



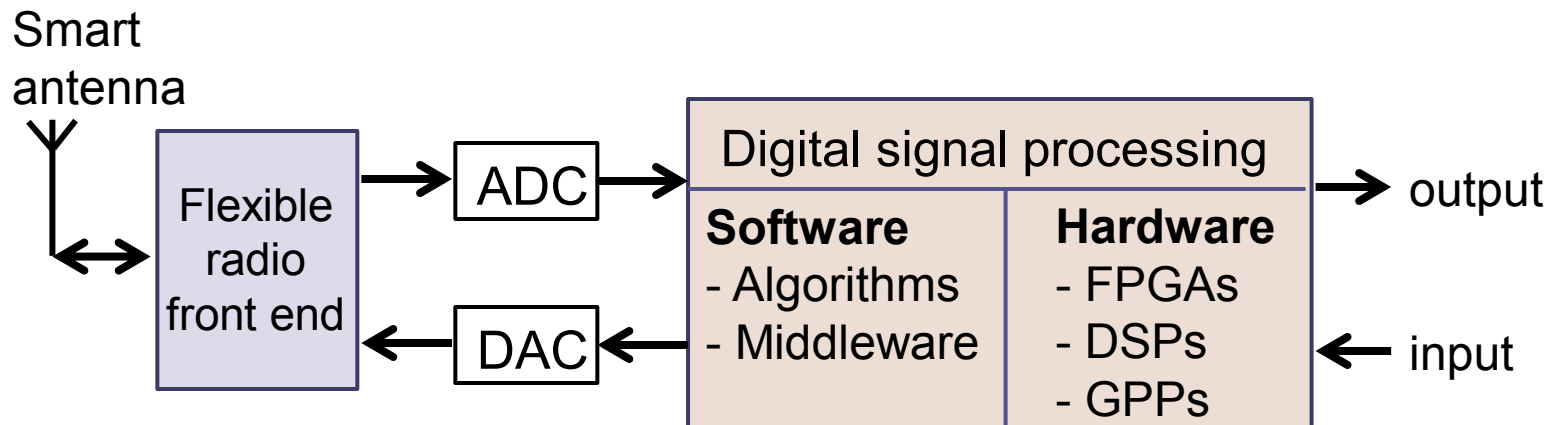
ALOE Framework and Waveform Design Workshop

Vuk Marojevic
Ismael Gomez
Antoni Gelonch

Outline

- 1. Context**
- 2. ALOE Concepts and Framework**
- 3. Computing Resource Management**
- 4. Waveform Development and Deployment**
- 5. Conclusions**

SDR Model



The software radio or SDR provides a flexible radio architecture that allows changing the radio personality, possibly in real-time

J. H. Reed, *Software Radio – A Modern Approach To Radio Engineering*, 2002 Prentice Hall PTR, Upper Saddle River, NJ.


Software

- **Waveforms (SDR applications)**
 - DSP algorithms
 - Radio's physical layer behavior
- **Middleware (SDR framework)**
 - Software layer between applications and hardware
 - Execution environment for waveforms
 - Individual hardware and software development
 - Waveform loading and unloading → reconfiguration
 - Portability and reuse of components

Hardware

TODAY

multicores




small clusters
(heterogeneous)




TOMORROW

many-cores



ALOE Context



UMTS




LTE



?




ALOE



SDR Computing



- 1. Multiprocessing
- 2. Lightweight
- 3. Platform Independence
- 4. Computing Resource Management



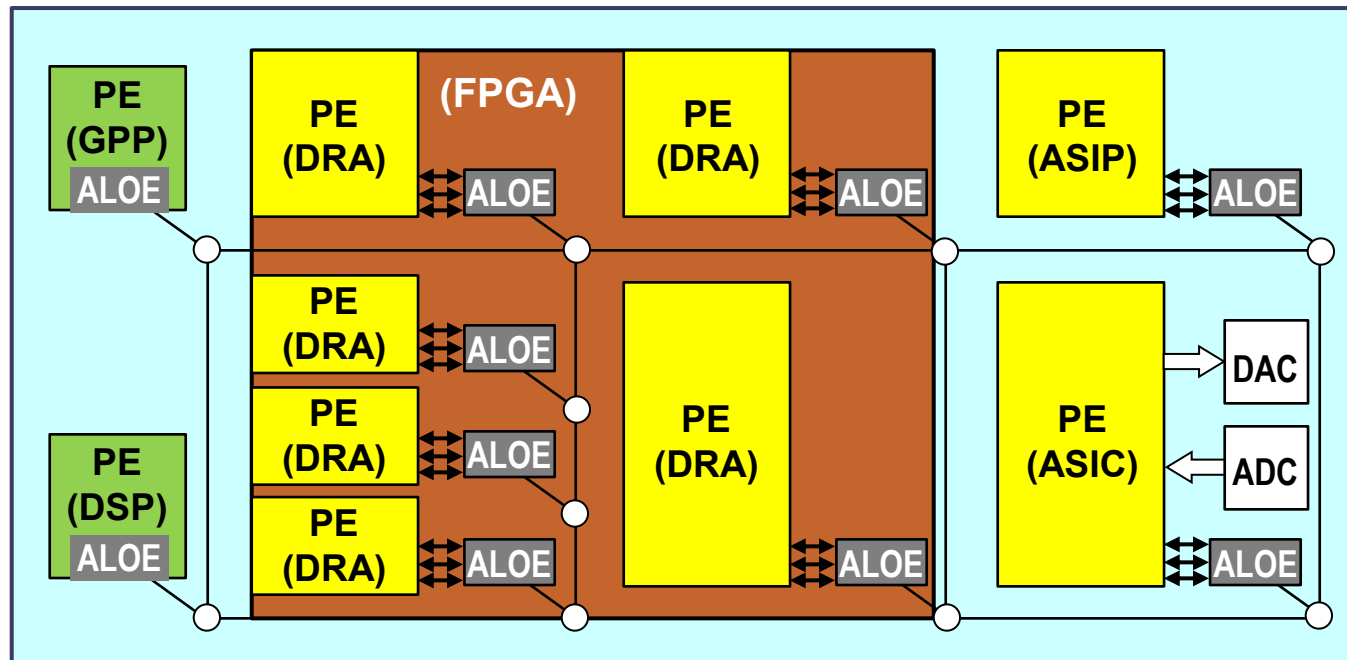
2004...

ALOE
Tools

Open source

Heterogeneous Multiprocessing

SDR Platform




PE with OS
 PE without OS

ALOE services
 Platform services

PE: processing element

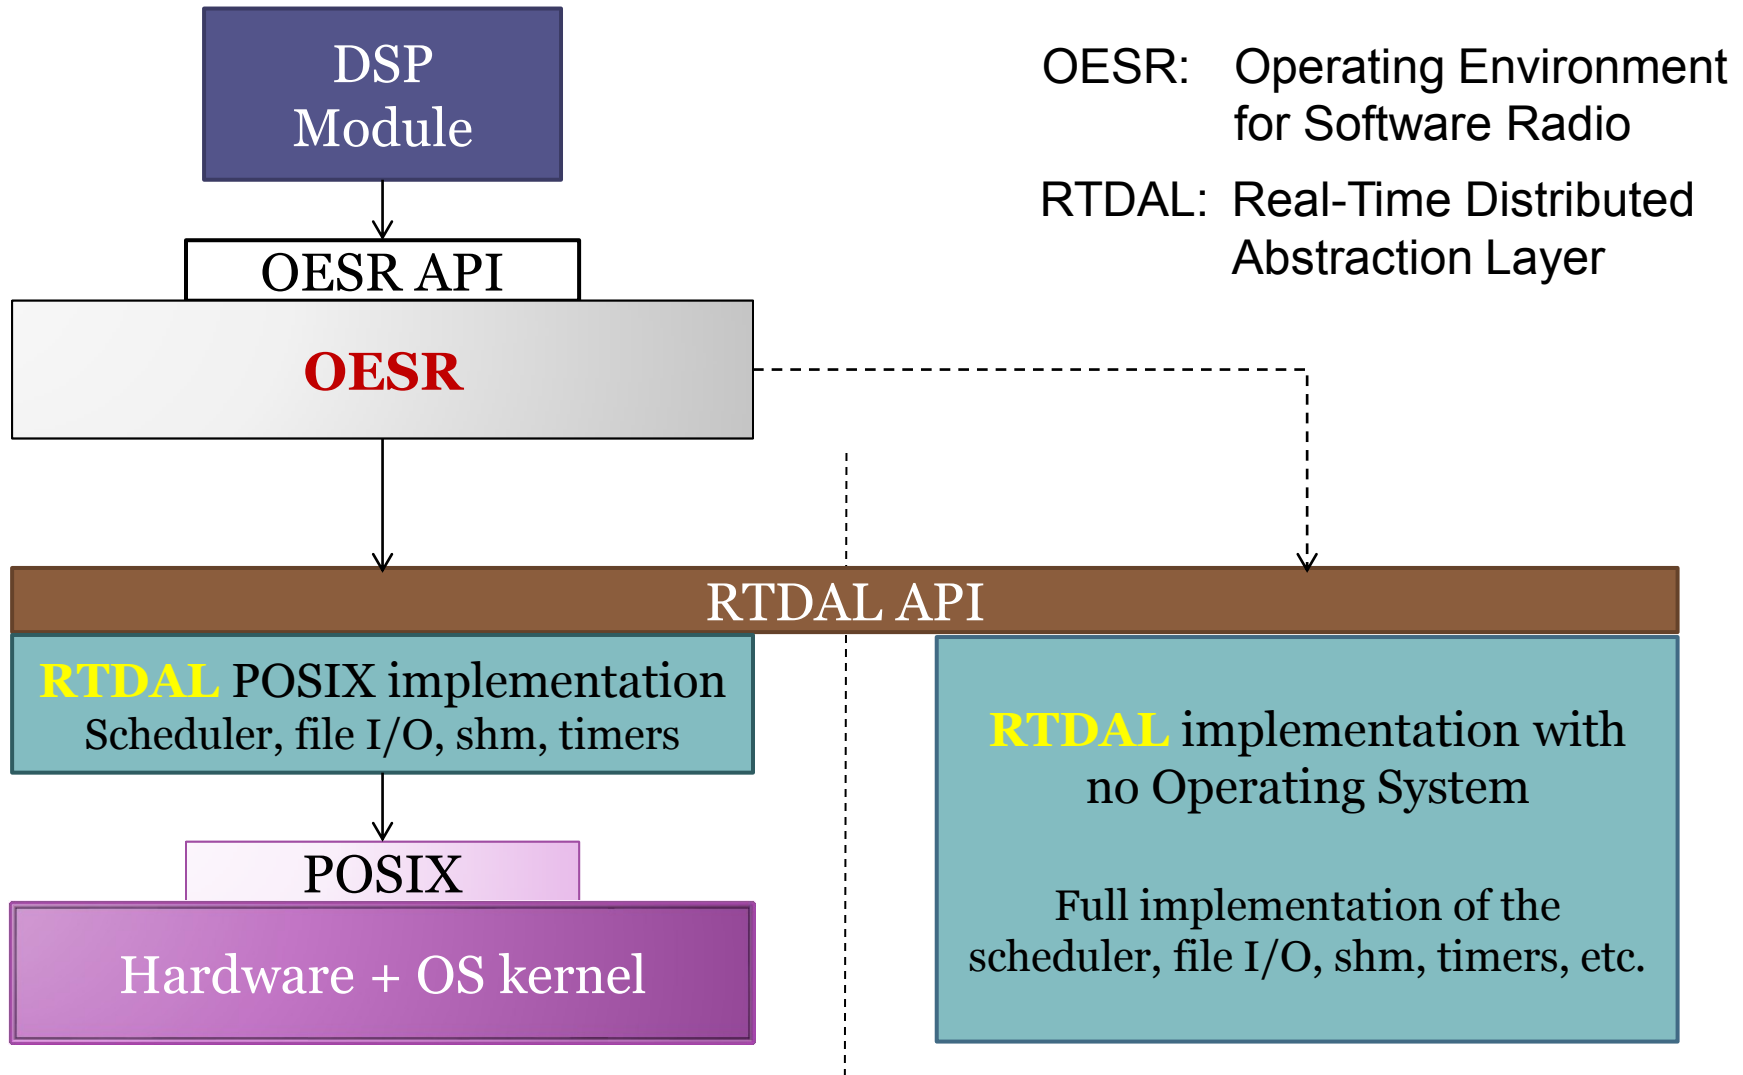
DRA: dynamically reconfigurable area

ALOE Layers



PE: Processing Element

ALOE Architecture



Real-Time Distributed Abstraction Layer (RTDAL)

- Interprocessor communication
- Synchronization
- Scheduling
 - pipelined execution, partitioned scheduling
 - 1 thread per processing core
- RTDAL API
 - Task creation and management
 - Interfaces
 - ADC/DAC abstraction
 - Time functions

Operating Environment for Software Radio (OESR)

- Automatic mapping of waveforms
- Location-transparent inter-module communications
- Configuration and visualization of variables & parameters
- Logs, counters, ...

Computing Resource Management

Wireless Communications Characteristics

- Continuous data transmission and reception
- Real-time services → real-time processing
- **RAT/mode/QoS target → processing demands**
- Heterogeneous multiprocessor platforms
- Limited computing resources
- Dynamic reconfigurations

Computing Resource Management

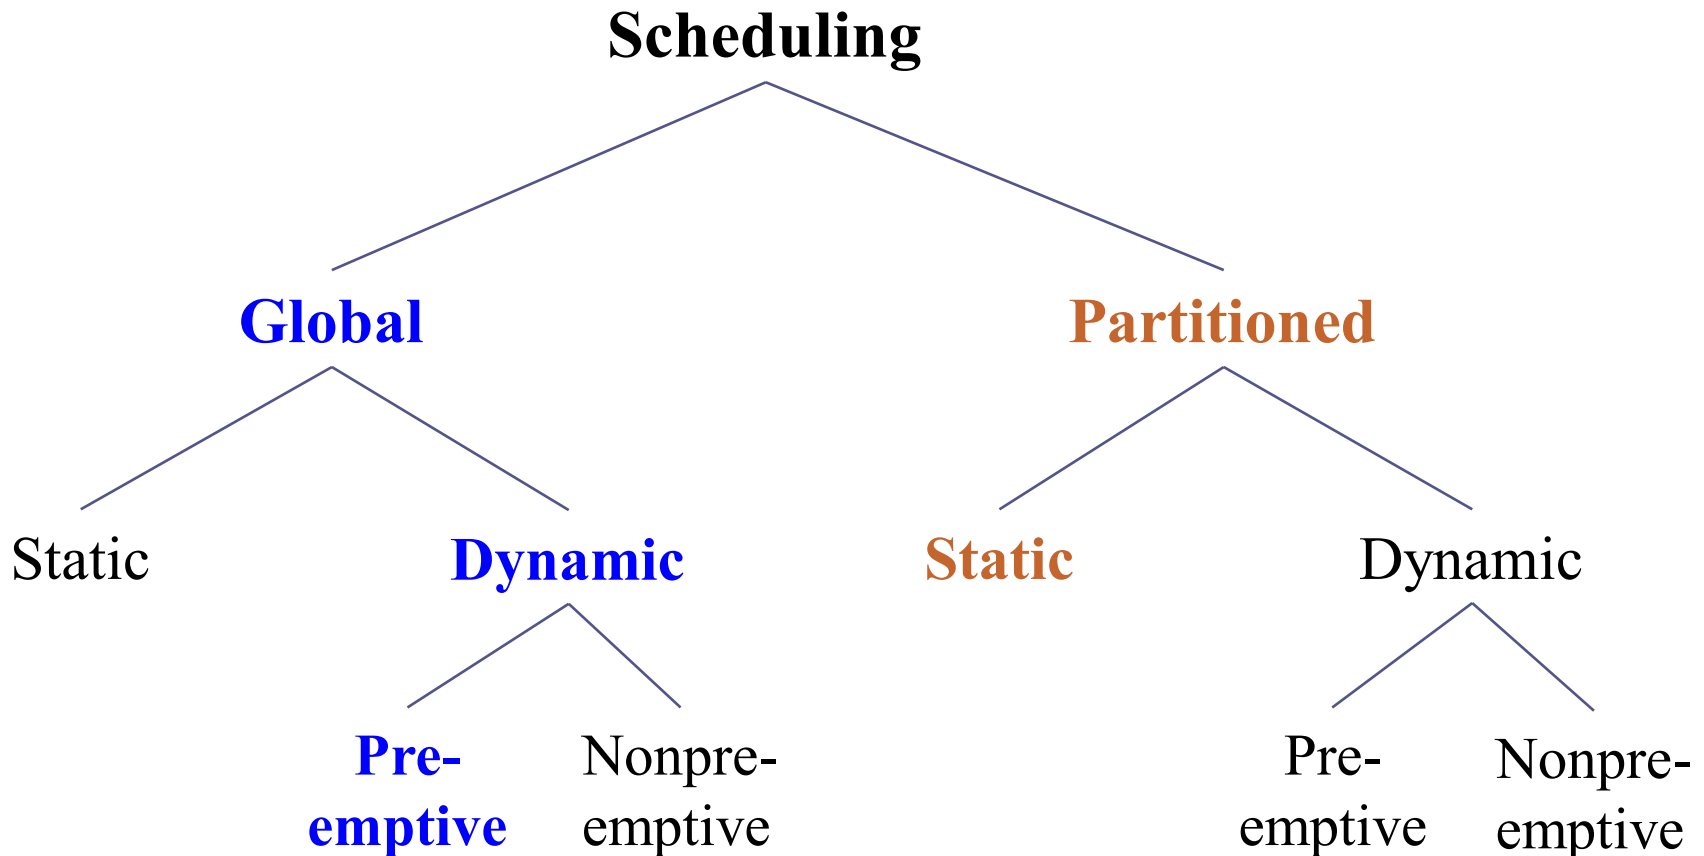
“Provide sufficient computing resource to waveforms for real-time processing”

Real-time constraints:

- *Minimum **throughput***
- *Maximum **latency***

Scheduling

Scheduling is the method by which threads, processes or data flows are given access to system resources. [Wikipedia]



Static vs. Dynamic Scheduling (I)

Static	Dynamic
Offline (compile time), before execution	Online (runtime), at execution
Deterministic performance	Nondeterministic performance
Avoid migrations → less overhead & fewer cache misses	Migrations → overhead & cache misses
Avoid task locks → less system calls → less overhead	Task locks → more system calls → more overhead
Regular, periodic tasks with a priori information	Irregular, aperiodic tasks with unknown characteristics a priori
Runtime rescheduling costly	Easy to add new task at runtime


Static vs. Dynamic Scheduling (II)

- Scheduling overhead increases...
 - ...with the waveform granularity
 - ...with the number of processing elements
 - ...inversely to task execution time
- *Global-dynamic-preemptive* scheduling
+ flexible
- may incur significant resource overhead
- *Partitioned-static* scheduling scales better


V. Marojevic, I. Gomez, A. Gelonch, “Evaluation of computing resource management methods for SDR clouds,” *SDR-WInnComm Conf.*, Washington DC, 2013.

ALOE Resource Management

- **Pipelining**
- **Partitioned scheduling:** static and cooperative
 - Low overhead
 - Easy to implement
 - Scalable
- **Heterogeneous platforms** (w/o shared memory)
- Requires task-to-processor **mapping**



Pipelining




Data packet: x samples

Packet arrival rate: $1/T$ Hz \rightsquigarrow **Throughput requirement:** x/T samples/s

Process 1 input data packet every T seconds


Pipelining

$t = 0$




Pipelining

$t = T$




Pipelining

$t = 2T$

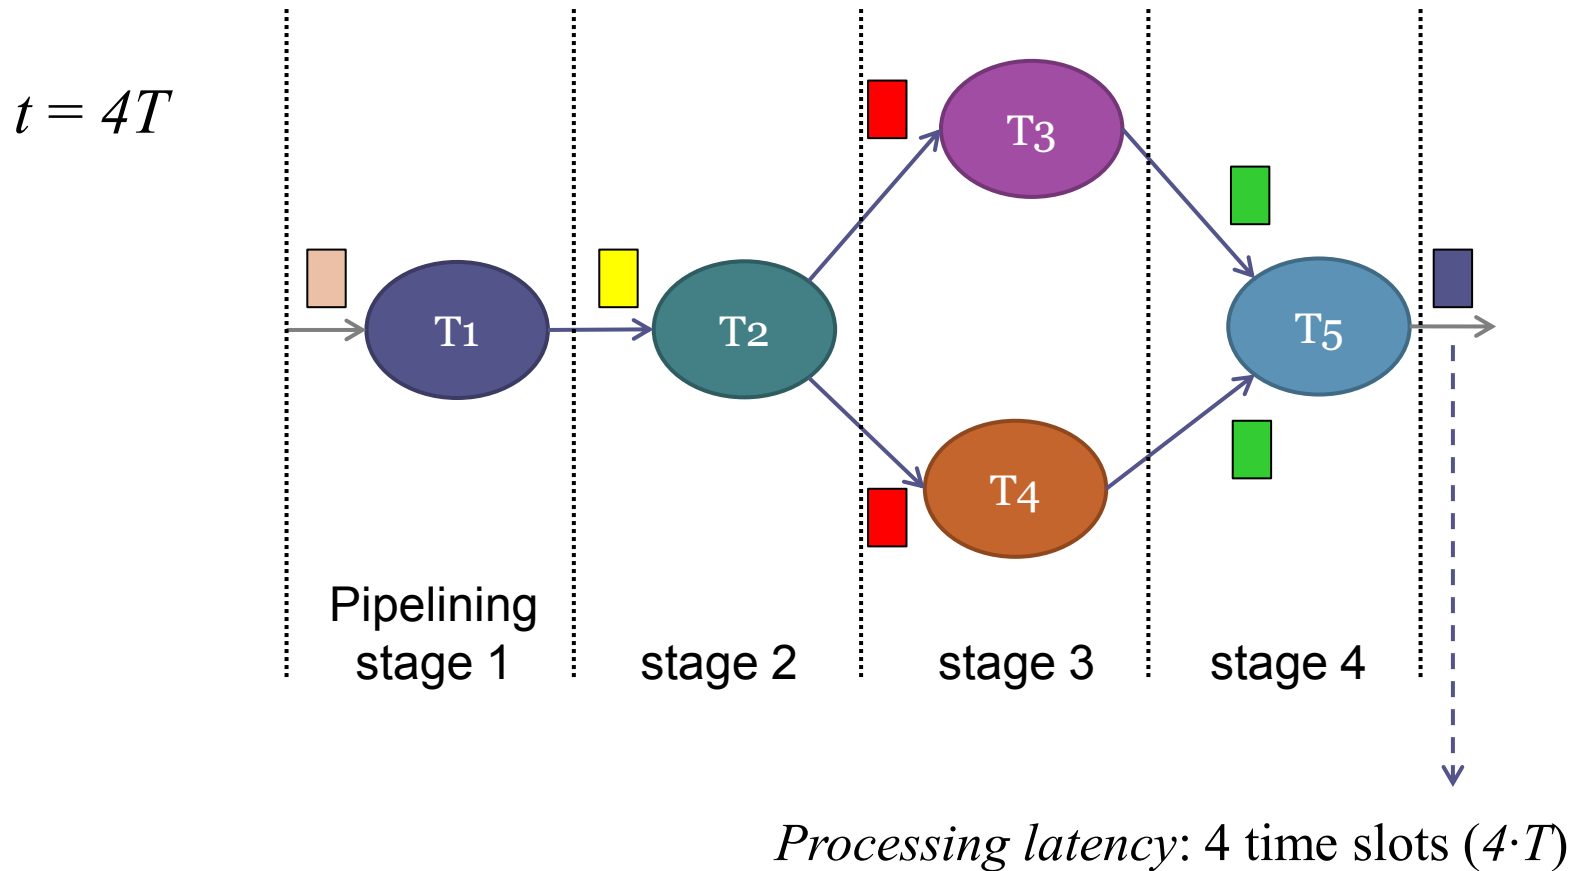


Pipelining

$t = 3T$



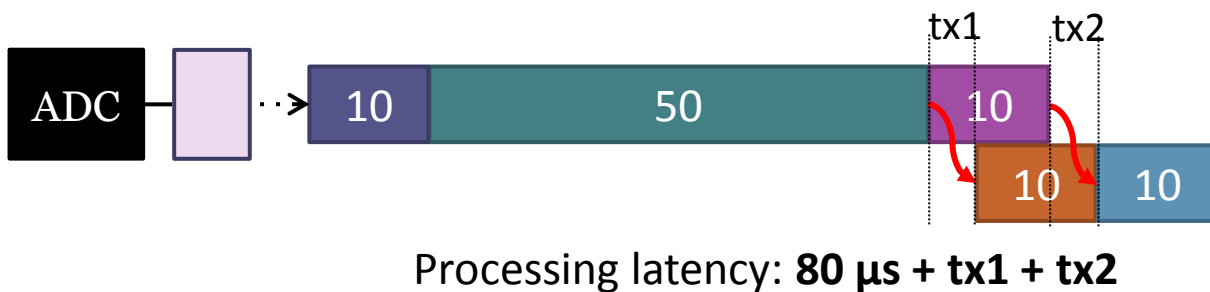
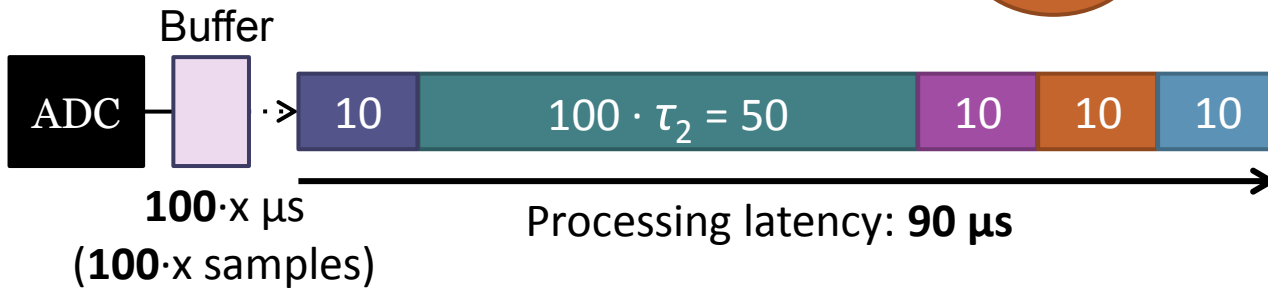
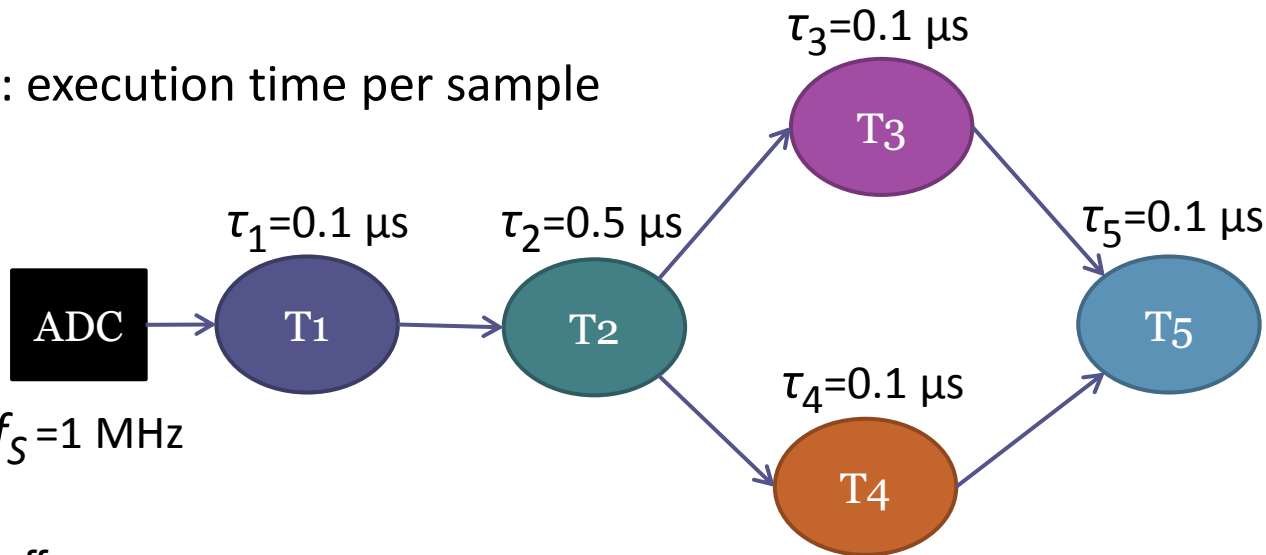
Pipelining



Processing throughput: 1 packet every T seconds (x/T input samples/s)

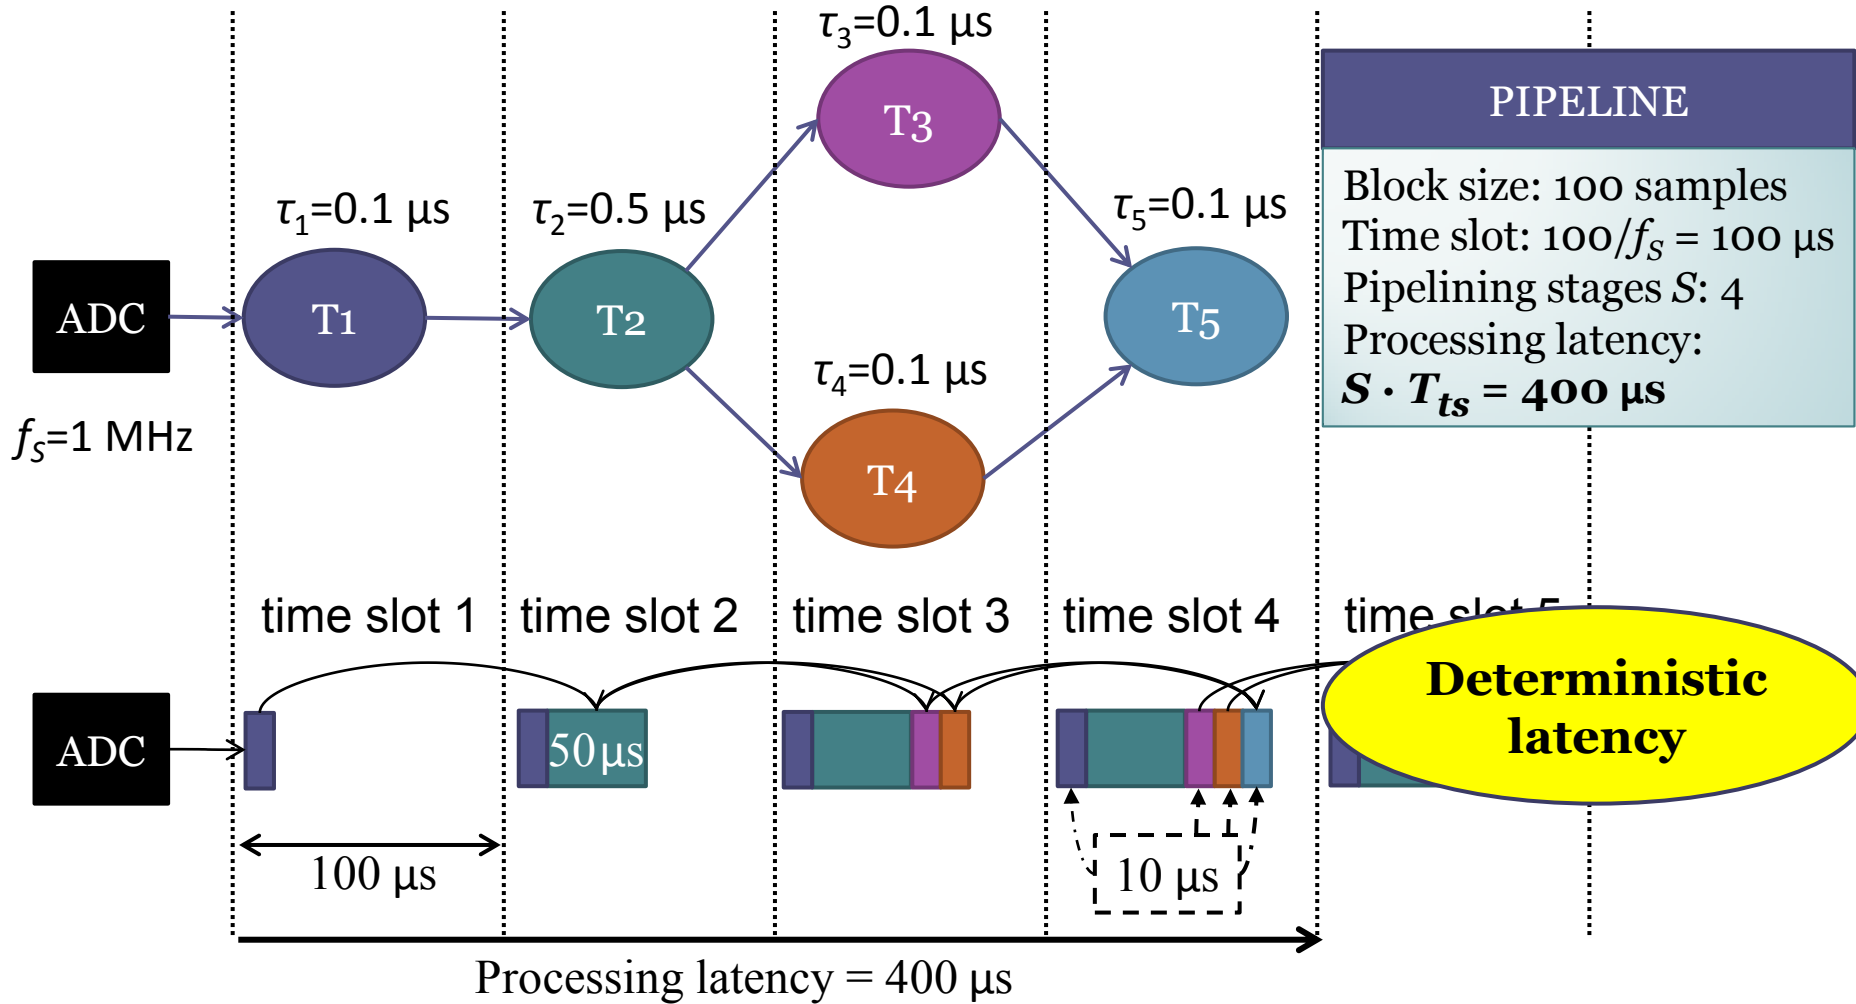
Scheduling Example w/o Pipeline

τ_i : execution time per sample



Latency:
 $f(\text{platform, mapping})$

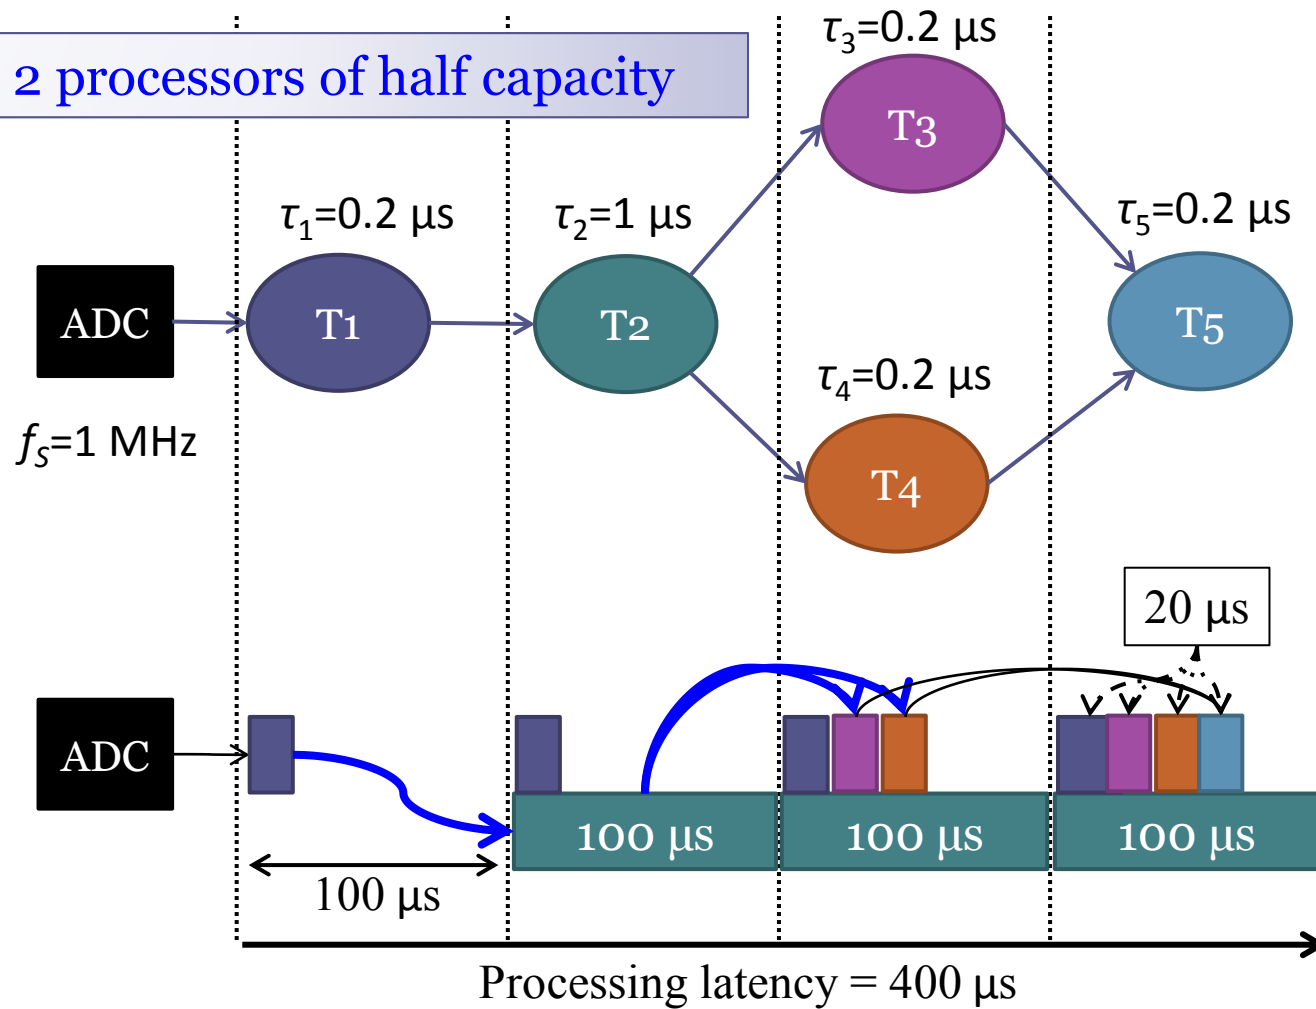
Pipelined Scheduling (I)



Pipelined execution: removes precedence constraints, simplifies scheduling

Pipelined Scheduling (II)

2 processors of half capacity



PIPELINE

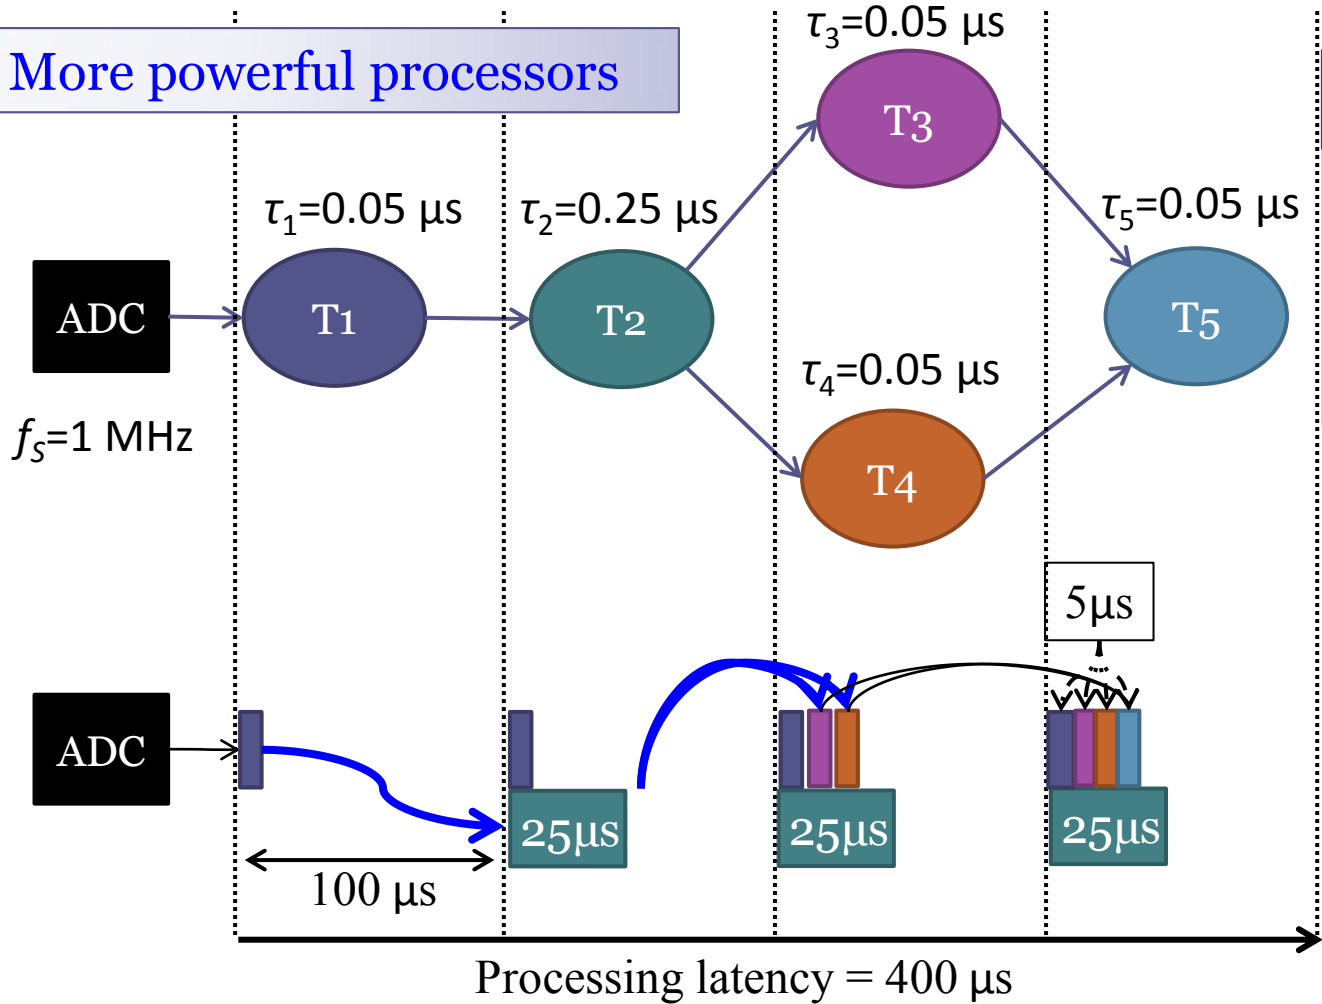
Block size: 100 samples
 Time slot: $100/f_s = 100 \mu\text{s}$
 Pipelining stages S : 4
 Processing latency:
 $S \cdot T_{ts} = 400 \mu\text{s}$

Equal latency and throughput

Scheduling performance: platform-independent

Pipelined Scheduling (III)

More powerful processors



PIPELINE

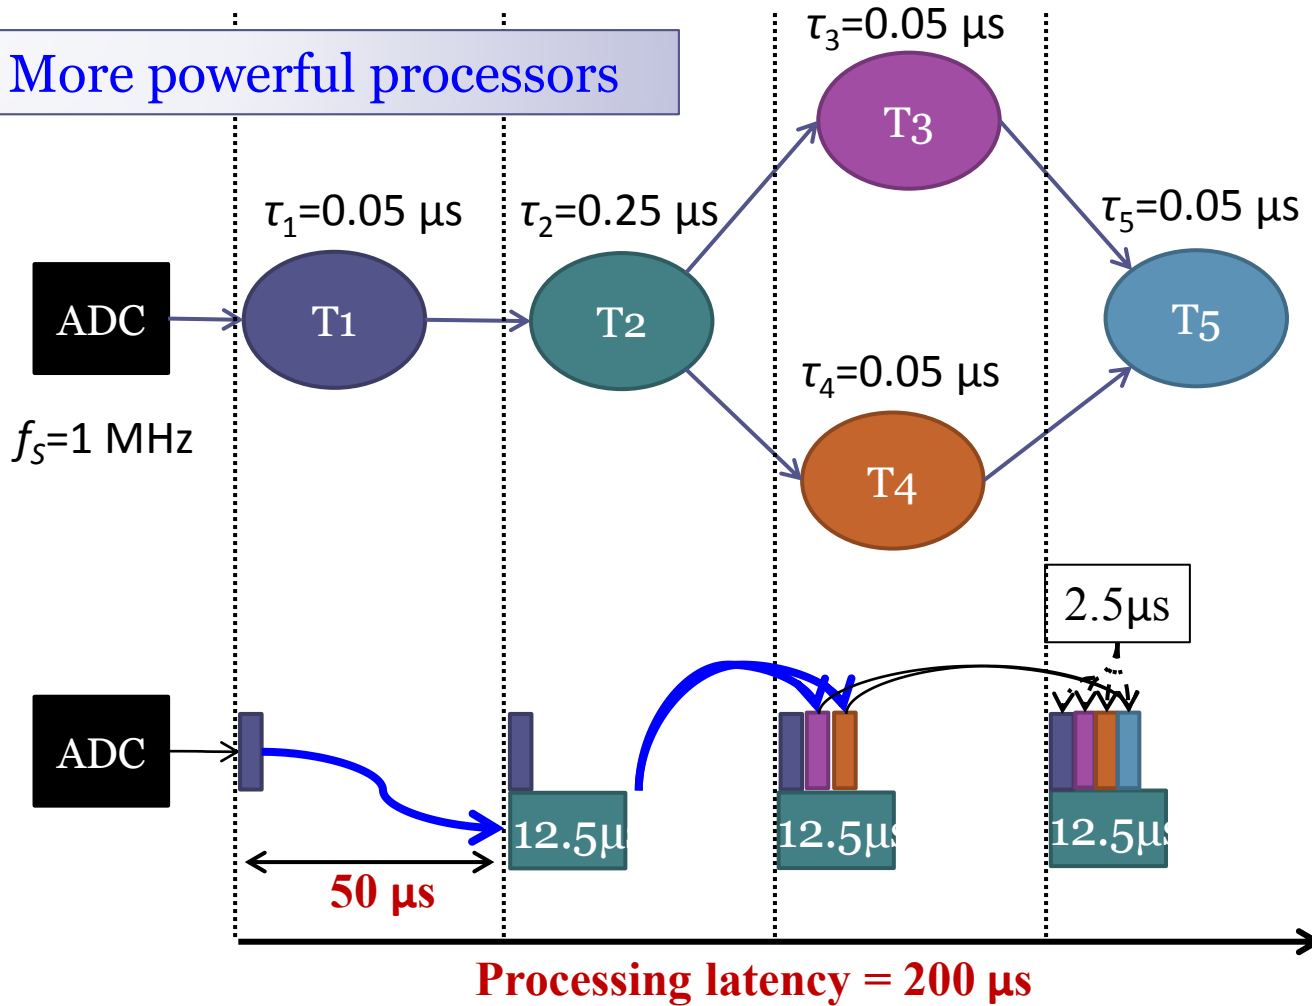
Block size: 100 samples
 Time slot: $100/f_s = 100 \mu\text{s}$
 Pipelining stages S : 4
 Processing latency:
 $S \cdot T_{ts} = 400 \mu\text{s}$

Equal latency and throughput

Scheduling performance: platform-independent

Pipelined Scheduling: Latency Control (I)

More powerful processors



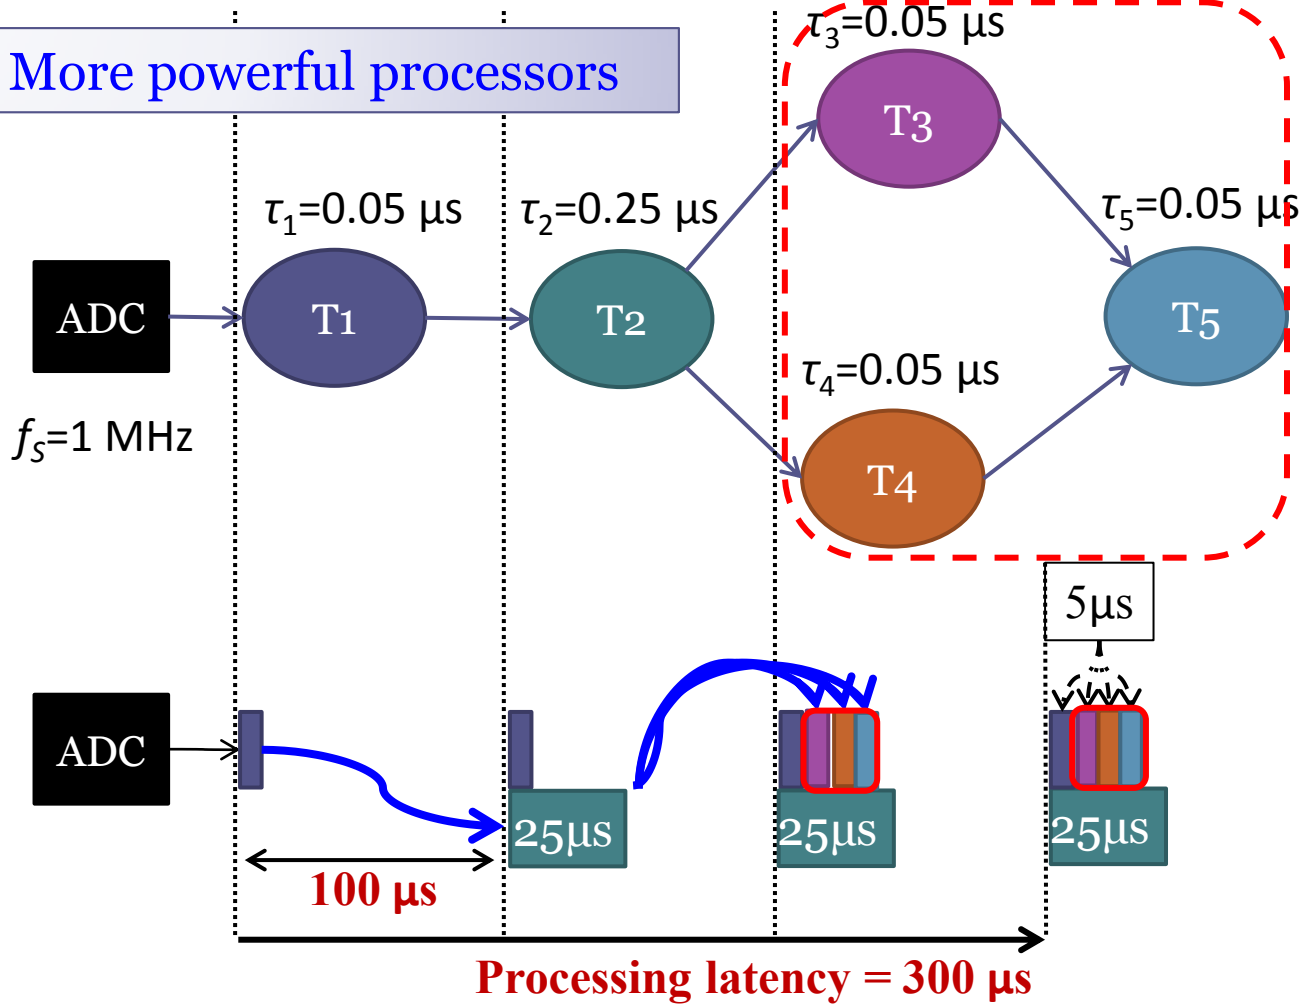
PIPELINE

Block size: **50 samples**
 Time slot: $50/f_s = \mathbf{50 \mu\text{s}}$
 Pipelining stages S : 4
 Processing latency:
 $S \cdot T_{ts} = \mathbf{200 \mu\text{s}}$

Scheduling performance: platform-independent

Pipelined Scheduling: Latency Control (II)

More powerful processors

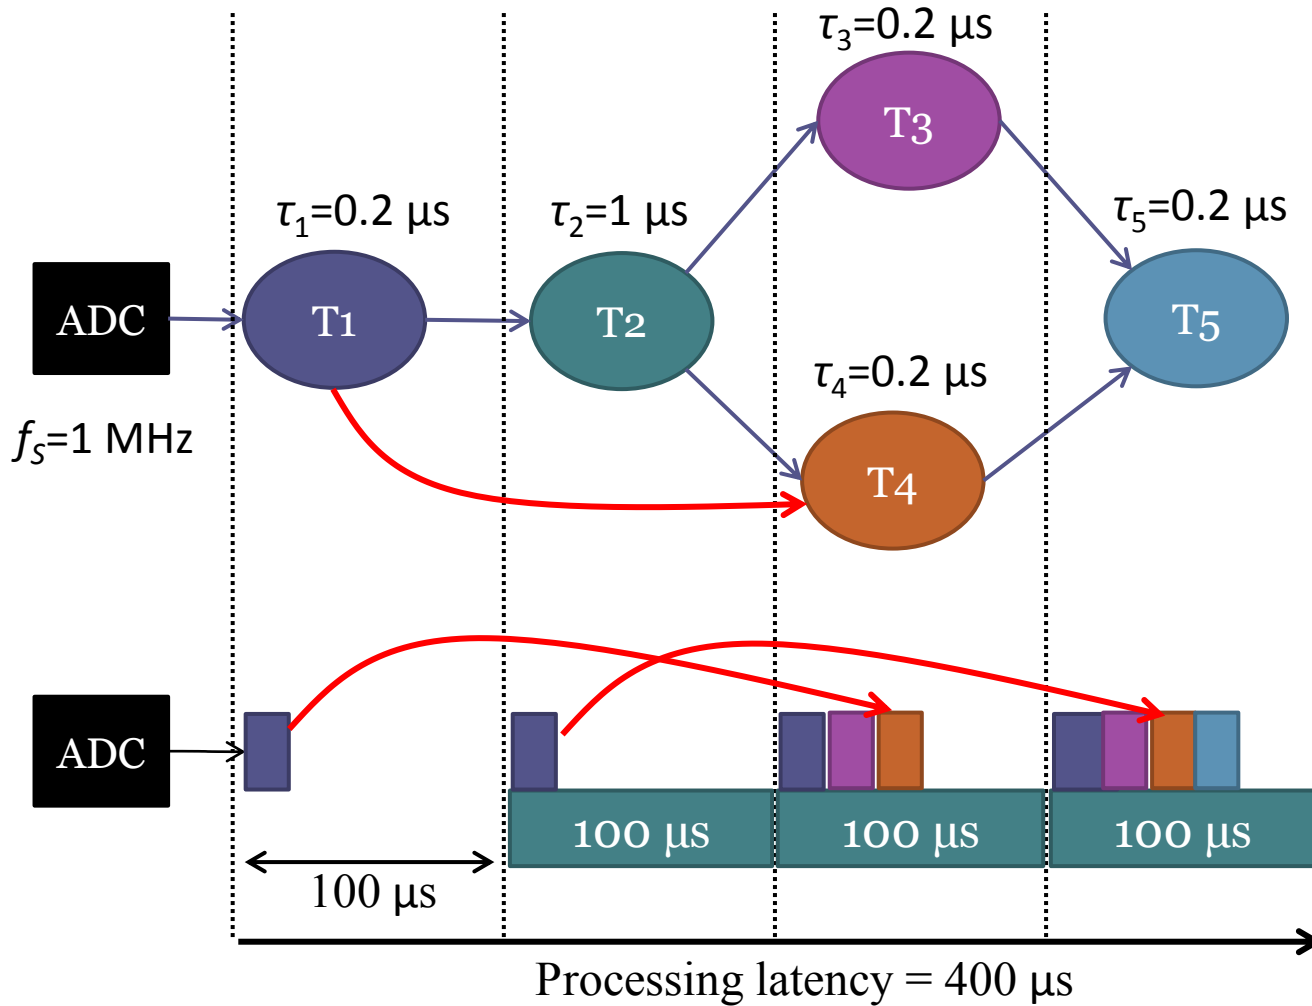


PIPELINE

Block size: **50 samples**
 Time slot: $50/f_s = \mathbf{50 \mu s}$
 Pipelining stages S : 3
 Processing latency:
 $S \cdot T_{ts} = \mathbf{300 \mu s}$

Scheduling performance: platform-independent

Pipelined Scheduling: Control Flow



PIPELINE
Block size: 100 samples
Period T : $1/f_s = 100 \mu\text{s}$
Pipelining stages S : 4
Processing latency: $S \cdot T = 400 \mu\text{s}$

...

Scheduling performance: platform-independent

Mapping

- Application and platform models
- Any mapping algorithm
- Two general-purpose algorithms:
 - t_w -mapping: $O(m \cdot n^{w+1})$
 - g_w -mapping: $O(m \cdot n^w)$
- Cost function





$$\text{Cost} = \frac{\text{processing requirement}}{\text{available processing power}} + \frac{\text{bandwidth requirement}}{\text{available bandwidth}}$$

balance processing load *minimize data flows*


DEMO 1: Computing Resource Management

- LTE 1.4 MHz
- Time slot: 1 ms
- Sampling frequency: 1.92 MHz
- Each time slot, 1920 complex samples are sent to/ received from the USRP
- Receiver has 7 pipeline stages: 7 ms latency


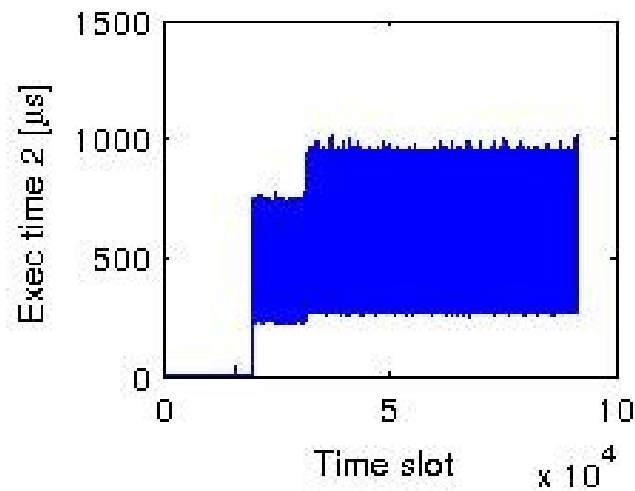

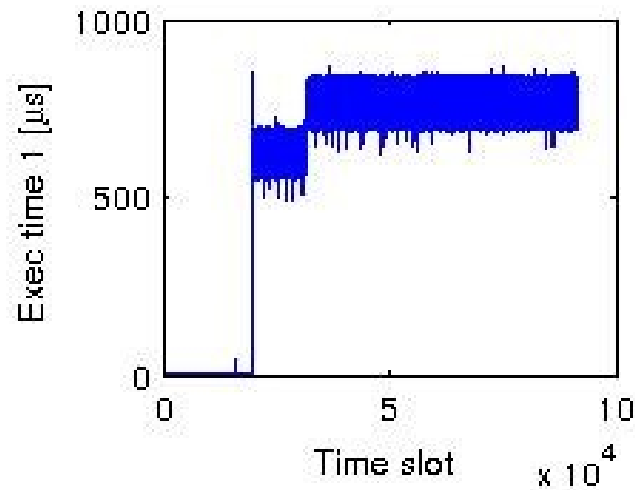
Soft Real-Time: Execution Trace




Soft Real-Time: Execution Trace



Hard Real-Time: Execution Trace




Hard Real-Time: Latency



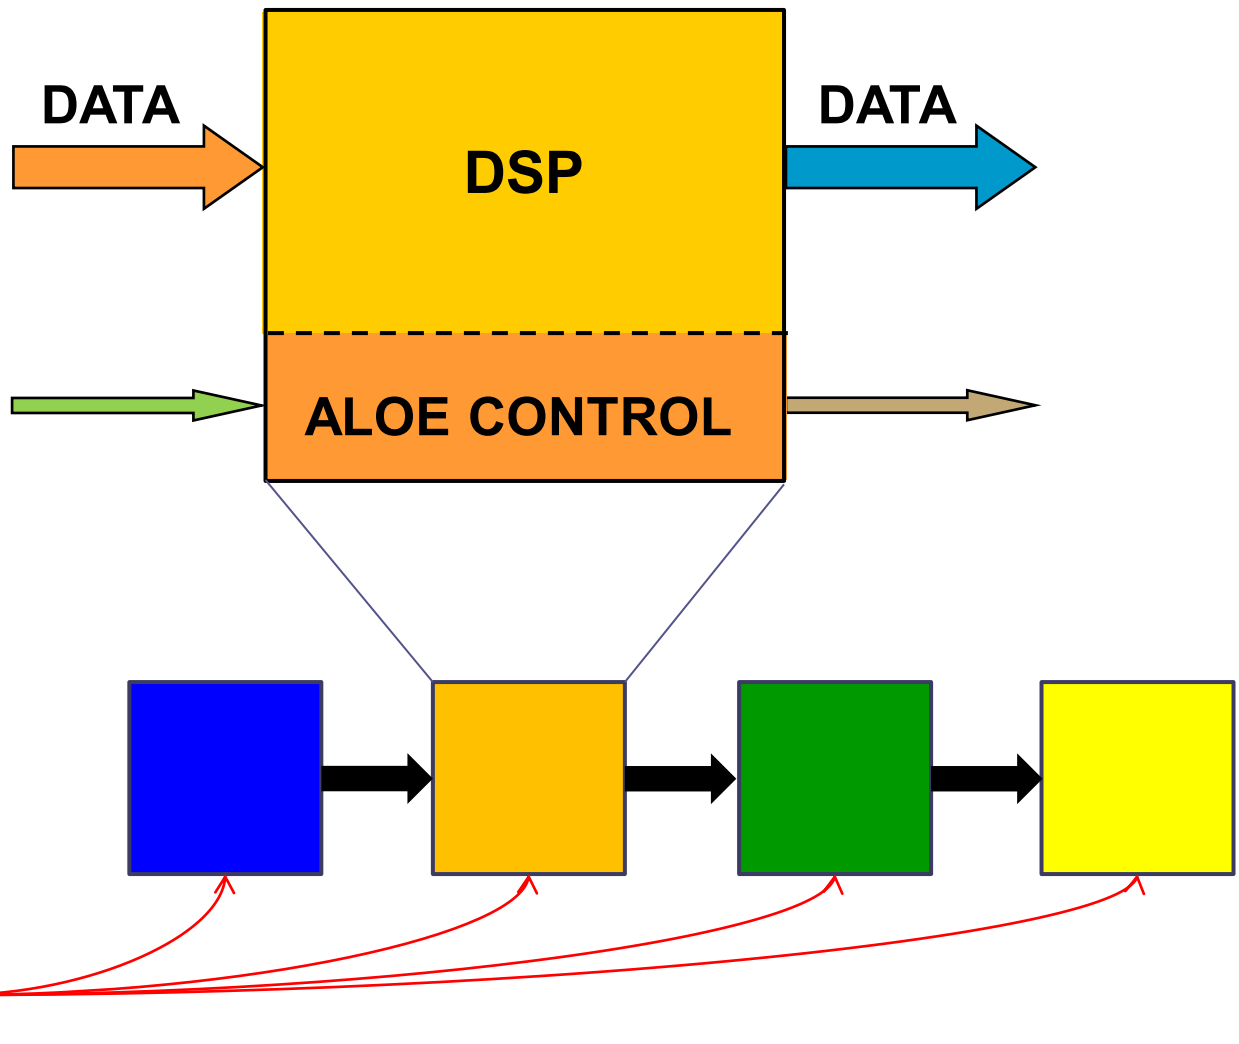
Waveform Design and Deployment

ALOE Waveform

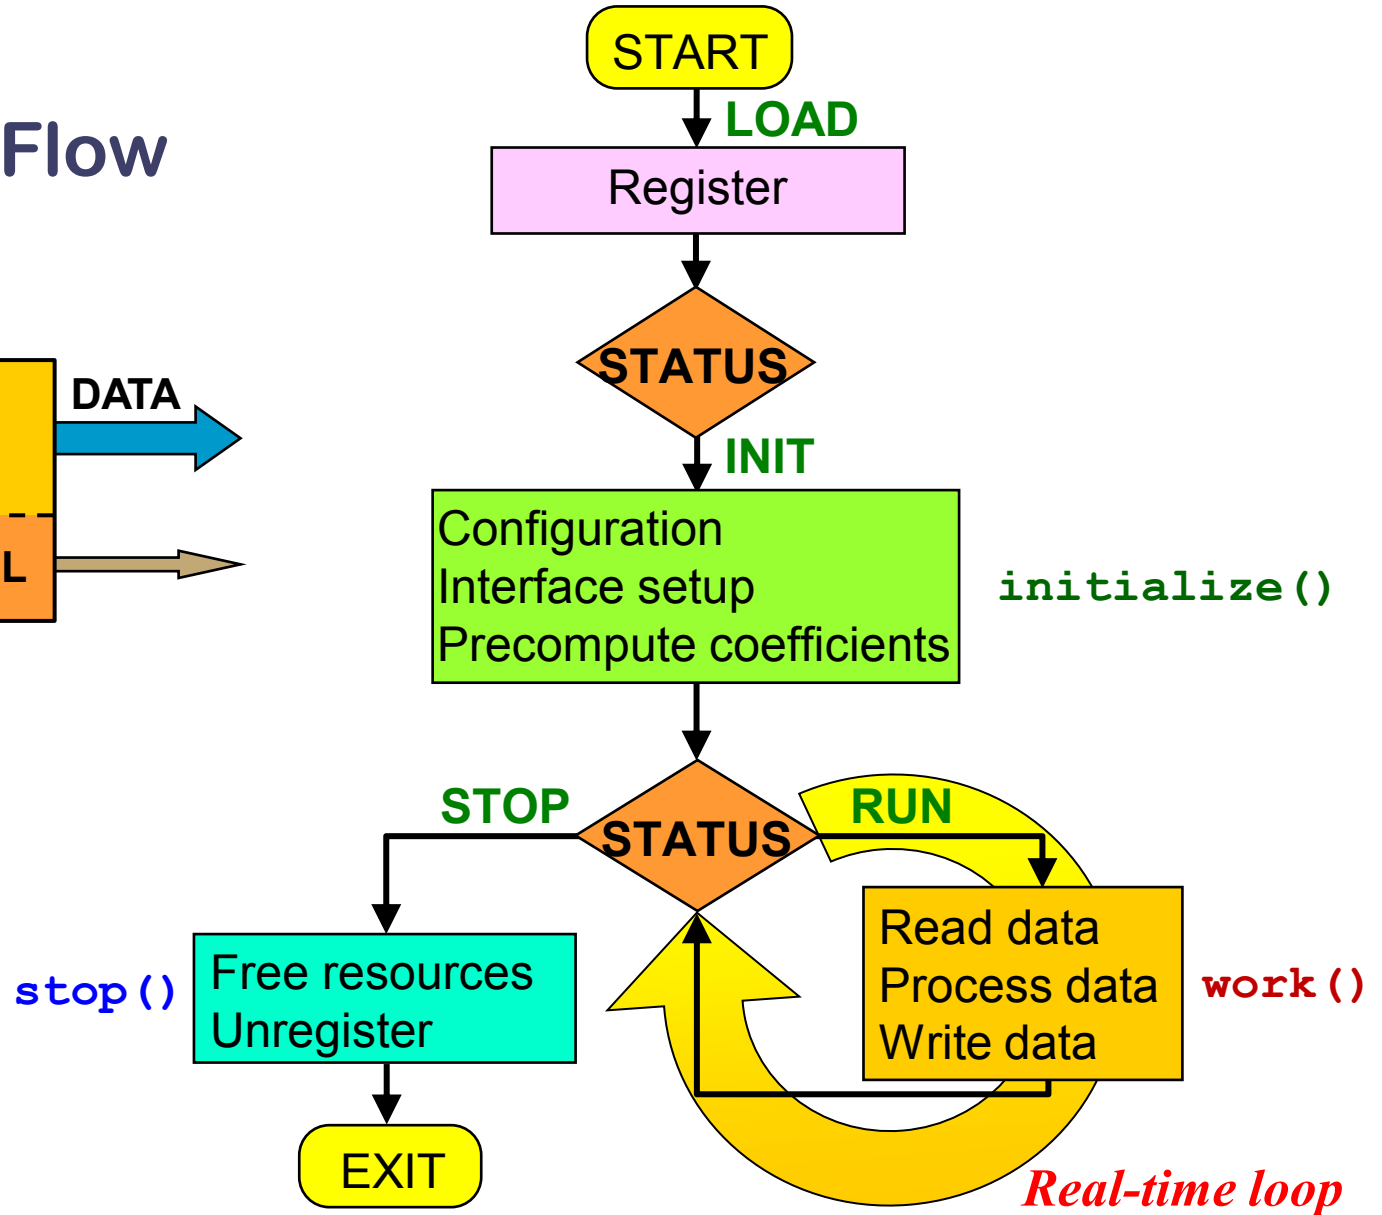



- **ALOE waveform**
 - Processing modules
 - Connections
 - Parameters
- **Module**
 - Computing requirements
 - Configuration parameters

Waveform Module



Module Execution Flow



Waveform Design and Deployment

Development

Implementation of DSP algorithms

CRC

Turbo
Coder


Code
Blk
Segm.

Code
Blk
Concat.

Rate
Matching

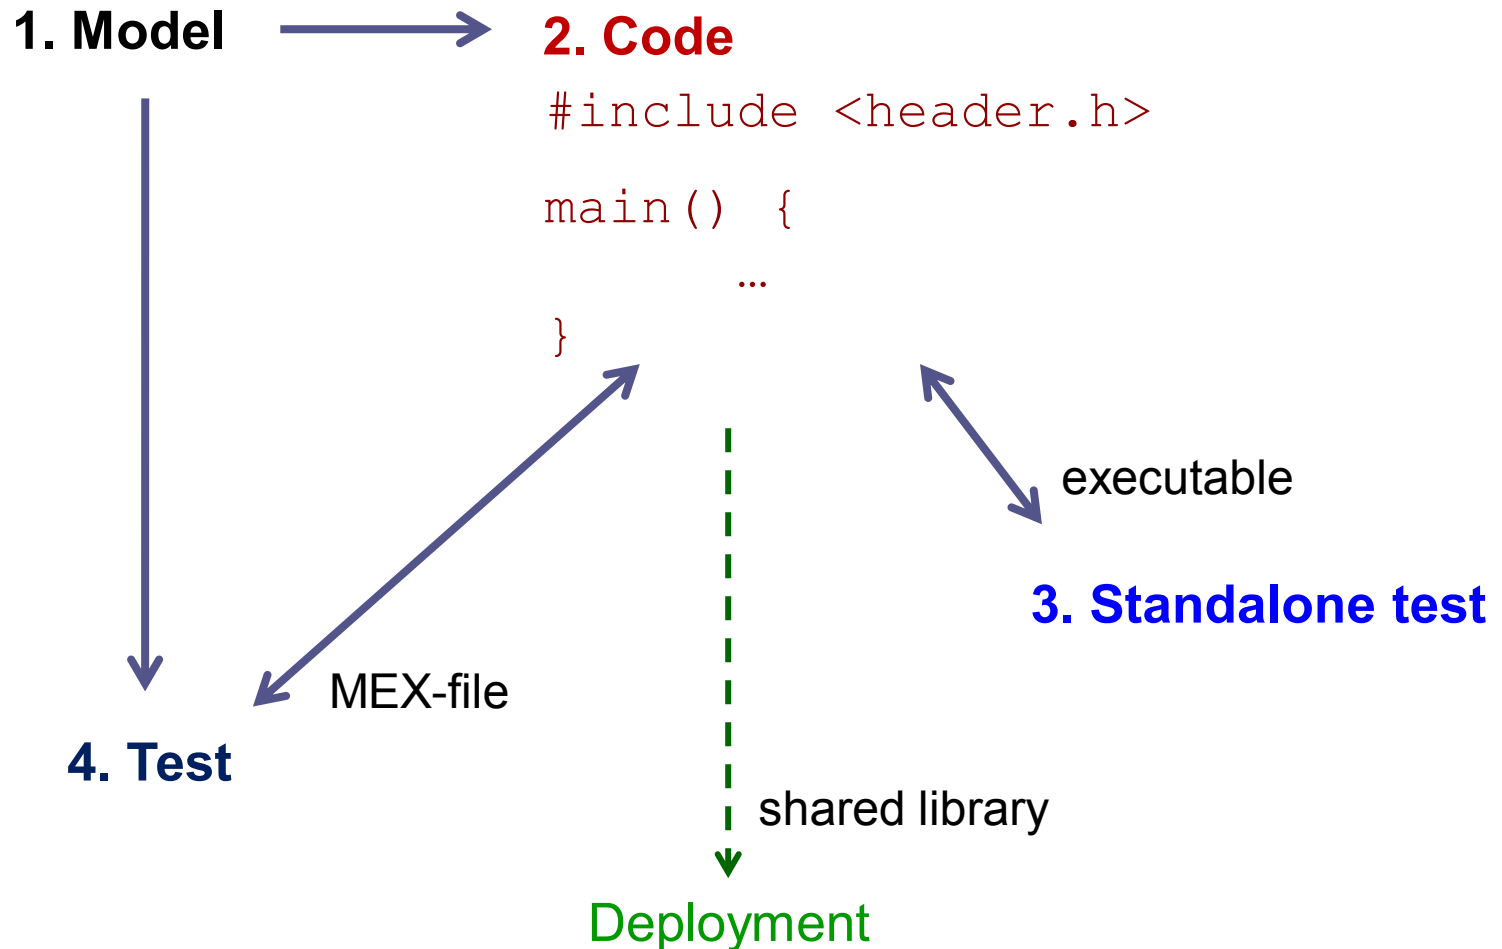
Deployment

Waveform creation and execution



- Parameters
- Execution time slot
- Pipelining stages
- ...

Module Development



Module Template

initialize()

```
int initialize() {  
    ...  
    param_get_int_name("dft_len",&dft_len);  
    dft_plan = compute_dft_plan(dft_len);  
    return 0;  
}
```

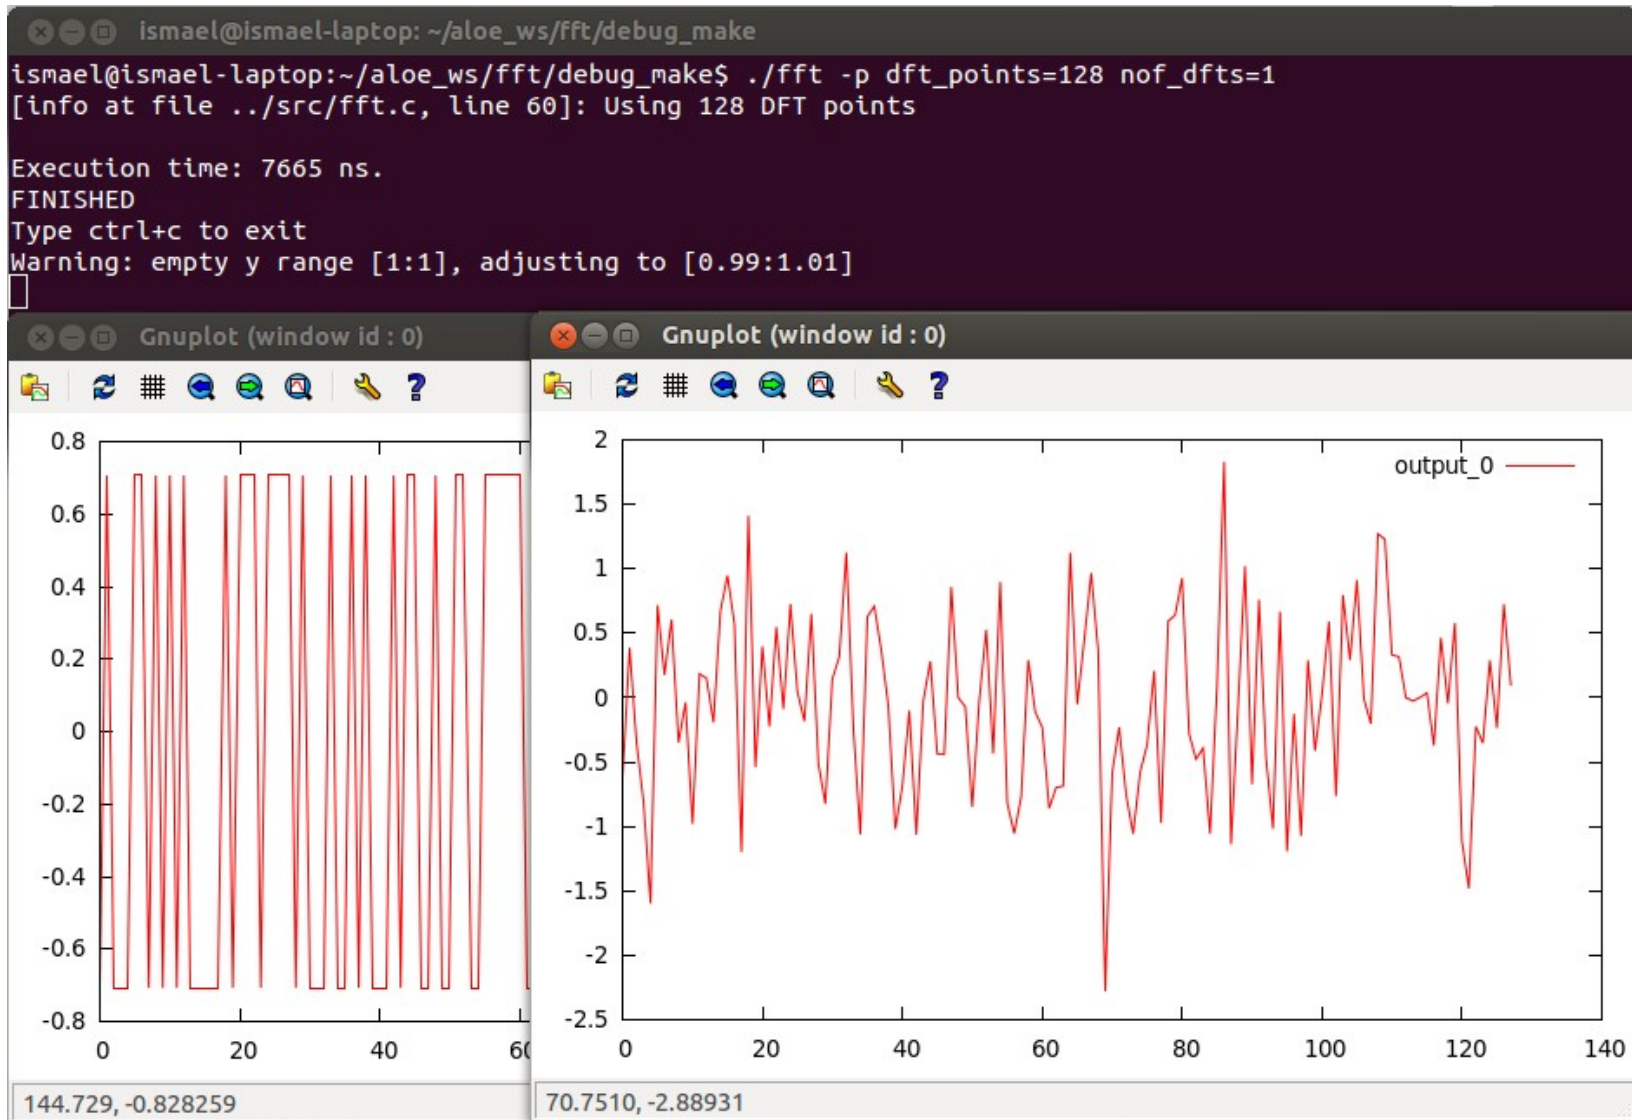
work()

```
int work(void **inp, void **out) {  
    complex_t *input = inp[0];  
    complex_t *output = out[0];  
    int nsamples = get_input_samples(0);  
    param_get_int_name("runtime_param",&my_param);  
    run_dft(dft_plan, input, output);  
    for (int i=0;i<nsamples;i++) {  
        ...  
    }  
    return 0;  
}
```

stop()

```
int stop() {  
    destroy_dft_plan(dft_plan);  
}
```

Standalone Execution: Debugging



Standalone Execution: Profiling (e.g. Valgrind)

KCachegrind Mem: 19% | CPU: 2% | ↑ 0 KiB/s ↓ 1 KiB/s temp 165.0°C 1:17 2:20 PM

```
ismael@ismael-laptop:~/aloe_ws/eclipse_aloe/modrep_osld/lte_sss_synch$ valgrind --tool=callgrind ./lte_sss_synch input_len=1920
==29784== Callgrind, a call-graph generating cache profiler
==29784== Copyright (C) 2002-2011, and GNU GPL'd, by Josef Weidendorfer et al.
==29784== Using Valgrind-3.7.0 and LibVEX; rerun with -h for copyright info
==29784== Command: ./lte_sss_synch input_len=1920
==29784==
==29784== For interactive control, run 'callgrind_control -h'.
```

callgrind.out.29784 [./lte_sss_synch input_len=1920]

Flat Profile

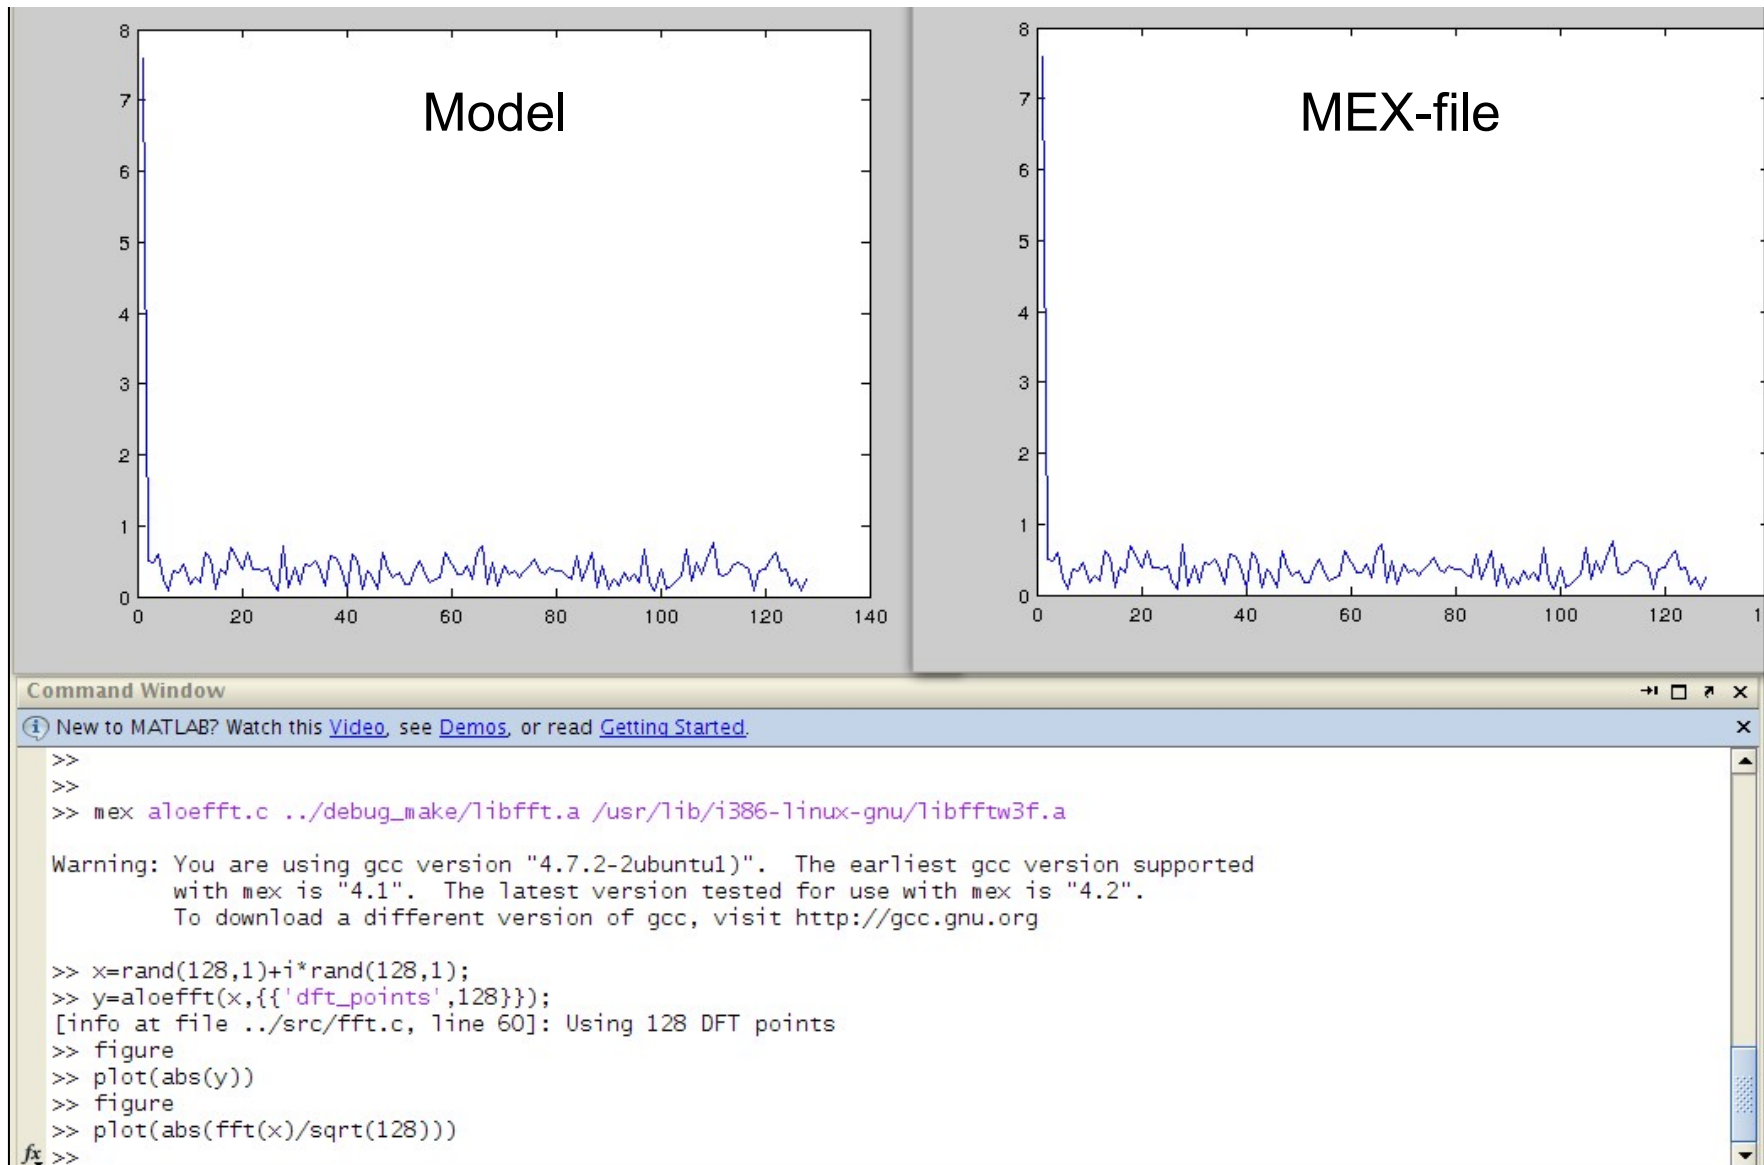
Search: (No Grouping)

Incl.	Self	Called	Function	Location
639 892	227 316	205	fftwf_twiddle_awake	lte_sss_synch
630 668	13 662	594	fftwf_mkplan	lte_sss_synch
616 224	5 720	440	fftwf_rdfc_solve	lte_sss_synch
612 320	74	1	work	lte_sss_synch:lte_sss_synch.c
606 083	161	1	get_m0m1	lte_sss_synch:find_sss.c
593 613	45 630	4 563	fftwf_ifree	lte_sss_synch
556 213	8 723	317	zero	lte_sss_synch
548 768	548 768	488	fftwf_cpy2d_pair	lte_sss_synch
539 696	18 656	424	fftwf_cpy2d_pair_ci	lte_sss_synch
538 857	45 630	4 563	fftwf_kernel_free	lte_sss_synch
529 481	529 211	135	malloc_consolidate	libc-2.15.so:malloc.c
503 841	1 407	67	vec_dot_prod	lte_sss_synch:vector.c
491 816	107 598	5 073	free	libc-2.15.so:malloc.c
487 727	30	1	fftwf_dft_conf_standard	lte_sss_synch
461 899	53 811	787	really_compress	lte_sss_synch
457 619	83 490	850	fftwf_mkstride	lte_sss_synch
451 431	38	1	fftwf_rdfc_conf_standard	lte_sss_synch
445 964	8 829	981	fftwf_solver_register	lte_sss_synch
441 504	53	1	get_volk_32fc_x2_multi...	libvolk.so.0.0.0:volk.c
437 135	101 662	981	register_solver	lte_sss_synch
433 516	12 753	981	fftwf_mksolver	lte_sss_synch
424 163	2 933	4	volk_rank_archs	libvolk.so.0.0.0:volk_rank_a...
417 856	11 880	360	copy	lte_sss_synch
410 825	6 546	1	load_preferences	libvolk.so.0.0.0:volk_prefs.c

work

Types Callers All Callers Callee Map Source Code

Verification



DEMO 2: Module Development

Waveform Design and Deployment

Development

Implementation of DSP algorithms

CRC

Turbo
Coder


Code
Blk
Segm.

Code
Blk
Concat.

Rate
Matching

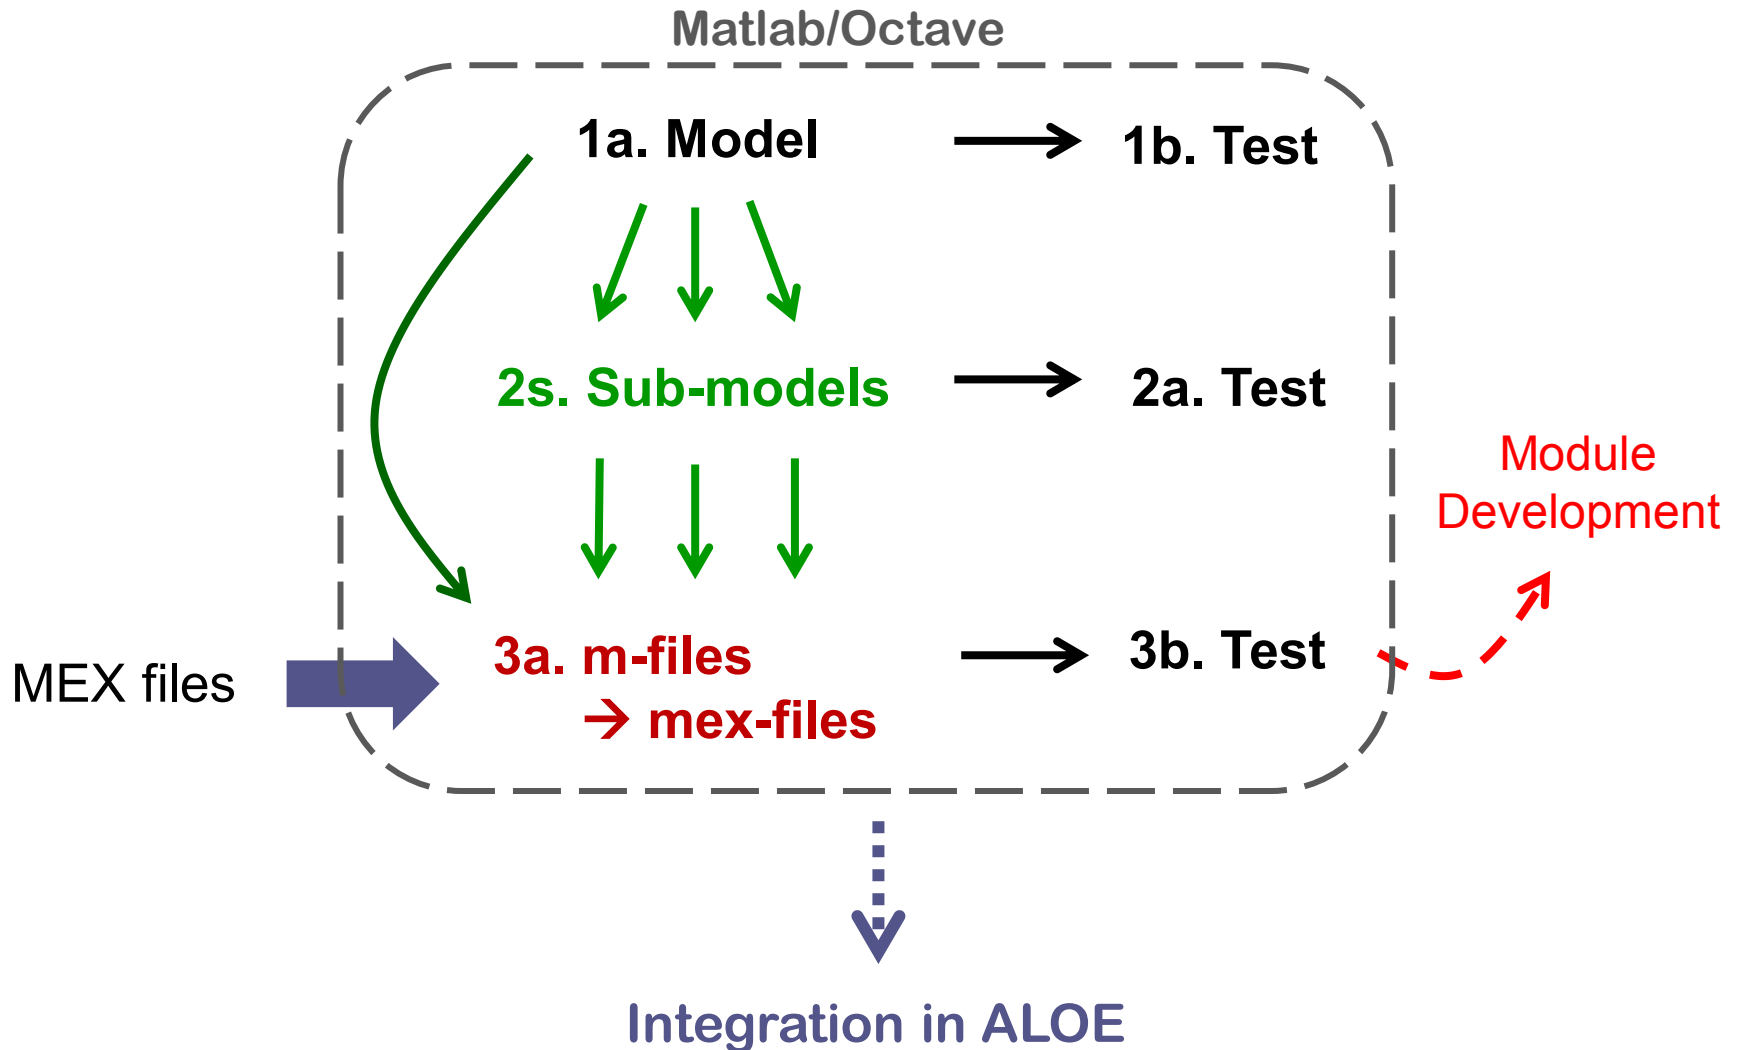
Deployment

Waveform creation and execution



- Parameters
- Execution time slot
- Pipelining stages
- ...

Waveform Definition and Testing

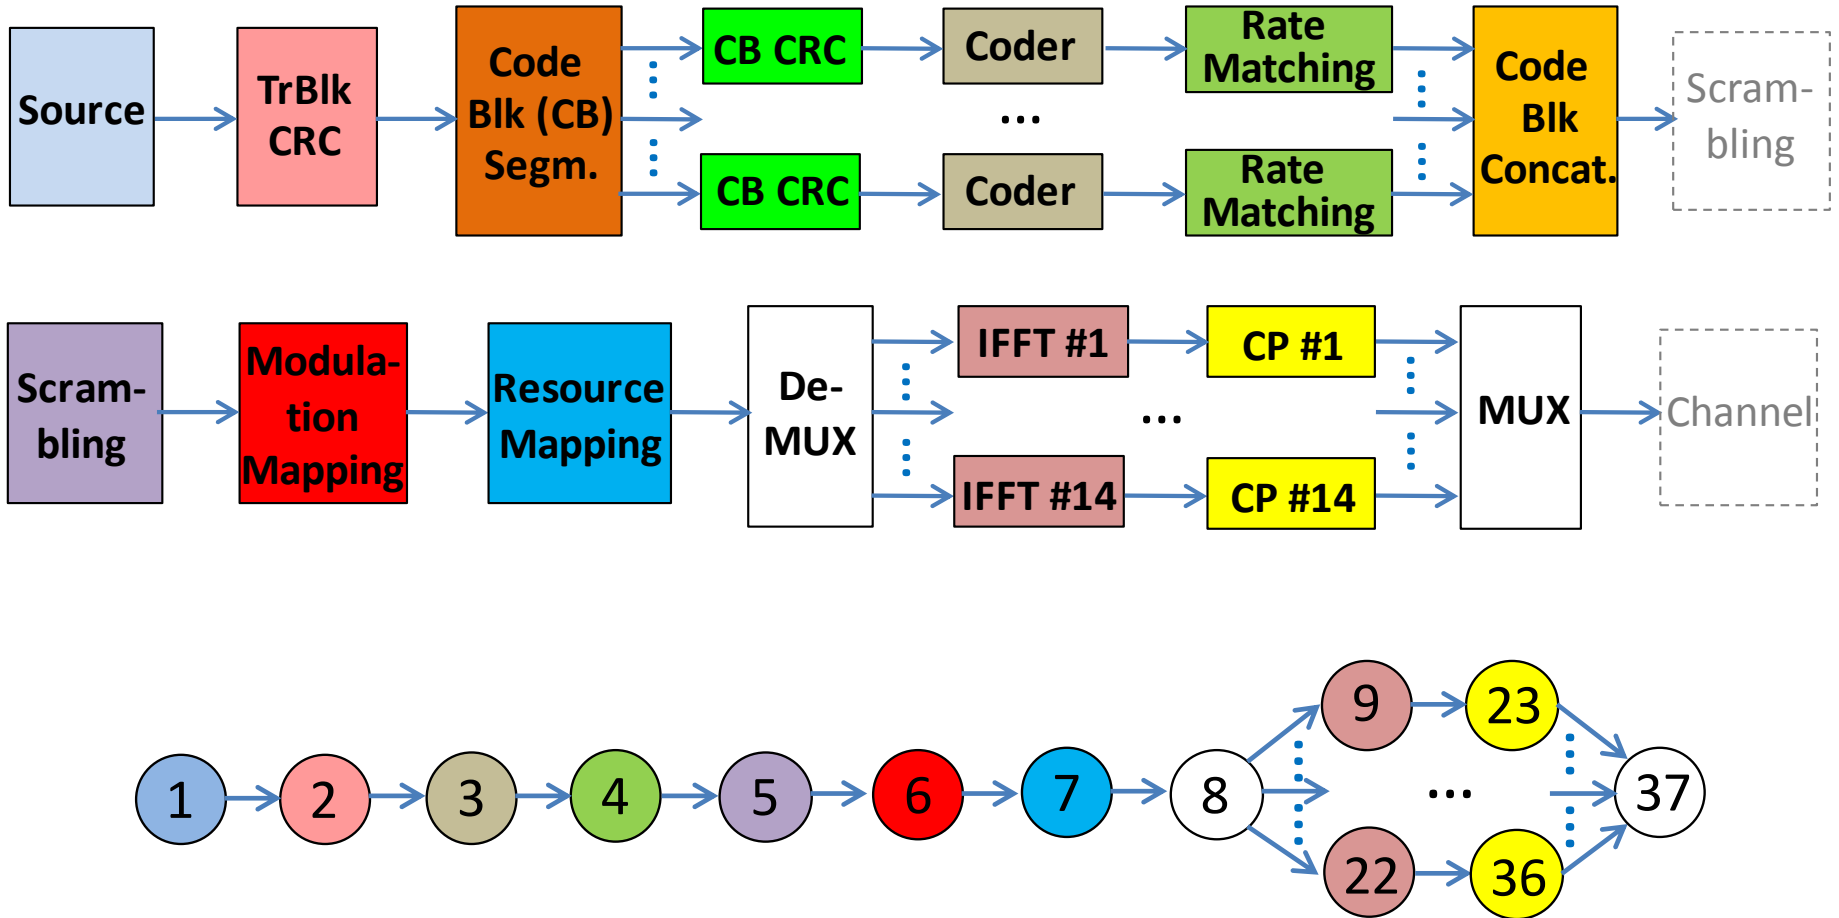


Application Description File (.app)

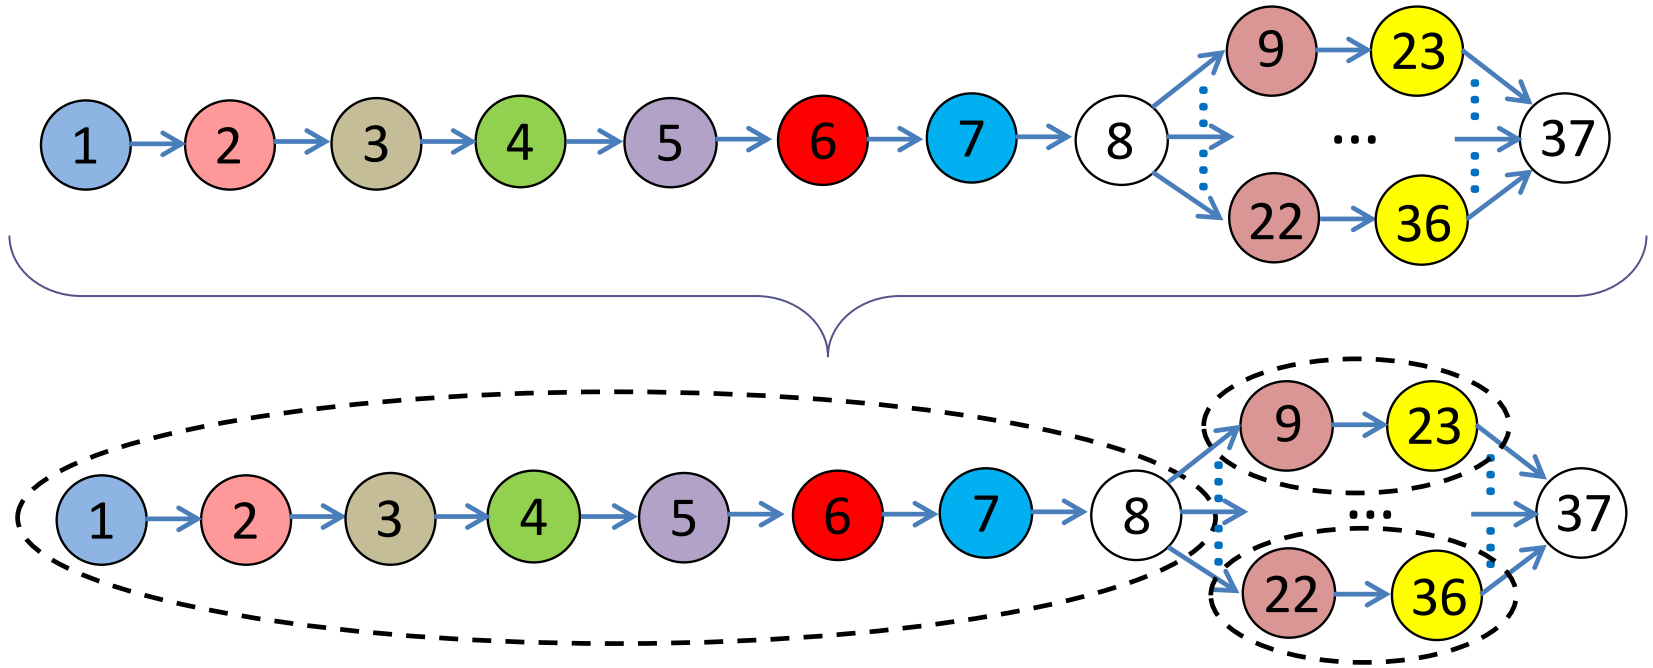
Field	Options	Description
modules	name	Unique ID
	binary	name and rel. location of the binary
	mopts	Processing requirements (for mapping)
	Variables	Configuration parameter values
interfaces	{src=<source>; dest=<destination>}, {...}, ...	Connection of modules: - <i>src</i> : output interface(s) of source module - <i>dest</i> : input interface(s) of destination module
join_stages	(<m1>,<m2>,...), (...)	Executes modules listed in each pair of brackets in a single pipelining stage

Pipelining Stages

PDSCH - Tx



Join Stages



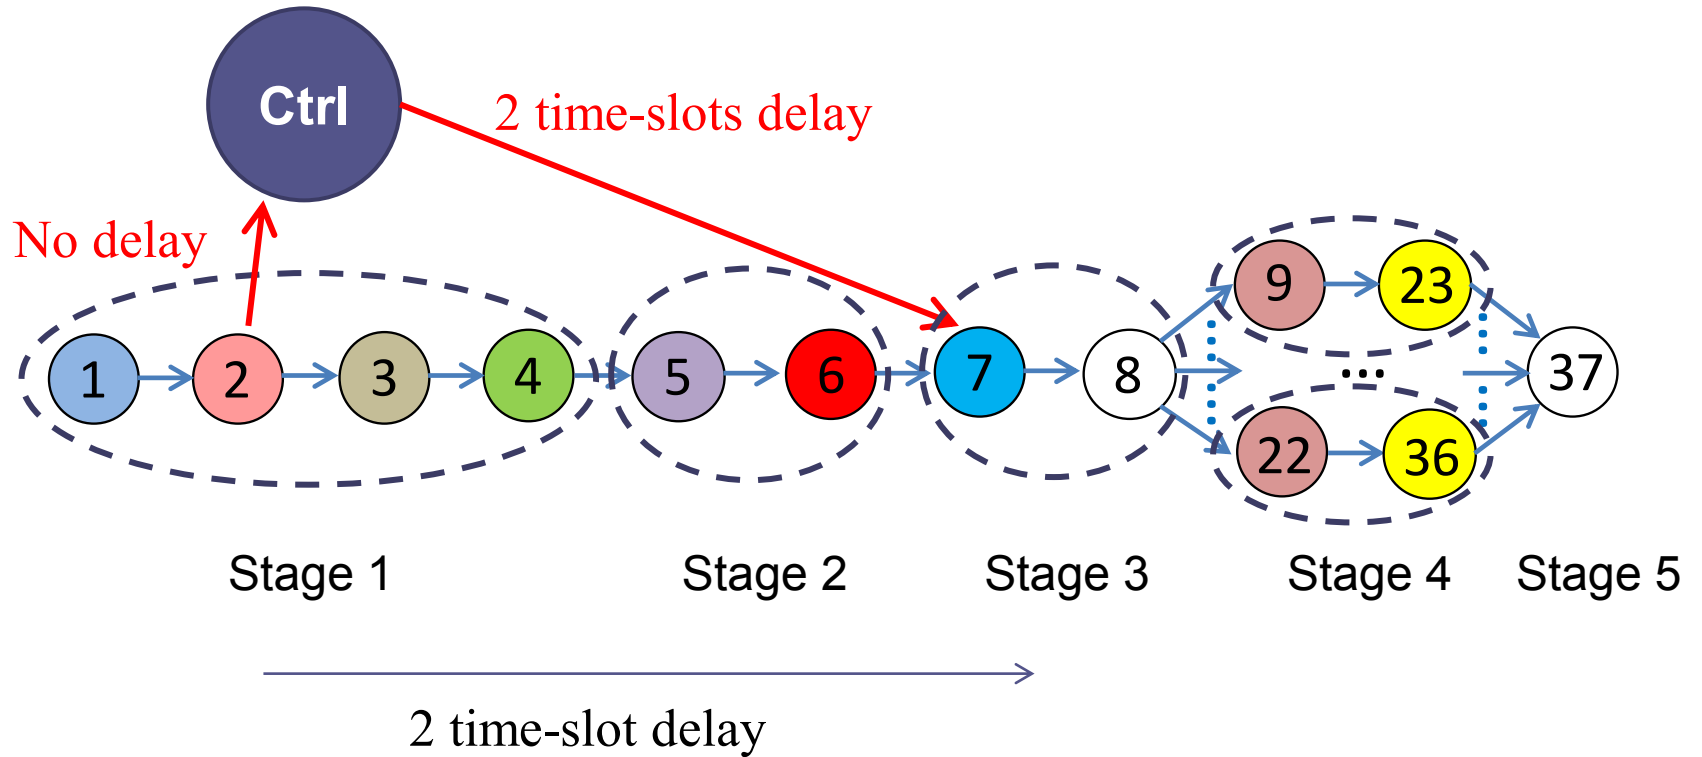
join stages=

```
(
    (M1, M2, M3, M4, M5, M6, M7, M8),
    (M9, M23),
    ...
    (M22, M36)
);
```

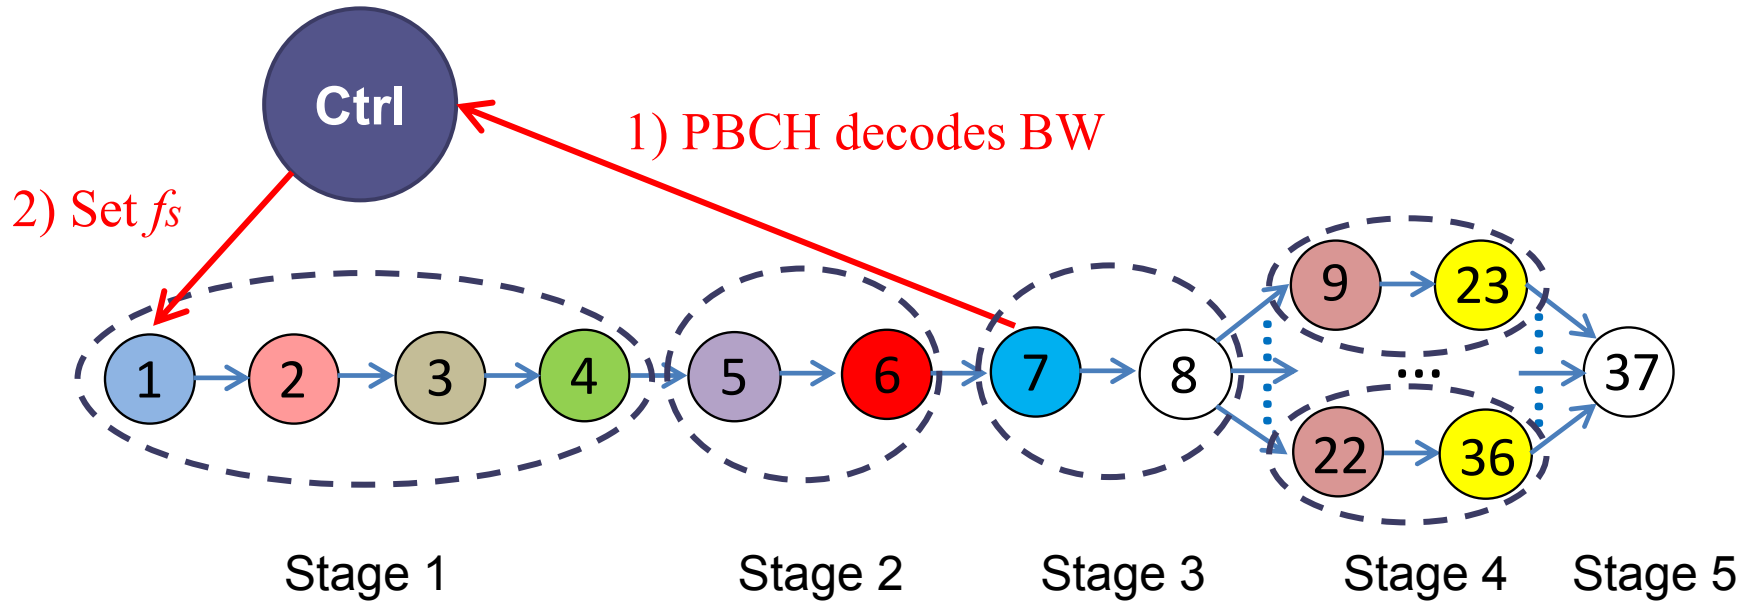
Waveform Deployment Tools

1. Waveform description file
 - Collection of sub-waveform description files
2. Decouple Tx-Rx
 - Tx writes to file, Rx reads from file
 - Modify file with Matlab (add channel noise/distortion)
3. Debug mode
 - Logging service
4. Real-time execution:
 - UHD support
 - Execution statistics

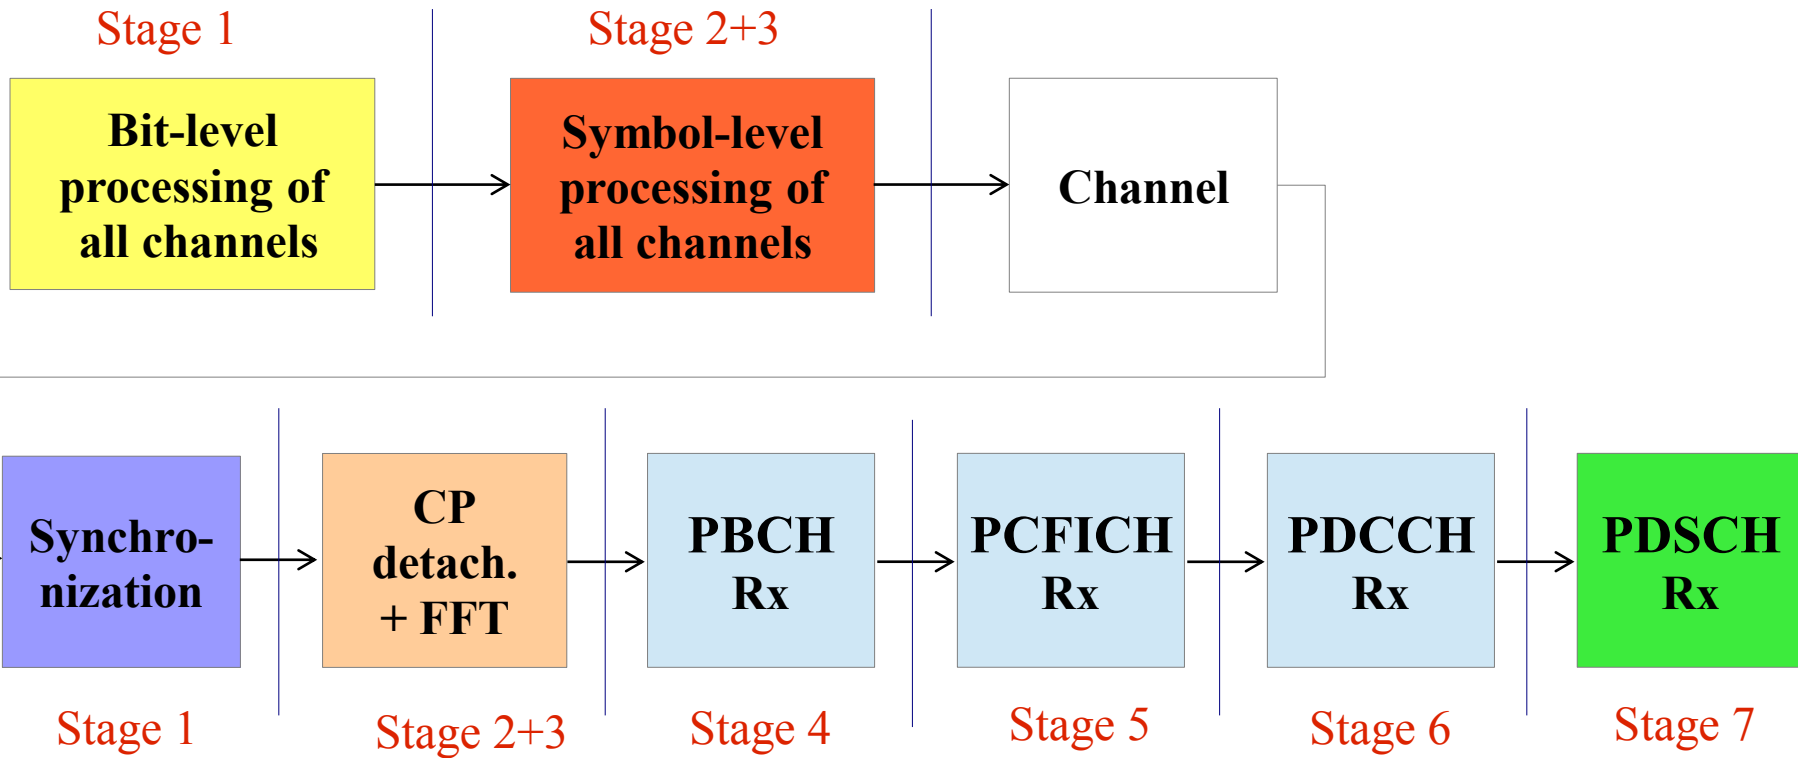
Control Module (I)

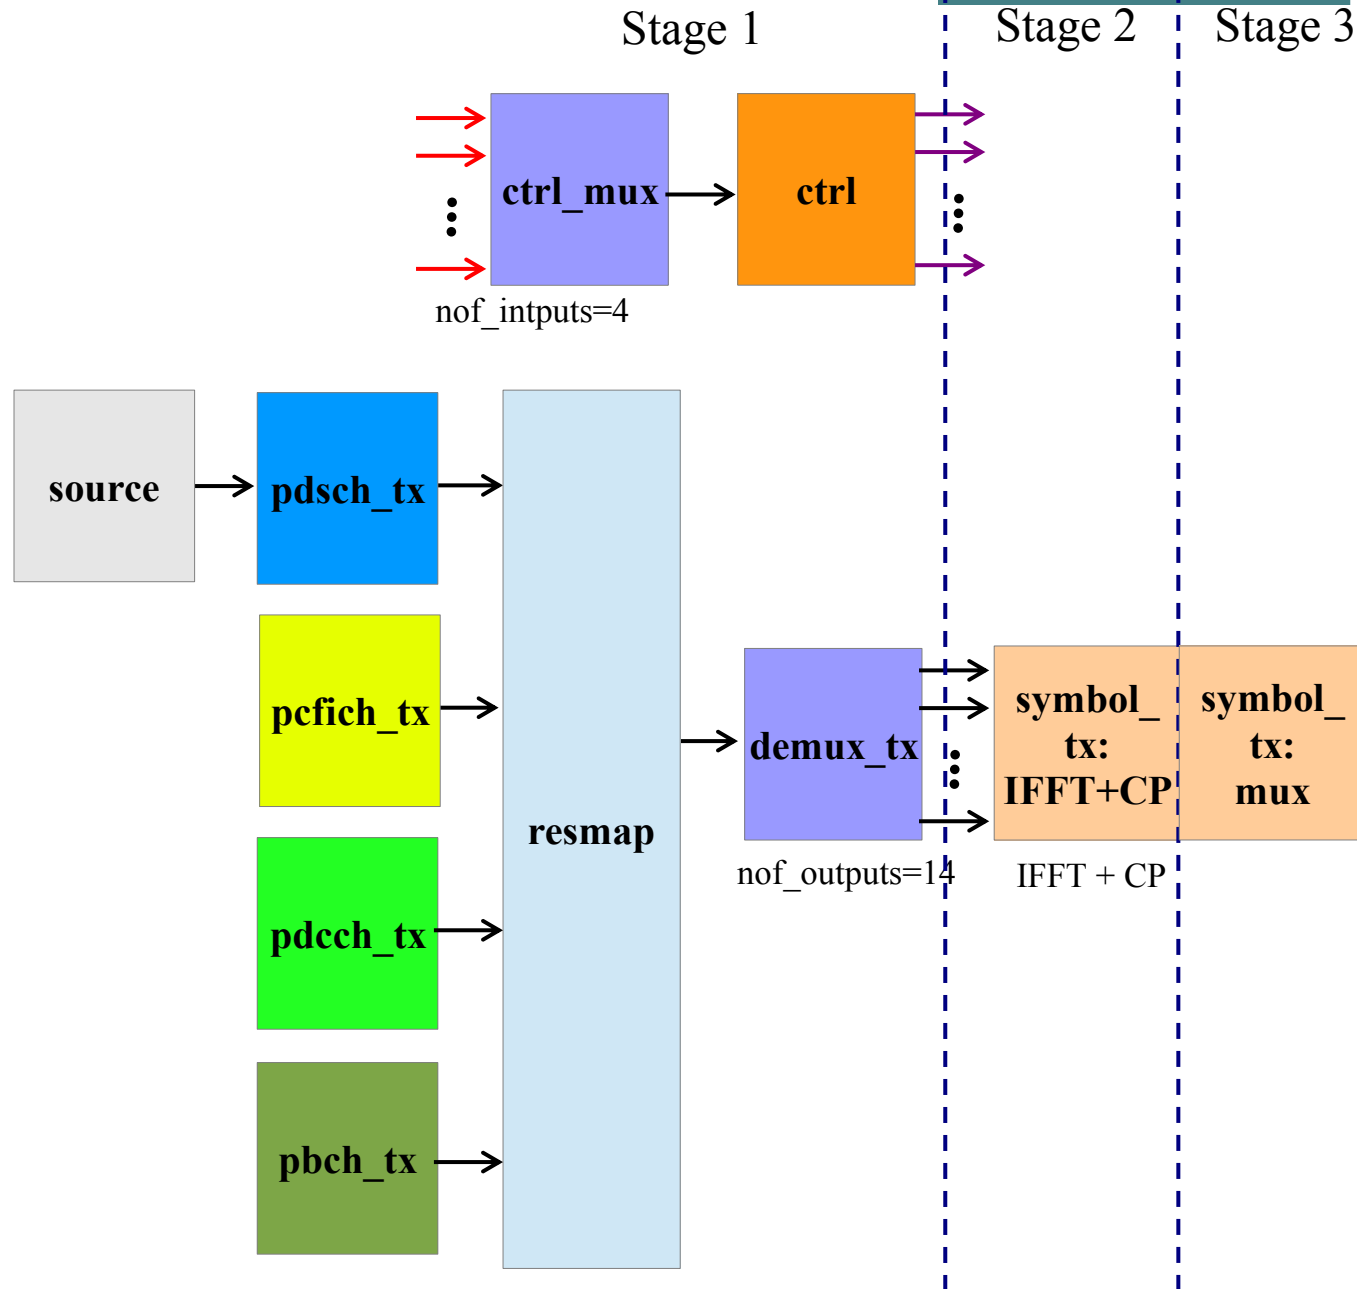



Control Module (II)

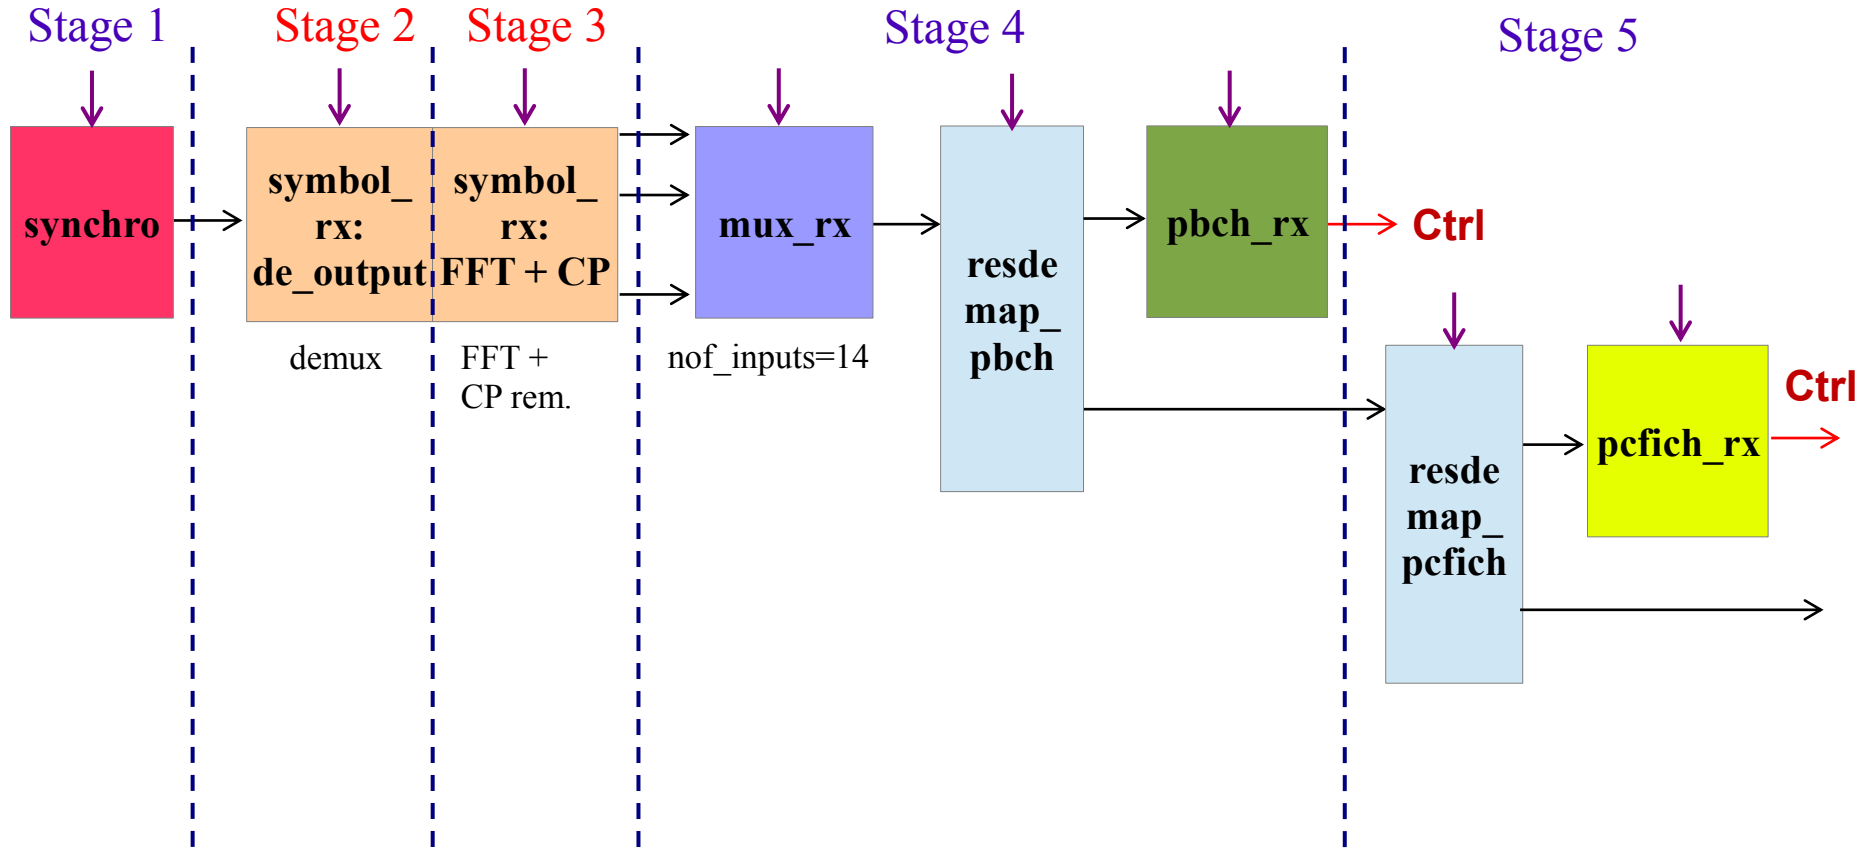


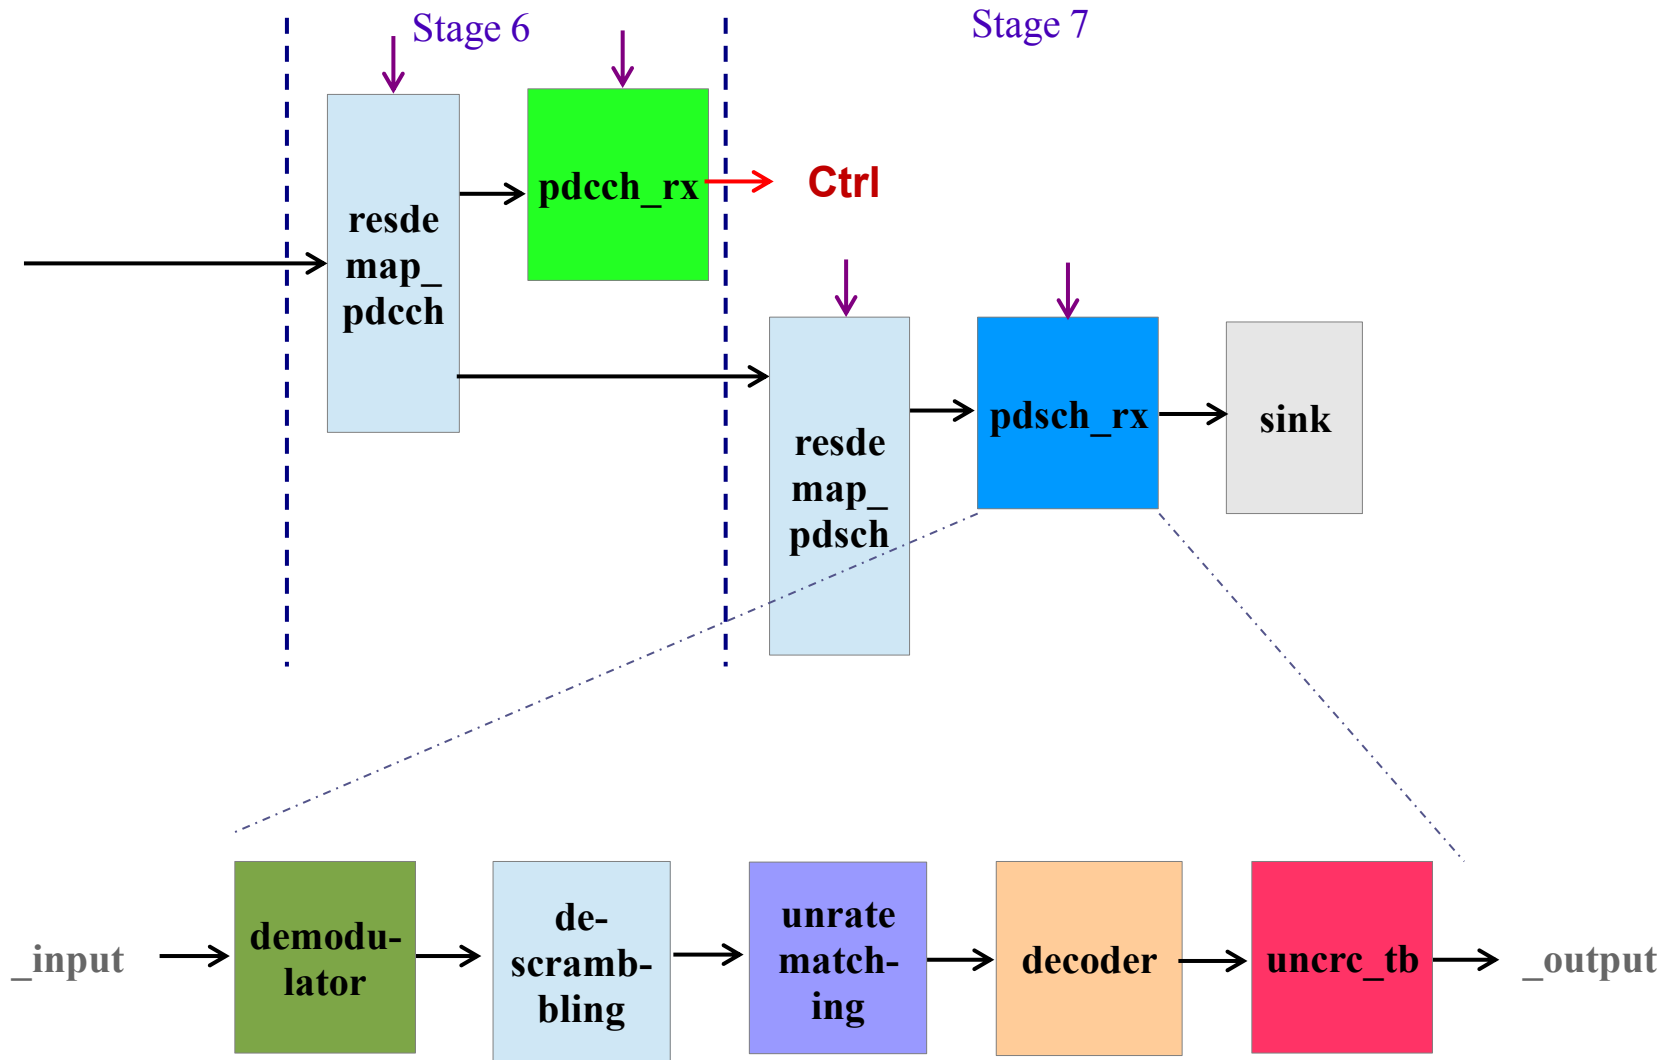
LTE DL Waveform



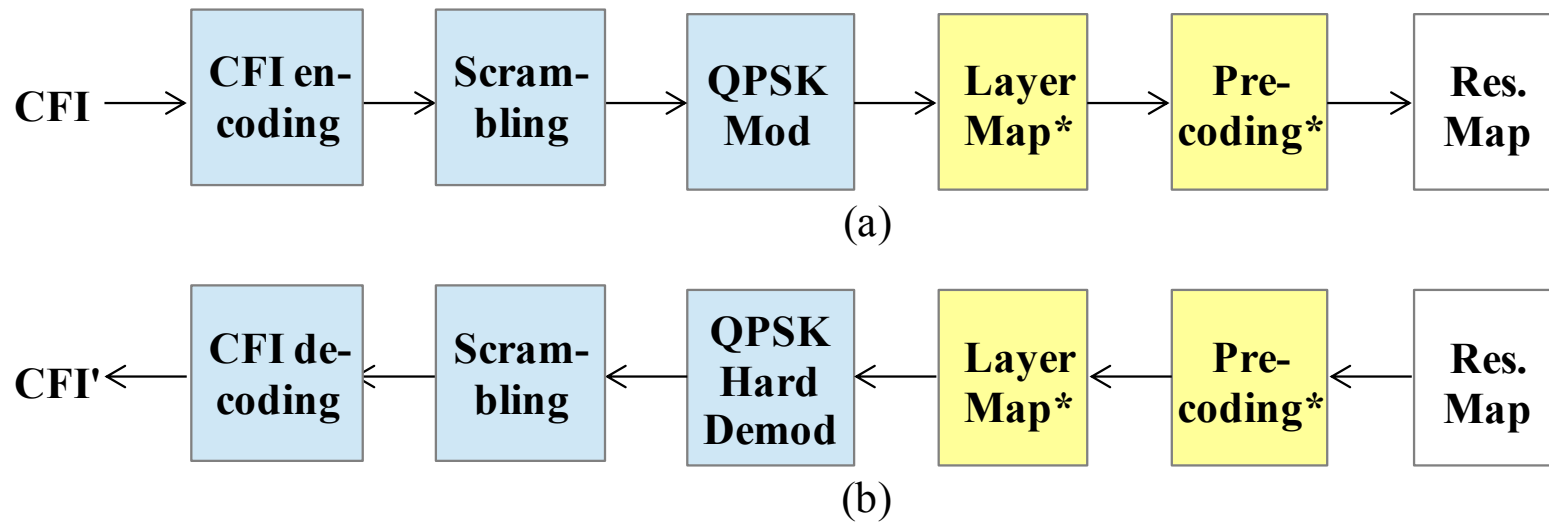


pdsch_tx

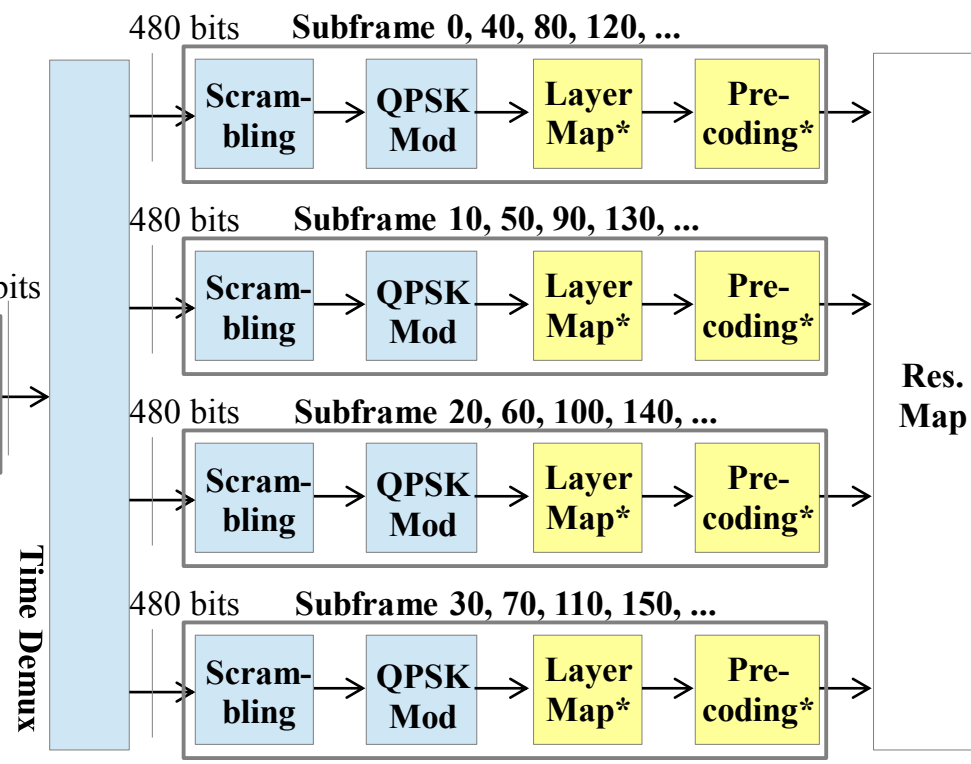
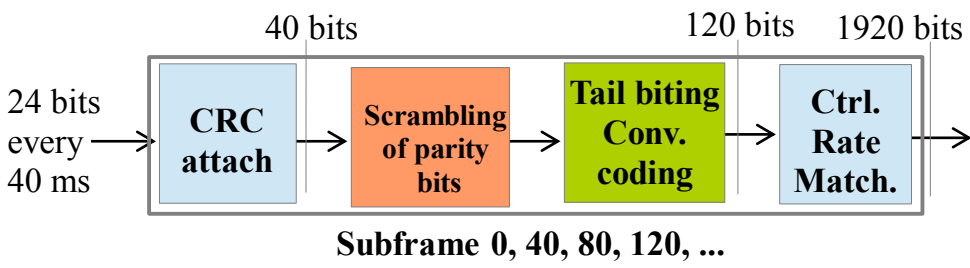





PCFICH



BCH

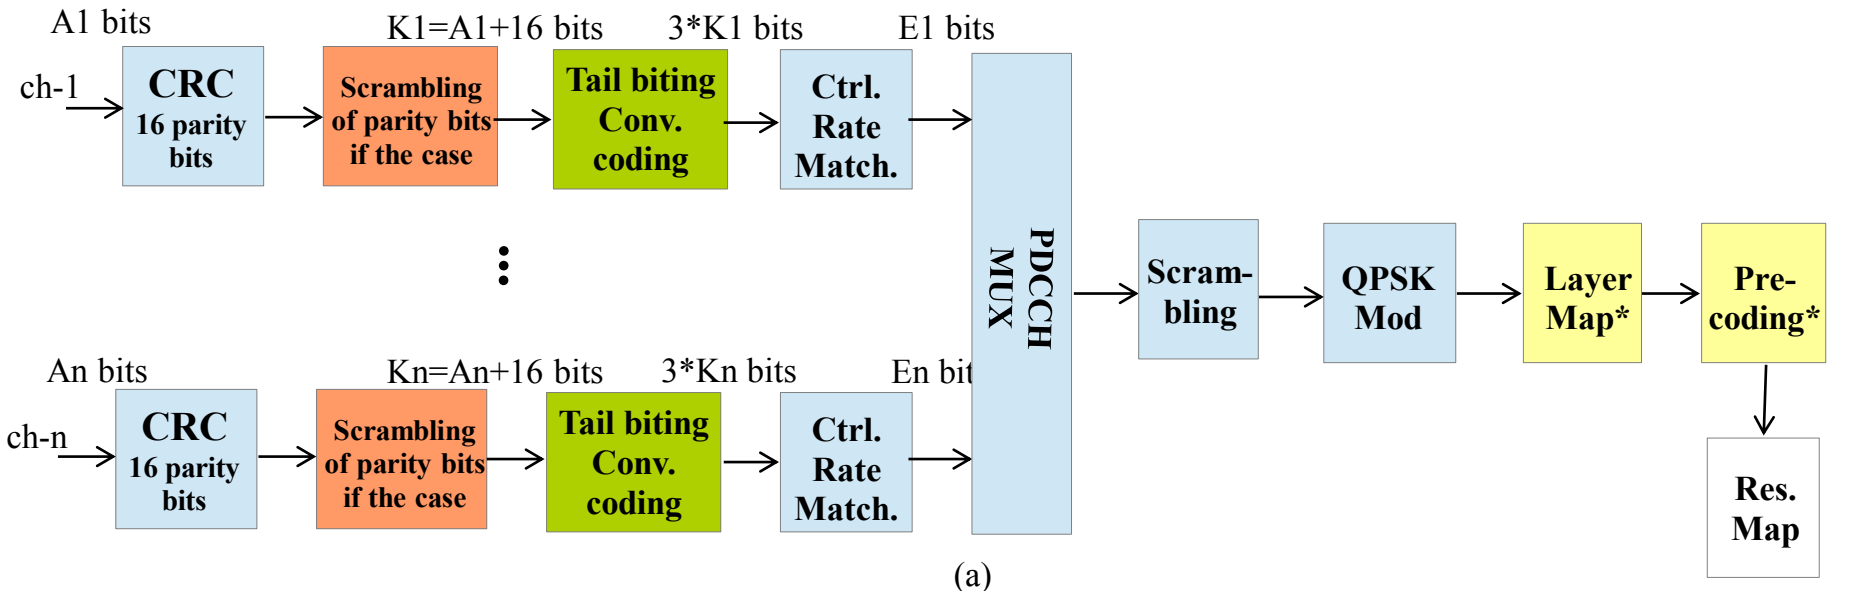


(a)




(b)

PDCCH



(a)



(b)

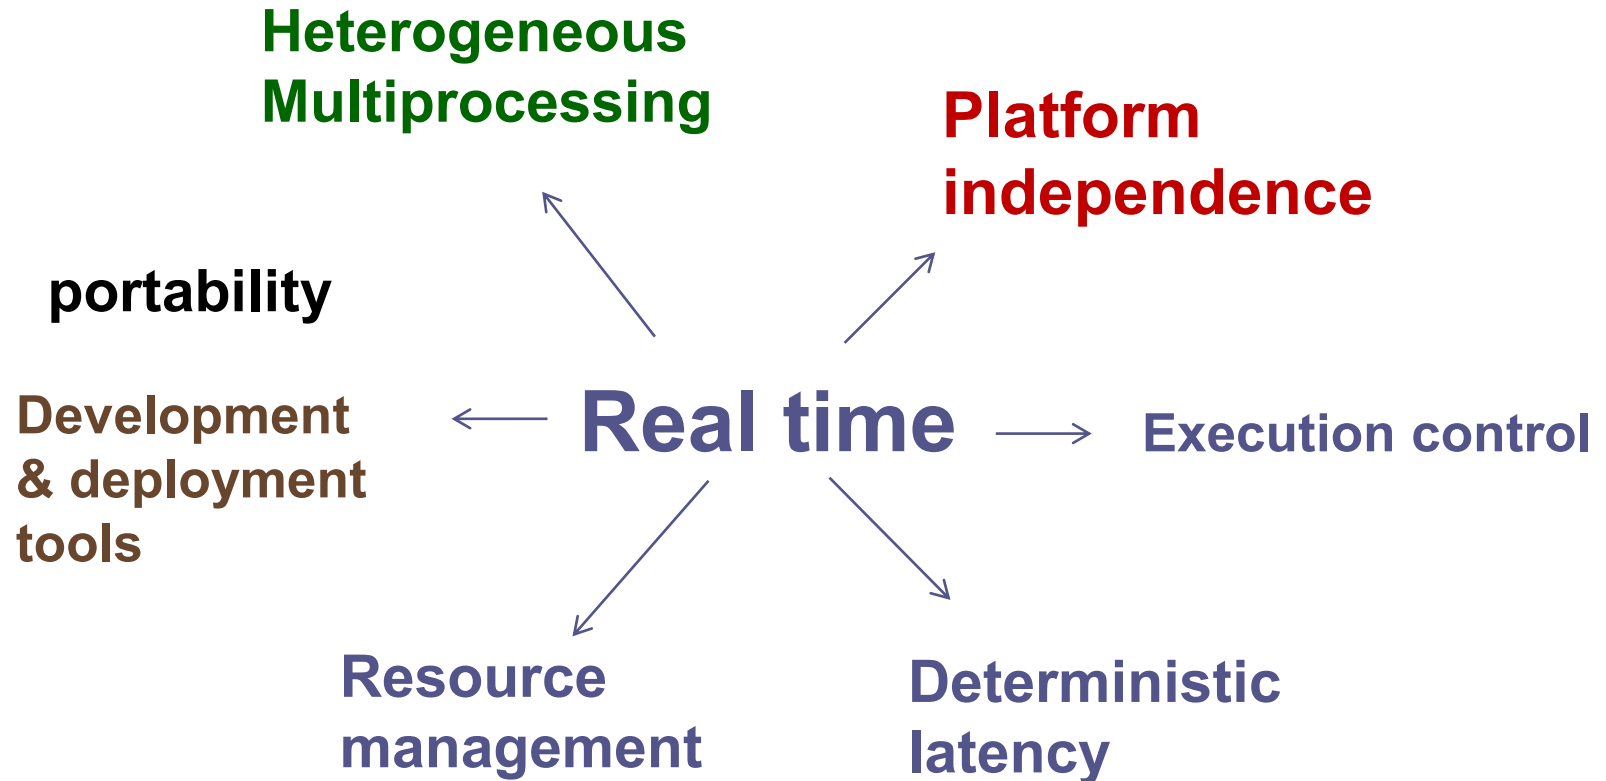
DEMO 3: Waveform Deployment

Conclusions

SDR Frameworks


- **SCA (Software Communication Architecture)**
 - Military
 - Research and education (e.g. OSSIE)
- **GNU Radio**
 - Research and education (PC, multicore)
- ...

ALOE Characteristics




Conclusions

multicores



many-cores



small clusters
(heterogeneous)



Future Work

- GNU Radio/ SCA compatibility
- Add and test new schedulers:
 - Dynamic, provided by the RTOS
 - Hybrid (static-dynamic)
- Tools
 - Waveform development and deployment
 - Graphical User Interface for ALOE++
 - ...

Call for Participation

- **FlexNets** (Flexible Wireless Communications Systems & Networks)
 - <http://flexnets.upc.edu/trac/>
 - ALOE releases
 - Computing resource management framework
 - Waveforms
 - Educational material
- **OSLD** (Open source LTE deployment)
 - <https://sites.google.com/site/osldproject/home>
 - <https://github.com/flexnets>
 - ALOE++
 - DSP modules library
 - Development Tools
- **Mailing lists:** <https://groups.google.com/group/flexnets>

