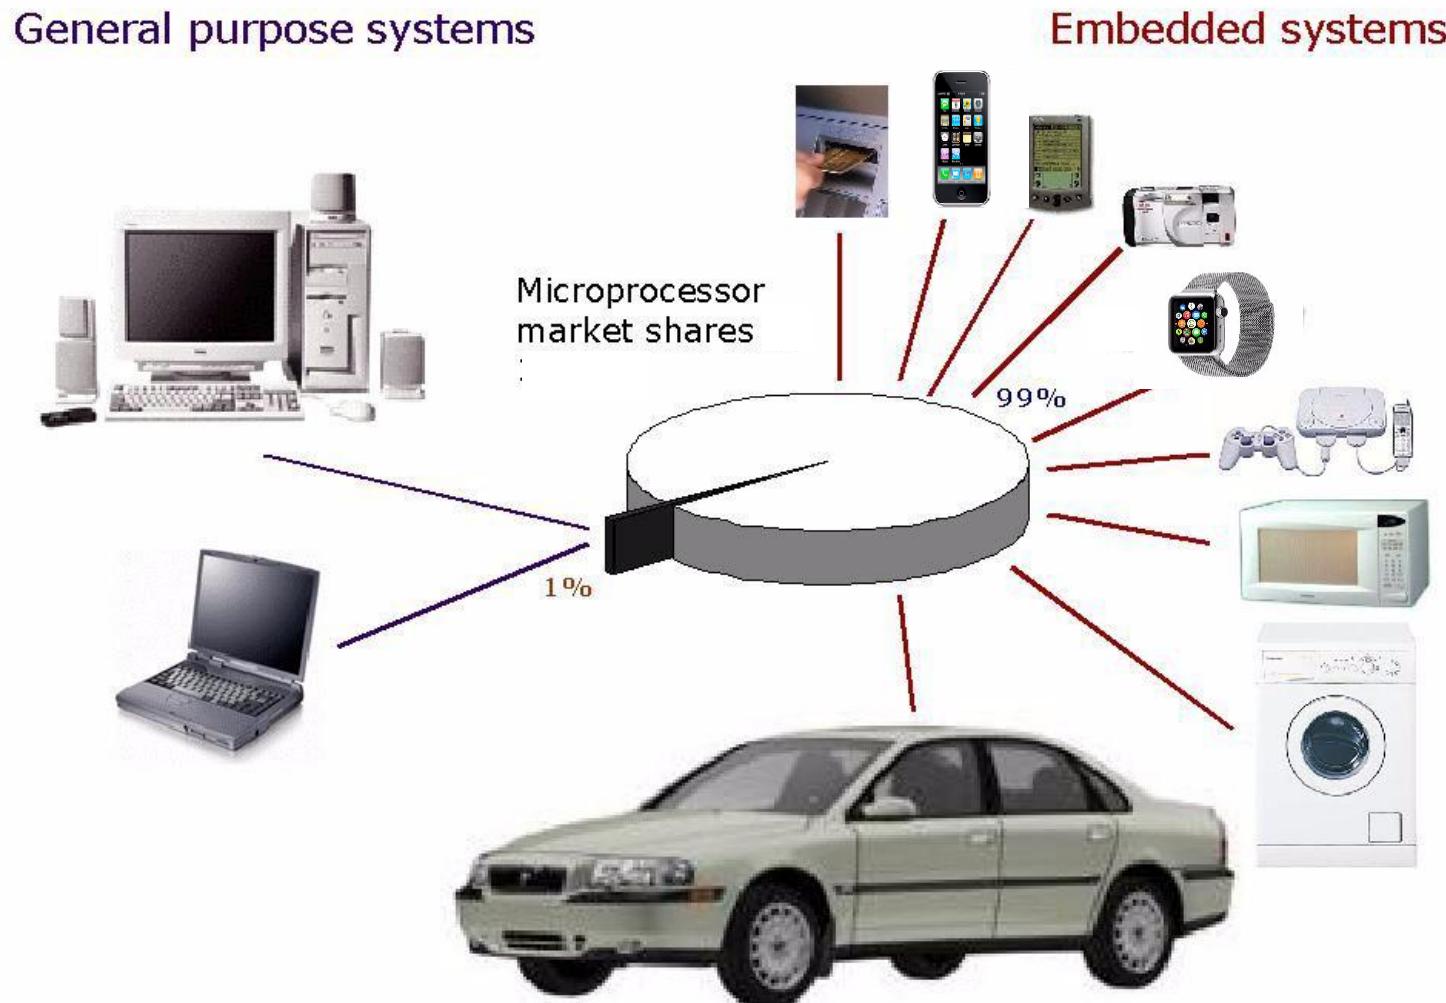
Lessons&Labs:

Xiaopeng Teng
Institutionen för datavetenskap (IDA)
email: xiaopeng.teng@liu.se

EMBEDDED SYSTEMS AND THEIR DESIGN

1. What is an Embedded System
2. Characteristics of Embedded Applications
3. The Traditional Design Flow
4. An Example
5. A New Design Flow
6. The System Level
7. Course Topics

That's how we use microprocessors



What is an Embedded System?

There are several definitions around:

- Some definitions highlight what it is (not) used for:

“An embedded system is any sort of device which includes a programmable component but itself is not intended to be a general purpose computer.”

What is an Embedded System?

There are several definitions around:

- Some definitions highlight what it is (not) used for:

“An embedded system is any sort of device which includes a programmable component but itself is not intended to be a general purpose computer.”
- Some focus on what it is built from:

“An embedded system is a collection of programmable parts surrounded by ASICs and other standard components, that interact continuously with an environment through sensors and actuators.”

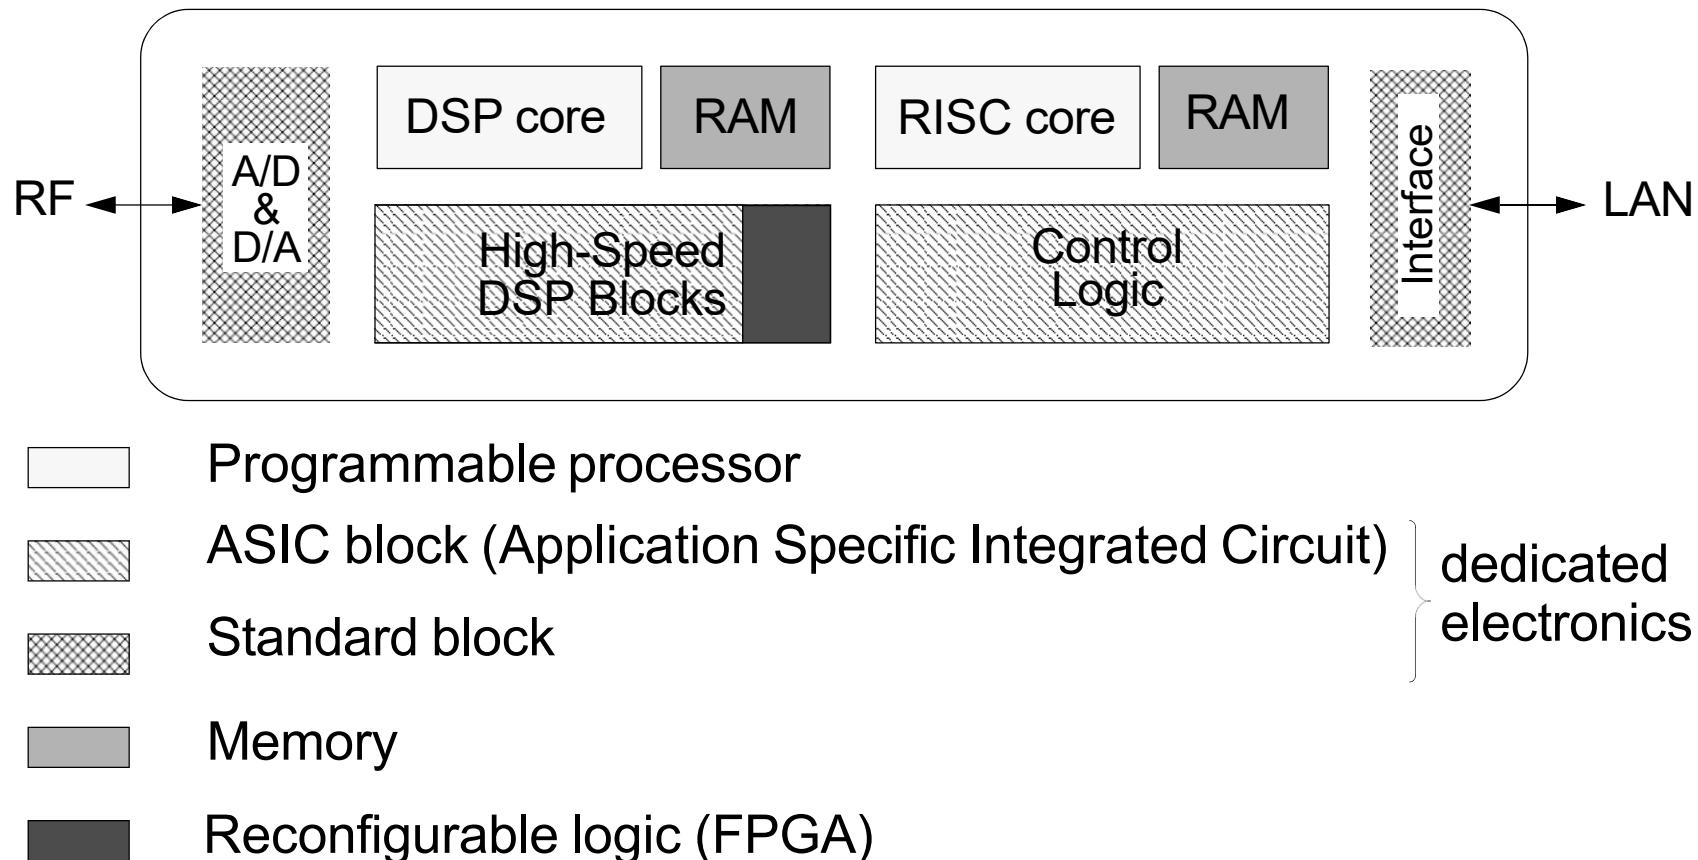
What is an Embedded System?

Some of the main characteristics:

- Dedicated (not *general purpose*)
- Contains a programmable component
- Interacts (continuously) with the environment

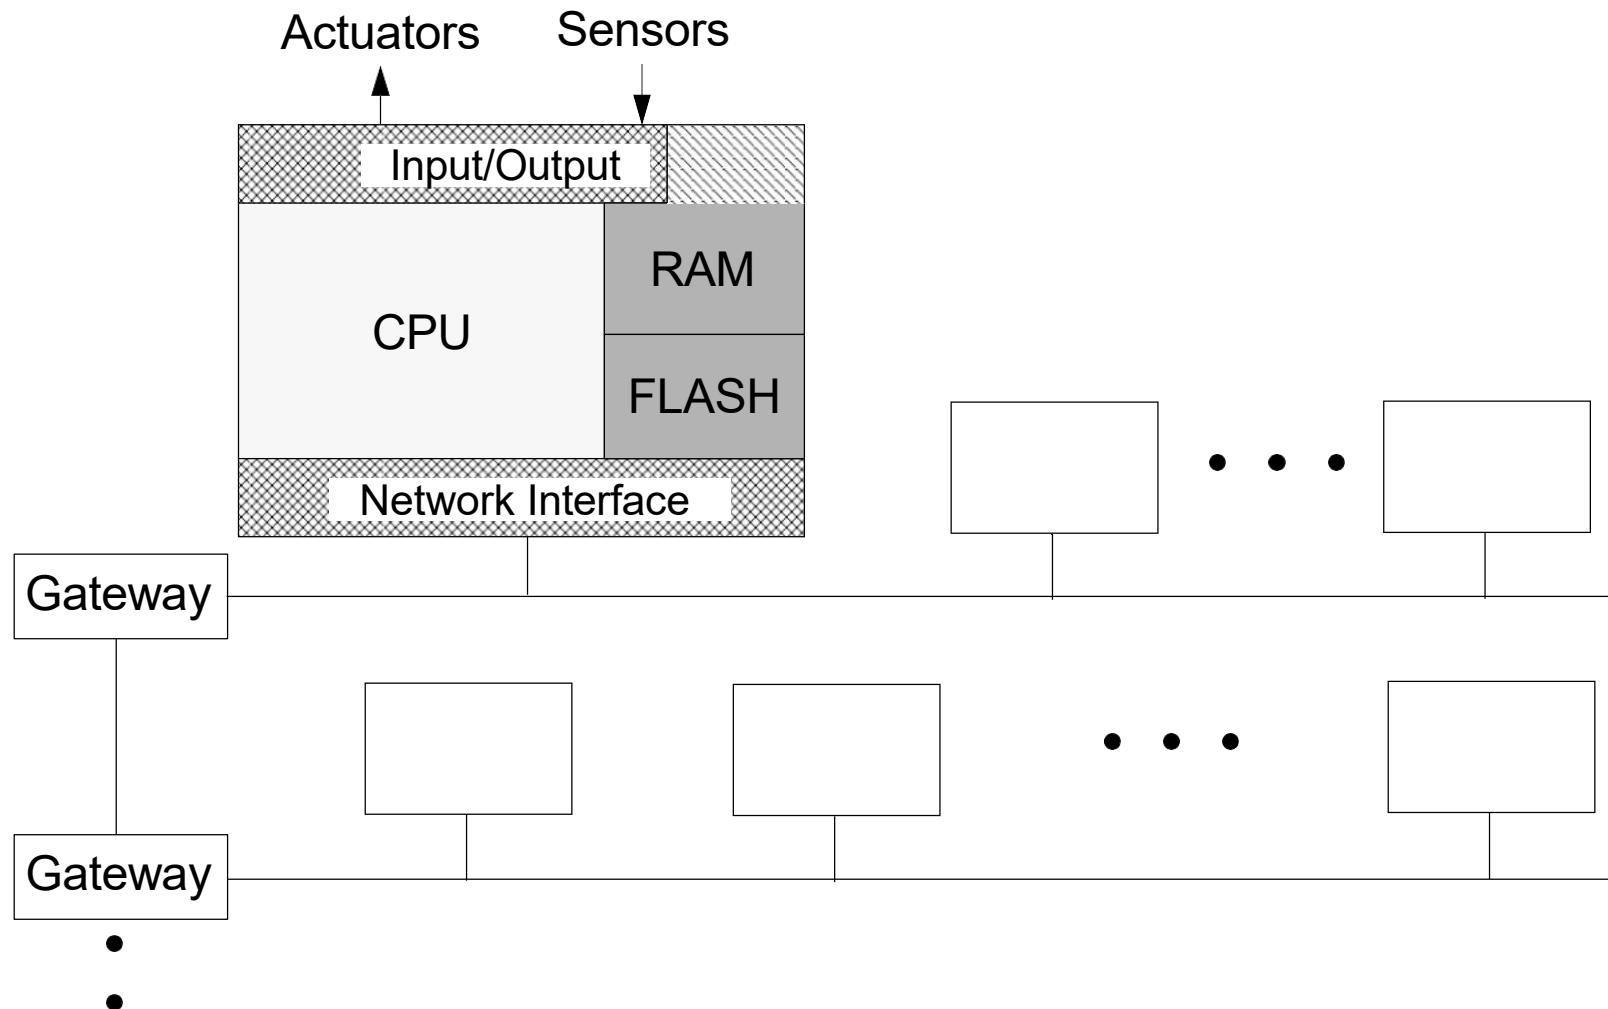
Two Typical Implementation Architectures

Telecommunication System on Chip



Two Typical Implementation Architectures

Distributed Embedded System (automotive application)



The Software Component

Software running on the programmable processors:

- Application tasks
- Real-Time Operating System
- I/O drivers, Network protocols, Middleware

Characteristics of Embedded Applications

What makes them special?

- Like with “ordinary” applications, functionality and user interfaces are often *very complex*.

But, in addition to this:

- Time constraints
- Power constraints
- Cost constraints
- Safety
- Time to market

Time constraints

- Embedded systems have to perform in real-time: if data is not ready by a certain deadline, the system fails to perform correctly.
 - *Hard deadline*: failure to meet leads to major hazards.
 - *Soft deadline*: failure to meet is tolerated but affects quality of service.

Power constraints

Cost constraints

- Embedded systems are very often mass products in highly competitive markets and have to be shipped at a low cost.

What we are interested in:

- Manufacturing cost
- Design cost
- Material cost (Bill of Material)
- Warranty cost

Safety

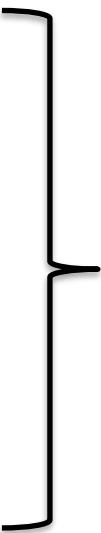
- Embedded systems are often used in life critical applications: avionics, automotive electronics, nuclear plants, medical applications, military applications, etc.
 - *Reliability* and *safety* are major requirements. In order to guarantee safety during design:
 - Formal verification: mathematics-based methods to verify certain properties of the designed system.
 - Automatic synthesis: certain design steps are automatically performed by design tools.

Short time to market

- In highly competitive markets it is critical to catch the *market window*: a short delay with the product on the market can have catastrophic financial consequences (even if the quality of the product is excellent).
 - Design time has to be reduced!
 - Good design methodologies.
 - Efficient design tools.
 - Reuse of previously designed and verified (hardw&softw) blocks.
 - *Good designers who understand both software and hardware!*

Why is Design of Embedded Systems Difficult?

- High Complexity
- Strong time&power constraints
- Low cost
- Short time to market
- Safety critical systems



In order to achieve these requirements, systems have to be highly optimized.

Why is Design of Embedded Systems Difficult?

- High Complexity
- Strong time&power constraints
- Low cost
- Short time to market
- Safety critical systems

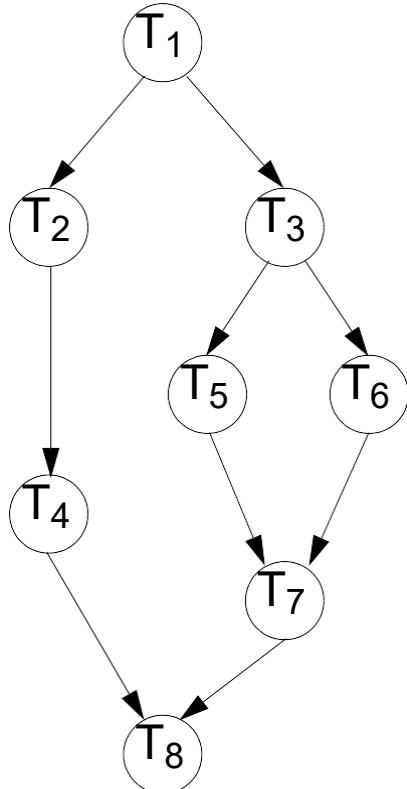


In order to achieve these requirements, systems have to be highly optimized.



Both hardware and software aspects have to be considered simultaneously!

An Example



The system to be implemented is modelled as a *task graph*:

- a node represents a *task* (a unit of functionality activated as response to a certain input and which generates a certain output).
- an edge represents a precedence constraint and data dependency between two tasks.

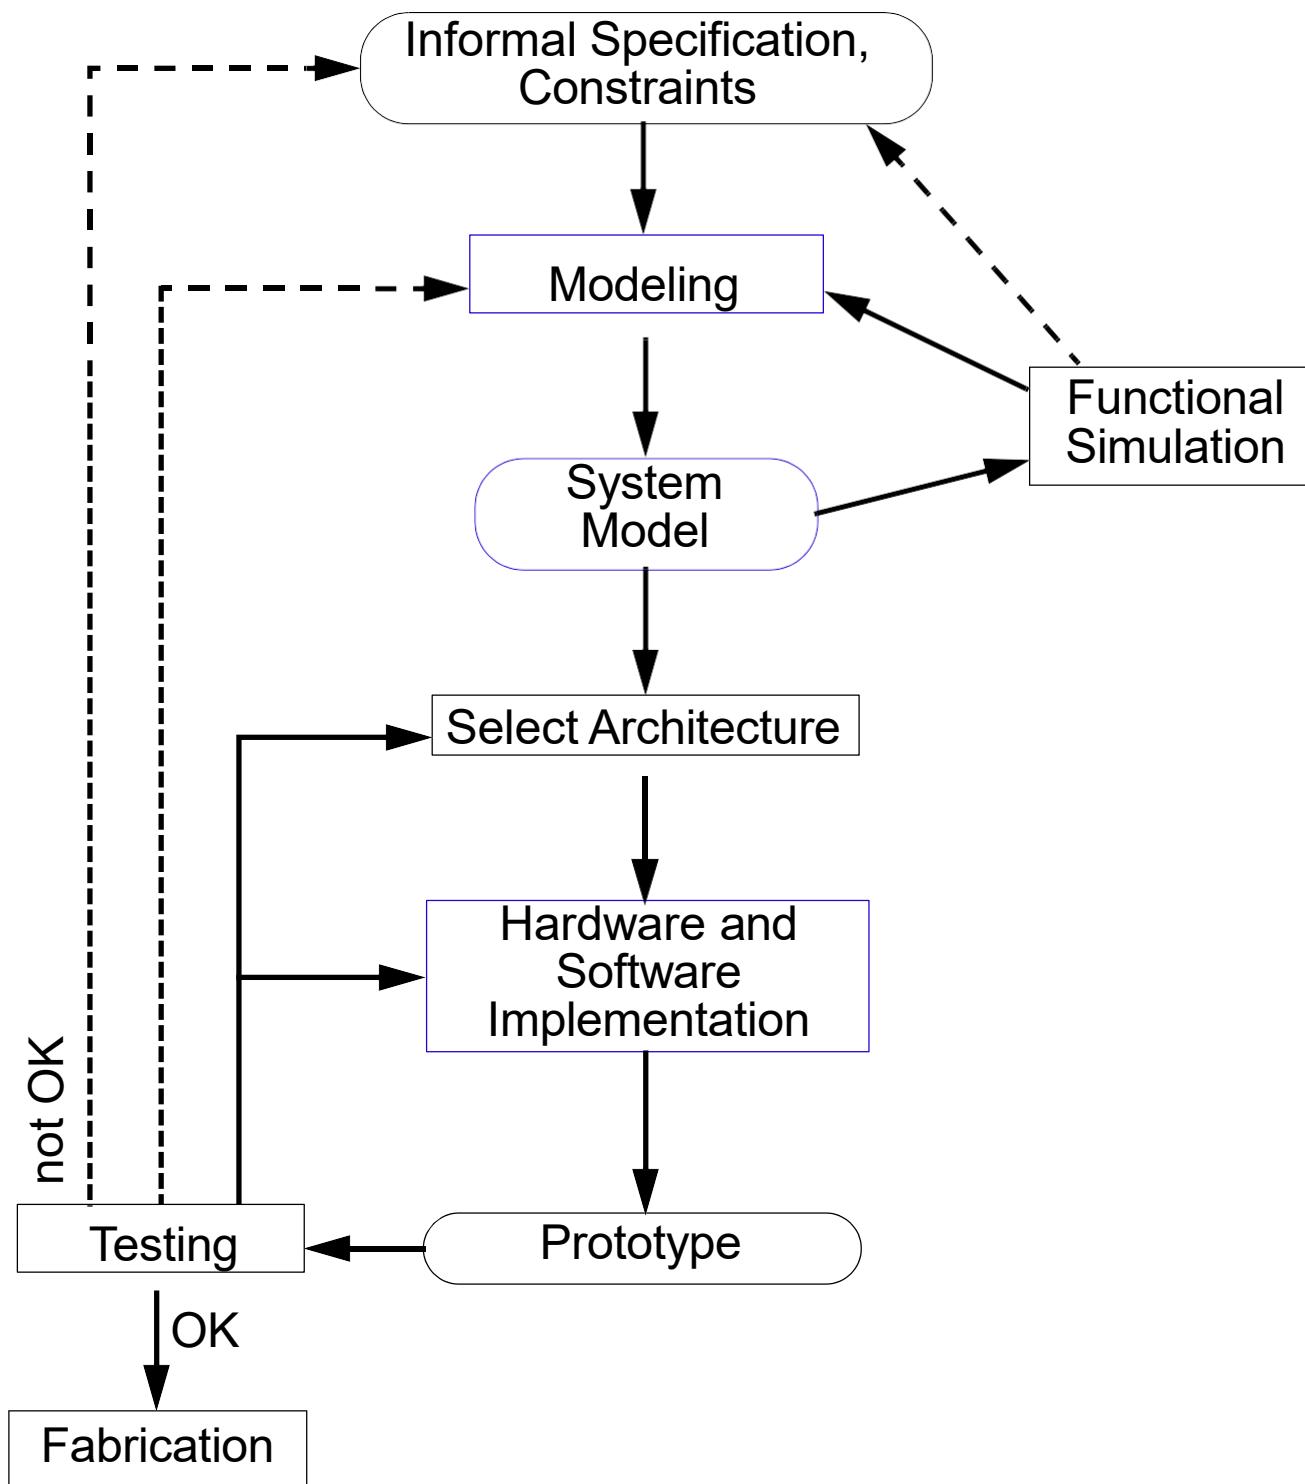
Period: 42 time units

- The task graph is activated every 42 time units \Rightarrow an activation has to terminate in time less than 42.

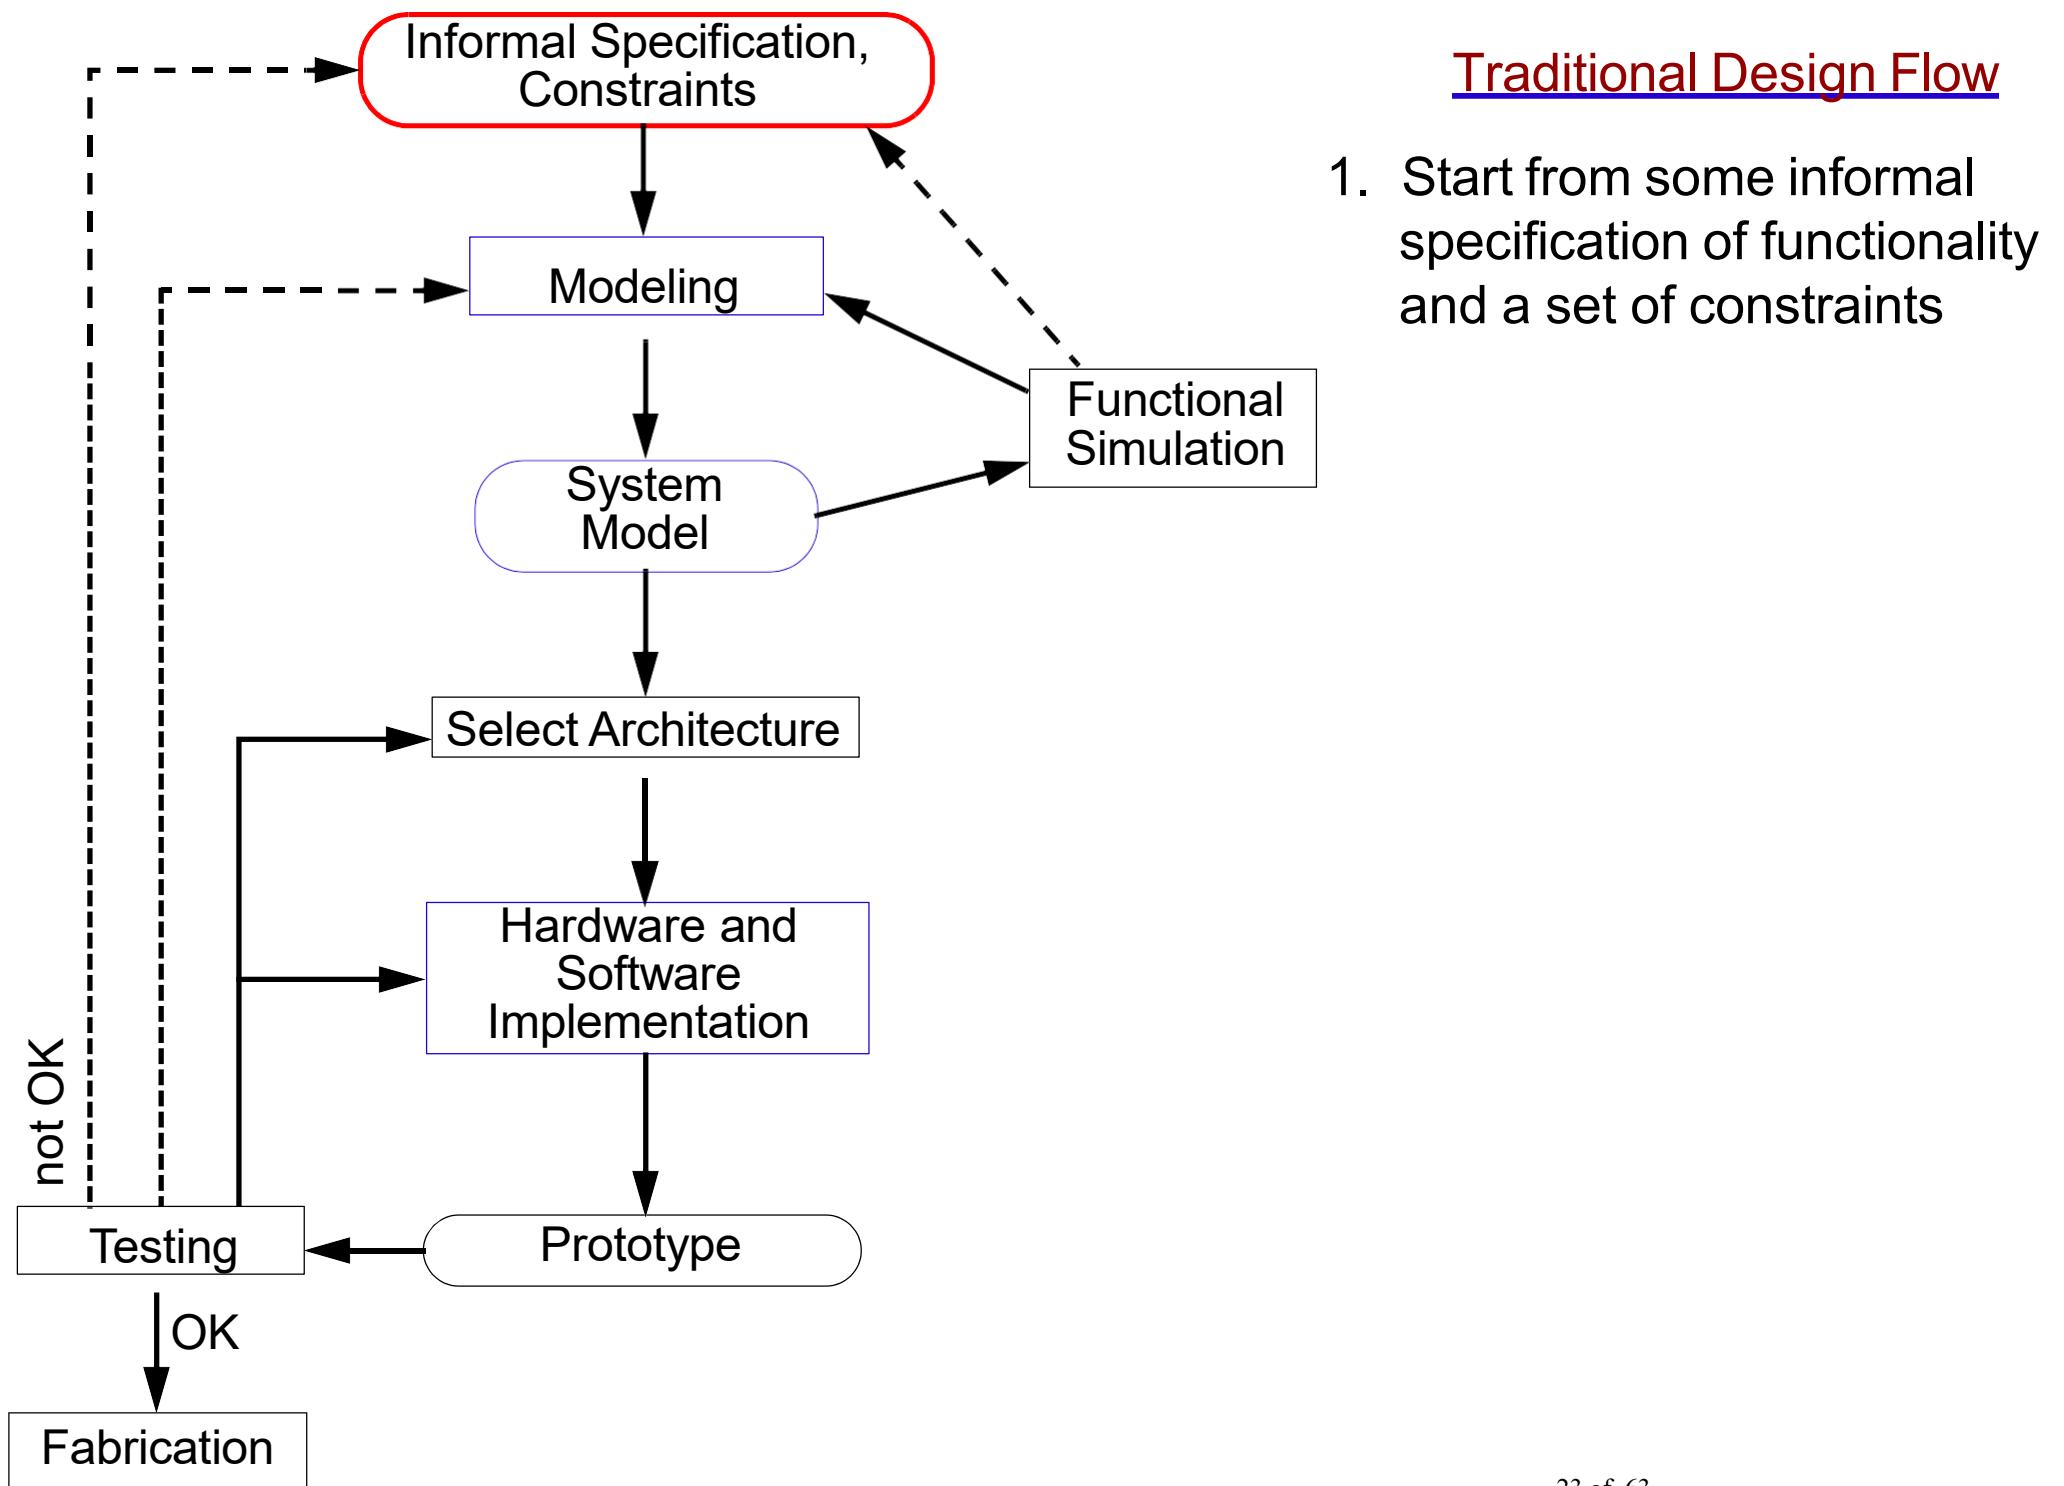
Cost limit: 8

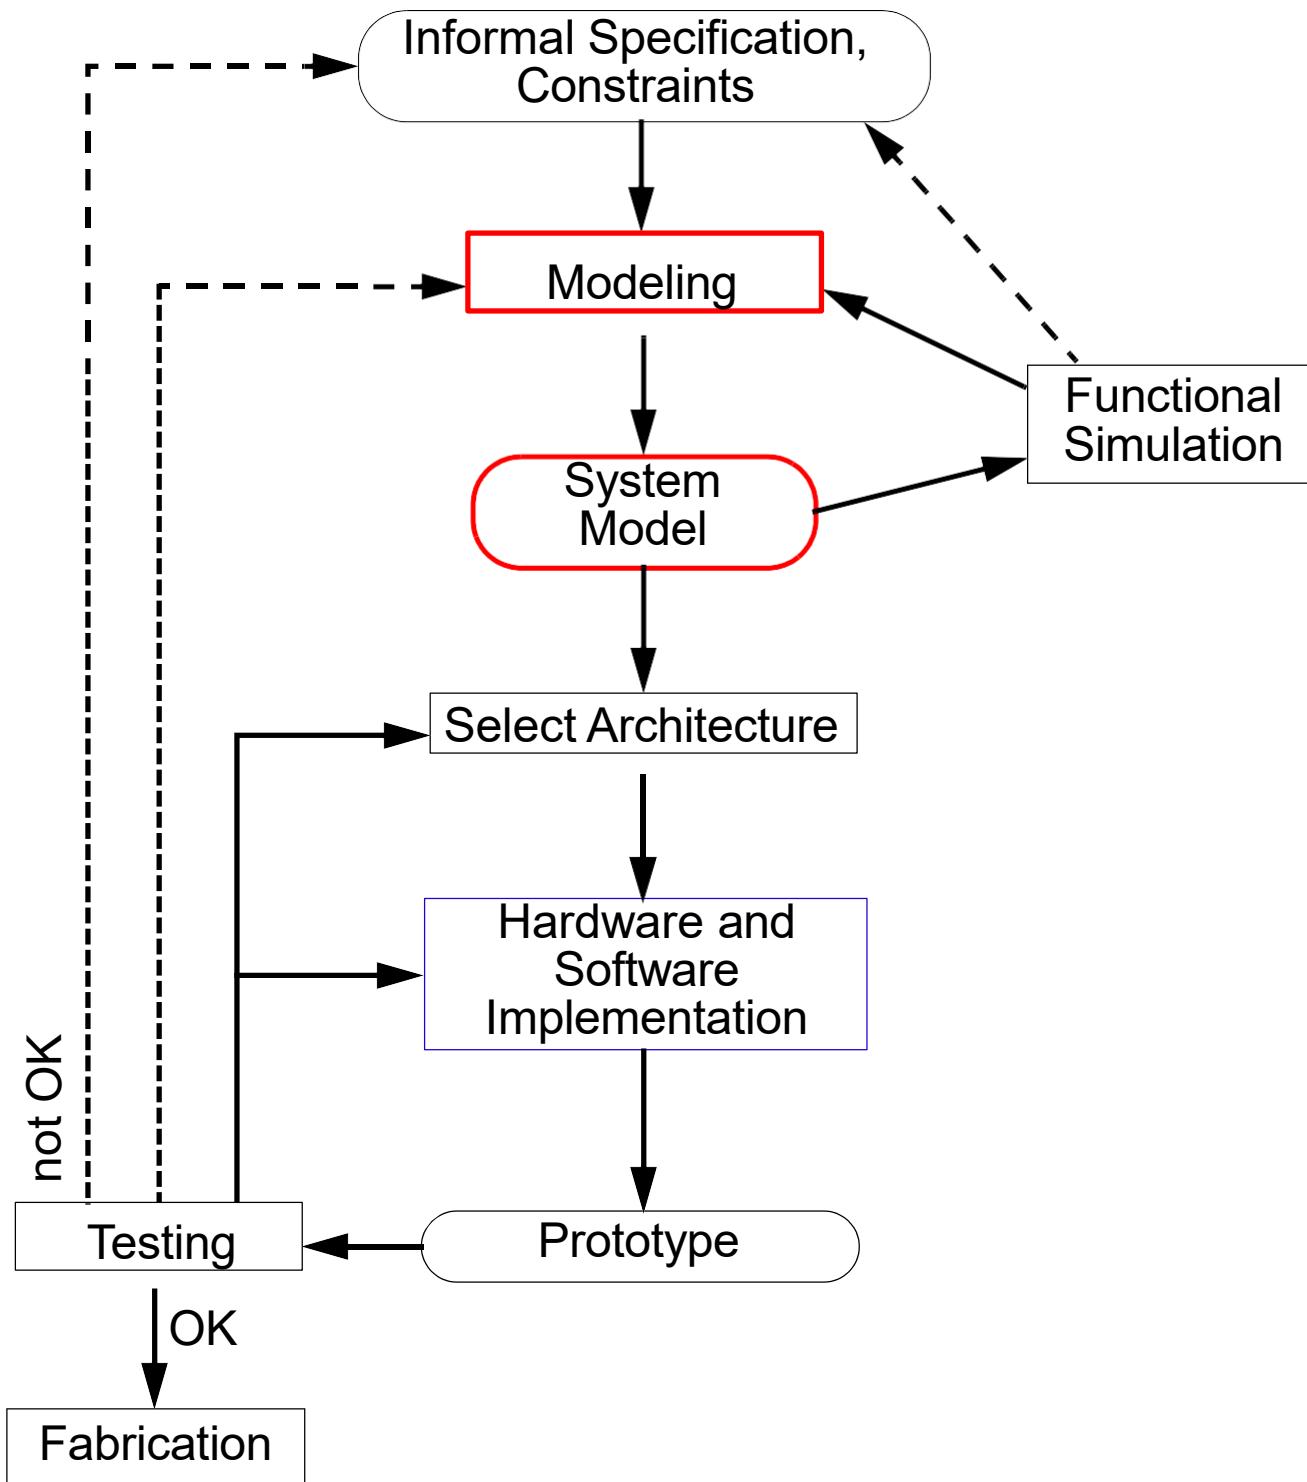
- The total cost of the implemented system has to be less than 8.

Traditional Design Flow



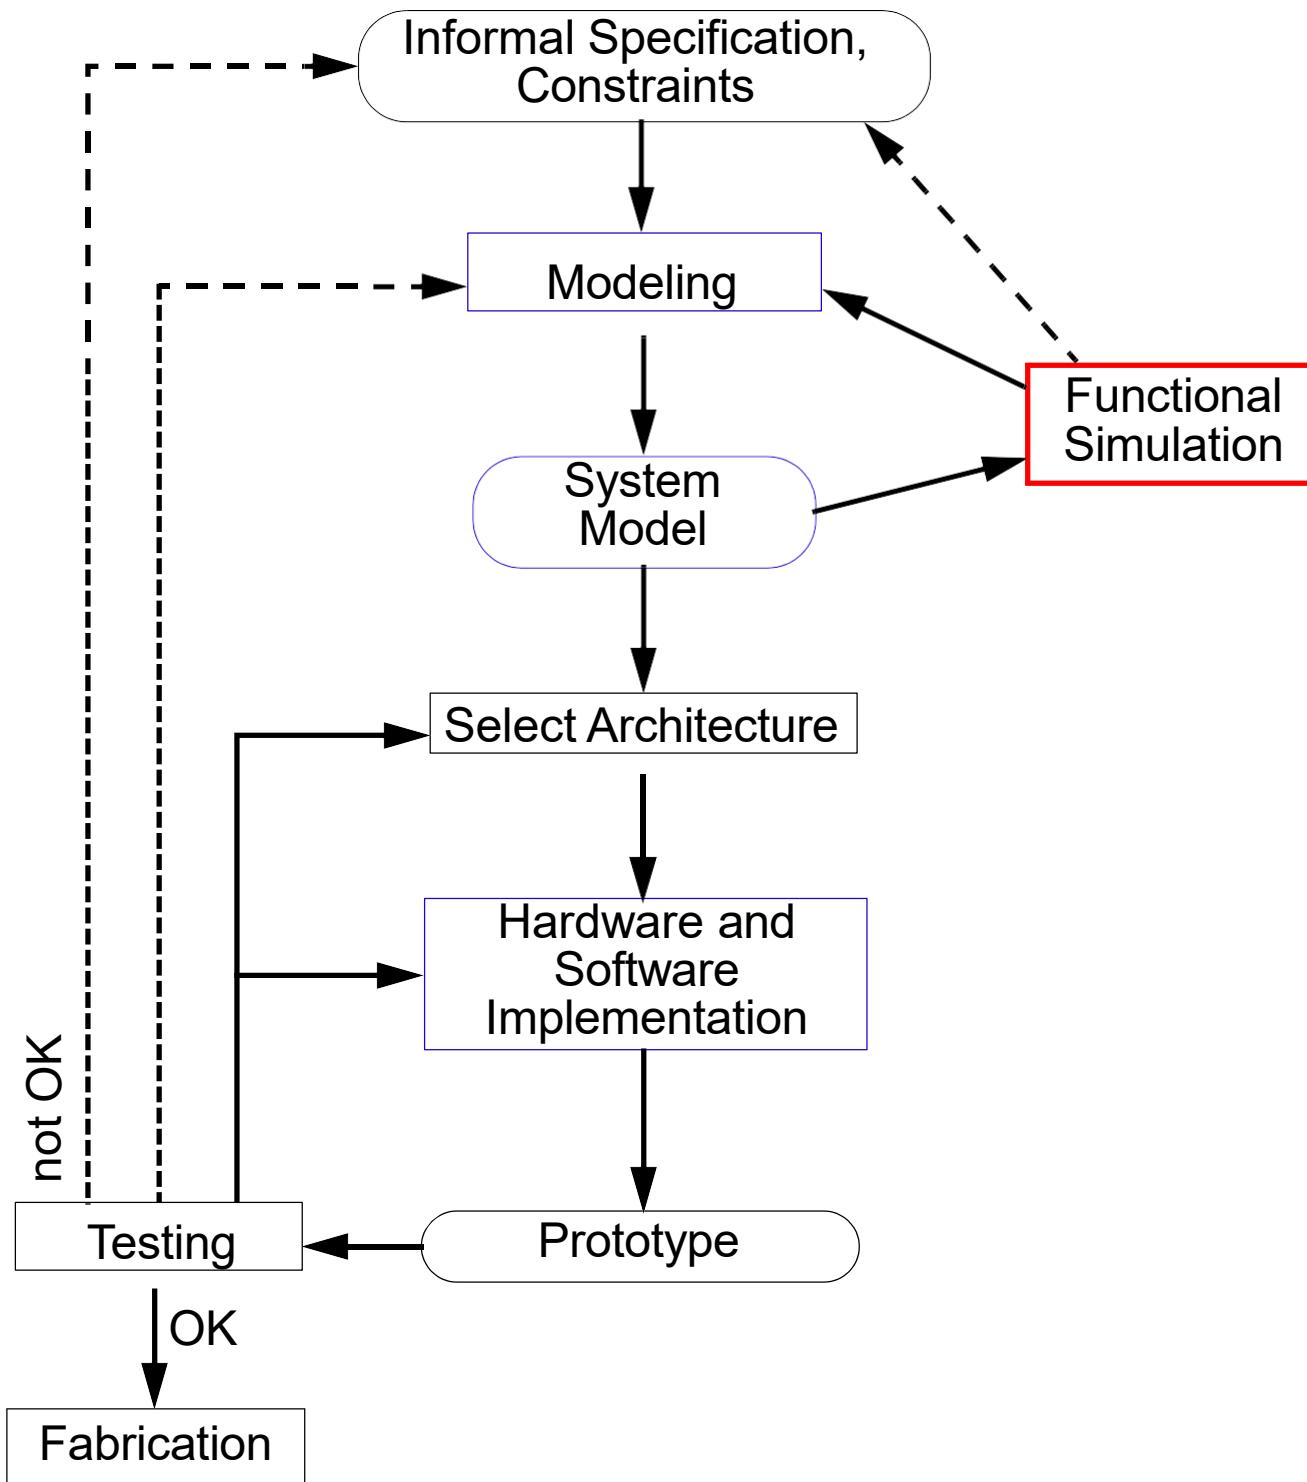
Traditional Design Flow





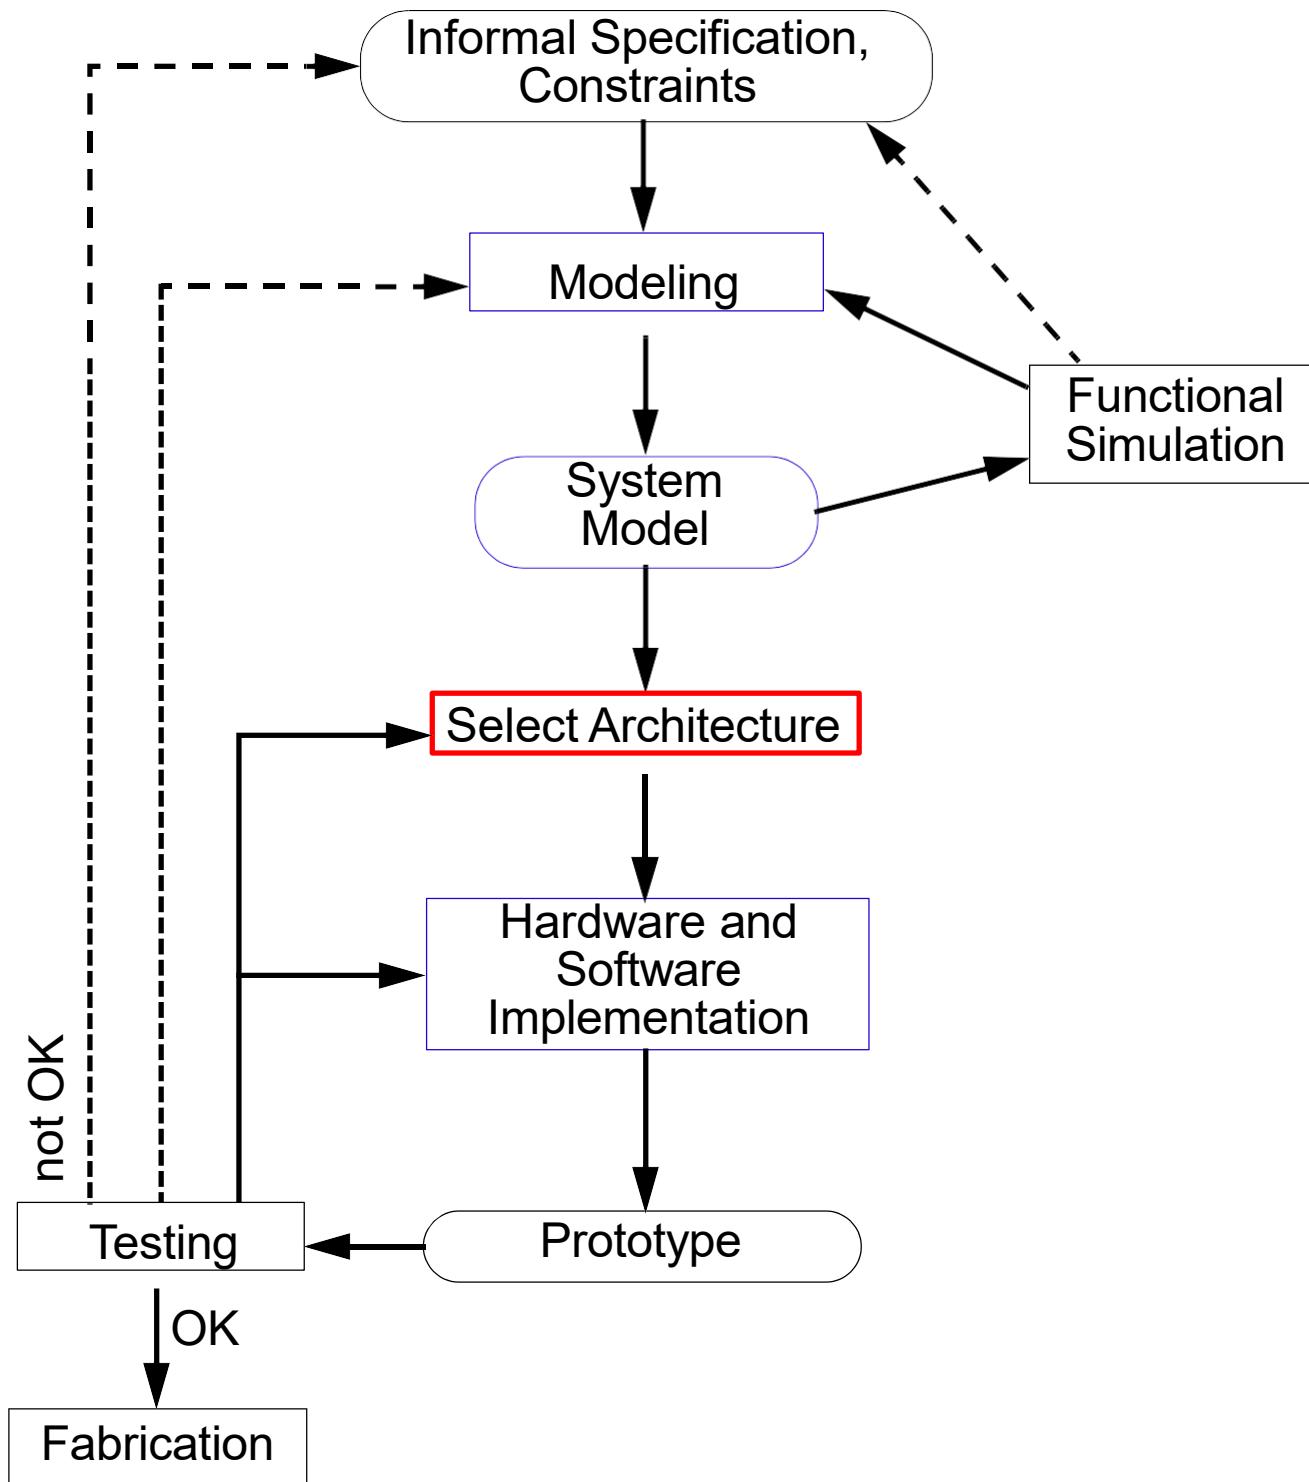
Traditional Design Flow

1. Start from some informal specification of functionality and a set of constraints
2. Generate a more formal model of the functionality, based on some modeling concept. Such model is our task graph



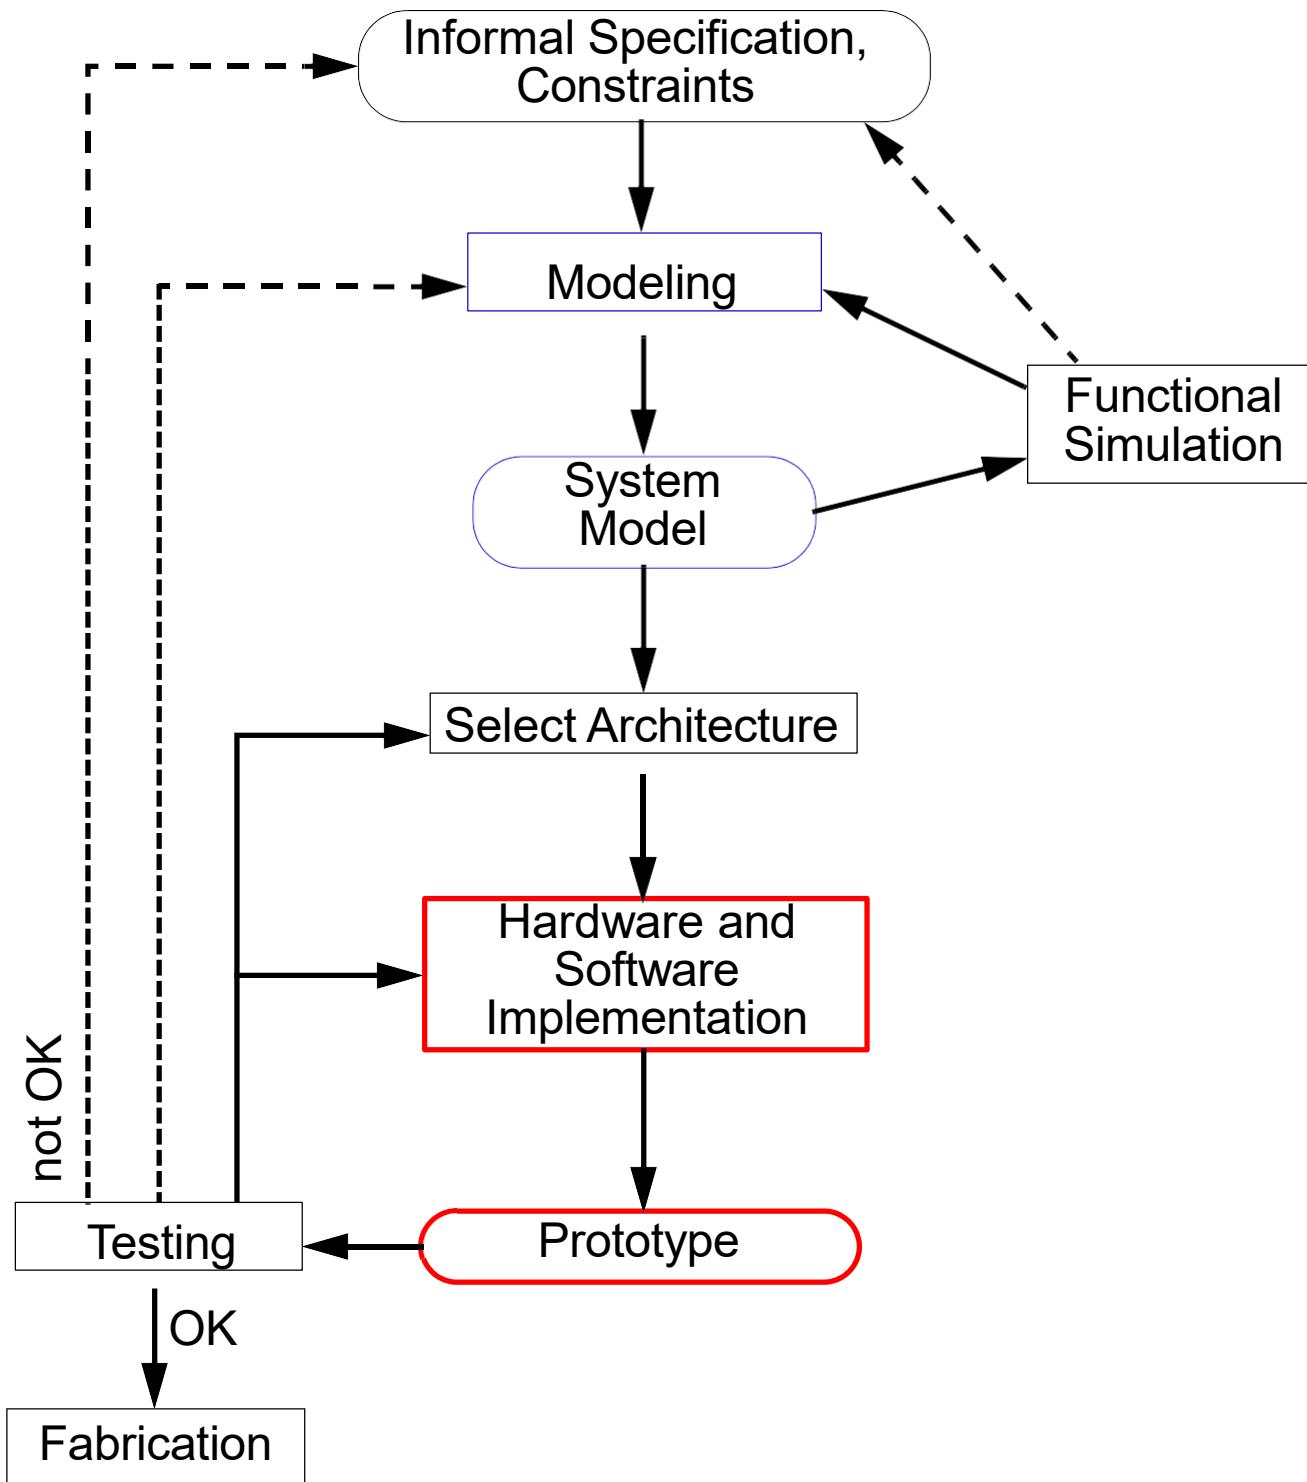
Traditional Design Flow

1. Start from some informal specification of functionality and a set of constraints
2. Generate a more formal model of the functionality, based on some modeling concept. Such model is our task graph
3. Simulate the model in order to check the functionality. If needed make adjustments.



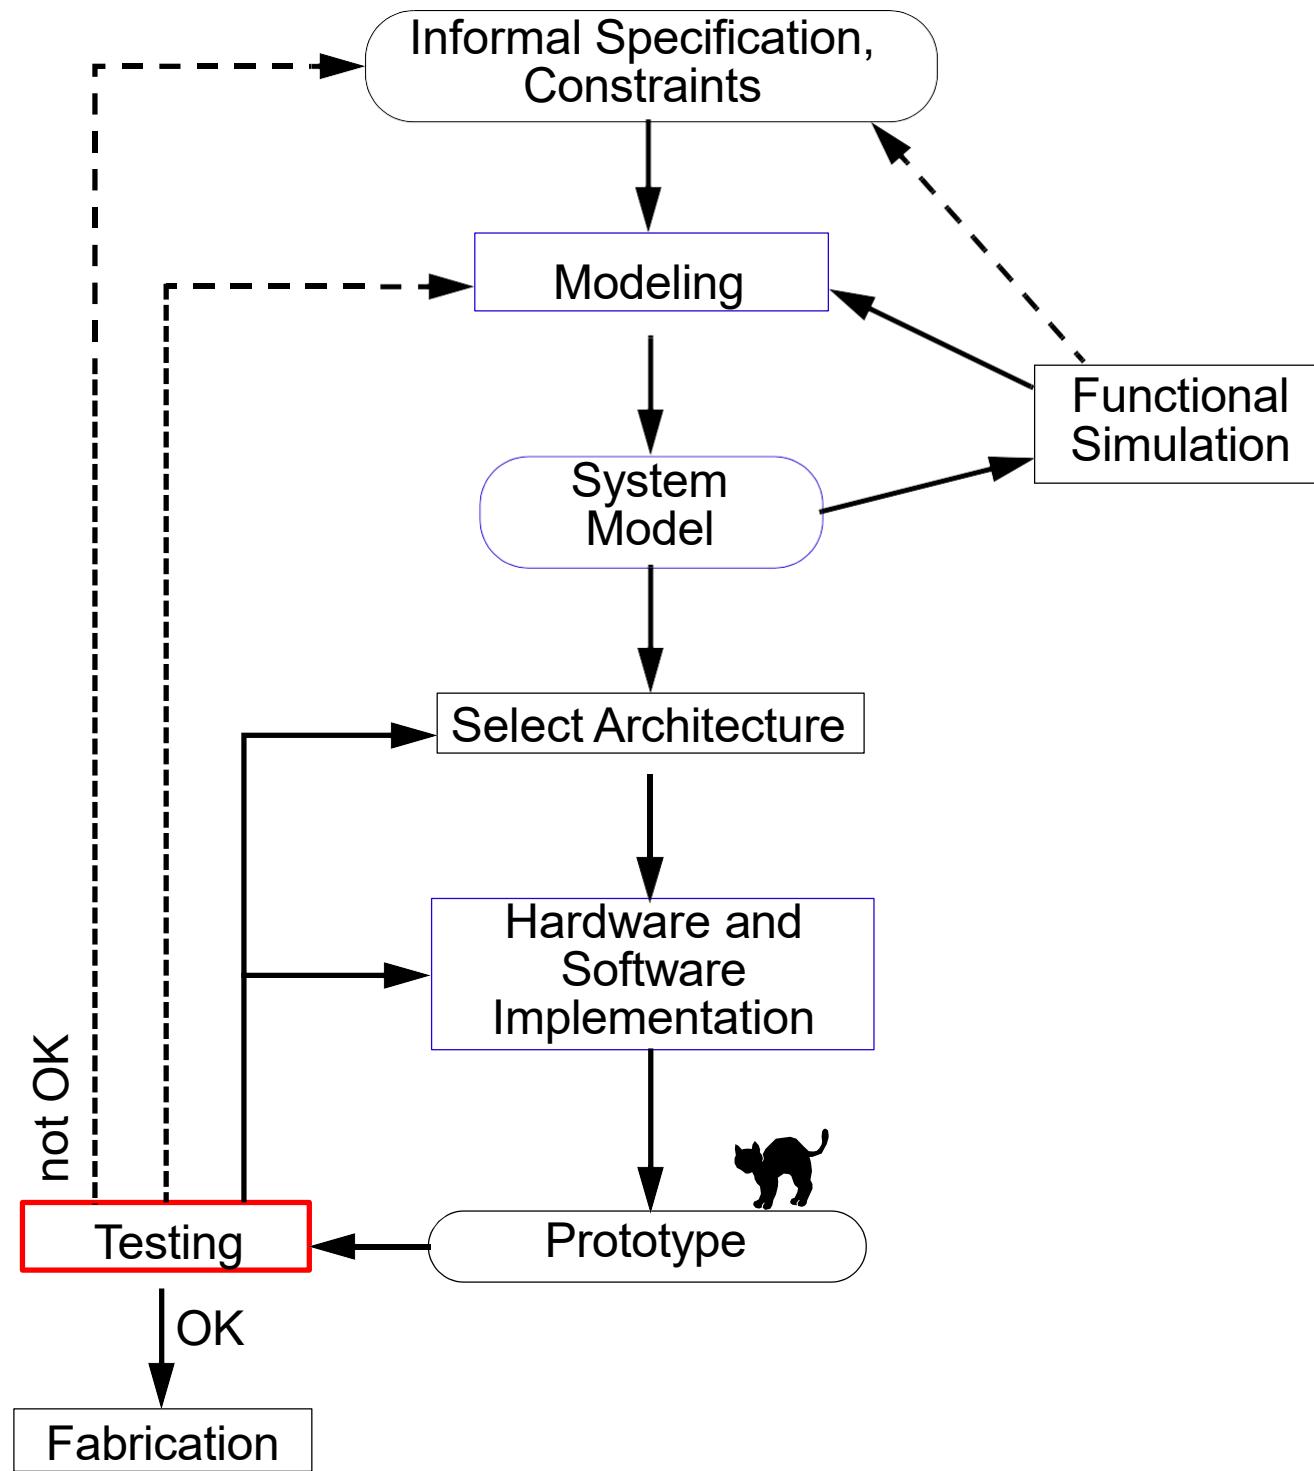
Traditional Design Flow

1. Start from some informal specification of functionality and a set of constraints
2. Generate a more formal model of the functionality, based on some modeling concept. Such model is our task graph
3. Simulate the model in order to check the functionality. If needed make adjustments.
4. Choose an architecture (μ processor, buses, etc.) such that cost limits are satisfied and, you hope, time and power constraints are fulfilled.



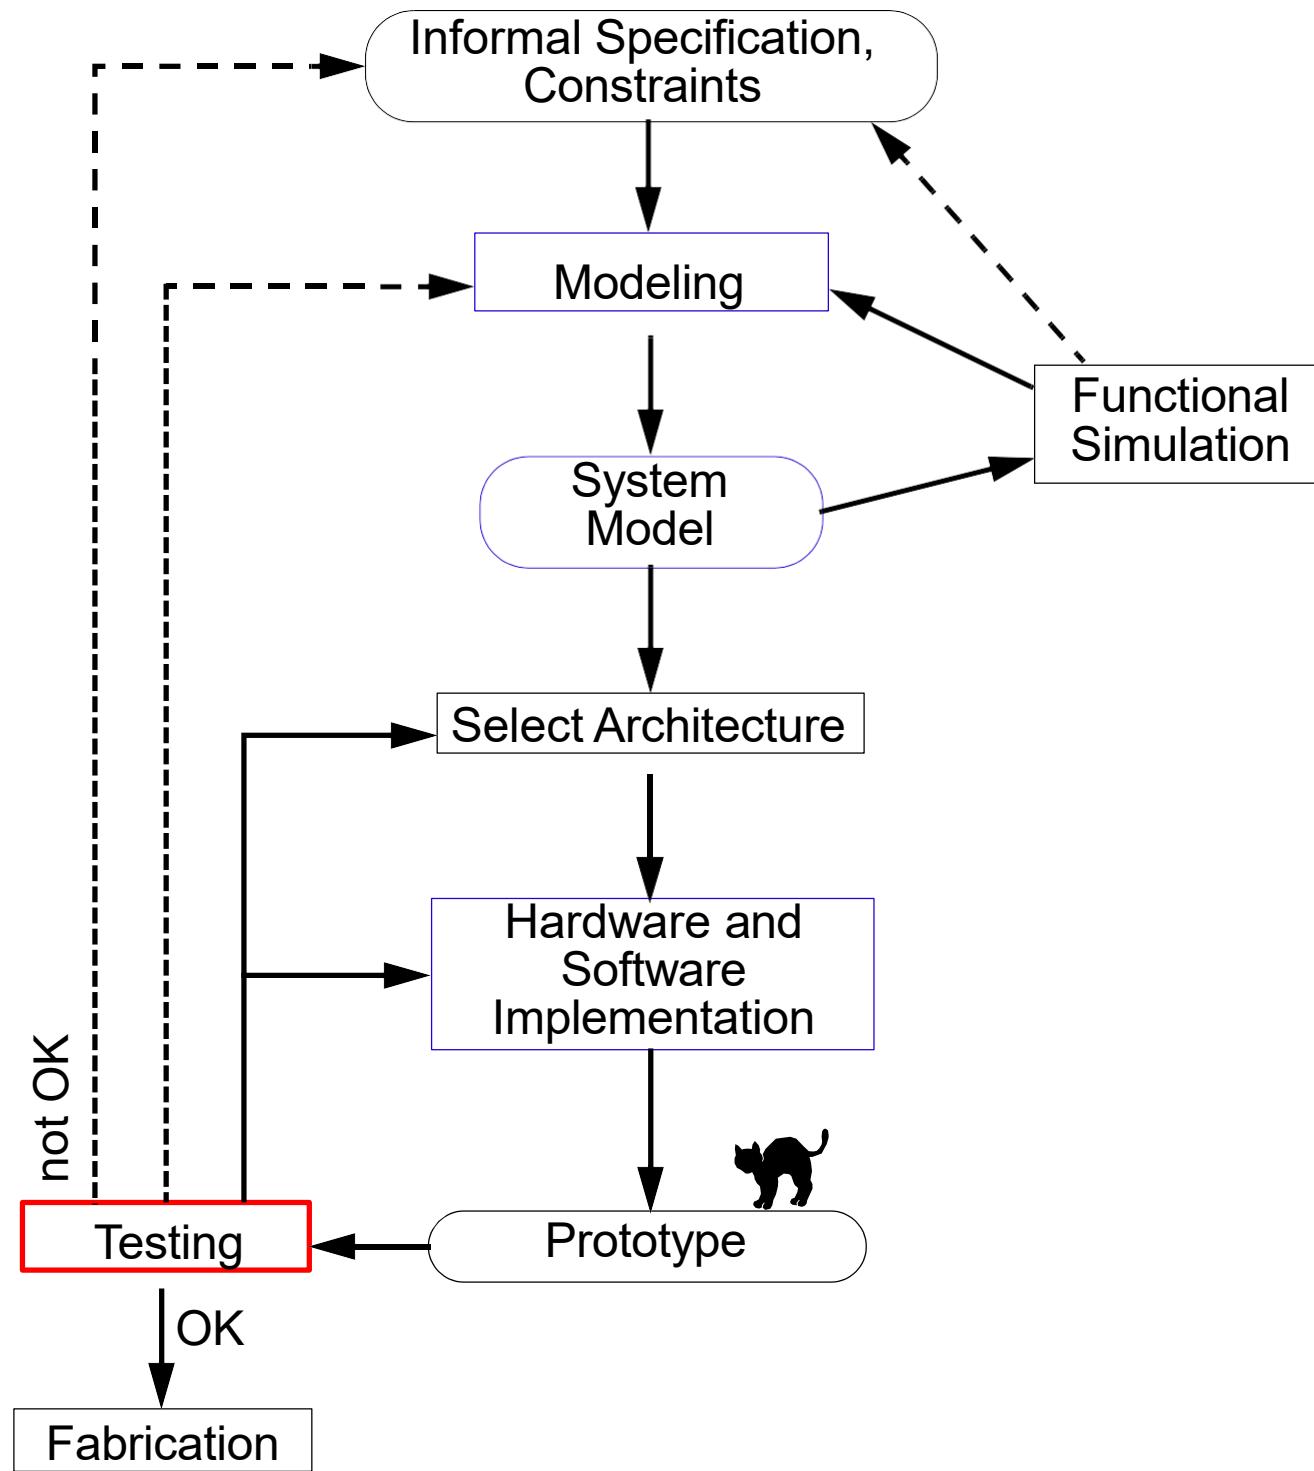
Traditional Design Flow

1. Start from some informal specification of functionality and a set of constraints
2. Generate a more formal model of the functionality, based on some modeling concept. Such model is our task graph
3. Simulate the model in order to check the functionality. If needed make adjustments.
4. Choose an architecture (μ processor, buses, etc.) such that cost limits are satisfied and, you hope, time and power constraints are fulfilled.
5. Build a prototype and implement the system.



Traditional Design Flow

1. Start from some informal specification of functionality and a set of constraints
2. Generate a more formal model of the functionality, based on some modeling concept. Such model is our task graph
3. Simulate the model in order to check the functionality. If needed make adjustments.
4. Choose an architecture (μ processor, buses, etc.) such that cost limits are satisfied and, you hope, time and power constraints are fulfilled.
5. Build a prototype and implement the system.
6. Verify the system: neither time nor power constraints are satisfied!!!



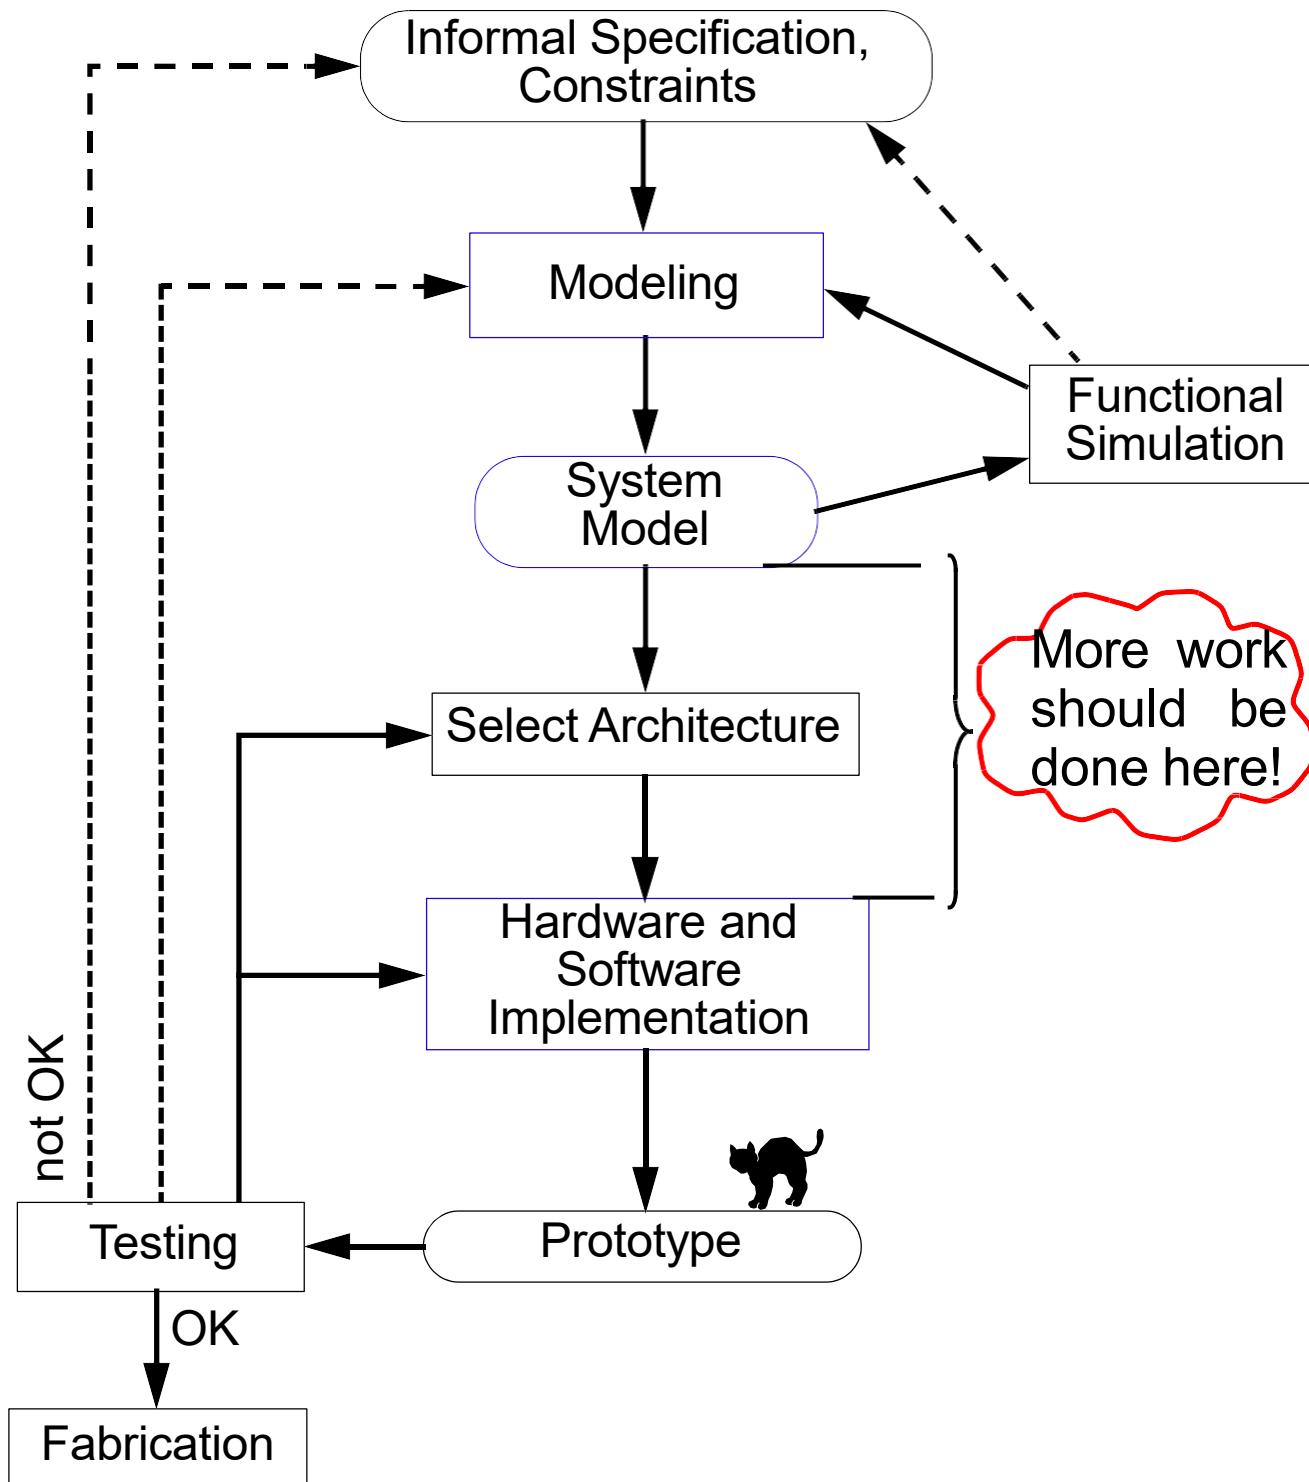
Traditional Design Flow

Now you are in great trouble: you have spent a lot of time and money and nothing works!

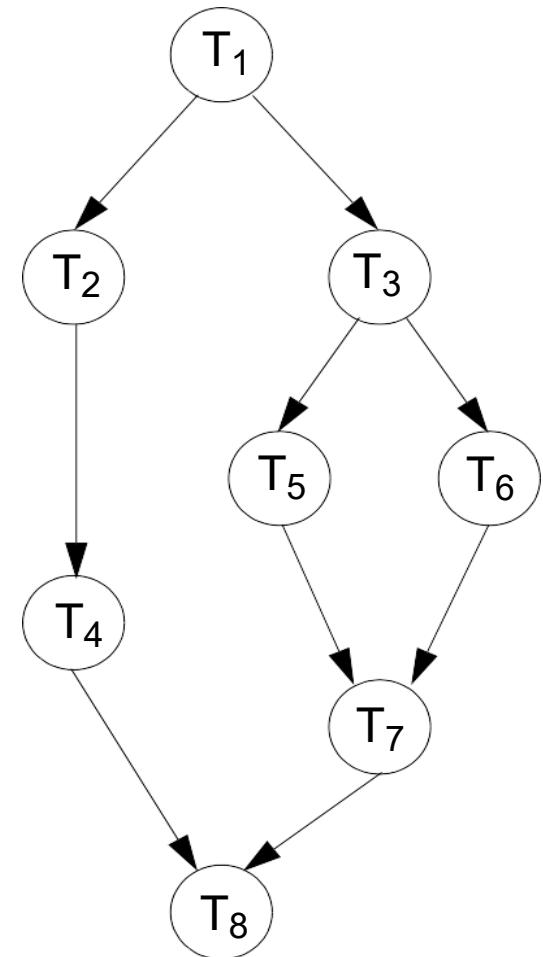
- Go back to 4, choose a new architecture and start a new implementation.
- Or negotiate with the customer on the constraints.

The Traditional Design Flow

- The consequences:
 - Delays in the design process
 - Increased design cost
 - Delays in time to market \Rightarrow missed market window
 - High cost of failed prototypes
 - Bad design decisions taken under time pressure
 - Low quality, high cost products

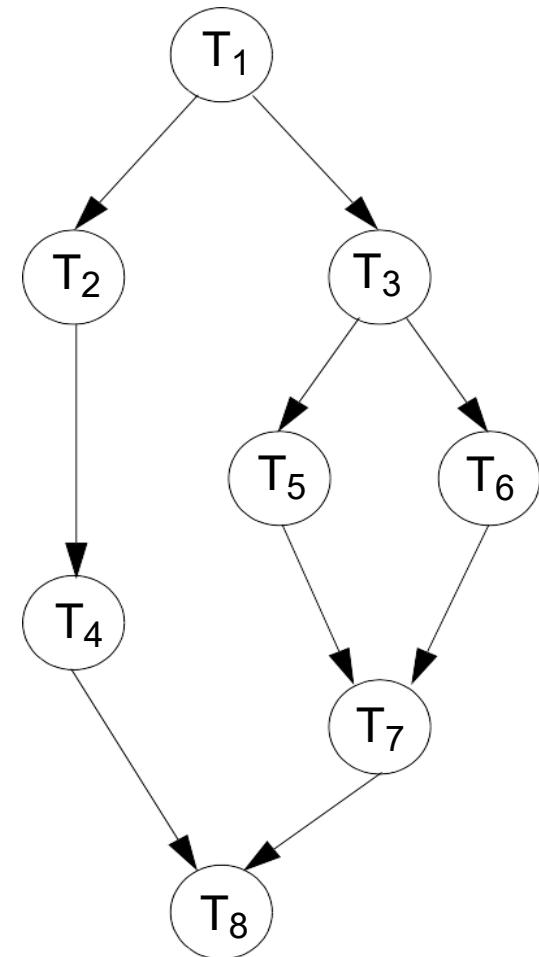


Example

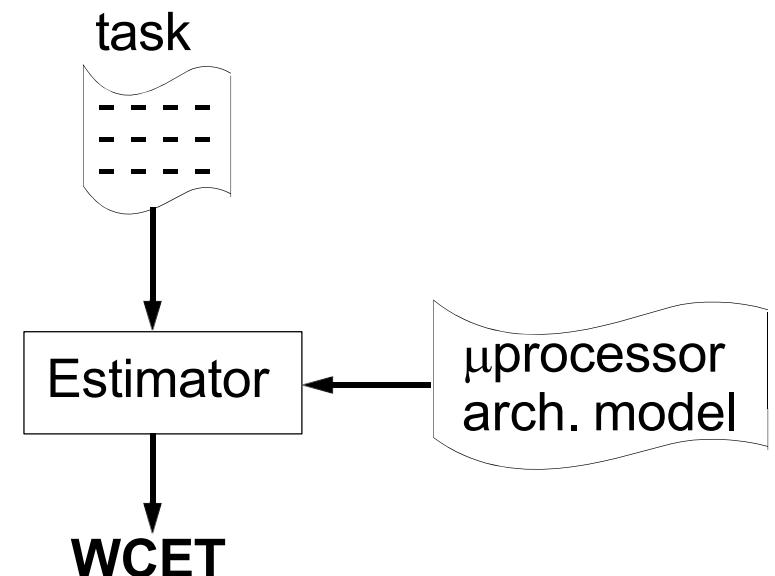


- We have the system model (task graph) which has been validated by simulation.
- We decide on a certain μ processor $\mu p1$, with cost 6.
- For each task the worst-case execution time (WCET) when run on $\mu p1$ is *estimated*.

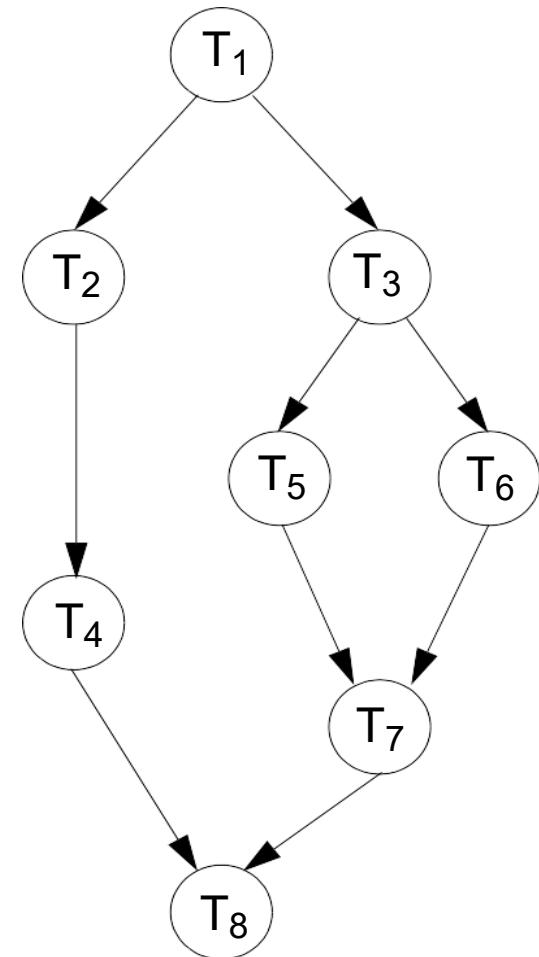
Example



- We have the system model (task graph) which has been validated by simulation.
- We decide on a certain μ processor $\mu p1$, with cost 6.
- For each task the worst-case execution time (WCET) when run on $\mu p1$ is *estimated*.

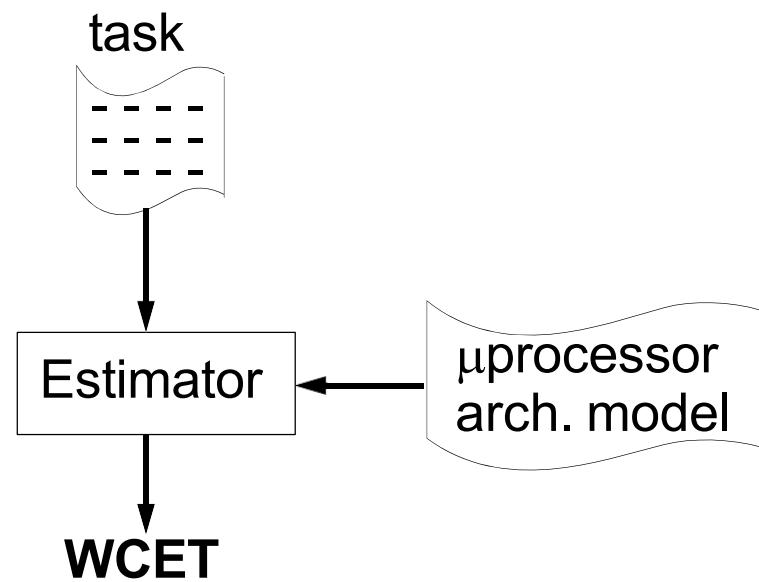


Example



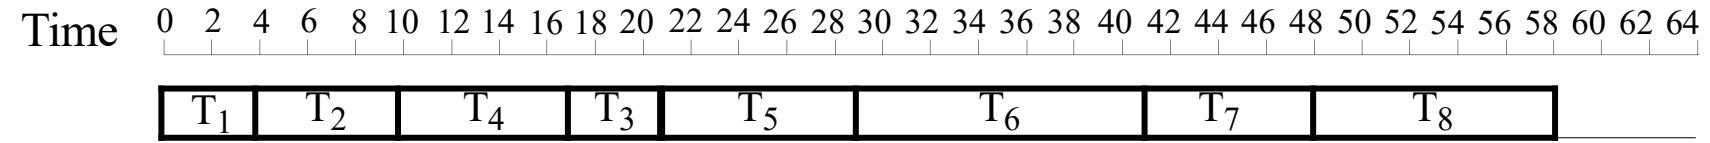
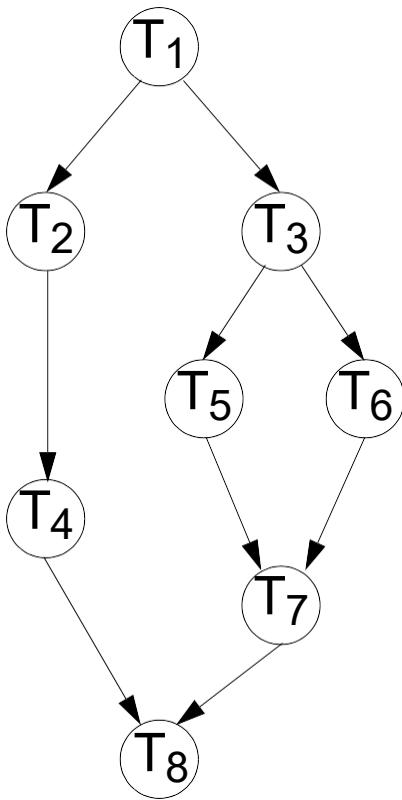
- We have the system model (task graph) which has been validated by simulation.
- We decide on a certain μ processor $\mu p1$, with cost 6.
- For each task the worst-case execution time (WCET) when run on $\mu p1$ is *estimated*.

Task	WCET
T ₁	4
T ₂	6
T ₃	4
T ₄	7
T ₅	8
T ₆	12
T ₇	7
T ₈	10



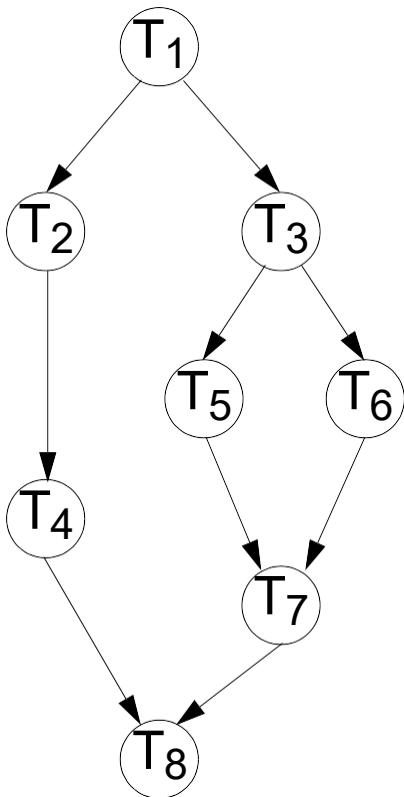
Example

We generate a schedule:

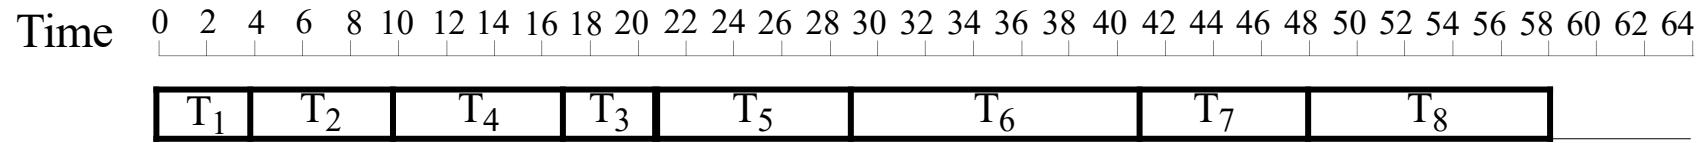


Task	WCET
T ₁	4
T ₂	6
T ₃	4
T ₄	7
T ₅	8
T ₆	12
T ₇	7
T ₈	10

Example



We generate a schedule:



Using the architecture with μ processor μ p1 we got a solution with:

- Execution time: $58 > 42$
- Cost: $6 < 8$

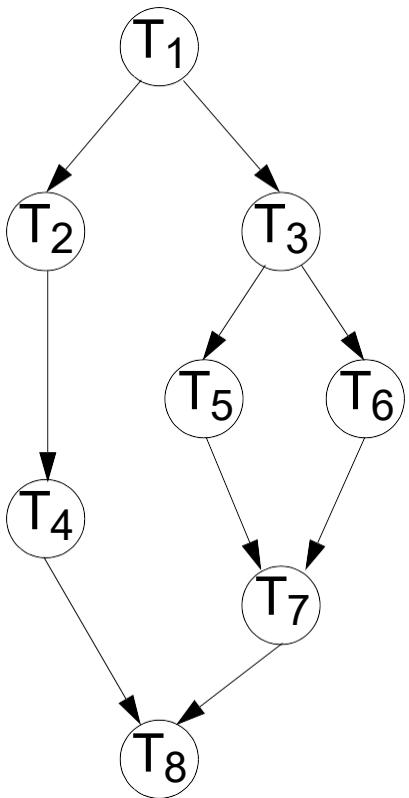


We have to try with another architecture!

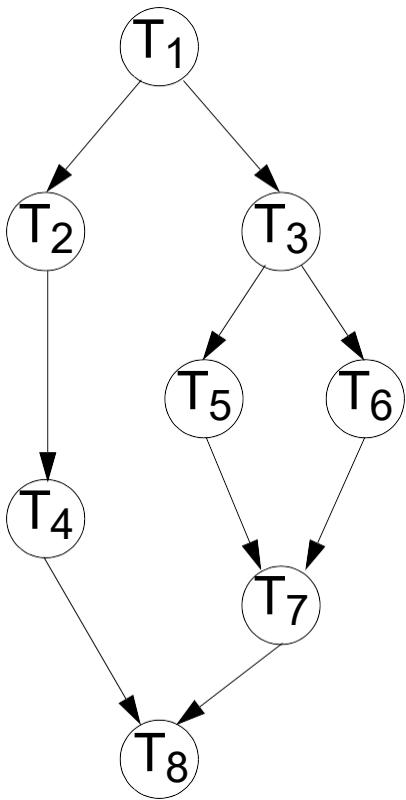
Task	WCET
T ₁	4
T ₂	6
T ₃	4
T ₄	7
T ₅	8
T ₆	12
T ₇	7
T ₈	10

Example

We look after a μ processor which is fast enough: μ p2



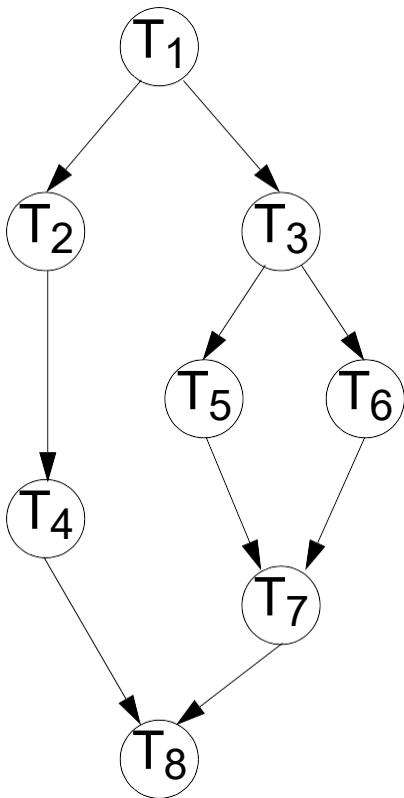
Example



We look after a μ processor which is fast enough: μ p2

For each task the WCET, when run on μ p2, is estimated.

Task	WCET
T ₁	2
T ₂	3
T ₃	2
T ₄	3
T ₅	4
T ₆	6
T ₇	3
T ₈	5



Example

We look after a μ processor which is fast enough: μ p2

For each task the WCET, when run on μ p2, is estimated.

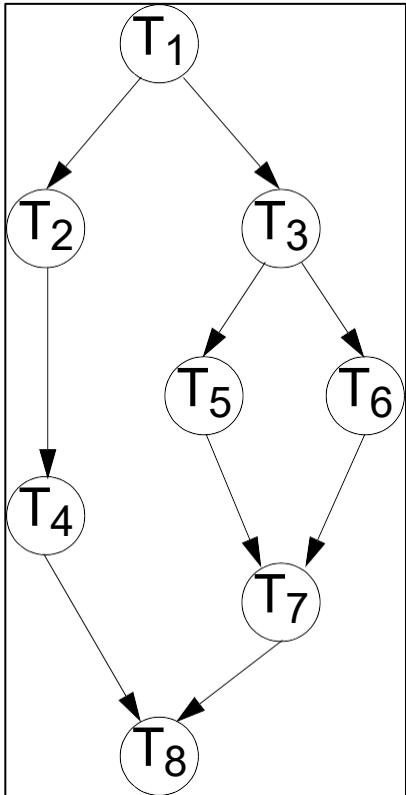
Using the architecture with μ processor μ p2, after generating a schedule, we got a solution with:

- Execution time: $28 < 42$
- Cost: $15 > 8$ 



We have to try with another architecture!

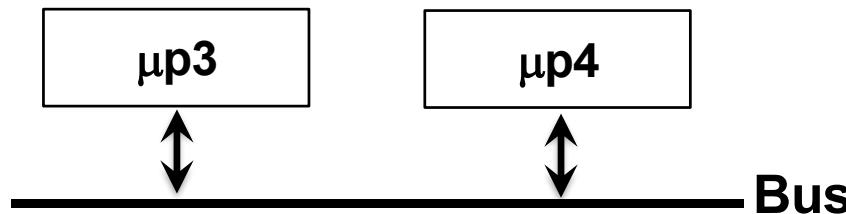
Task	WCET
T ₁	2
T ₂	3
T ₃	2
T ₄	3
T ₅	4
T ₆	6
T ₇	3
T ₈	5



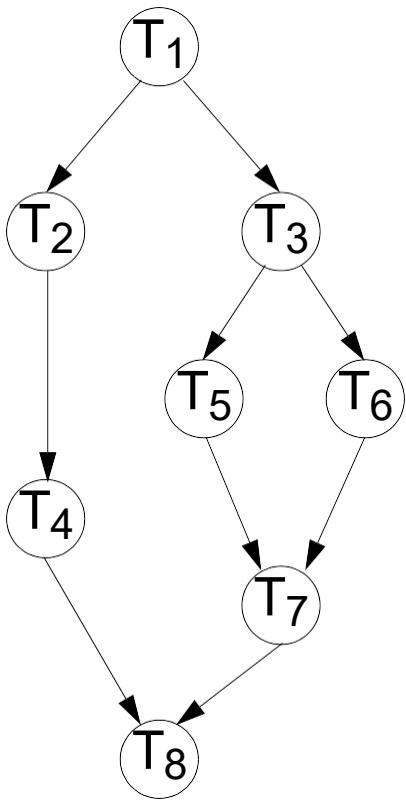
Example

We have to look for a multiprocessor solution

- In order to meet cost constraints try 2 cheap (and slow) µps:
 µp3: cost 3
 µp4: cost 2
 interconnection bus: cost 1

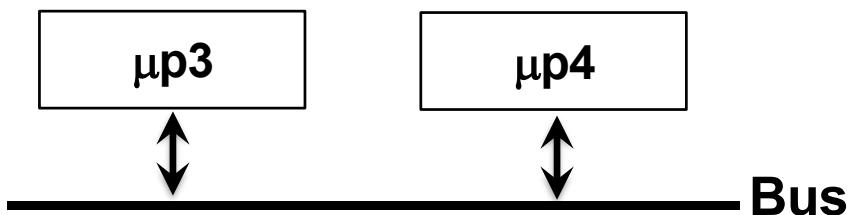


Example



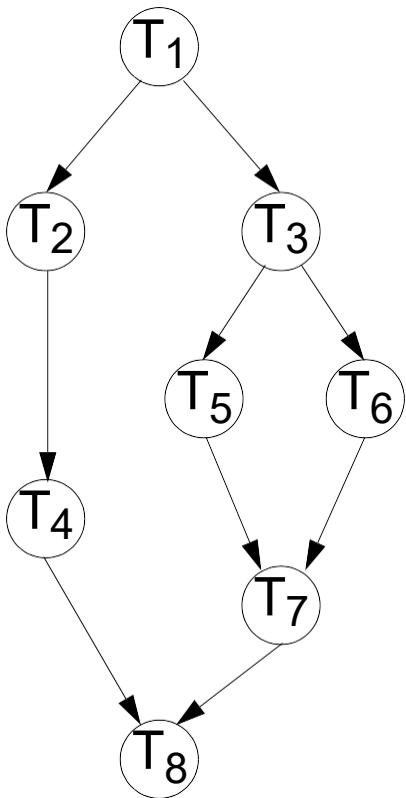
We have to look for a multiprocessor solution

- In order to meet cost constraints try 2 cheap (and slow) µps:
 - µp3: cost 3
 - µp4: cost 2
 - interconnection bus: cost 1



For each task the WCET, when run on µp3 and µp4, is estimated.

Task	WCET	
	µp3	µp4
T ₁	5	6
T ₂	7	9
T ₃	5	6
T ₄	8	10
T ₅	10	11
T ₆	17	21
T ₇	10	14
T ₈	15	19



Example

Now we have to *map* the tasks to processors:

$\mu p3: T_1, T_3, T_5, T_6, T_7, T_8$.

$\mu p4: T_2, T_4$.

If communicating tasks are mapped to different processors, they have to communicate over the bus.

Communication time has to be estimated; it depends on the amount of bits transferred between the tasks and on the speed of the bus.

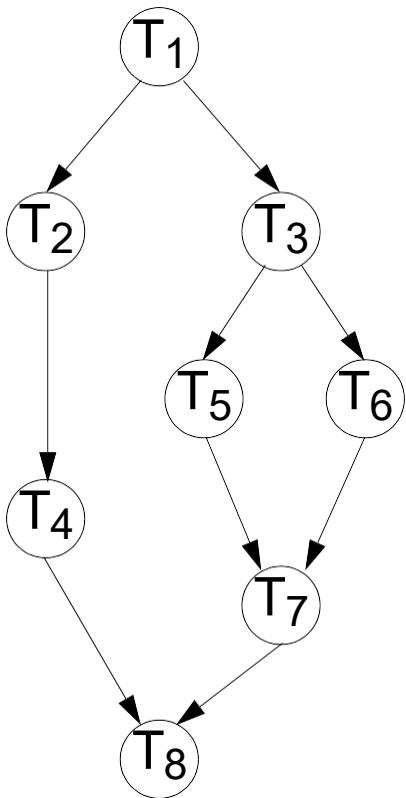
Task	WCET	
	$\mu p3$	$\mu p4$
T ₁	5	6
T ₂	7	9
T ₃	5	6
T ₄	8	10
T ₅	10	11
T ₆	17	21
T ₇	10	14
T ₈	15	19

Estimated communication times:

$C_{1-2}: 1$

$C_{4-8}: 1$

Example



$\mu p3: T_1, T_3, T_5, T_6, T_7, T_8$.

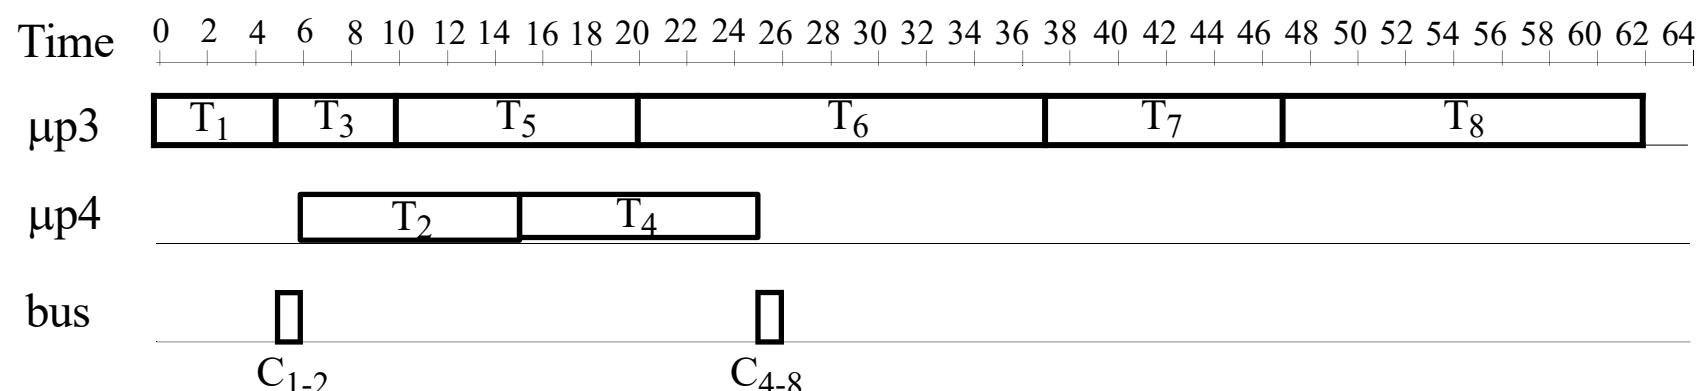
$\mu p4: T_2, T_4$.

Estimated communication times:

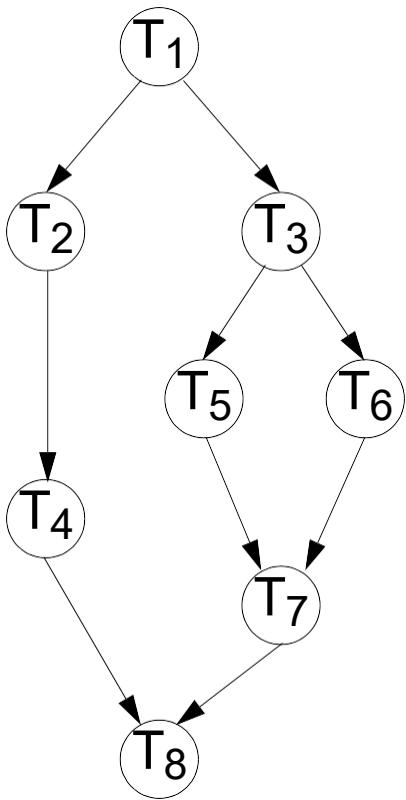
$$\begin{aligned} C_{1-2} &: 1, \\ C_{4-8} &: 1 \end{aligned}$$

We generate a schedule:

Task	WCET	
	$\mu p3$	$\mu p4$
T ₁	5	6
T ₂	7	9
T ₃	5	6
T ₄	8	10
T ₅	10	11
T ₆	17	21
T ₇	10	14
T ₈	15	19



Example



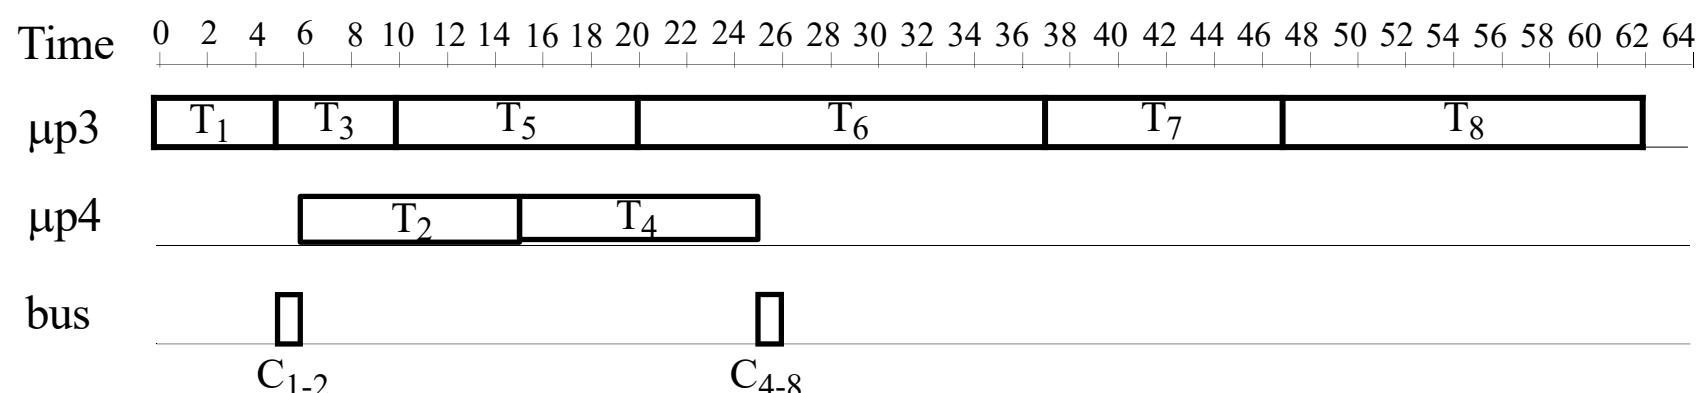
$\mu p3: T_1, T_3, T_5, T_6, T_7, T_8.$

$\mu p4: T_2, T_4.$

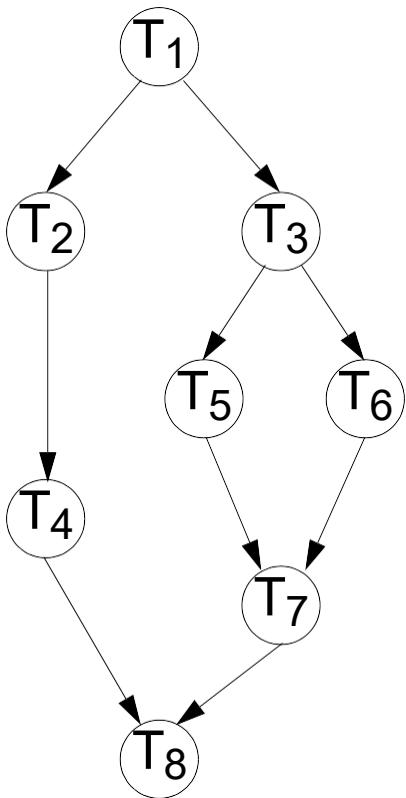
Estimated communication times:

$$\begin{aligned} C_{1-2} &: 1, \\ C_{4-8} &: 1 \end{aligned}$$

We generate a schedule:



We have exceeded the allowed execution time (42)!



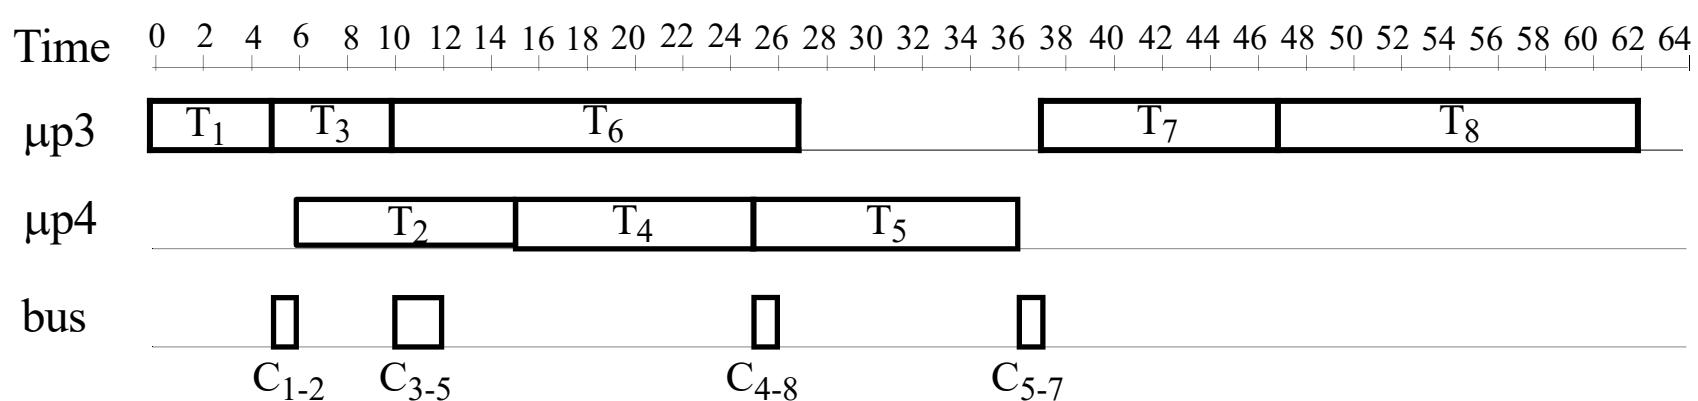
Example

Try a new mapping; T_5 to $\mu p4$, in order to increase parallelism. Two new communications are introduced, with estimated times:

C₃₋₅: 2

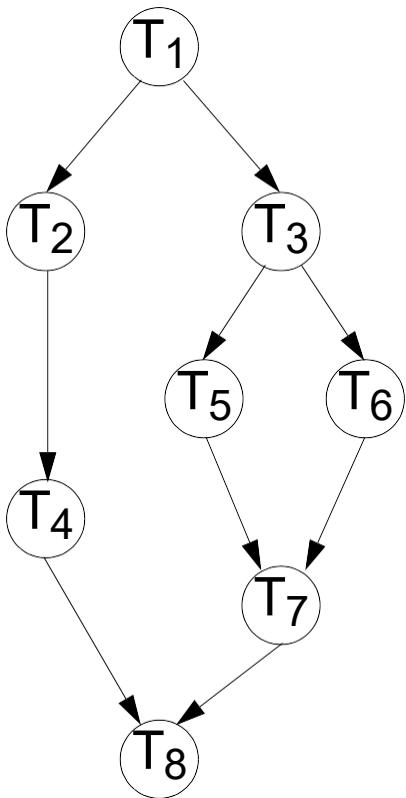
C₅₋₇: 1

We generate a schedule:



The execution time is still 62, as before!

Example



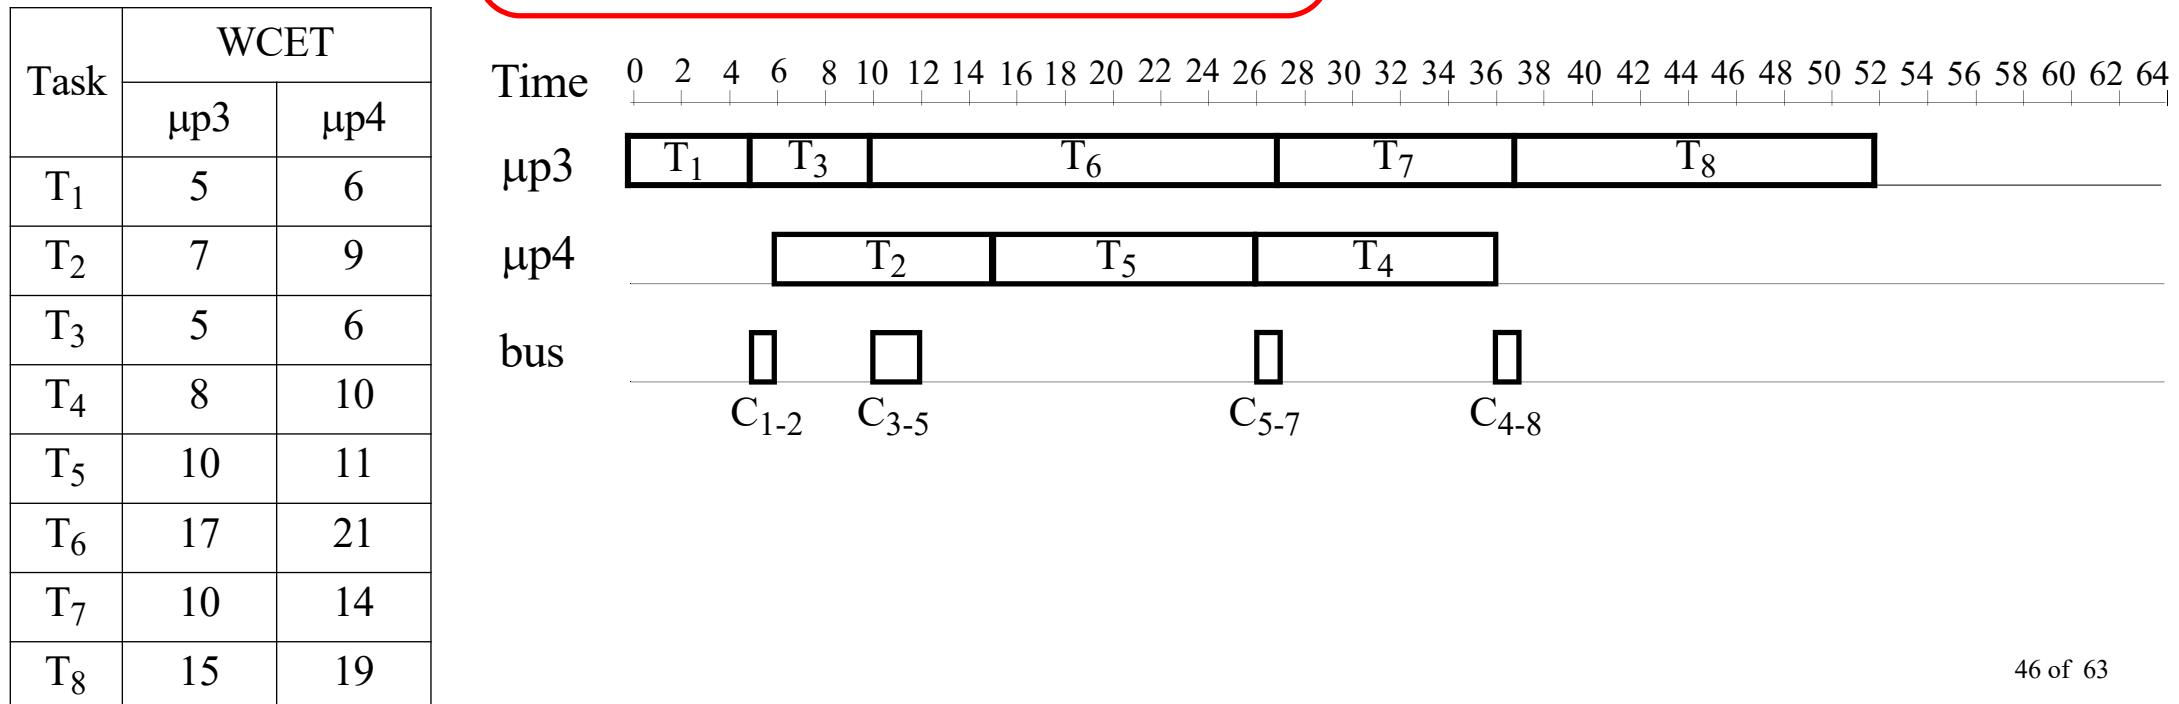
Try a new mapping; T_5 to $\mu p4$, in order to increase parallelism.

Two new communications are introduced, with estimated times:

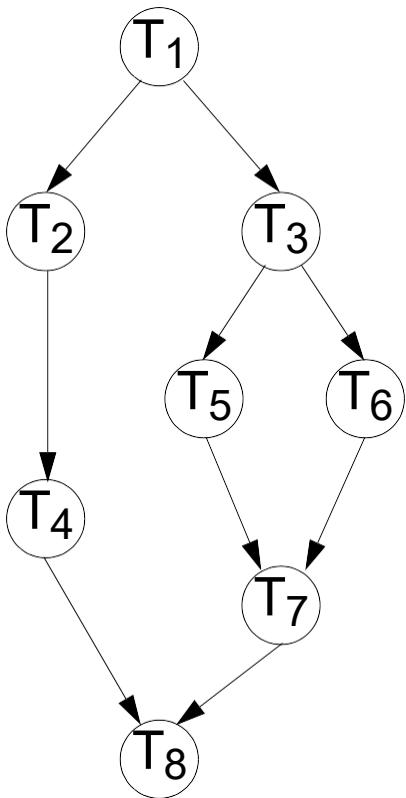
C_{3-5} : 2

C_{5-7} : 1

There exists a better schedule!



Example



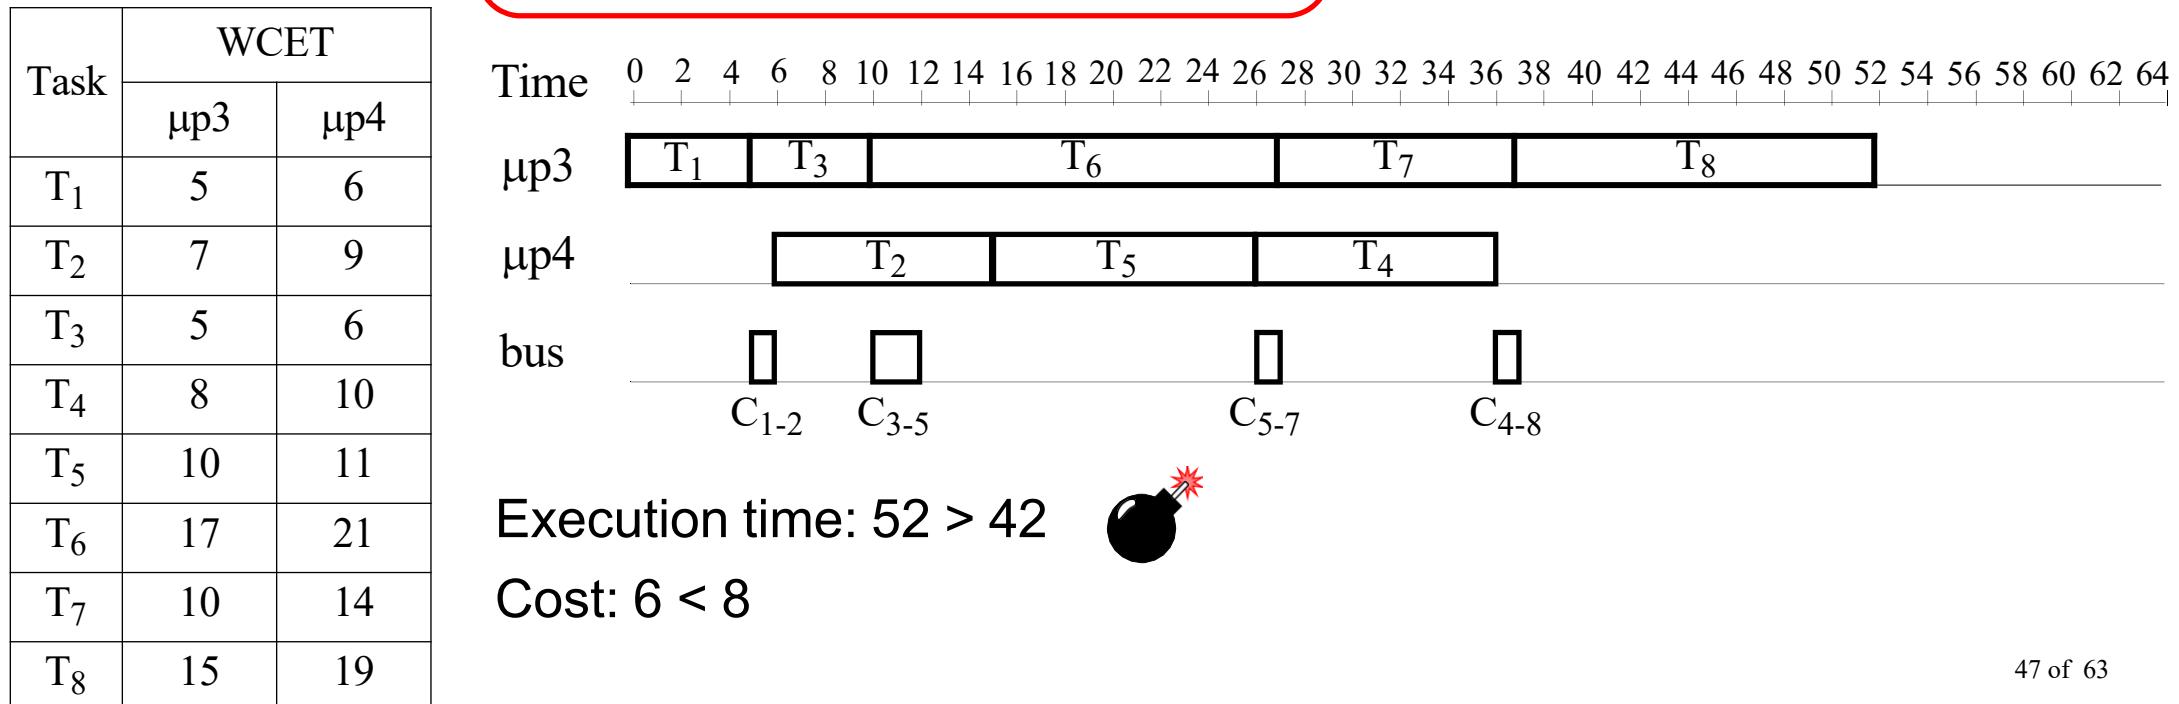
Try a new mapping; T_5 to $\mu p4$, in order to increase parallelism.

Two new communications are introduced, with estimated times:

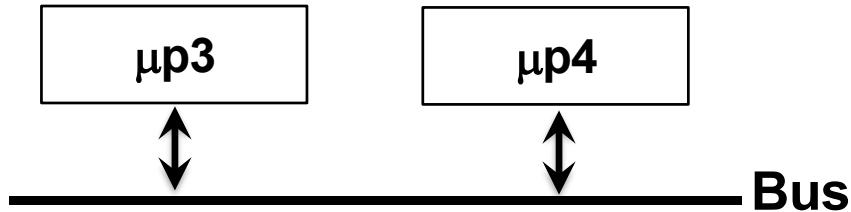
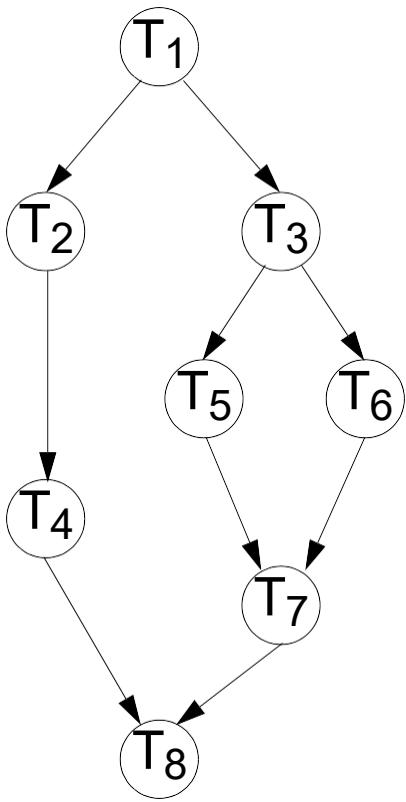
C_{3-5} : 2

C_{5-7} : 1

There exists a better schedule!



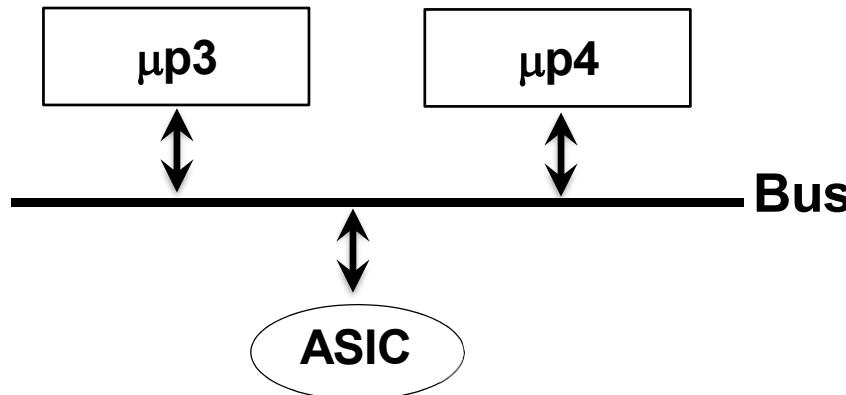
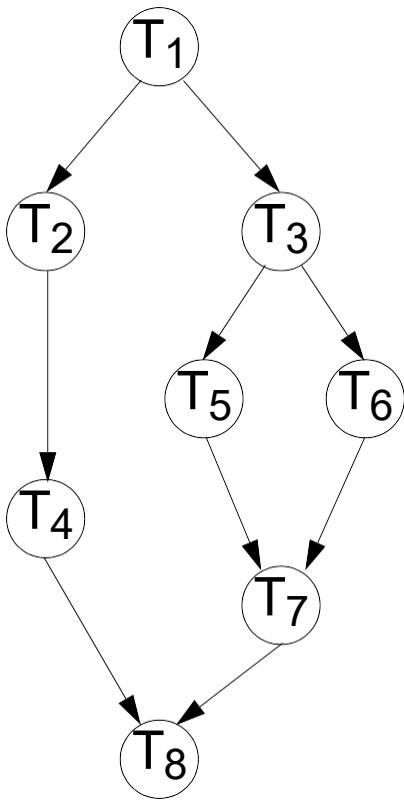
Example



- Possible solutions:
 - Change μproc. μp3 with faster one \Rightarrow cost limits exceeded

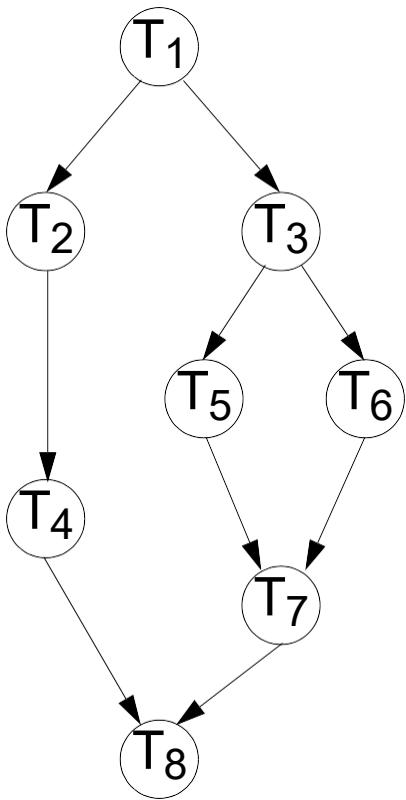
Task	WCET	
	μp3	μp4
T ₁	5	6
T ₂	7	9
T ₃	5	6
T ₄	8	10
T ₅	10	11
T ₆	17	21
T ₇	10	14
T ₈	15	19

Example

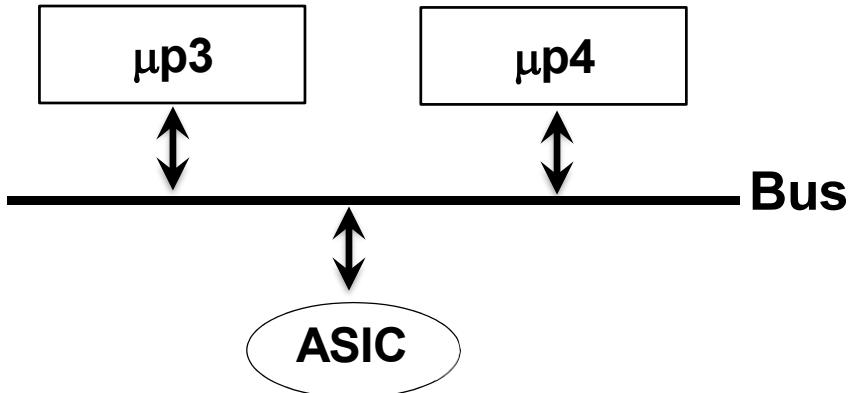


- Possible solutions:
 - Change μproc. μp3 with faster one \Rightarrow cost limits exceed
 - Implement part of the functionality in hardware as an ASIC Cost of ASIC: 1

Task	WCET	
	μp3	μp4
T ₁	5	6
T ₂	7	9
T ₃	5	6
T ₄	8	10
T ₅	10	11
T ₆	17	21
T ₇	10	14
T ₈	15	19



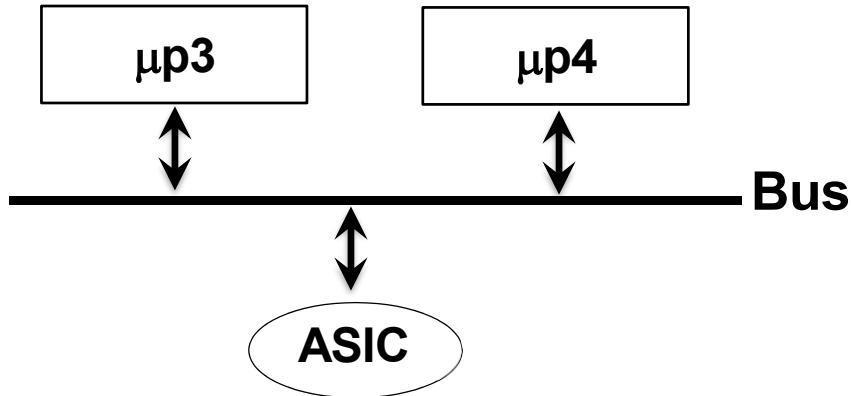
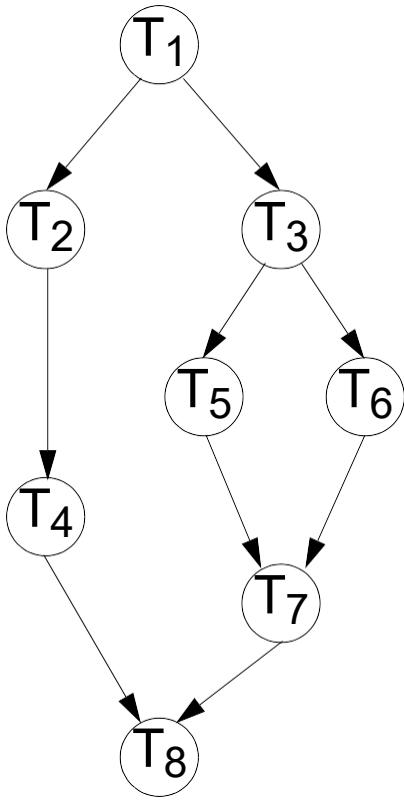
Example



- Possible solutions:
 - Change μproc. μp3 with faster one \Rightarrow cost limits exceed
 - Implement part of the functionality in hardware as an ASIC
- New architecture Cost of ASIC: 1
- Mapping
 - μp3: T₁, T₃, T₆, T₇.
 - μp4: T₂, T₄, T₅.
 - ASIC: T₈ with estimated WCET= 3
 - New communication, with estimated time: C₇₋₈: 1

Task	WCET	
	μp3	μp4
T ₁	5	6
T ₂	7	9
T ₃	5	6
T ₄	8	10
T ₅	10	11
T ₆	17	21
T ₇	10	14
T ₈	15	19

Example



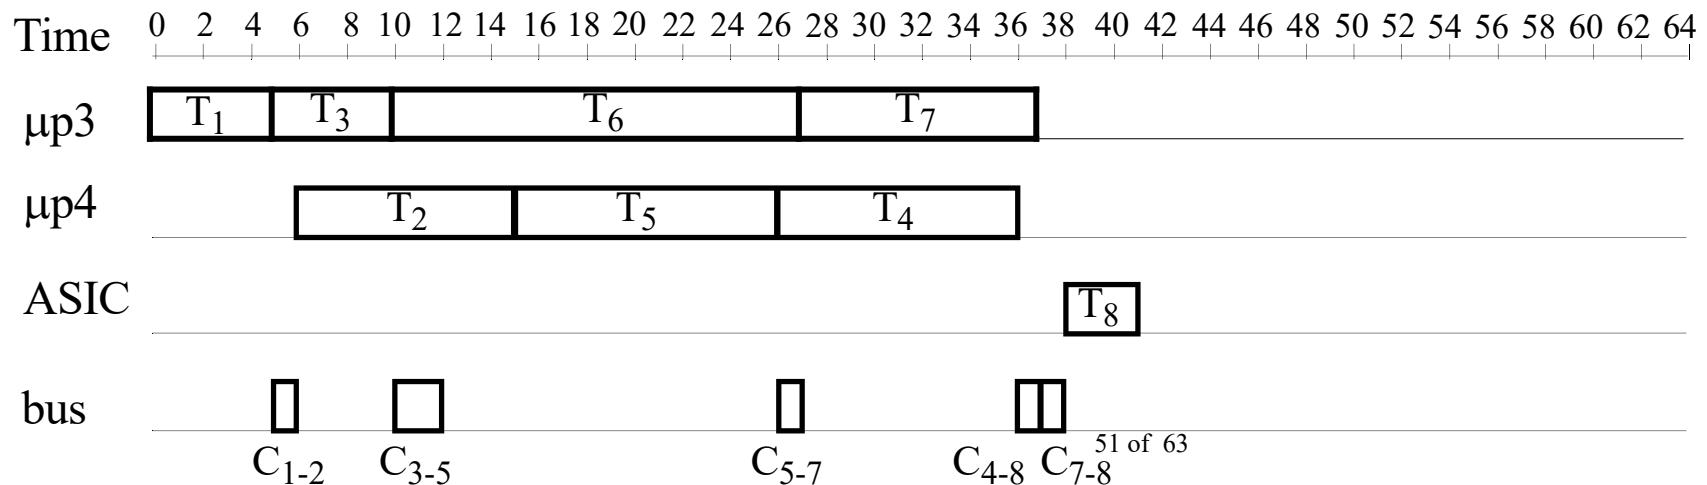
- Mapping

$\mu p3$: T_1, T_3, T_6, T_7 .

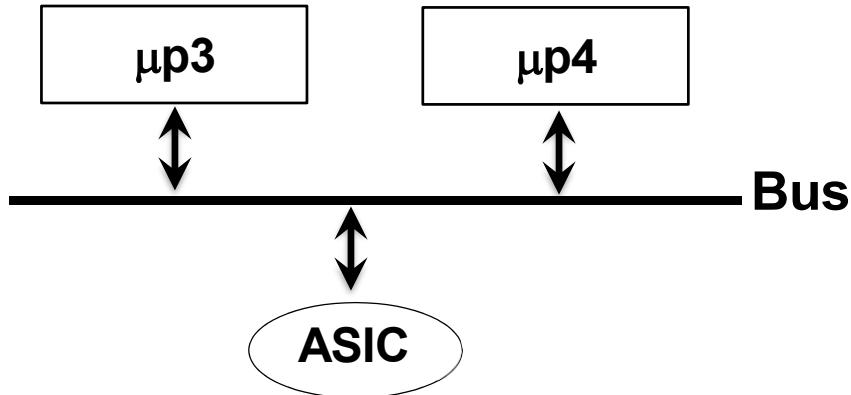
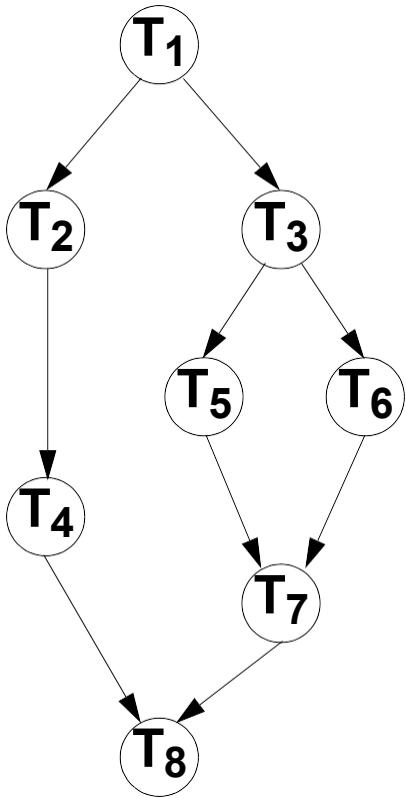
$\mu p4$: T_2, T_4, T_5 .

ASIC: T_8 with estimated WCET= 3

- New communication, with estimated time: C_{7-8} : 1

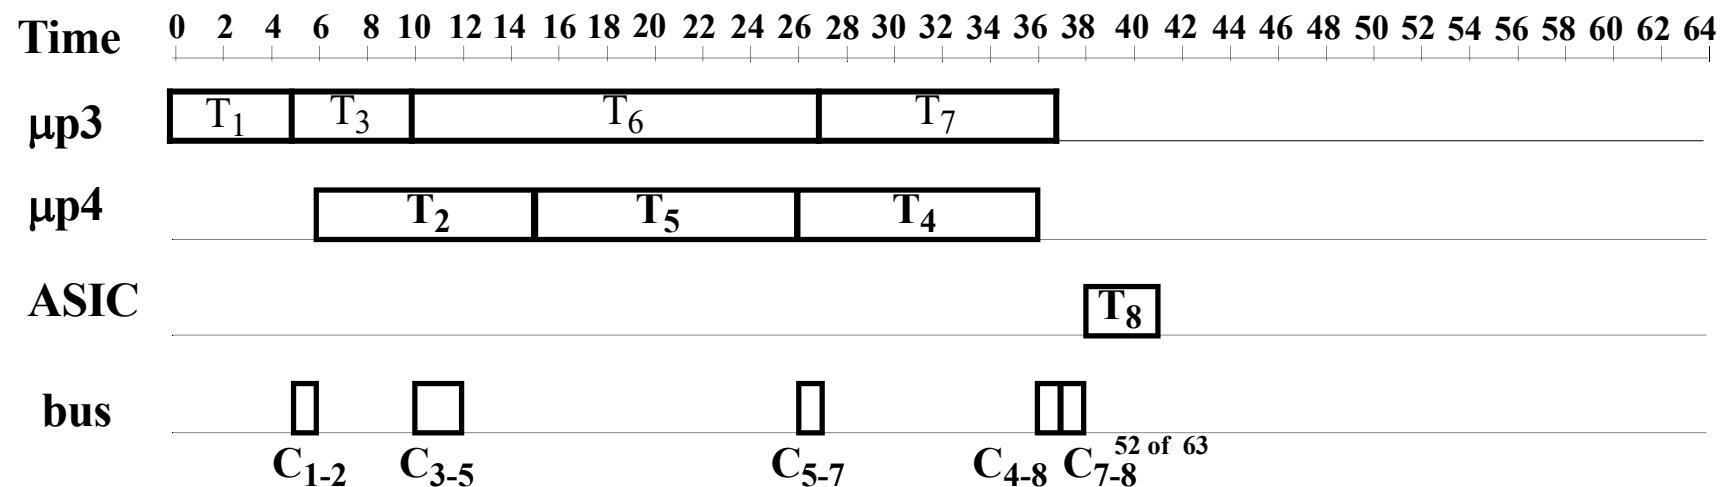
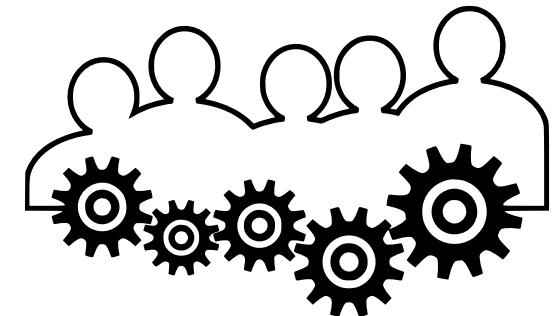


Example



Using this architecture we got a solution with:

- Execution time: $41 < 42$
- Cost: $7 < 8$



What did we achieve?

Example

- We have selected an architecture.
- We have mapped tasks to the processors and ASIC.
- We have elaborated a a schedule.

What did we achieve?

Example

- We have selected an architecture.
- We have mapped tasks to the processors and ASIC.
- We have elaborated a a schedule.

Extremely important!!!

Nothing has been built yet.

All decisions are based on simulation and estimation.

What did we achieve?

Example

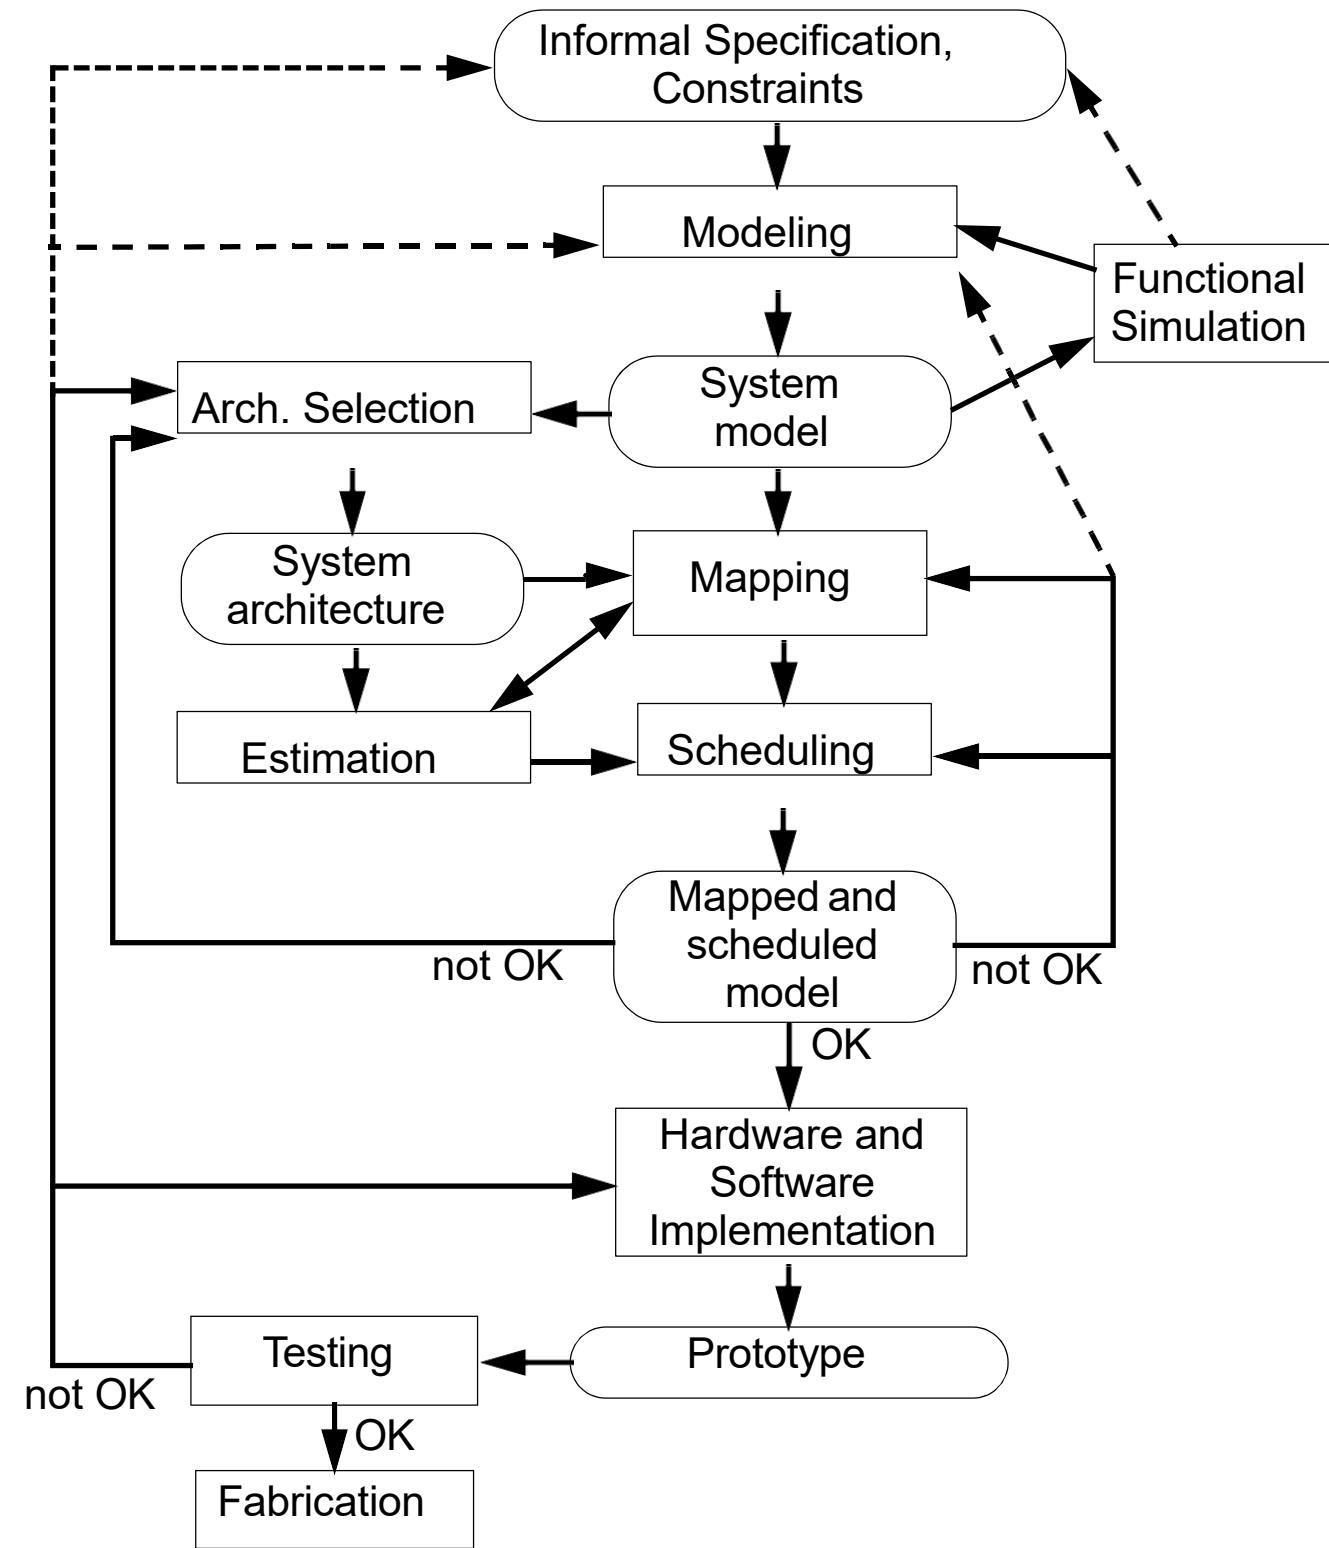
- We have selected an architecture.
- We have mapped tasks to the processors and ASIC.
- We have elaborated a a schedule.

Extremely important!!!

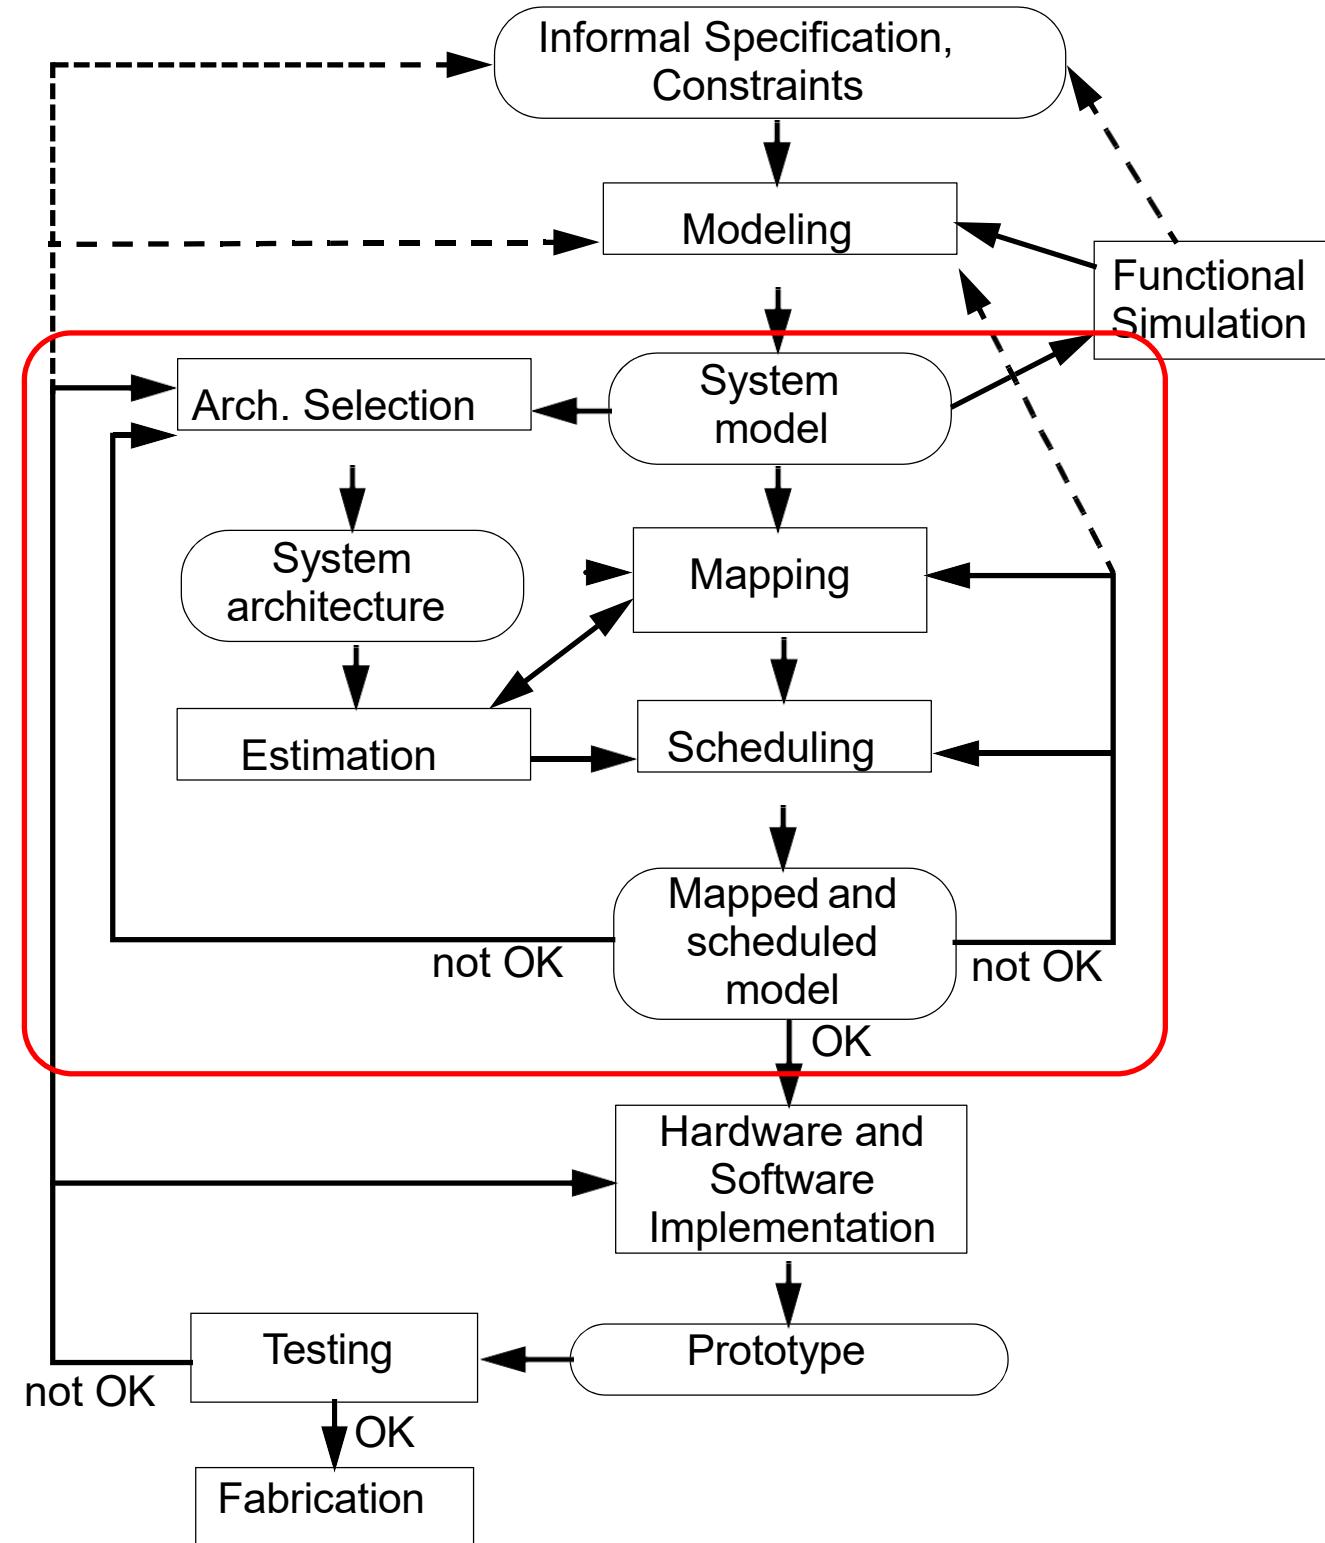
Nothing has been built yet.

All decisions are based on simulation and estimation.

- Now we can go and do the software and hardware implementation, with a high degree of confidence that we get a correct prototype.



What is the essential difference compared to the “traditional” design flow?

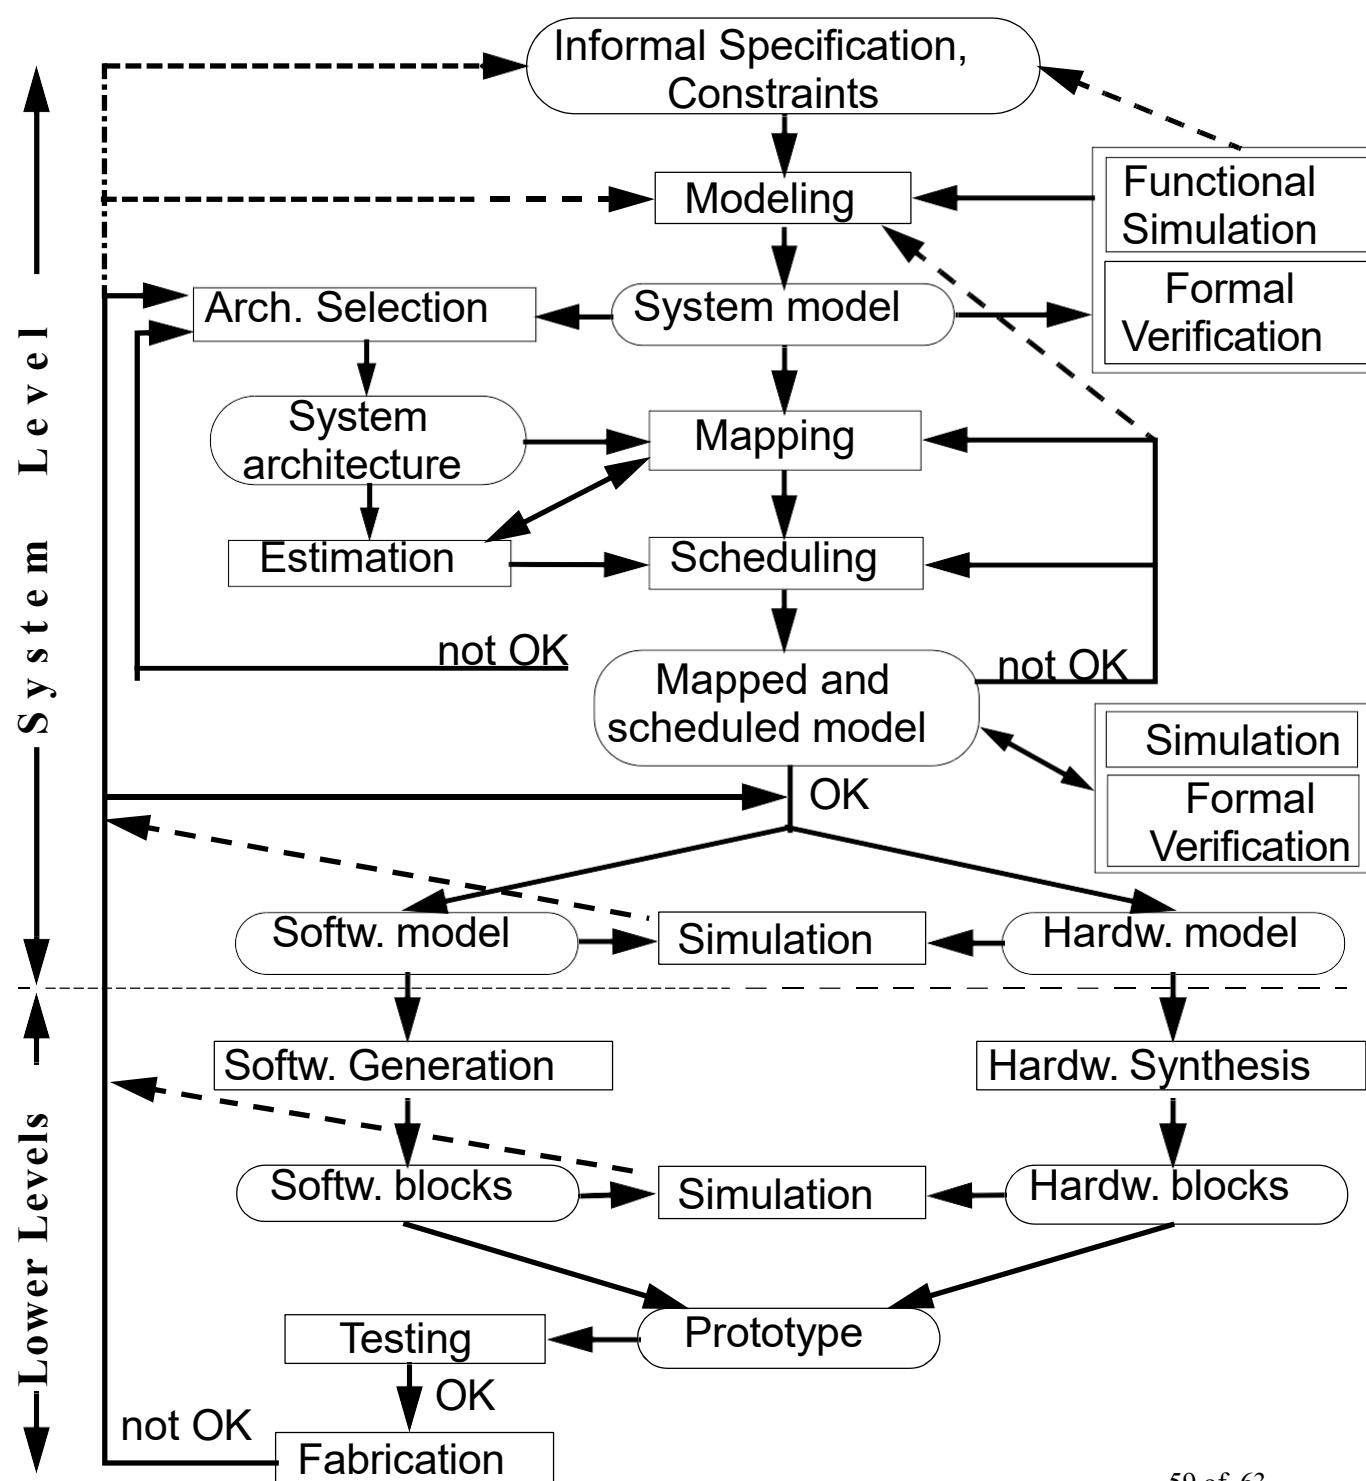


What is the essential difference compared to the “traditional” design flow?

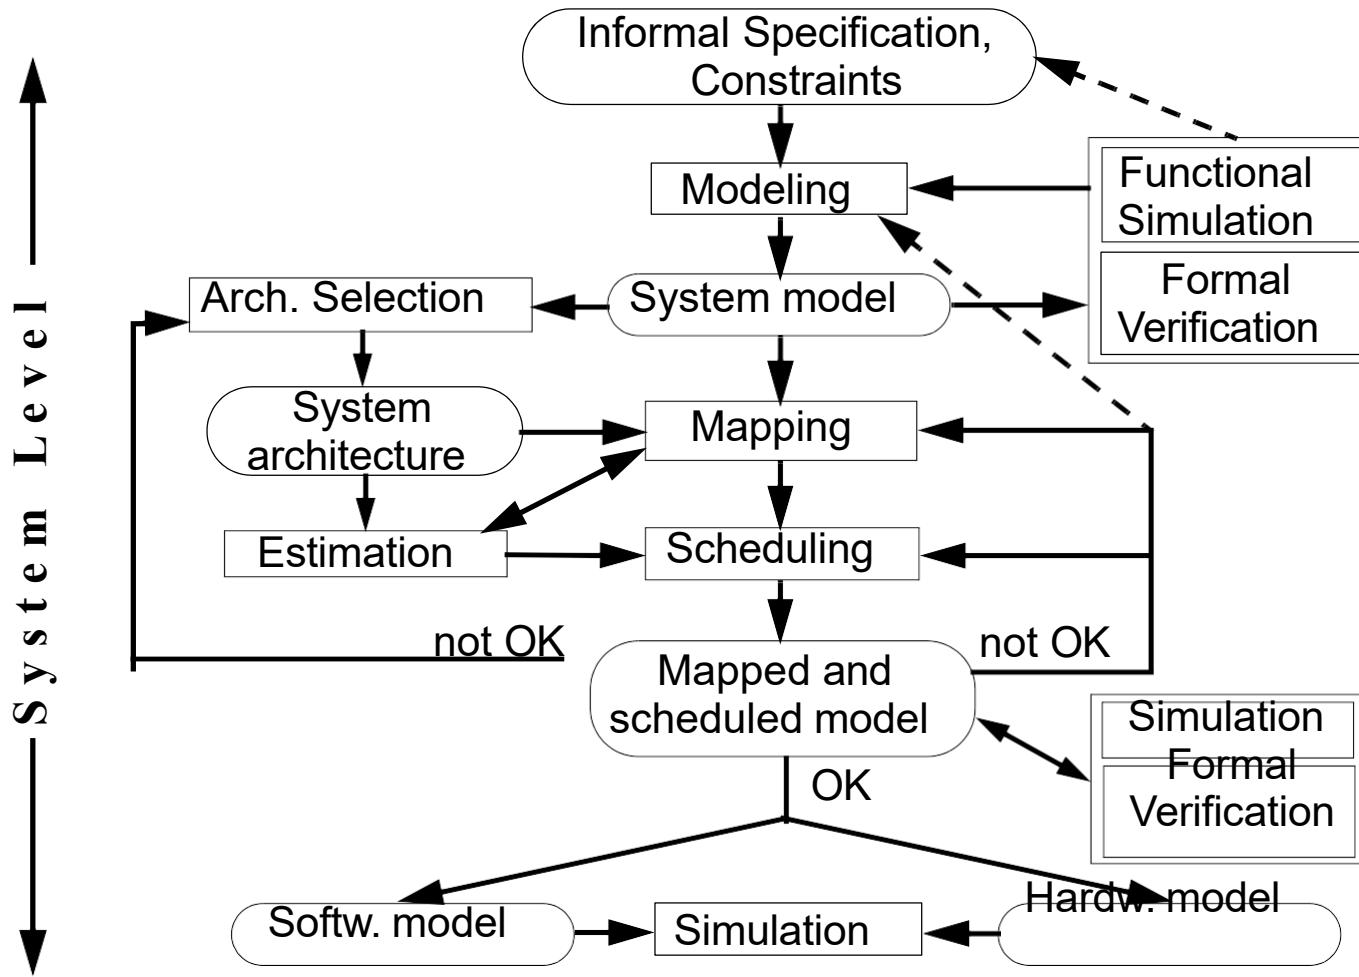
- The inner loop which is performed before the hardware/software implementation. This loop is performed several times as part of the *design space exploration*. Different architectures, mappings and schedules are explored, before the actual implementation and prototyping.
- We get highly optimized good quality solutions in short time. We have a good chance that the outer loop, including prototyping, is not repeated.

The Design Flow

- Formal verification
 - It is impossible to do an exhaustive verification by simulation! Especially for safety critical systems formal verification is needed.
- Hardware/Software codesign
 - During the mapping/scheduling step we also decide what is going to be executed on a programmable processor (software) and what is going into hardware (ASIC, FPGA).
 - During the implementation phase, hardware and software components have to be developed in a coordinated way, keeping care of their consistency (hardware/software cosimulation)



System Level Design Flow



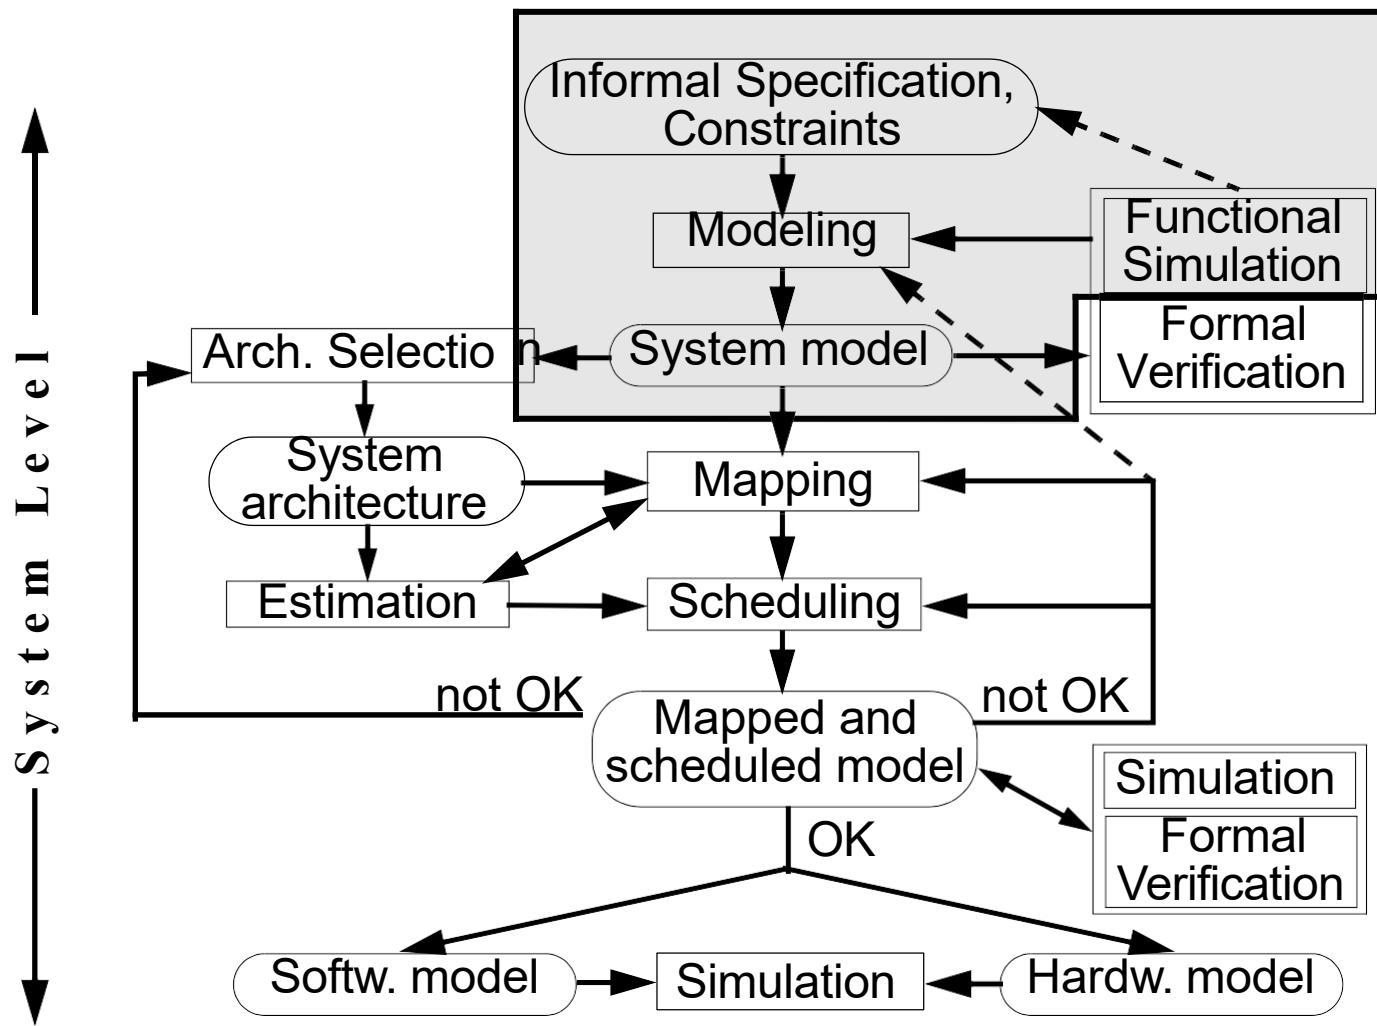
This is what we are interested in, in this course!

Course Topics at a Glance

- Introduction: Embedded Systems and Their Design (*just finished!*)
- Models of Computation and Specification Languages
 - Dataflow Models, Petri Nets, Discrete Event Models, Synchronous Finite State Machines & Synchronous Languages, Globally Asynchronous Locally Synchronous Systems, Timed Automata, Hybrid Automata.
- Architectures and Platforms for Embedded Systems Design
 - General Purpose vs. Application Specific Architectures, Component and Platform-based Design, Reconfigurable Systems, Functionality Mapping.
- Real-Time Embedded Systems
- System-Level Power/Energy Optimization

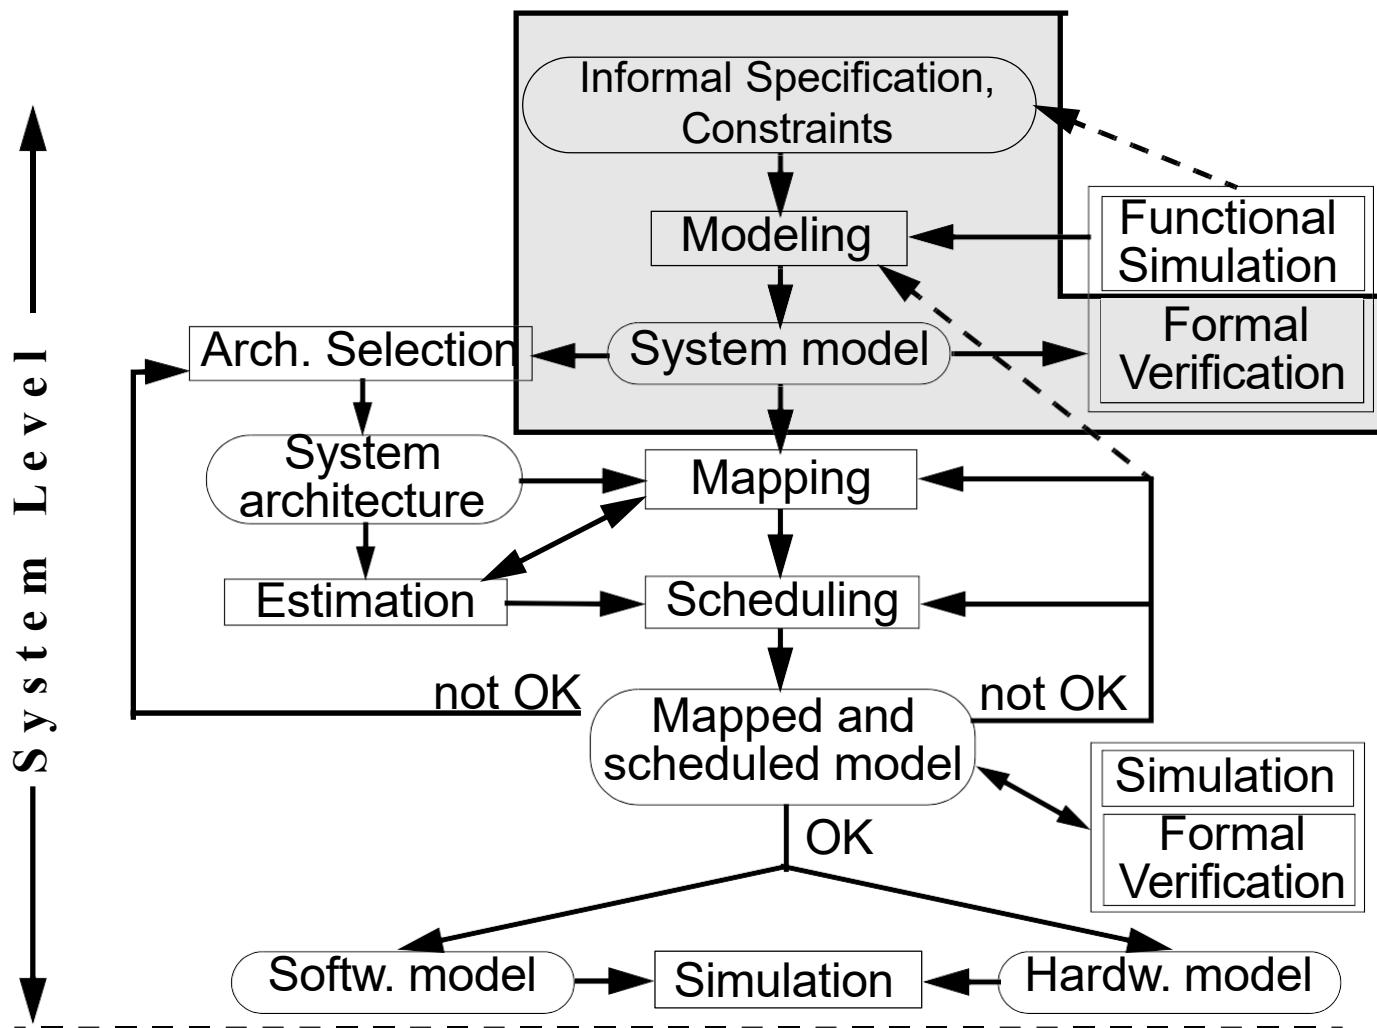
Lab Assignment 1

- Modeling and simulation with System C



Lab Assignment 2

- Formal verification with UPPAAL (TDTS07 only)



Lab Assignment 3

- Design space exploration with an MPARM simulator.

