Exploiting Dataflows for Reconfigurable Hardware Accelerators

Francesca Palumbo[1], Claudio Rubattu[1,2], Carlo Sau[3], Tiziana Fanni[3], Luigi Raffo[3]

[1]University of Sassari, PolComIng – Information Engineering Group

[2]University of Rennes, INSA Group

[3]University of Cagliari, Diee – Microelectronics and Bioengineering Group

Rennes, 12-14 December 2017

Who and Where

UNIVERSITY OF SASSARI

UNIVERSITY OF CAGLIARI

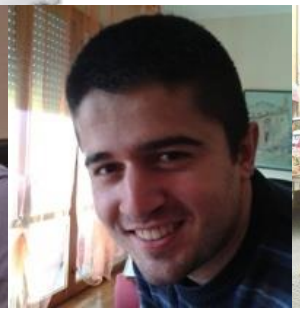

Who and Where

UNIVERSITY OF SASSARI

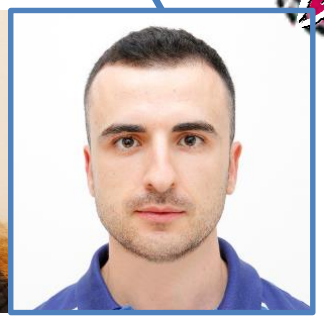

UNIVERSITY OF CAGLIARI

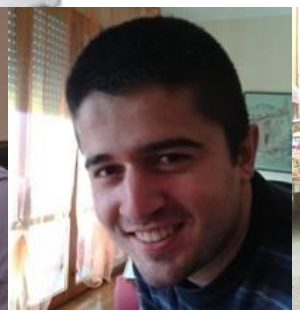

Outline

- The origins of our dataflow to hardware studies: the RPCT Project
  - Context
  - Target Technologies
  - Project Development
- The MDC tool
  - Approach
  - Baseline Functionality and Extensions
- Contexts of application
  - Neural Signal Decoding
  - HEVC Interpolation Filters
- Final Remarks

Outline

- The origins of our dataflow to hardware studies: the RPCT Project
  - Context
  - Target Technologies
  - Project Development
- The MDC tool
  - Approach
  - Baseline Functionality and Extensions
- Contexts of application
  - Neural Signal Decoding
  - HEVC Interpolation Filters
- Final Remarks

Modern Embedded Systems

Embedded Systems (*real-time* computing systems with a dedicated functionality) are pervasive (*98%* of computers are embedded) and may present *sensing* and *actuating* capabilities.

Modern Embedded Systems

Embedded Systems (*real-time* computing systems with a dedicated functionality) are pervasive (*98%* of computers are embedded) and may present *sensing* and *actuating* capabilities.

	Safety	Security	Certif.	Distrib.	HMI	Seamless	MPSoC	Energy
Automotive	x	x	x	x	x	x	x	
Aerospace	x	x	x	x	x		x	x
Healthcare	x	x	x	x	x	x	x	x
Consumer					x	x	x	

Complex functionalities.

Colliding technical requirements.

Multimedia Domain

HIGH PERFORMANCES

real time, portability, long battery life

UP-TO-DATE SOLUTIONS

last audio/video codecs, file formats...

MORE INTEGRATED FEATURES

MP3, Camera, Video, GPS...

MARKET DEMAND

convenient form factor, affordable price, fashion

Target & Technological Challenges

- **DATAFLOW MODEL OF COMPUTATION**
  - Modularity and parallelism → ***EASIER INTEGRATION AND FAVOURED RE-USABILITY***
- **COARSE-GRAINED RECONFIGURABILITY**
  - Flexibility and resource sharing → ***MULTI-APPLICATION PORTABLE DEVICES***

The RPCT project (2012-2015) has been funded by Sardinian Regional Government (L.R. 7/2007, CRP-18324).
http://sites.unica.it/rpct/

Target & Technological Challenges

- **DATAFLOW MODEL OF COMPUTATION**
  - Modularity and parallelism → ***EASIER INTEGRATION AND FAVOURED RE-USABILITY***
- **COARSE-GRAINED RECONFIGURABILITY**
  - Flexibility and resource sharing → ***MULTI-APPLICATION PORTABLE DEVICES***

Reconfigurable Platform Composer Tool Project

Automated **DESIGN FLOW** are fundamental to guarantee **SHORTER TIME-TO-MARKET**. Dealing with **APPLICATION SPECIFIC MULTI-CONTEXT** systems, in particular for **KERNEL ACCELERATORS**, state of the art still lacks in providing a broadly accepted solution.

The RPCT project (2012-2015) has been funded by Sardinian Regional Government (L.R. 7/2007, CRP-18324).
http://sites.unica.it/rpct/

Reasons for Coarser-Grain

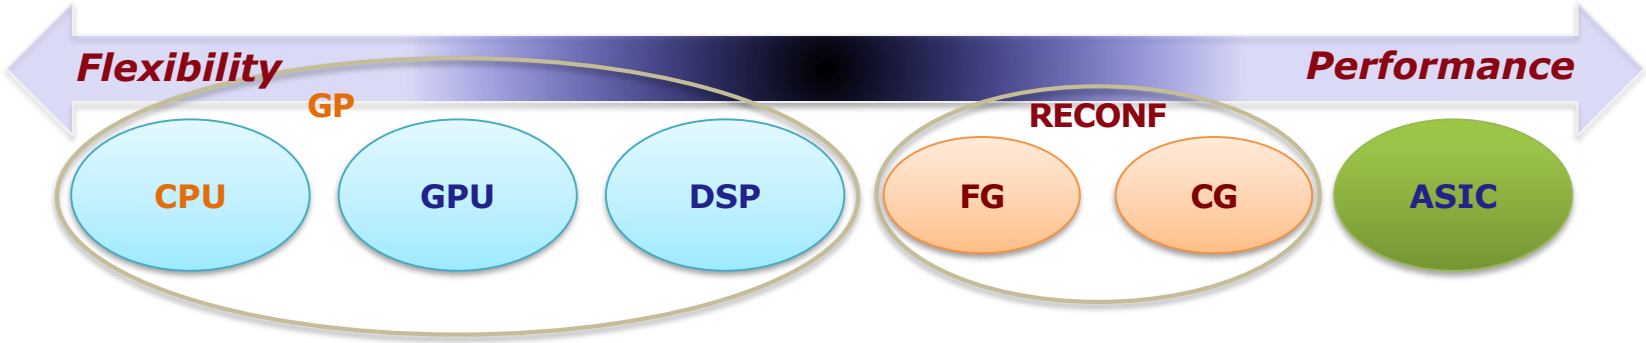

Reasons for Coarser-Grain

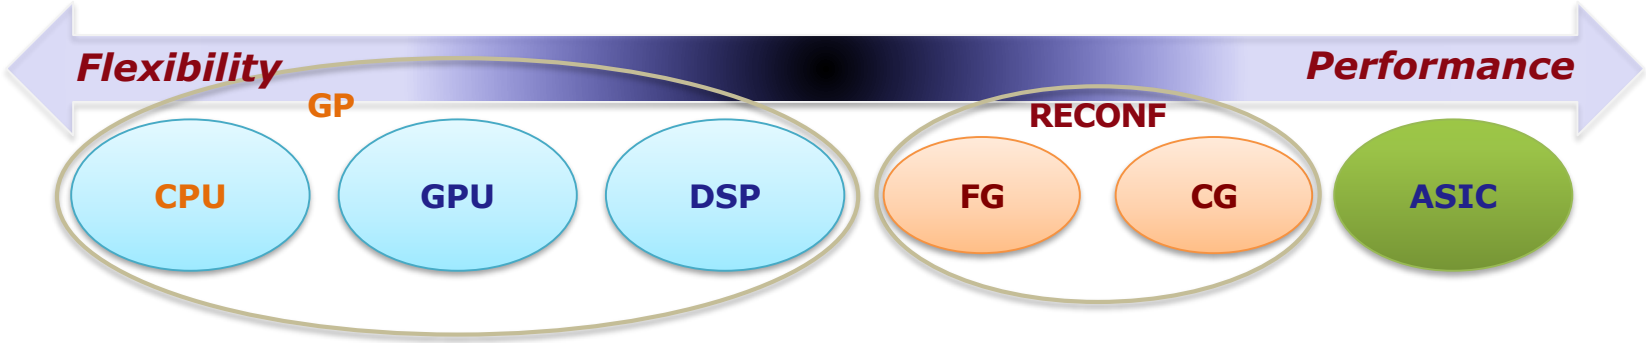

	Fine Grained	Coarse Grained
	bit-level	word-level
Flexibility	☺	😐
Speed	😐	☺
Memory	☹	😐

- **Coarse Grained (CG):**
  - both in ASIC and FPGA
  - 1 clock cycle switching, with dedicated switching blocks.
- **Fine Grained (FG):**
  - FPGA only
  - switching requires a new bit-stream

Framework Development

2010 | 2011 | 2012 | 2013 | 2014 | 2015 | 2016

Baseline tool specification:
Multi-Dataflow Composer (MDC) tool


MPEG-RVC Framework Integration:
Orcc + MDC + Xronos + Turnus



Framework Development

2010 | 2011 | 2012 | 2013 | 2014 | 2015 | 2016

Baseline tool specification:
Multi-Dataflow Composer (MDC) tool


MPEG-RVC Framework Integration:
Orcc + MDC + Xronos + Turnus


MDC:
Structural Profiler


MDC:
Low-Power Extension


MDC:
Co-processor Generator

Framework Evaluation

2010 2011 2012 2013 2014 2015 2016

Reconfigurable Image/Video Coding: JPEG e H.264

Adaptive Filtering: HEVC Encoding

Framework Evaluation

2010 2011 2012 2013 2014 2015 2016

Reconfigurable Image/Video Coding: JPEG e H.264

Neural Signal Decoding

Adaptive Filtering: HEVC Encoding

Cryptographic Systems

Outline

- The origins of our dataflow to hardware studies: the RPCT Project
  - Context
  - Target Technologies
  - Project Development
- The MDC tool
  - Approach
  - Baseline Functionality and Extensions
- Contexts of application
  - Neural Signal Decoding
  - HEVC Interpolation Filters
- Final Remarks

Design Suite & Targeted Challenges

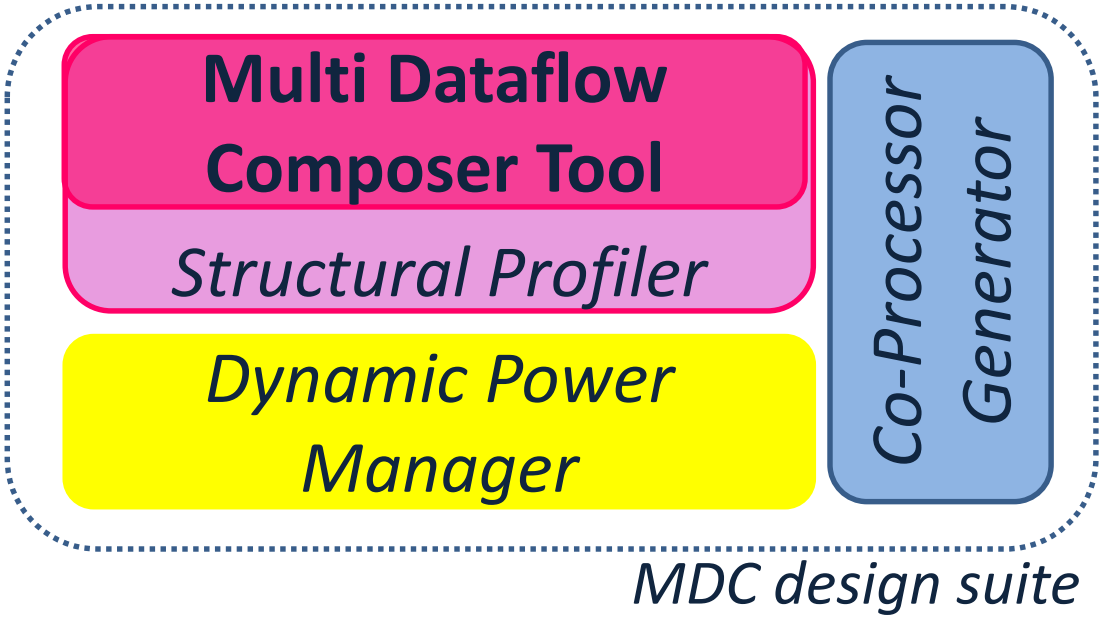

http://sites.unica.it/rpct/

Design Suite & Targeted Challenges

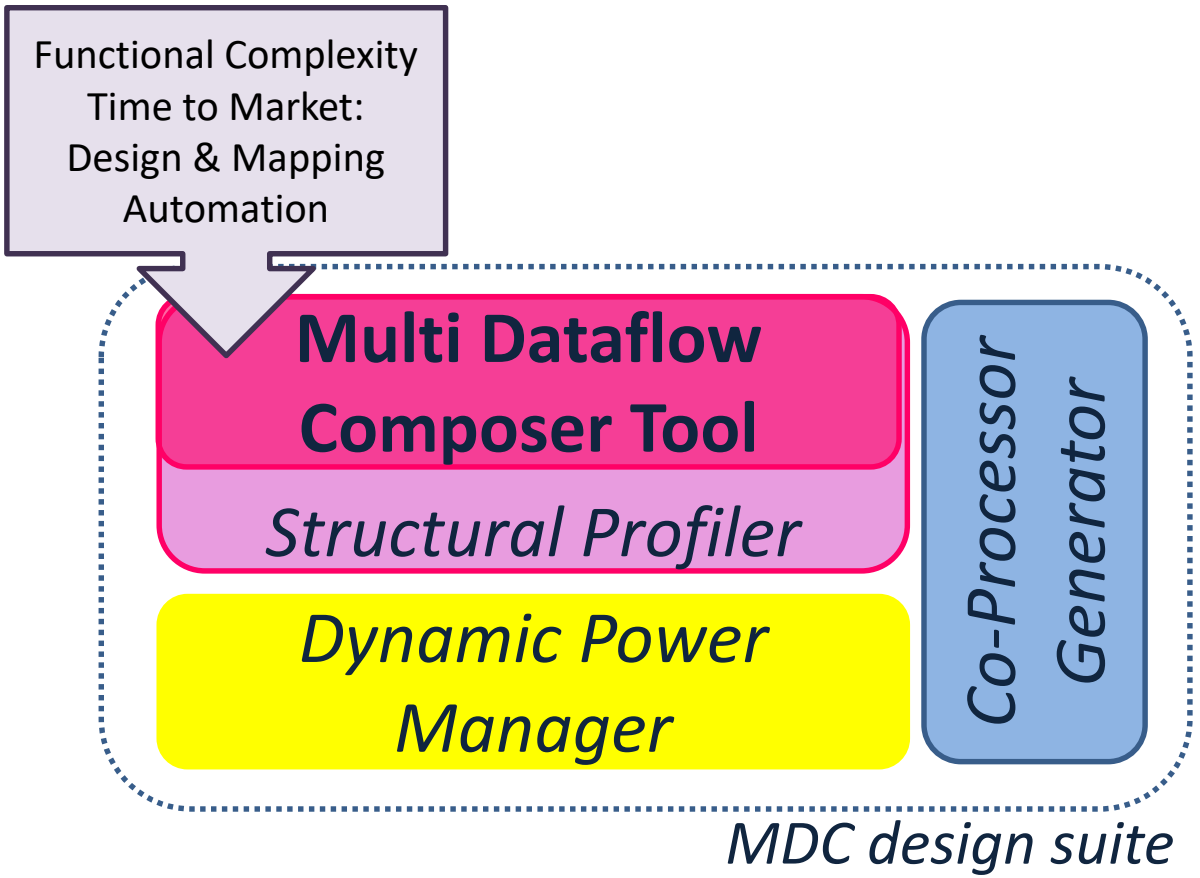

http://sites.unica.it/rpct/

Design Suite & Targeted Challenges

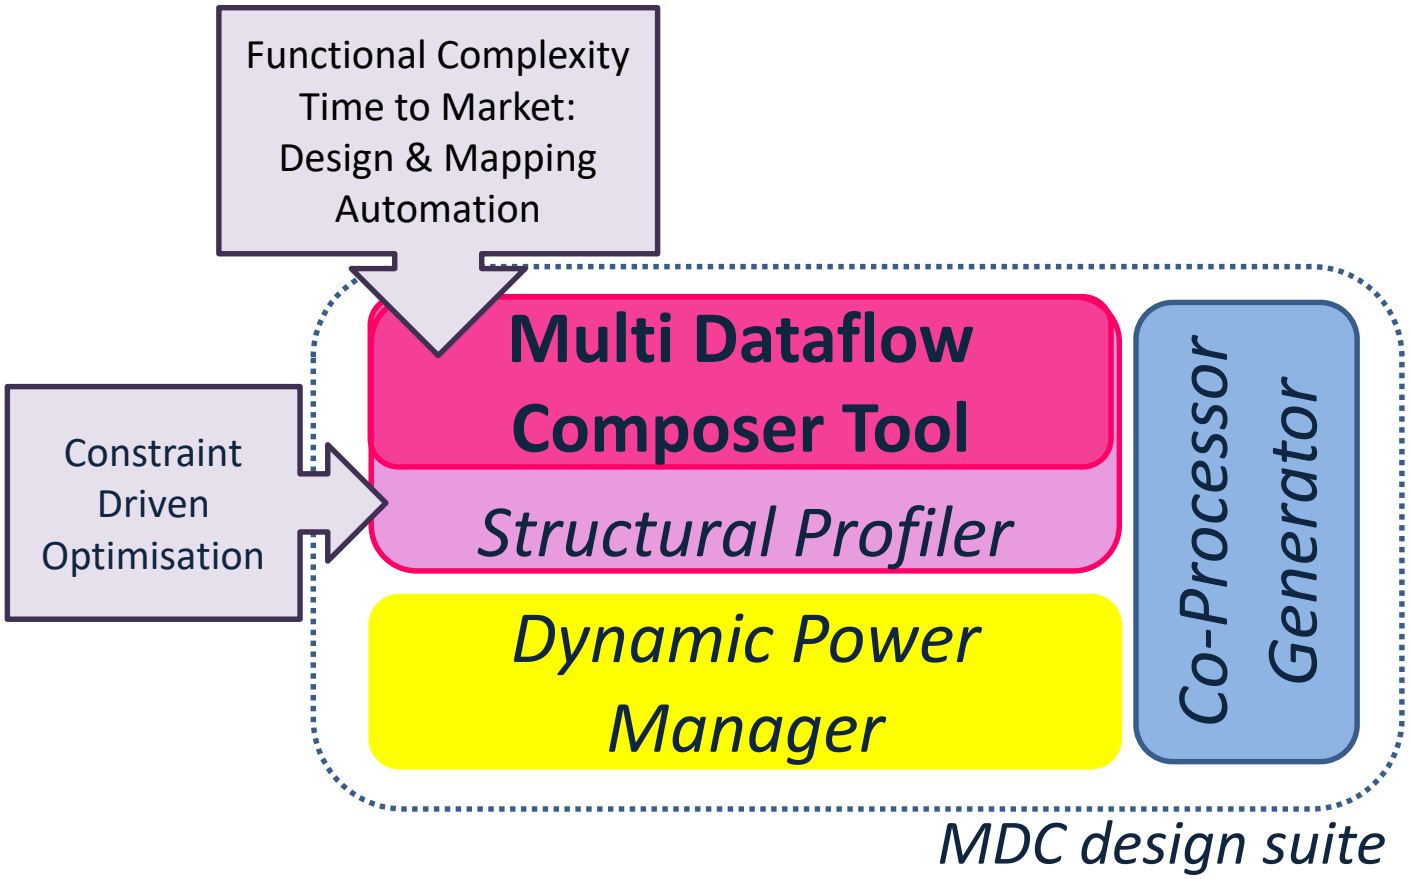

http://sites.unica.it/rpct/

Design Suite & Targeted Challenges

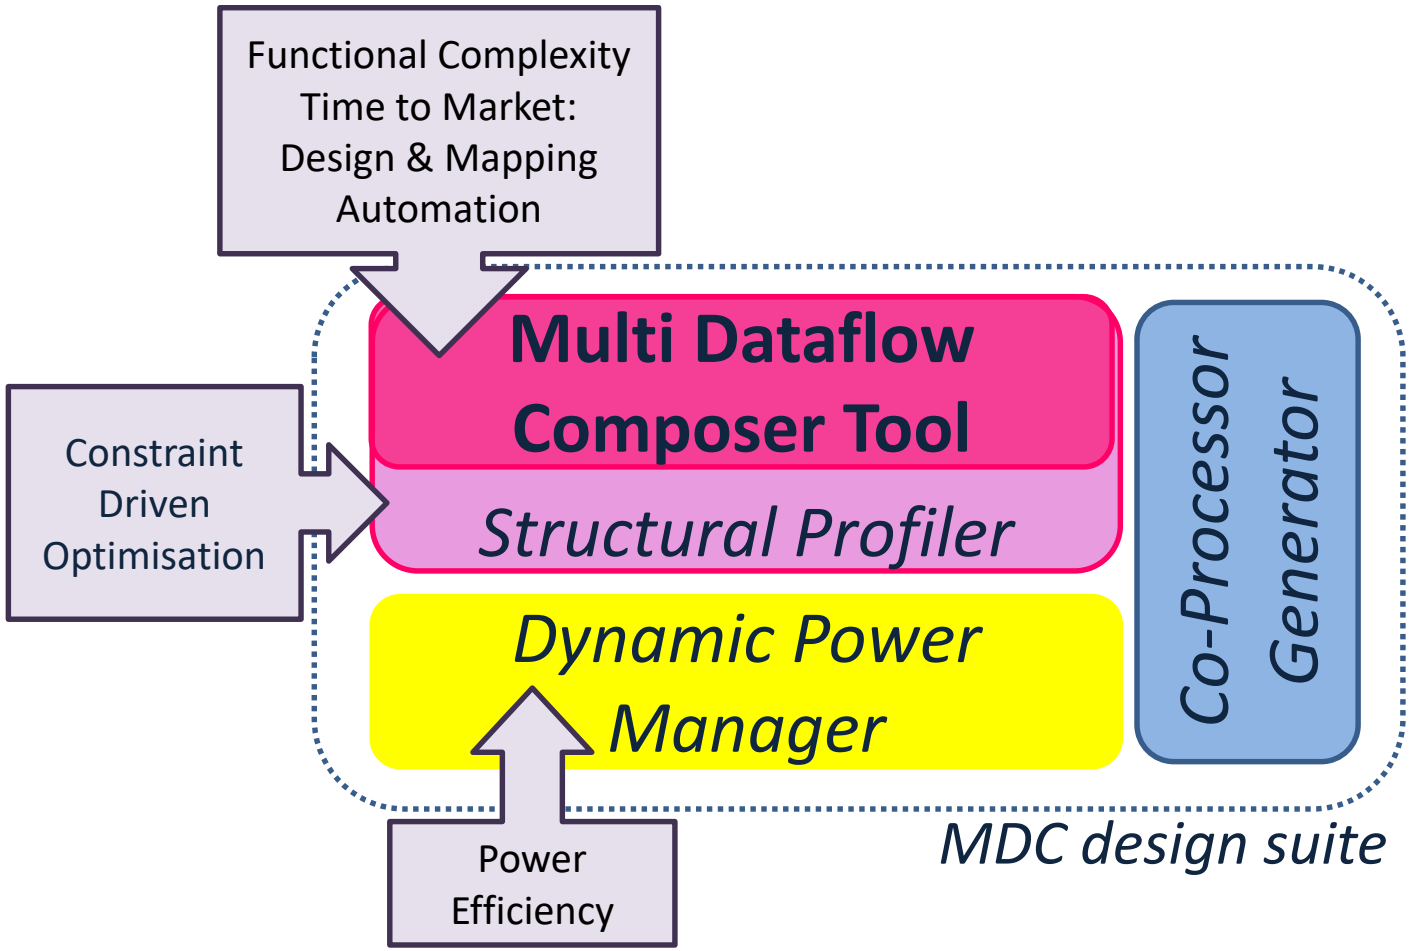

http://sites.unica.it/rpct/

Design Suite & Targeted Challenges

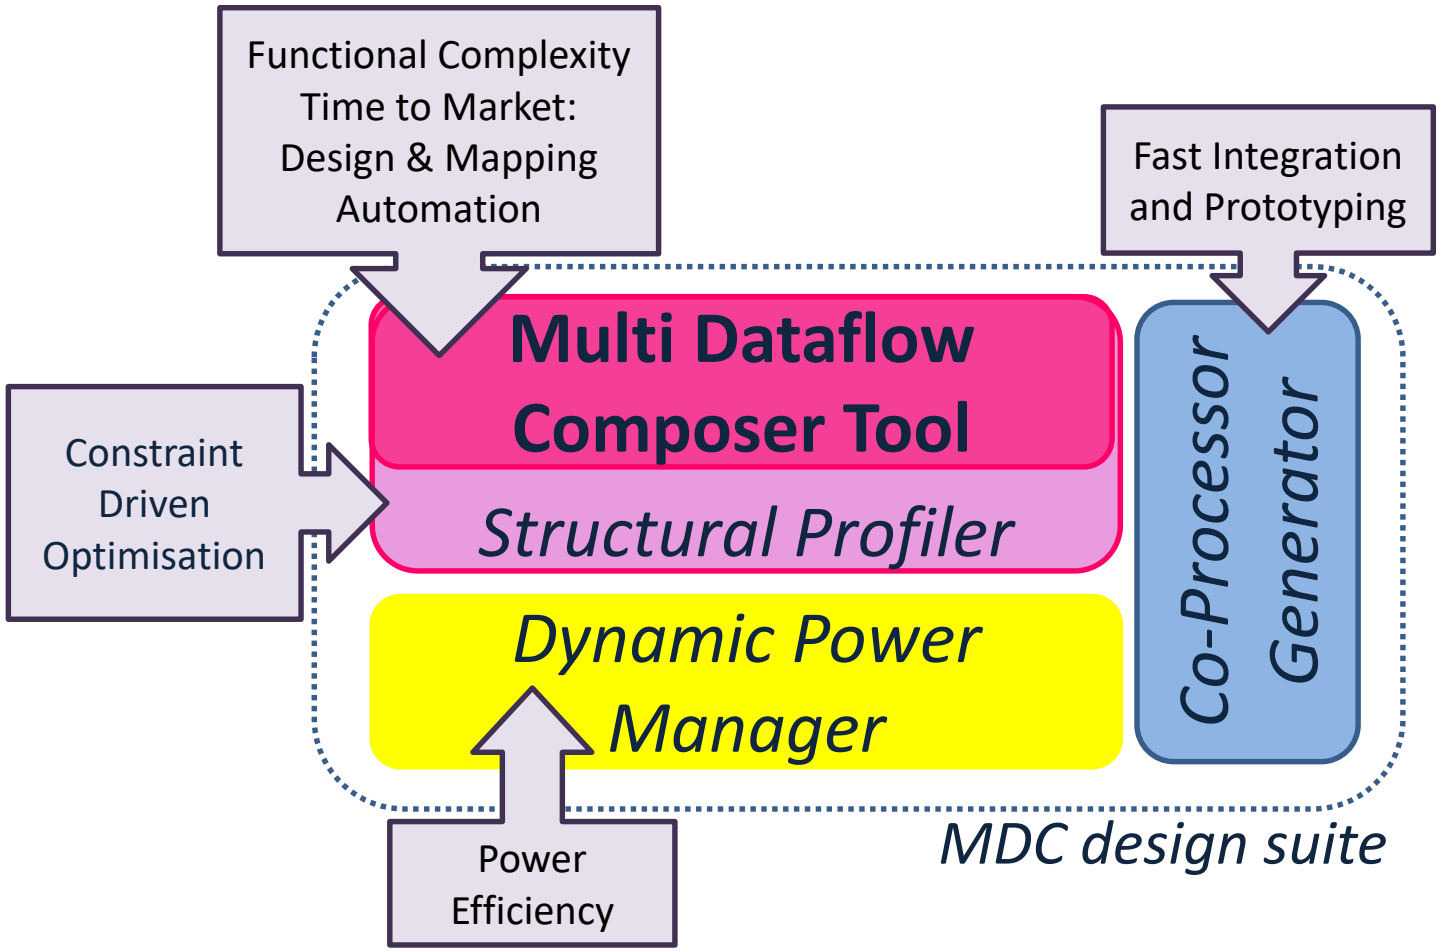

http://sites.unica.it/rpct/

Baseline: Dataflow to HW

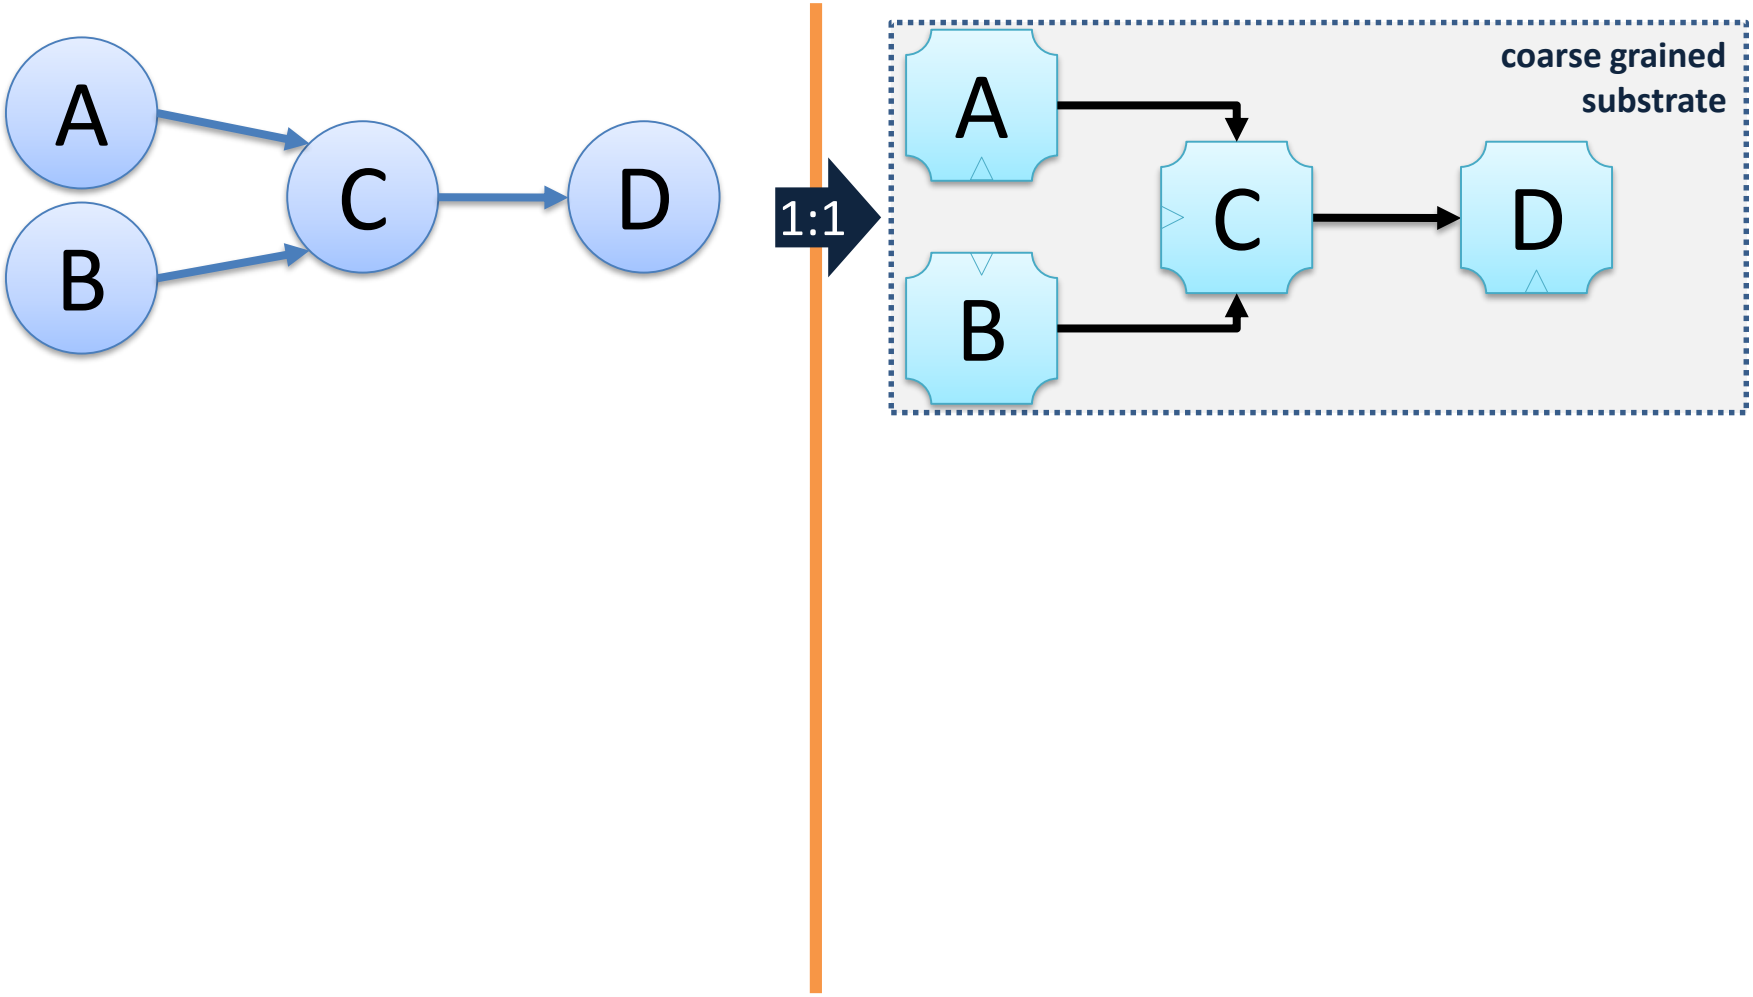

Baseline: Dataflow to HW

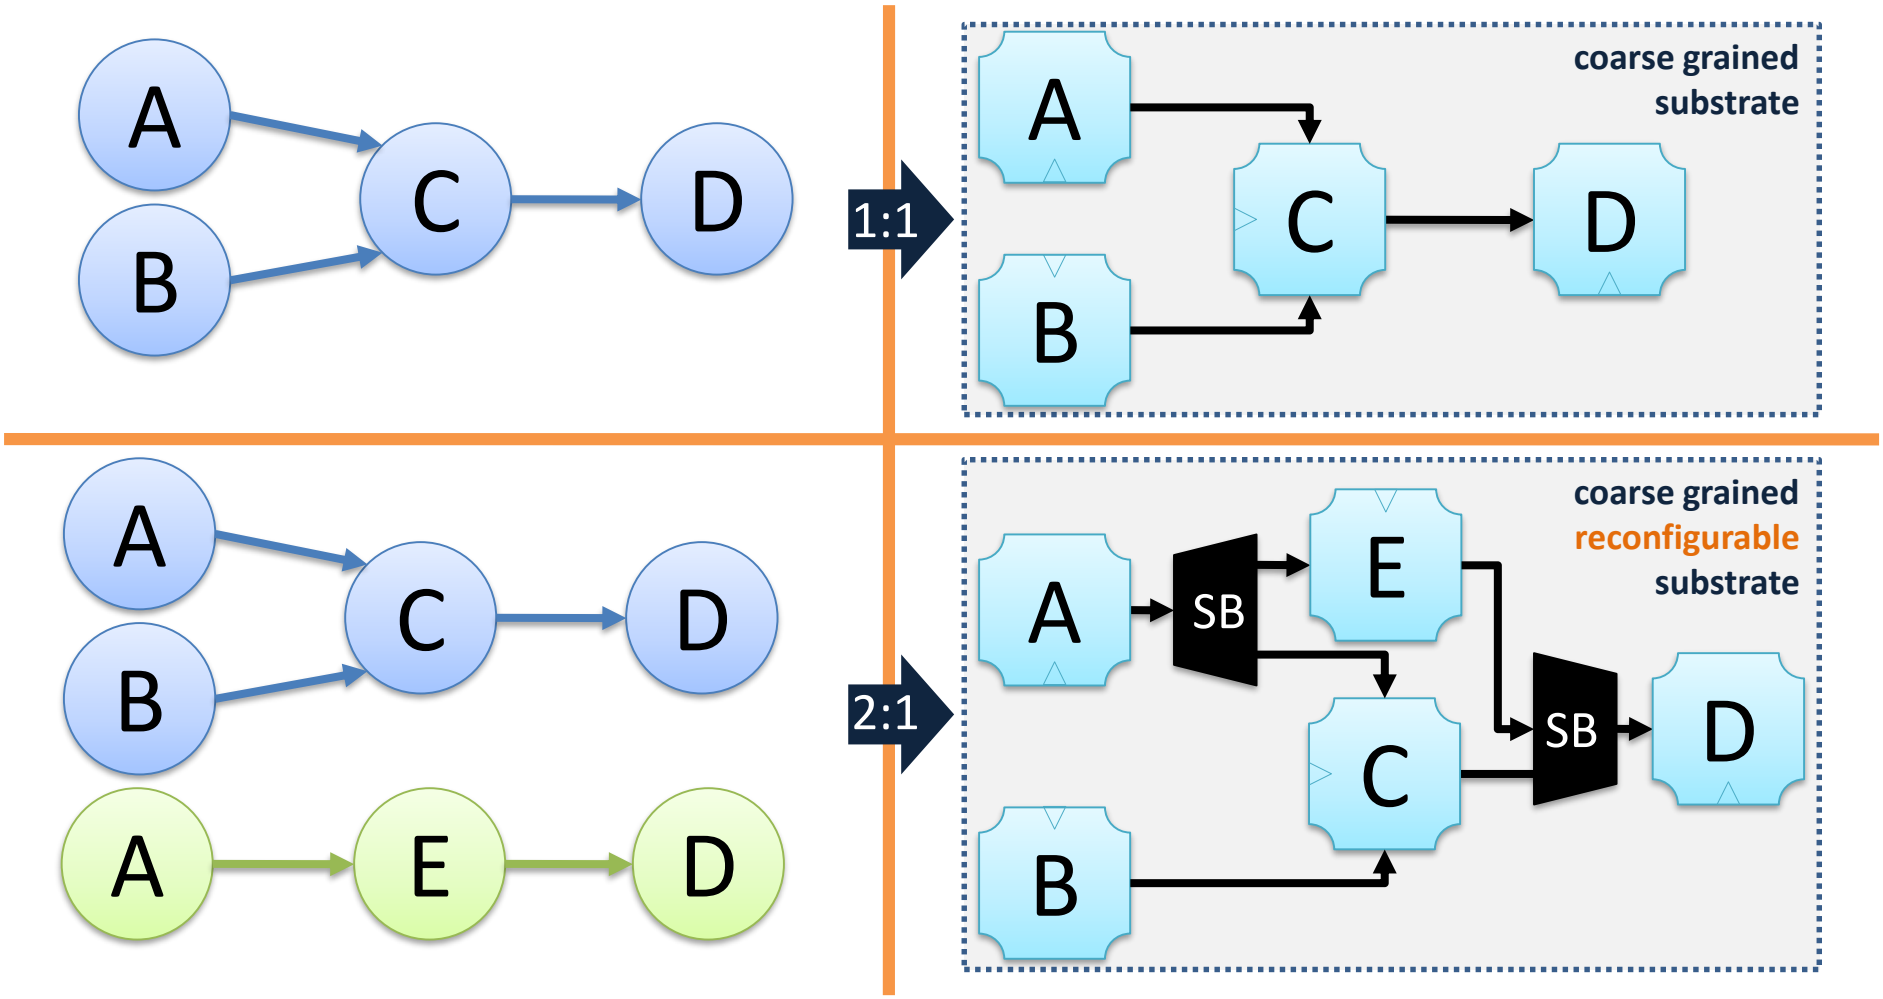

MDC Front-End:
Multi-Dataflow Generator

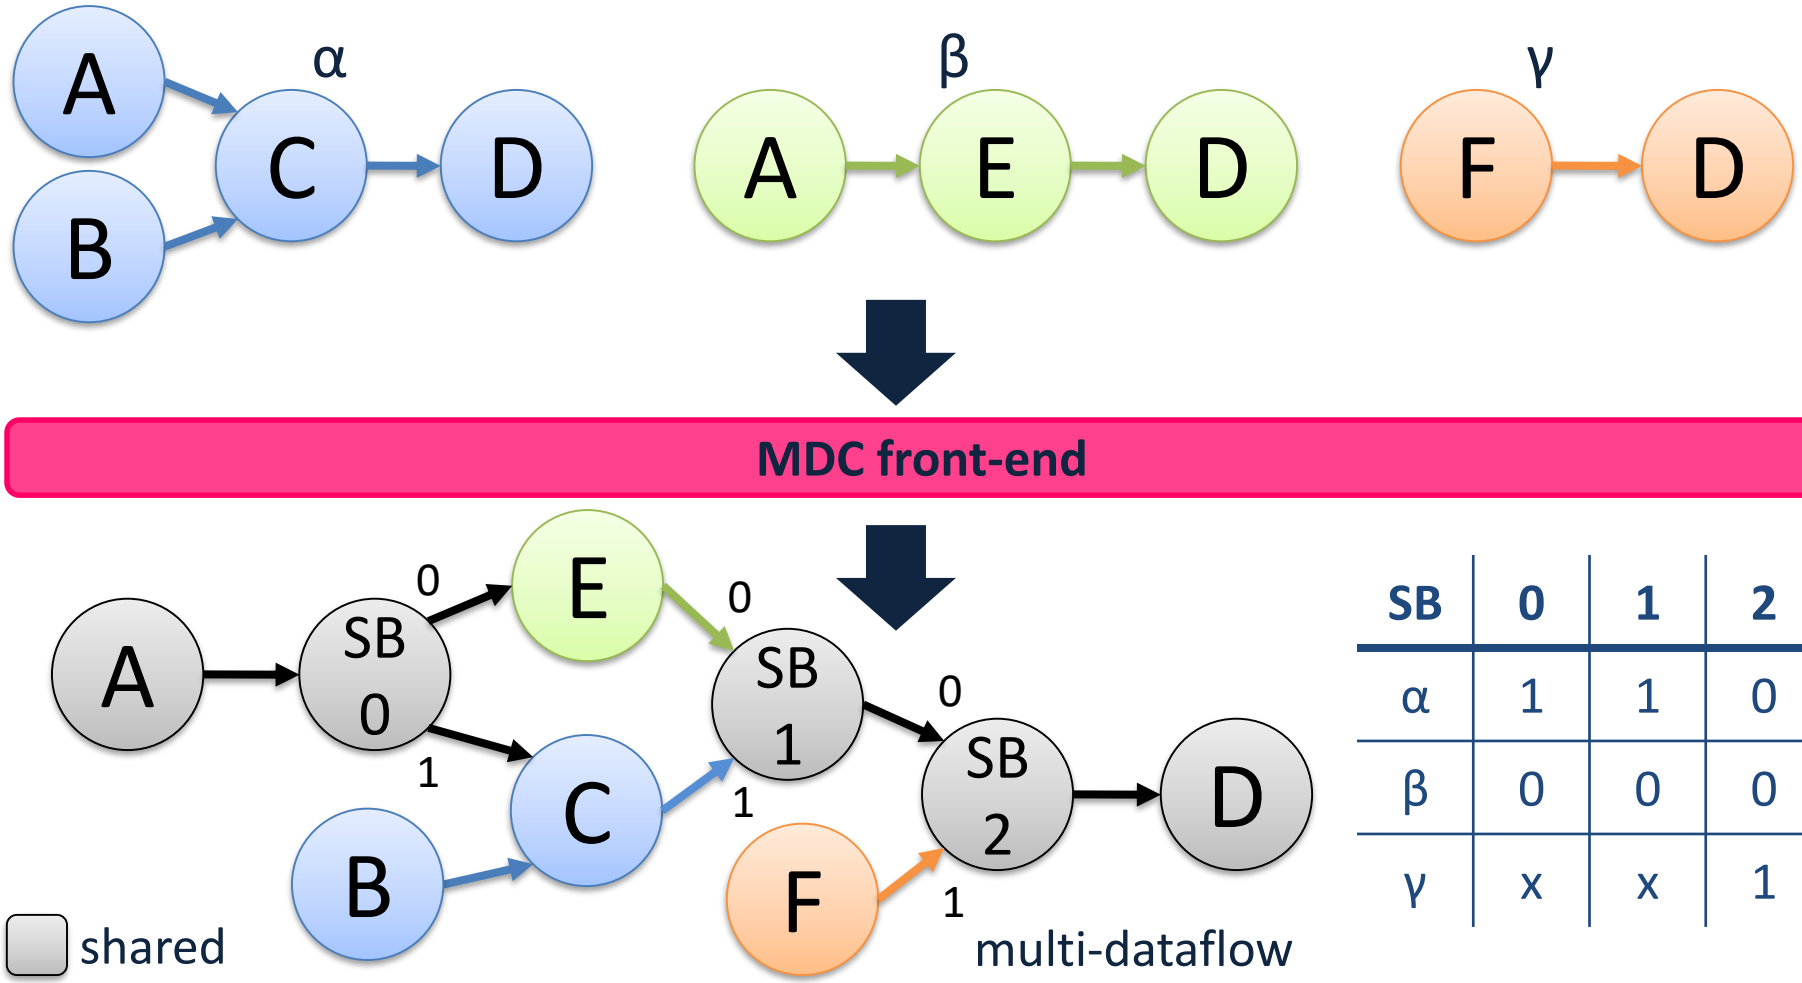

SB	0	1	2
α	1	1	0
β	0	0	0
γ	x	x	1

Datapath Merging Problem:
Graph Model

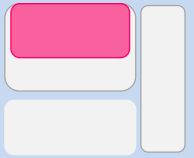

GRAPHS

$G_i = (V_i, E_i)$

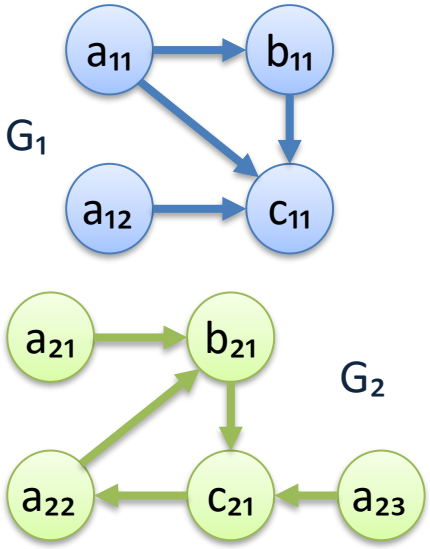

Datapath Merging Problem:
Graph Model

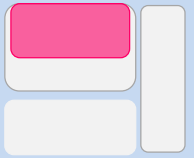

GRAPHS

$G_i = (V_i, E_i)$

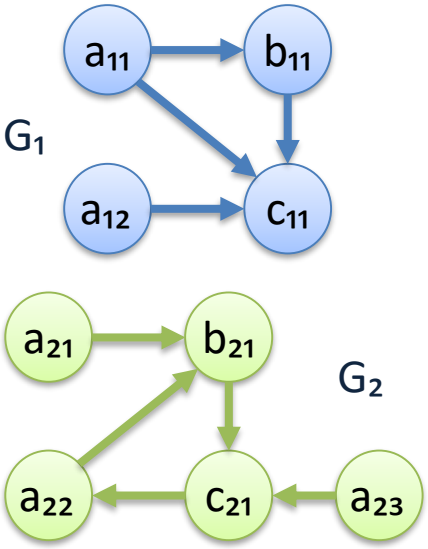

LABELING

$\pi_i : V_i \rightarrow T$

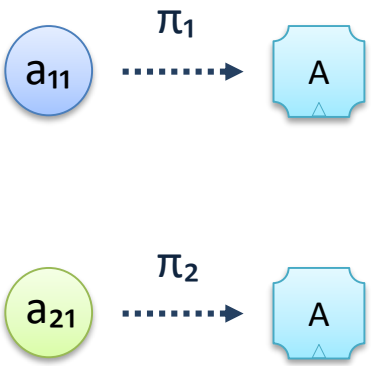

Datapath Merging Problem:
Graph Model

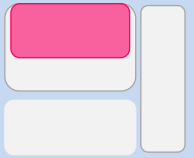

GRAPHS

$G_i = (V_i, E_i)$

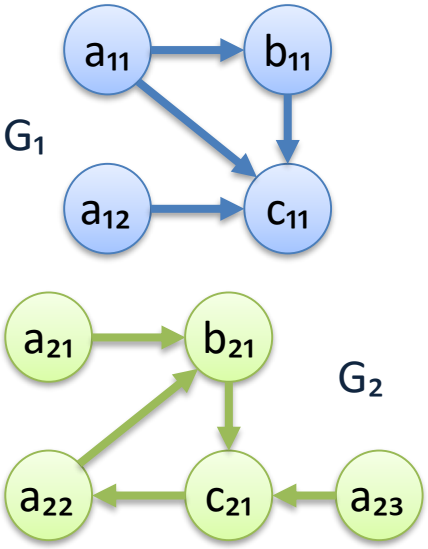

LABELING

$\pi_i : V_i \rightarrow T$

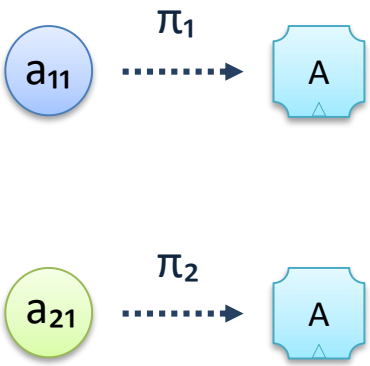

MAPPING

$\mu_i(v) = u,$
$(v \in V_i, u \in V)$
$\downarrow$
$\pi_i(v) = \pi(u)$

$e(v_i, v_i') \in E_i$
$\downarrow$
$e(\mu_i(v_i), \mu_i(v_i')) \in E$

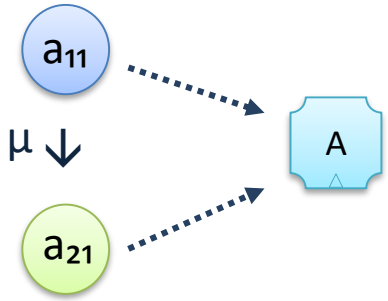

Datapath Merging Problem:
Graph Model

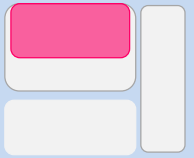

GRAPHS

$G_i = (V_i, E_i)$

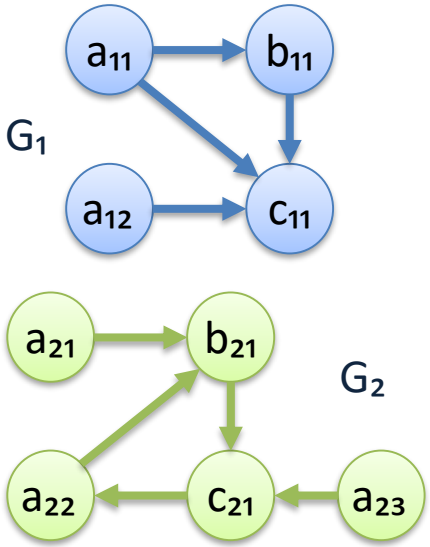

LABELING

$\pi_i : V_i \rightarrow T$

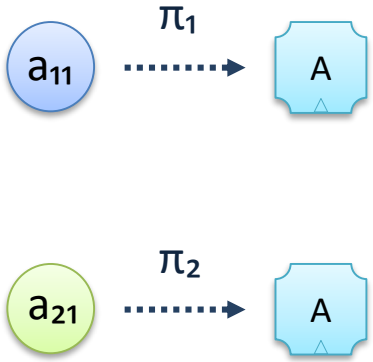

MAPPING

$\mu_i(v) = u,$
$(v \in V_i, u \in V)$
$\downarrow$
$\pi_i(v) = \pi(u)$

$e(v_i, v_i') \in E_i$
$\downarrow$
$e(\mu_i(v_i), \mu_i(v_i')) \in E$

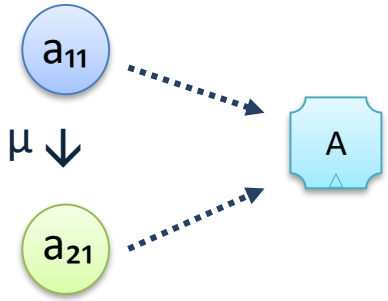

PROBLEM STATEMENT: *find a* ***Reconfigurable Graph*** ***G*** *(V,E) with the minimum costs (****min\|V\|*** *and* ***min \|E\|****)*

Datapath Merging Problem:
Graph Model

GRAPHS

$G_i = (V_i, E_i)$

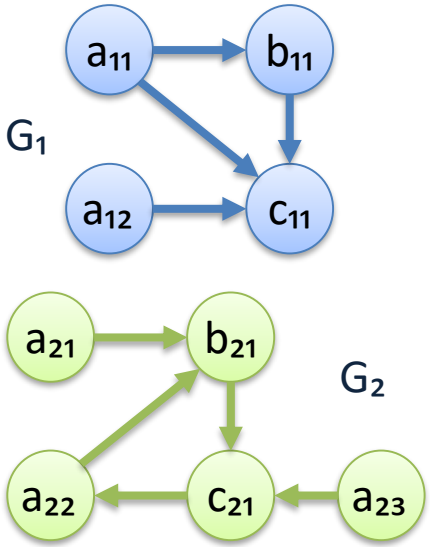

LABELING

$\pi_i : V_i \rightarrow \mathbf{T}$

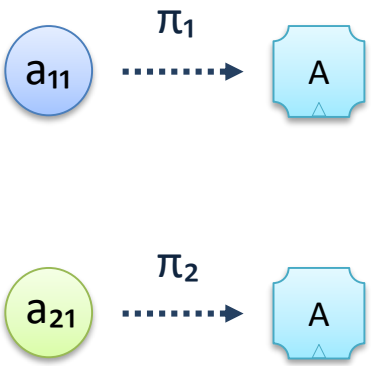

MAPPING

$\mu_i(v) = u,$
$(v \in V_i, u \in V)$
$\downarrow$
$\pi_i(v) = \pi(u)$

$e(v_i, v_i') \in E_i$
$\downarrow$
$e(\mu_i(v_i), \mu_i(v_i')) \in E$

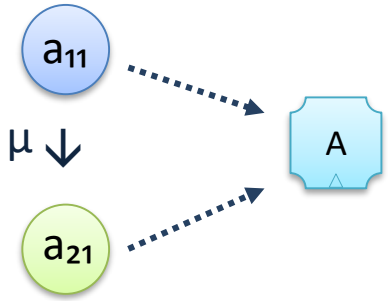

PROBLEM STATEMENT: *find a* ***Reconfigurable Graph*** ***G*** *(V,E) with the minimum costs (**min|V|** and **min |E|**)*

<u>feasible solution:</u>

$\forall T \in \mathbf{T}, V^T=\{v : \pi(v) = T\} \quad \rightarrow \quad |V^T| = \max |V_i^T|, V_i^T=\{v_i : \pi_i(v_i) = T\}$

Datapath Merging Problem:
Graph Model

GRAPHS

$G_i = (V_i, E_i)$

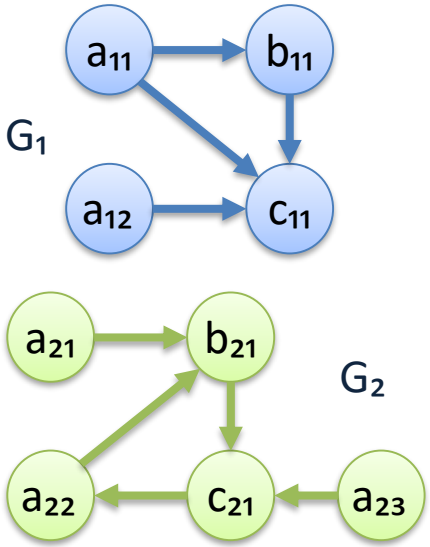

LABELING

$\pi_i : V_i \rightarrow \mathbf{T}$

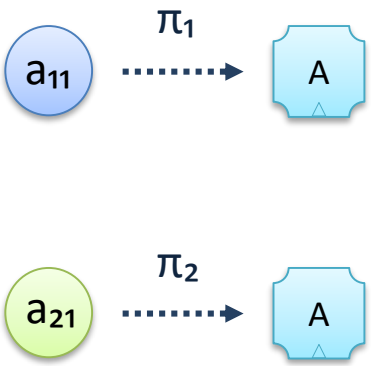

MAPPING

$\mu_i(v) = u,$
$(v \in V_i, u \in V)$
$\downarrow$
$\pi_i(v) = \pi(u)$

$e(v_i, v_i') \in E_i$
$\downarrow$
$e(\mu_i(v_i), \mu_i(v_i')) \in E$

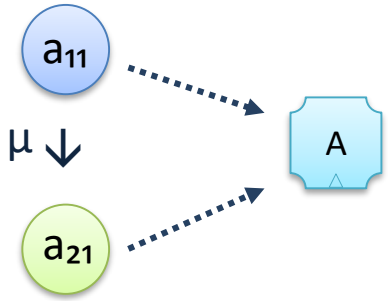

PROBLEM STATEMENT: *find a* ***Reconfigurable Graph*** ***G*** *(V,E) with the minimum costs (**min|V|** and **min |E|**)*

<u>feasible solution:</u>

$\forall T \in \mathbf{T}, V^T=\{v : \pi(v) = T\} \rightarrow |V^T| = \max |V_i^T|, V_i^T=\{v_i : \pi_i(v_i) = T\}$

<u>optimal solution:</u>

feasible solution with min|E|

Datapath Merging Problem:
Graph Model

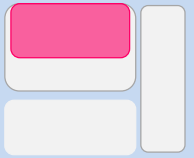

GRAPHS

$G_i = (V_i, E_i)$

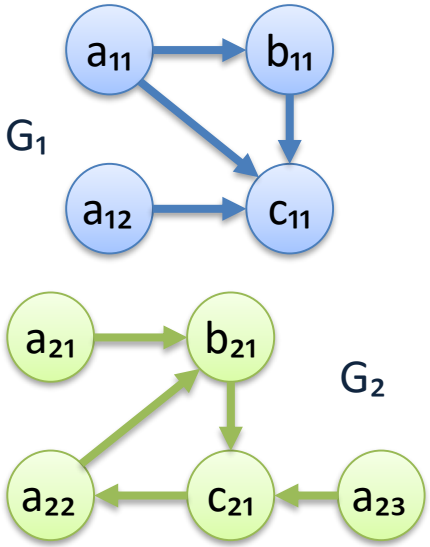

LABELING

$\pi_i : V_i \rightarrow T$

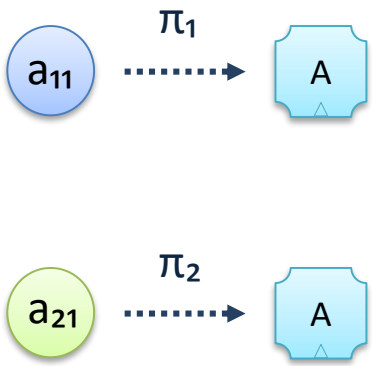

MAPPING

$\mu_i(v) = u,$
$(v \in V_i, u \in V)$
$\downarrow$
$\pi_i(v) = \pi(u)$

$e(v_i, v_i') \in E_i$
$\downarrow$
$e(\mu_i(v_i), \mu_i(v_i')) \in E$

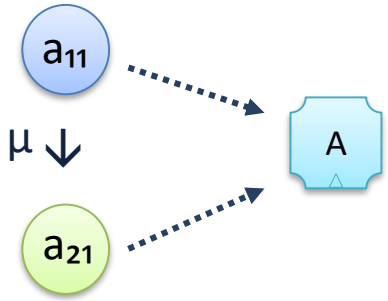

PROBLEM STATEMENT: *find a* ***Reconfigurable Graph*** ***G*** *(V,E) with the minimum costs (****min\|V\|*** *and* ***min \|E\|****)*

NP-complete problem: N. Moreano, et al., "*Datapath merging and interconnection sharing for reconfigurable architectures*", Symp. On *System Synthesis, 2002.*

MDC Back-End:

Platform Composer

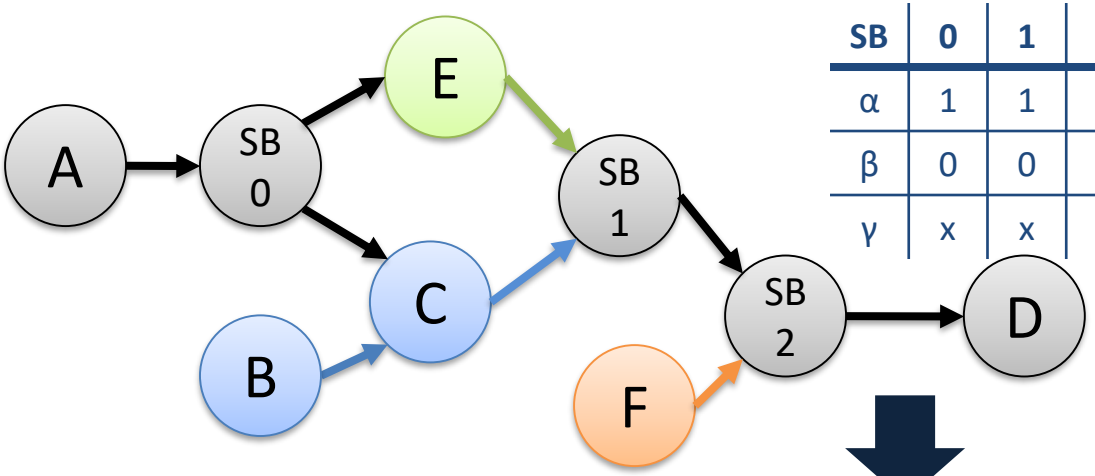

SB	0	1	2
α	1	1	0
β	0	0	0
γ	x	x	1

HDL components library

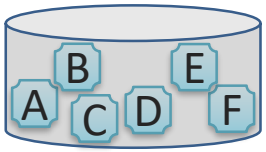

hardware communication protocol

MDC back-end

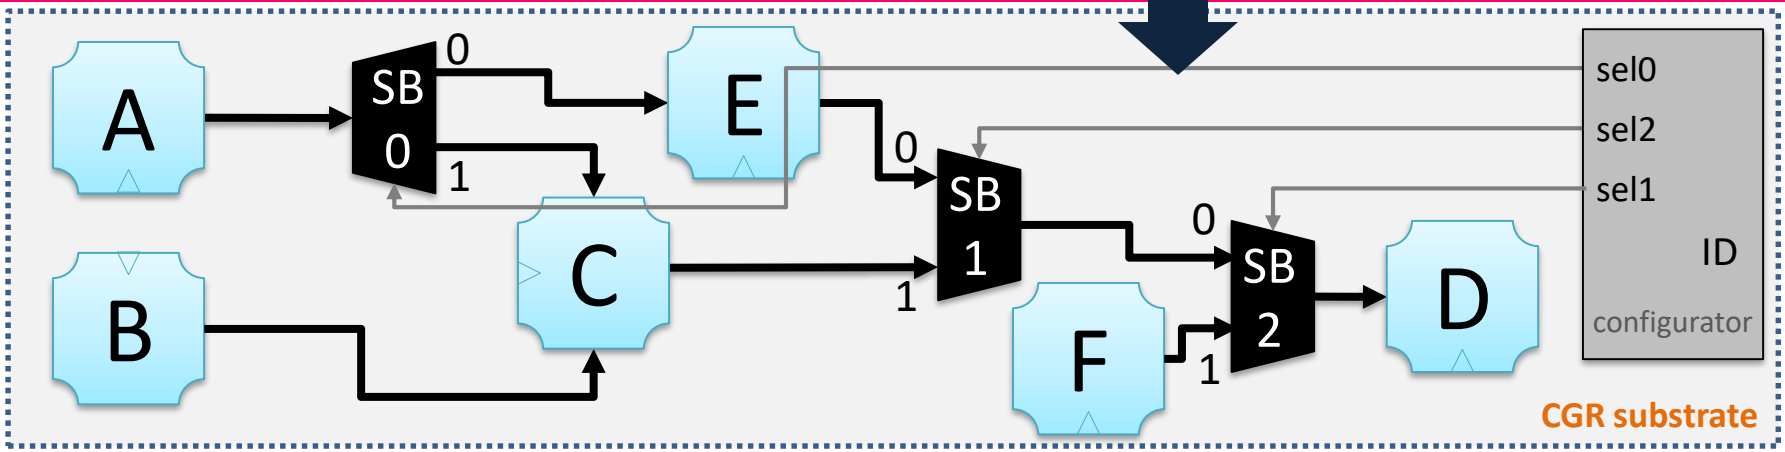

Integration within MPEG-RVC

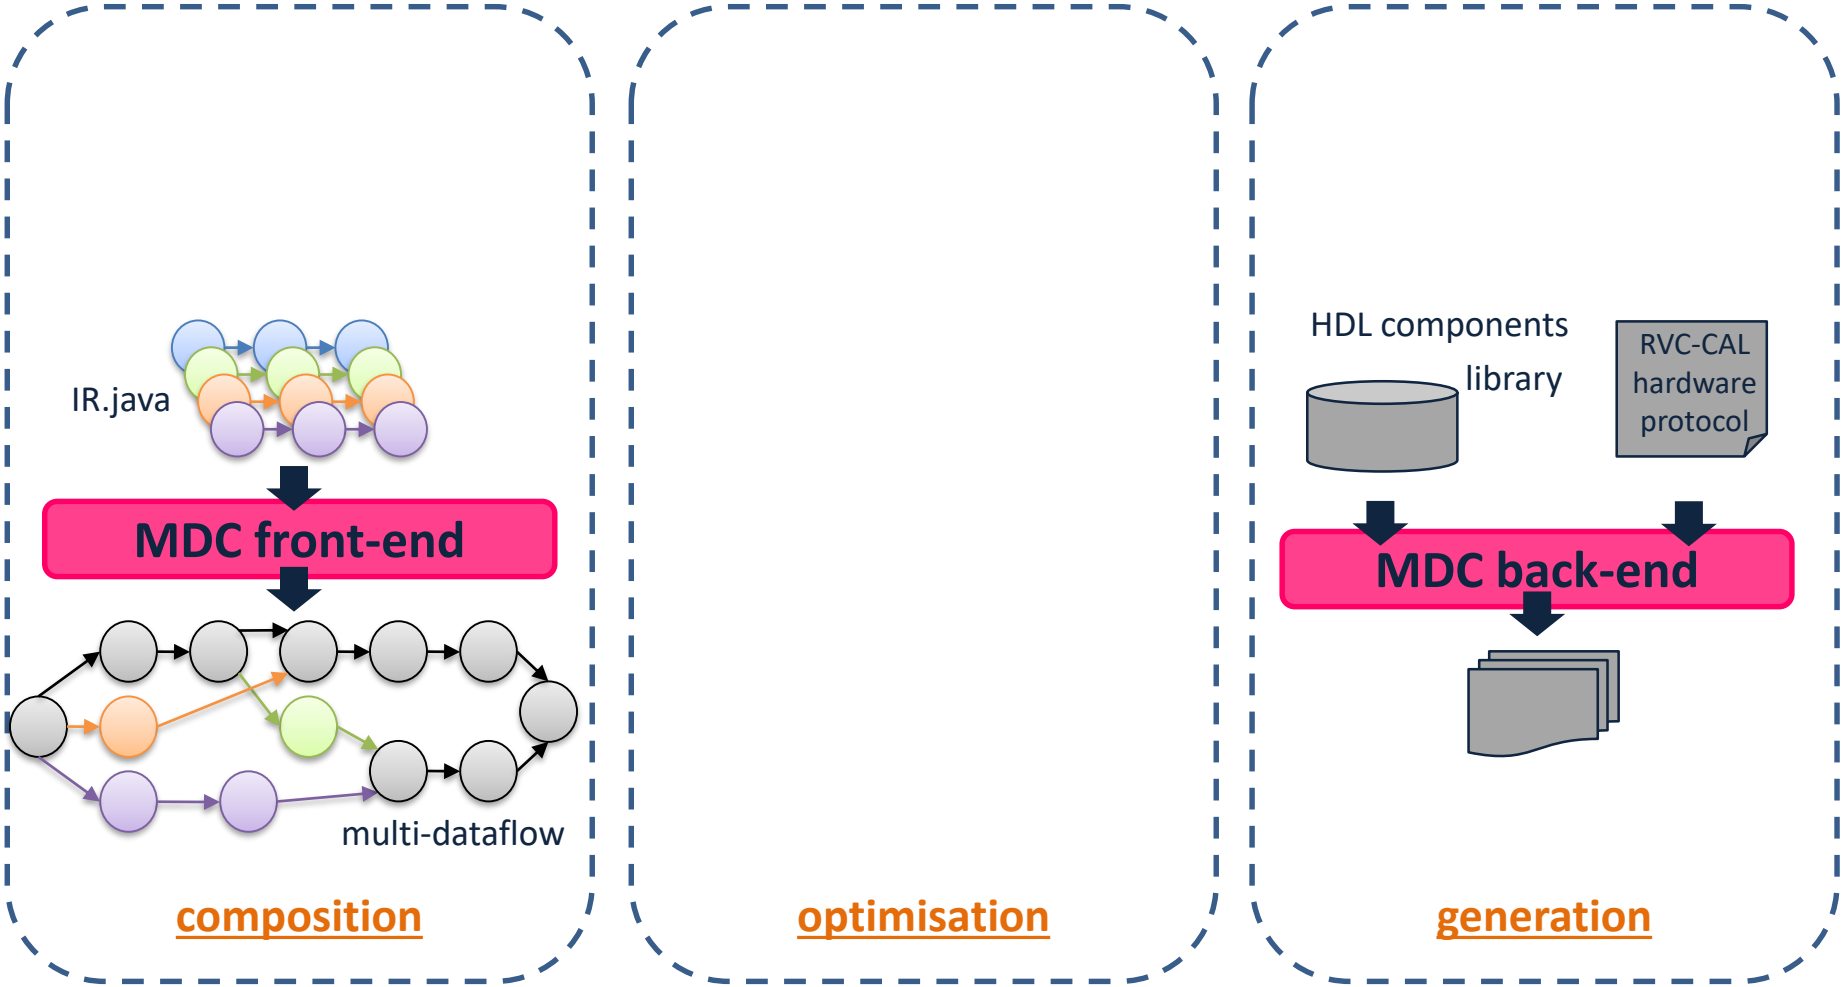

Integration within MPEG-RVC

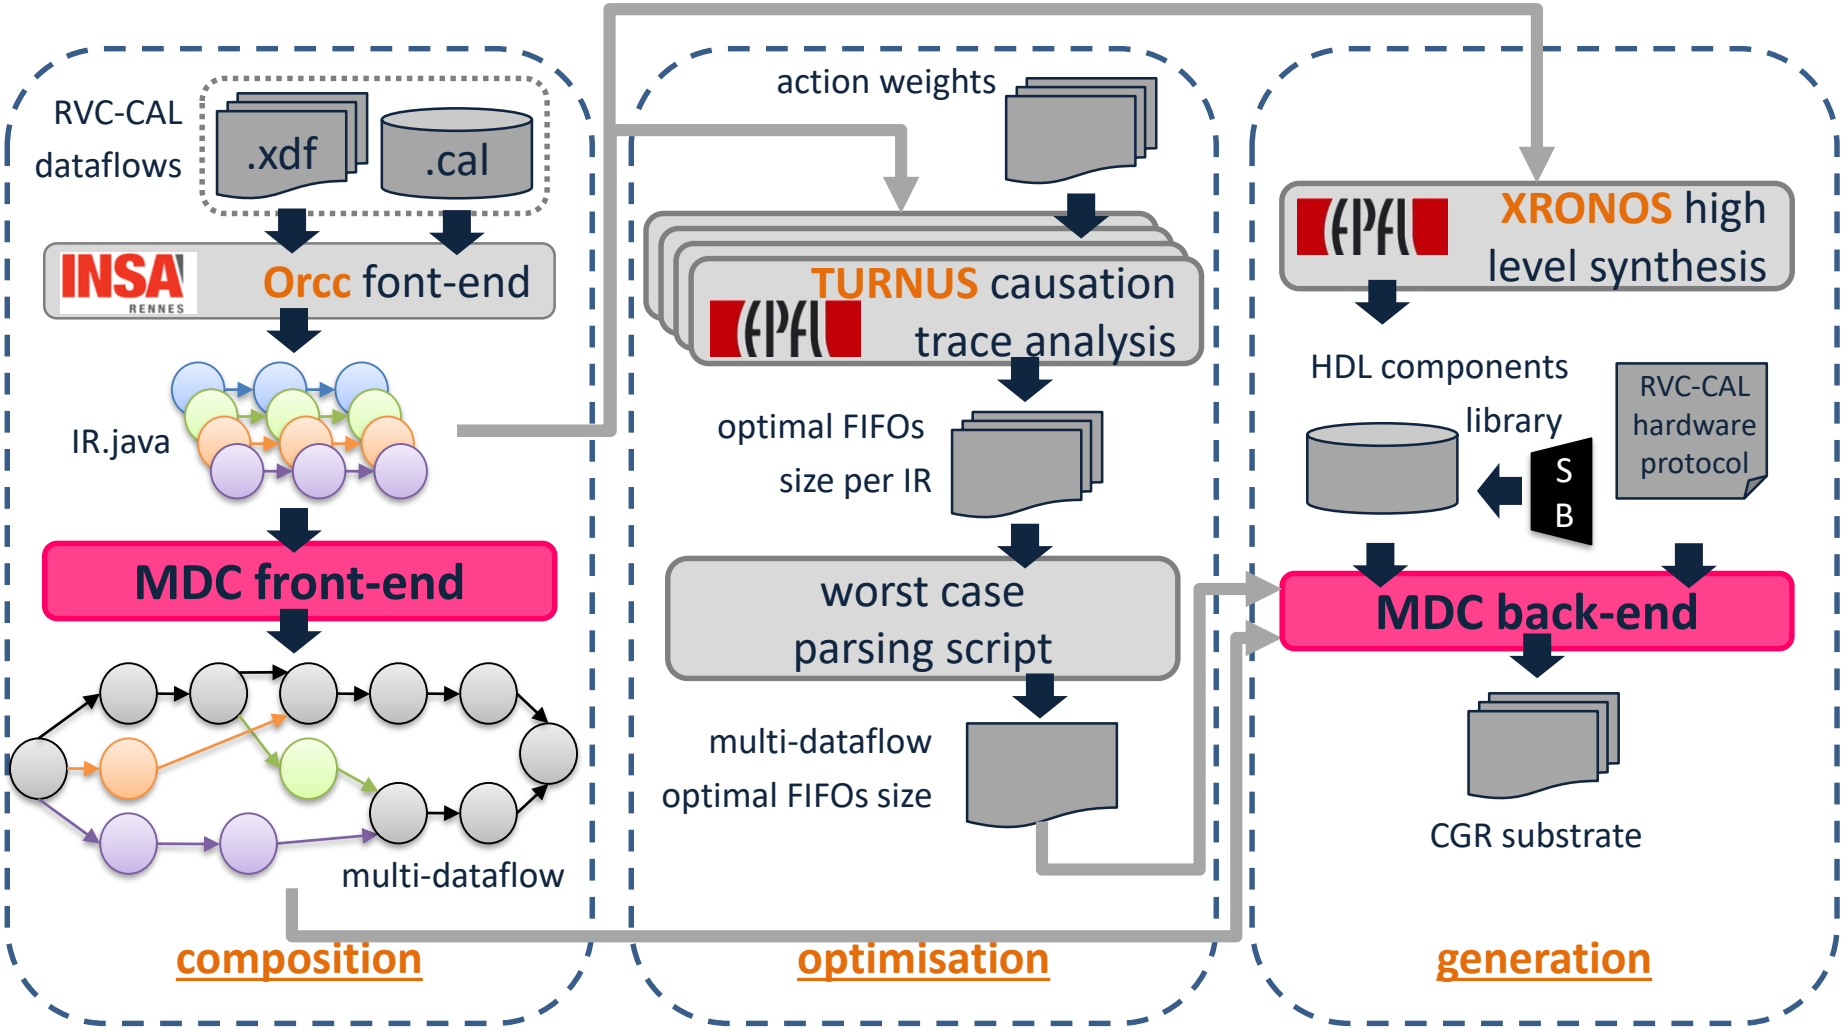

Structural Profiler

What are the topological characteristics impacting on the CGR substrate?

1. Number of merged dataflow specifications

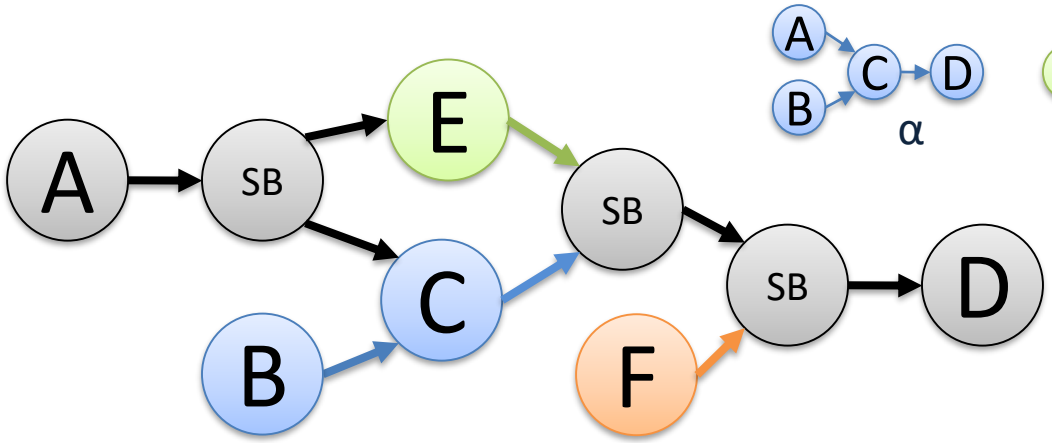

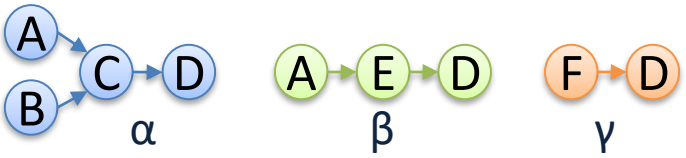

α+β+γ

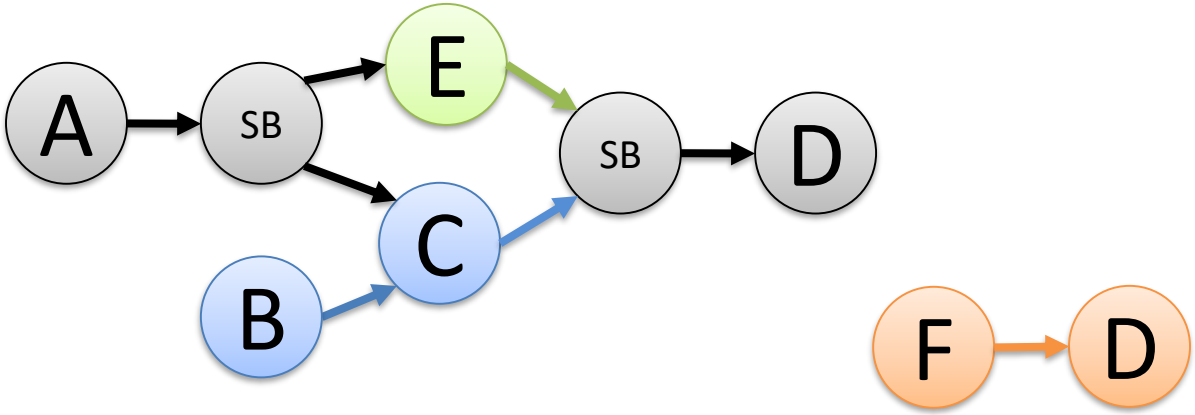

α+β|γ

Structural Profiler

What are the topological characteristics impacting on the CGR substrate?

1. Number of merged dataflow specifications

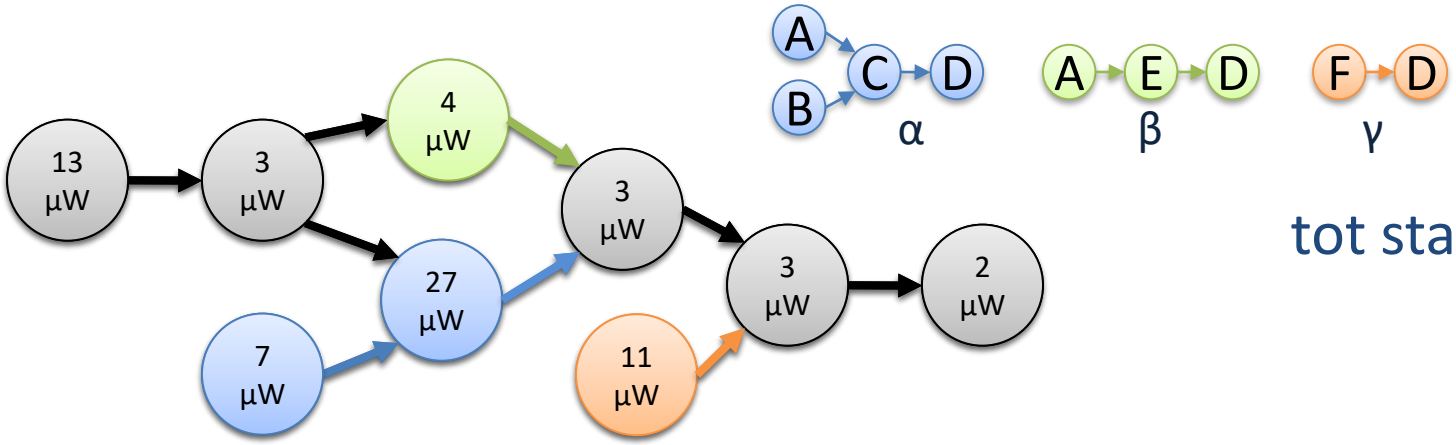

α+β+γ

tot static power

73 µW

☹

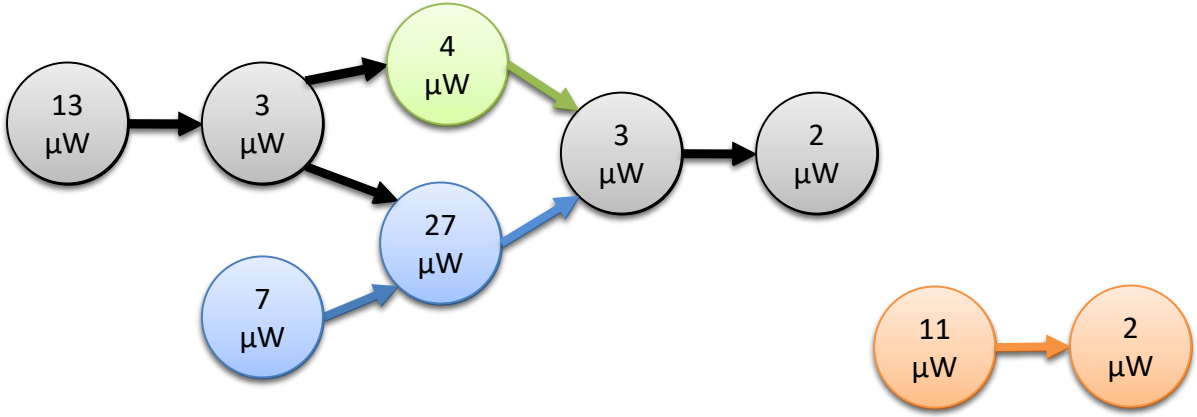

α+β|γ

tot static power

72 µW

☺

Structural Profiler

What are the topological characteristics impacting on the CGR substrate?

2. Merging order

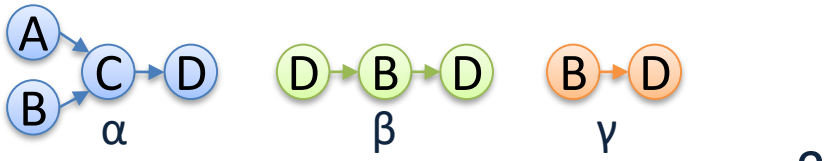

α+γ+β

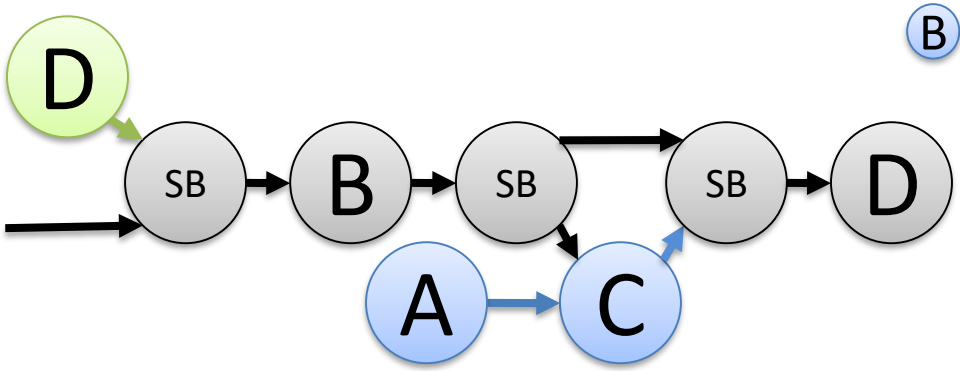

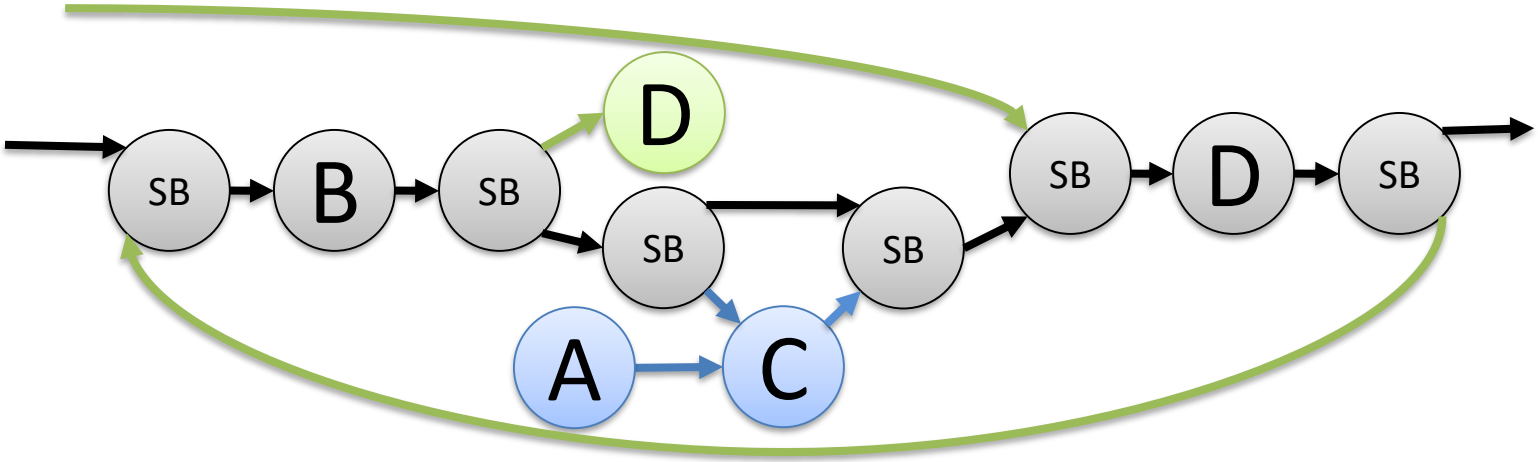

β+α+γ

Structural Profiler

What are the topological characteristics impacting on the CGR substrate?

2. Merging order

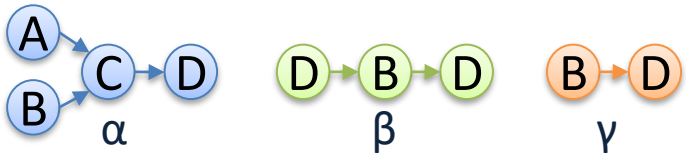

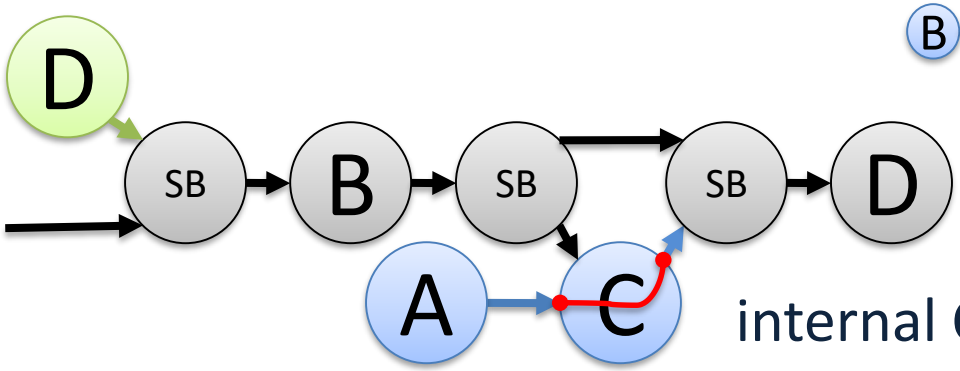

internal CP

α+γ+β
frequency
45 MHz
☺

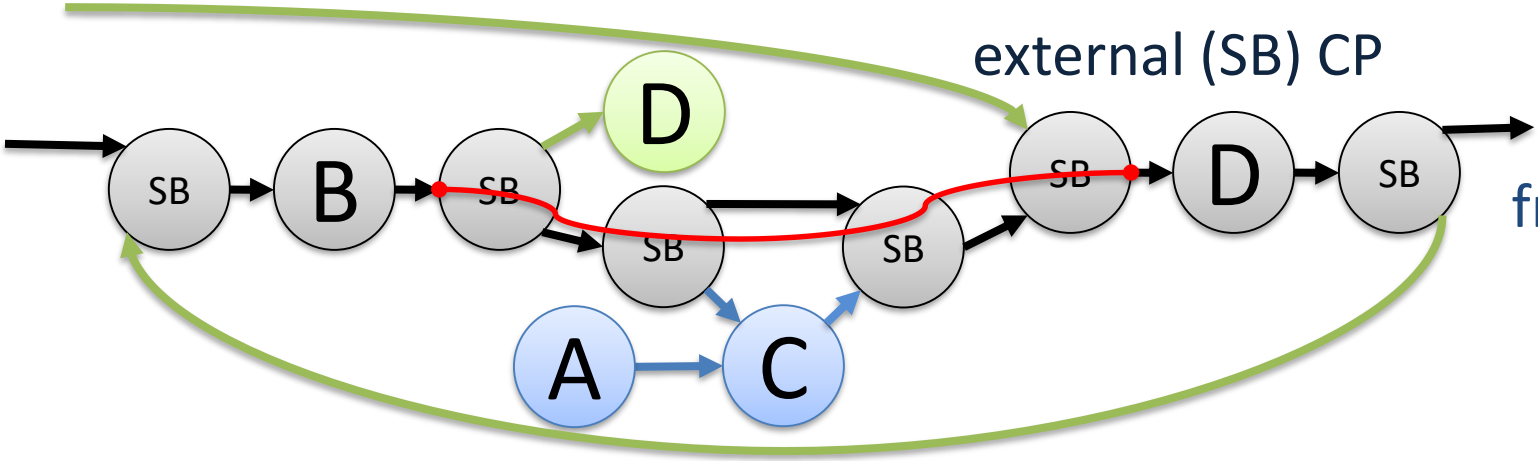

β+α+γ
frequency
42 MHz
☹

Structural Profiler

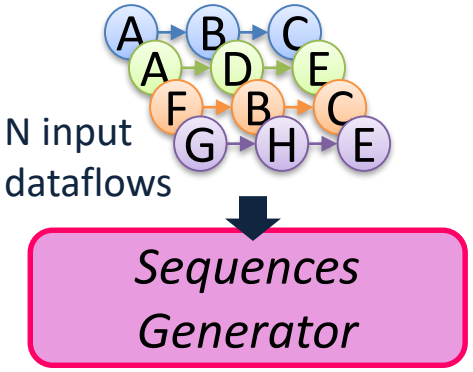

Structural Profiler

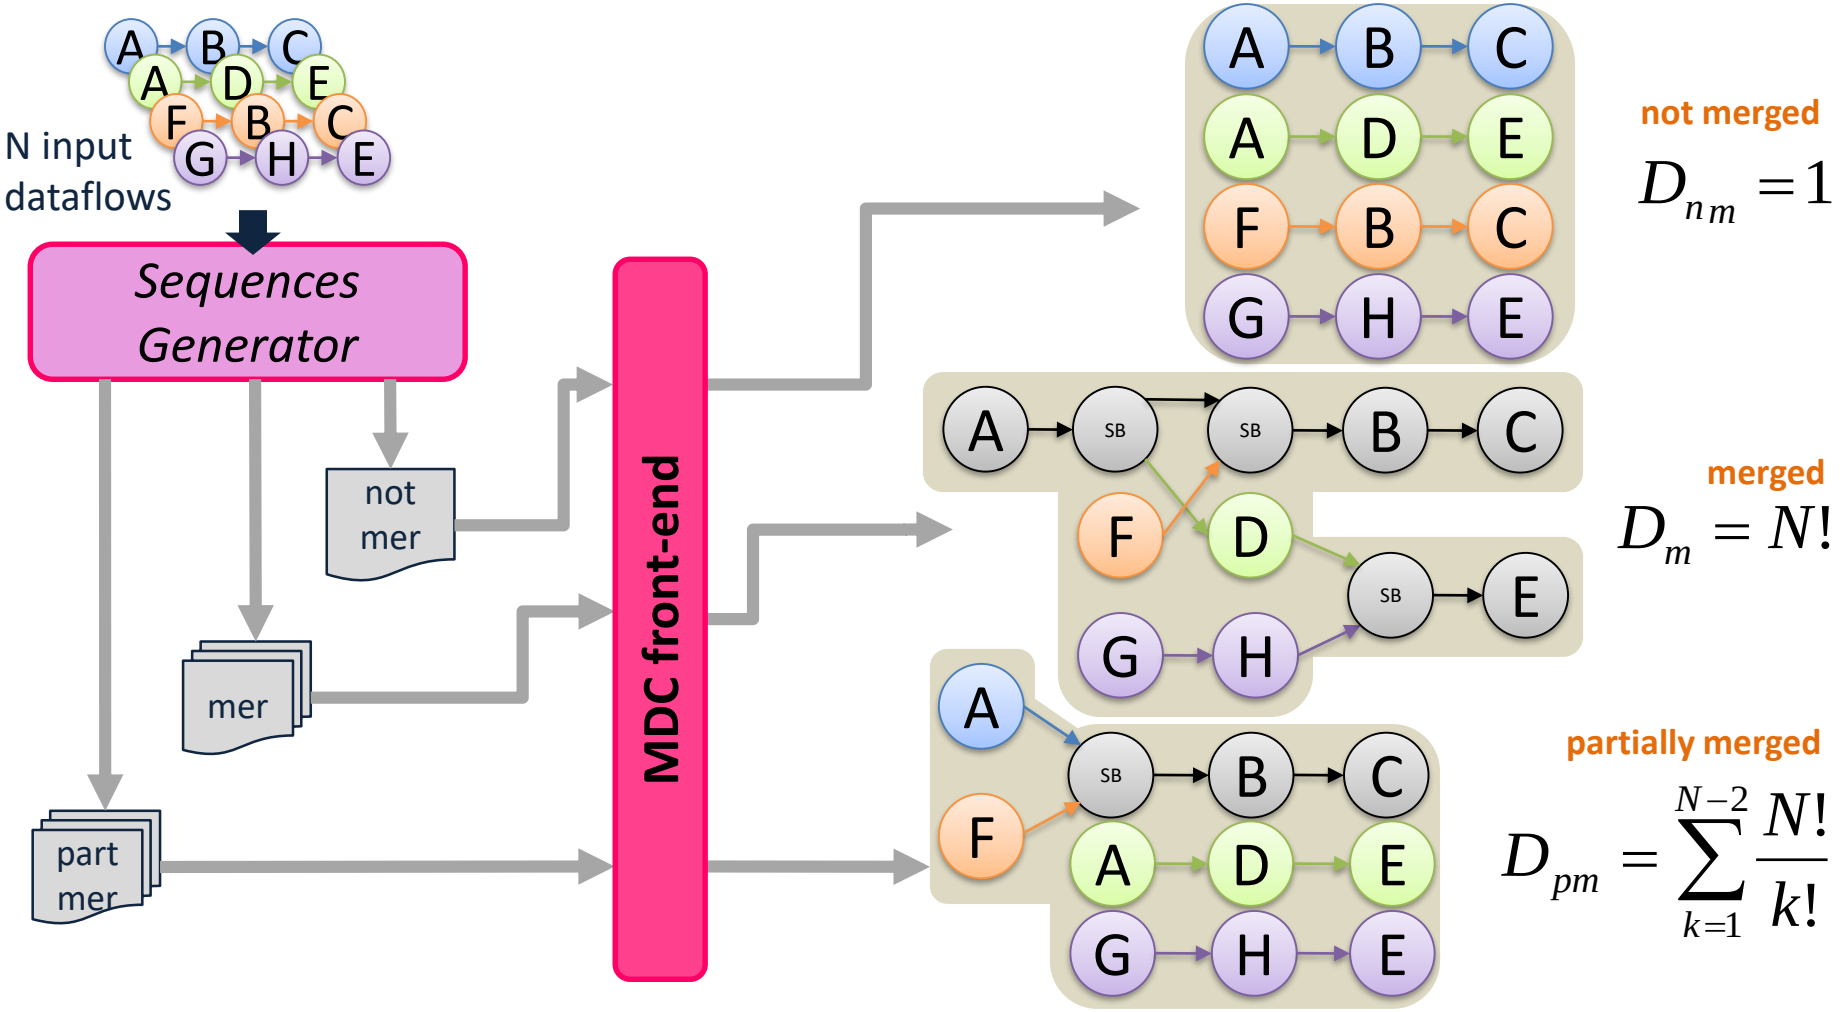

Structural Profiler

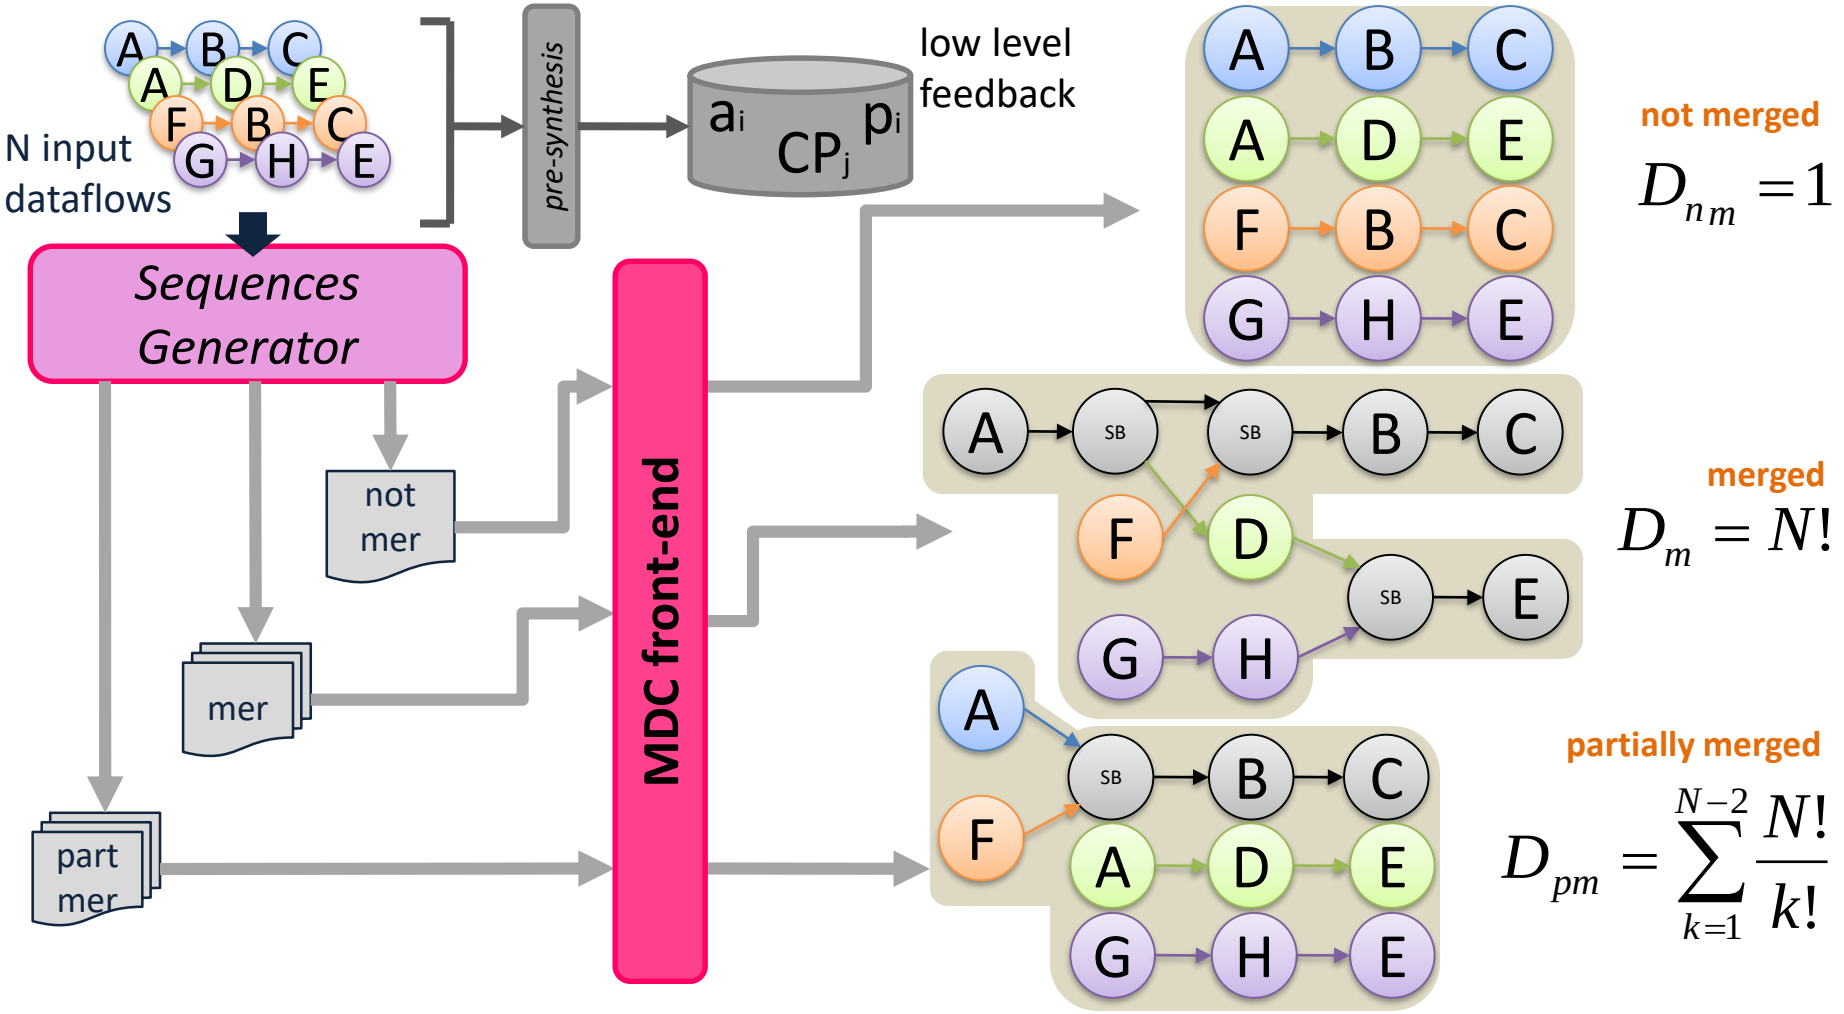

Structural Profiler

$$\textbf{Area} = \sum_{i=1}^{M} a_i$$

$$\textbf{Power} = \sum_{i=1}^{M} p_i$$

number of actors involved in the DP

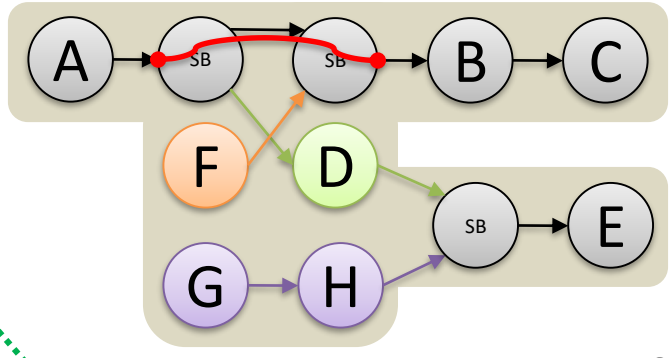

current design point (DP)

low level feedback

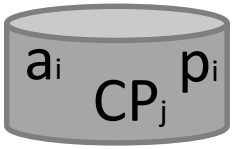

a_i/ p_i = actor area/power

CP_j = input dataflow critical path

$$\textbf{Frequency} = \frac{1}{CP} = \frac{1}{\max(CP_{in}, CP_{SB})}$$

longest SB chain within the DP

$$CP_{in} = \max(CP_j)$$

$$CP_{SB} = f(b) * \ln(N_{SB}) + g(b)$$

number of SBs in the DP chain

empirical functions of the SB size in bits *b*

Structural Profiler

Automated Pareto Analysis

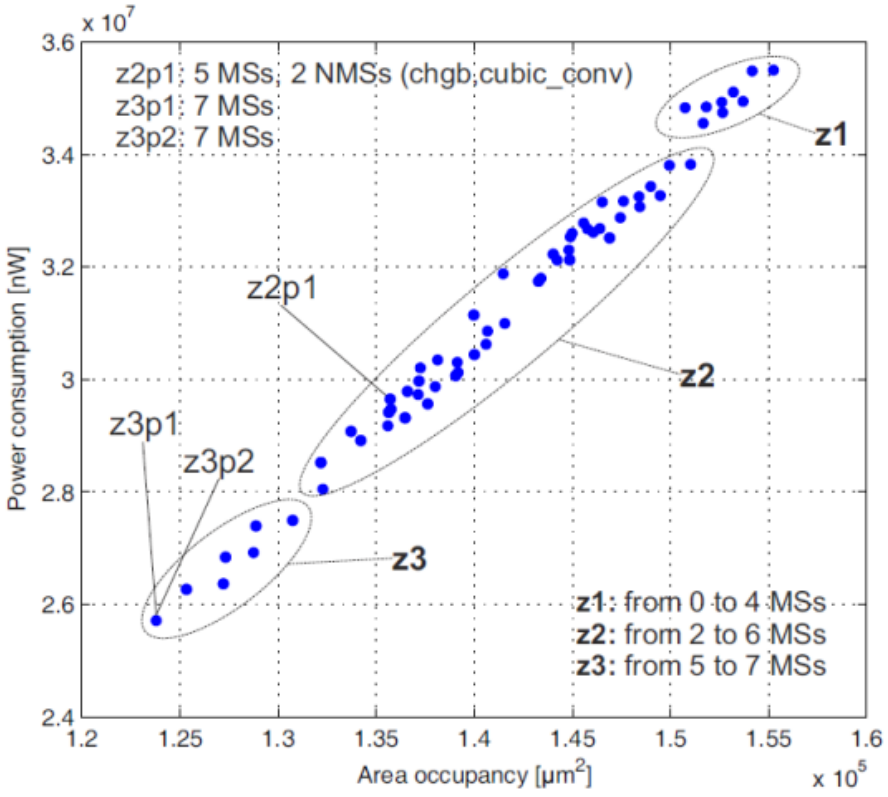

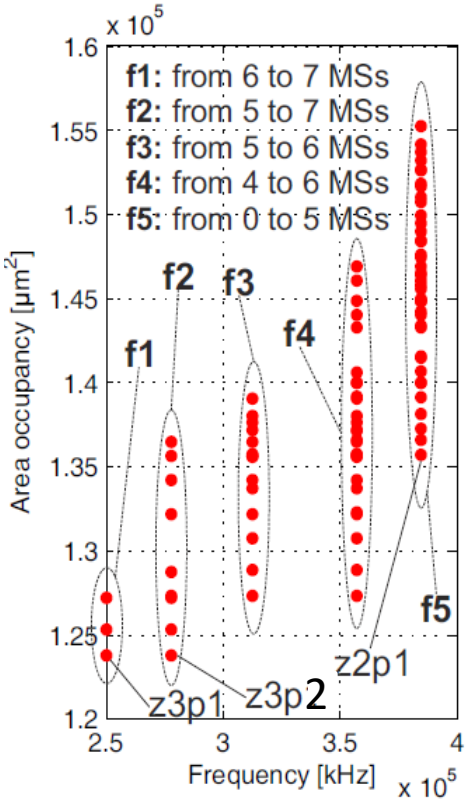

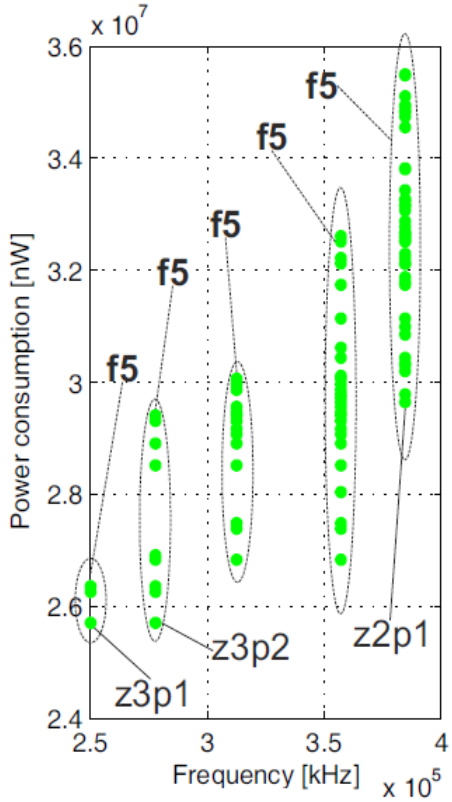

MSs = Merged dataflow Specifications (example with N=7)

Structural Profiler

Automated Pareto Analysis

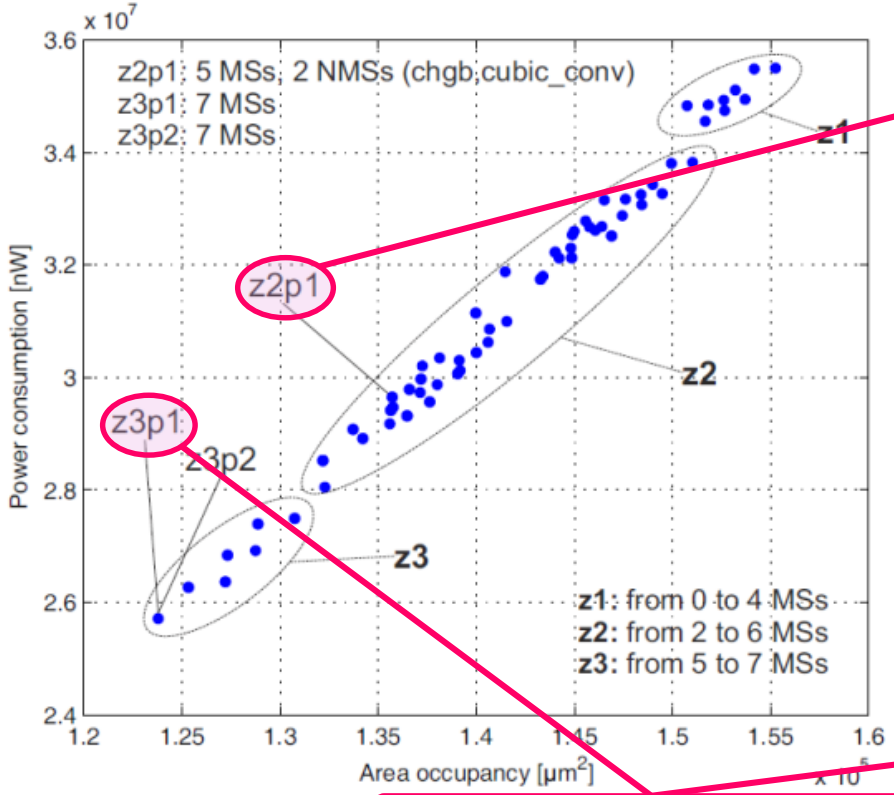

FREQ. OPTIMAL

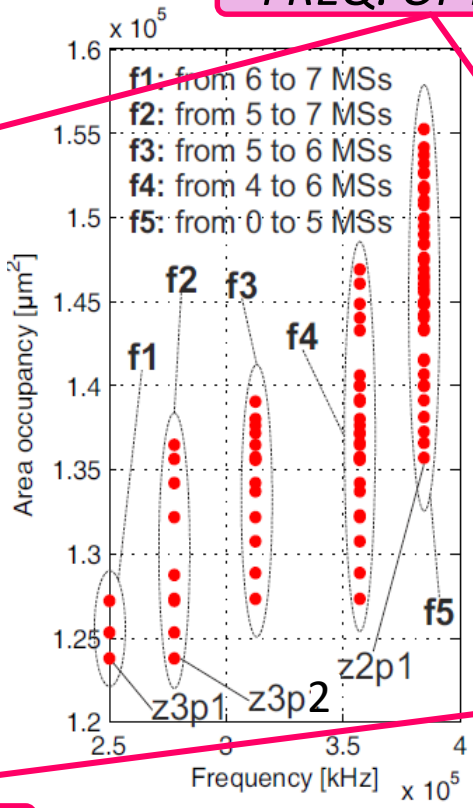

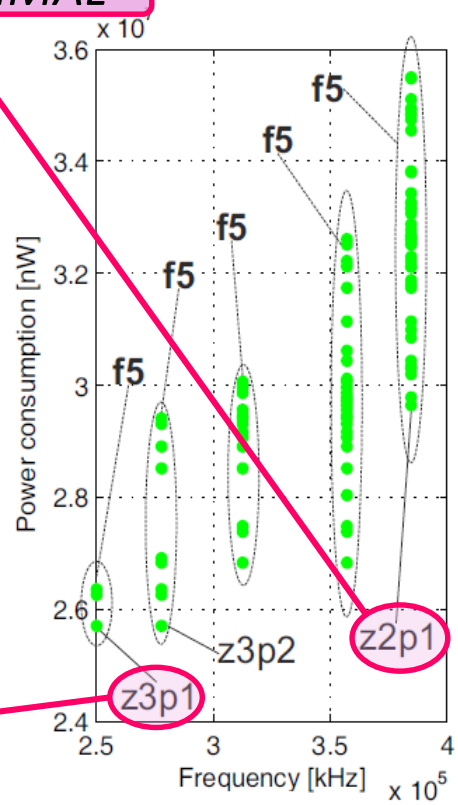

AREA/POWER OPTIMAL

MSs = Merged dataflow Specifications (example with N=7)

Dynamic Power Management

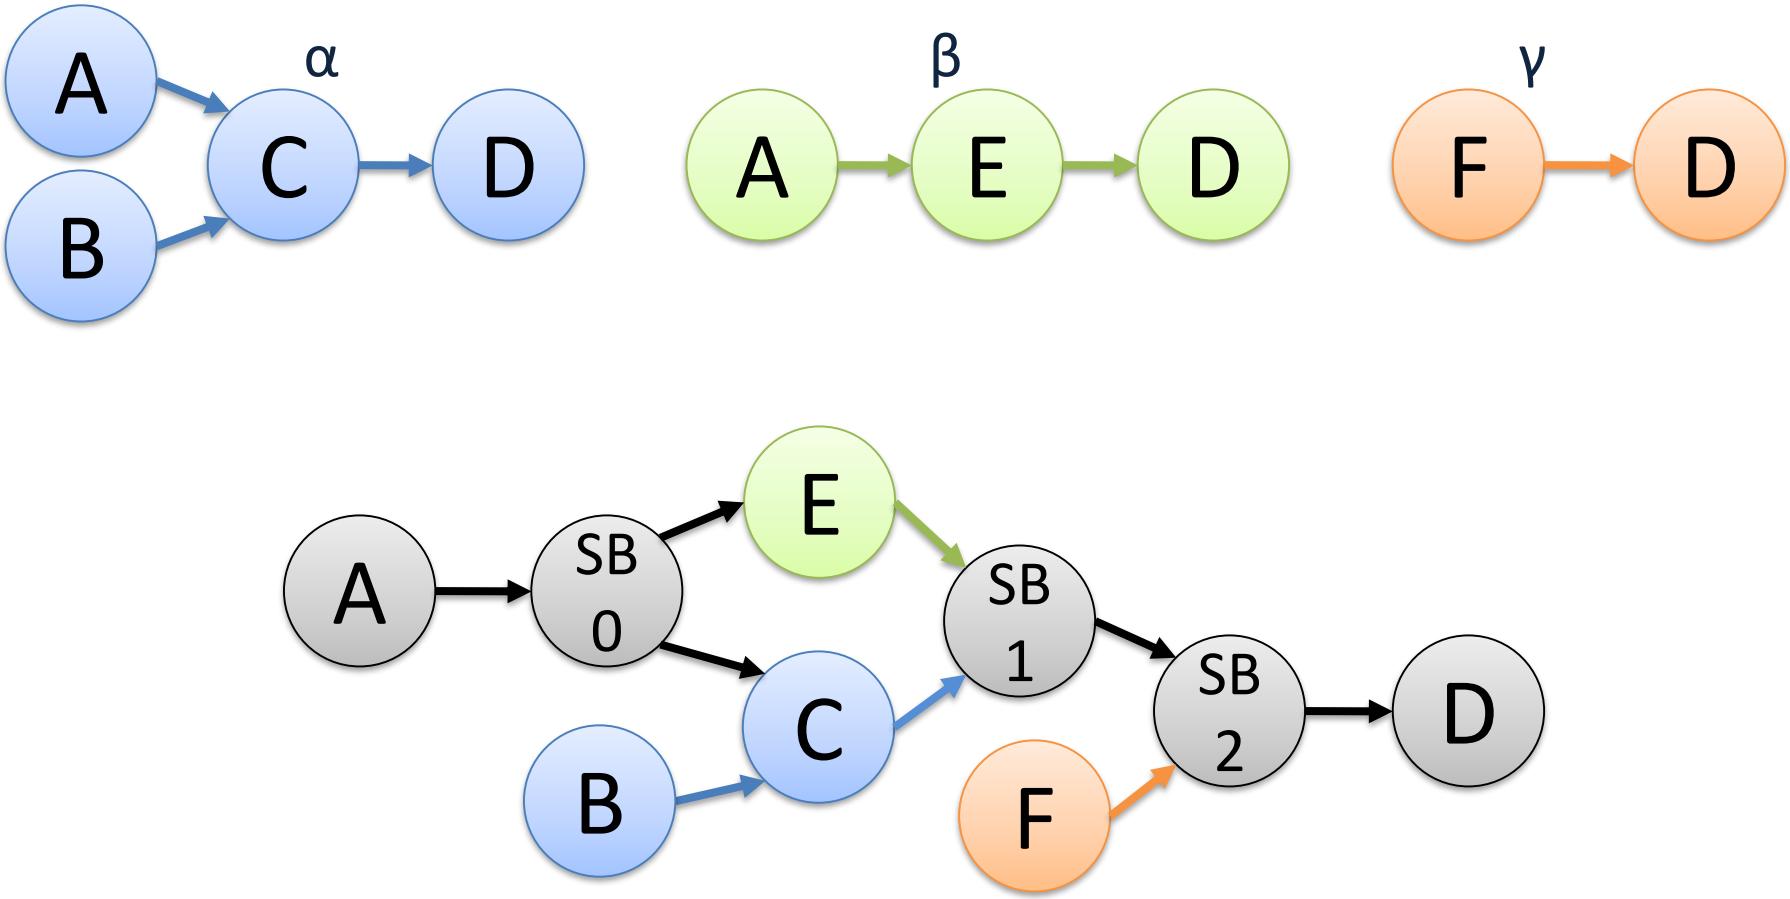

Dynamic Power Management

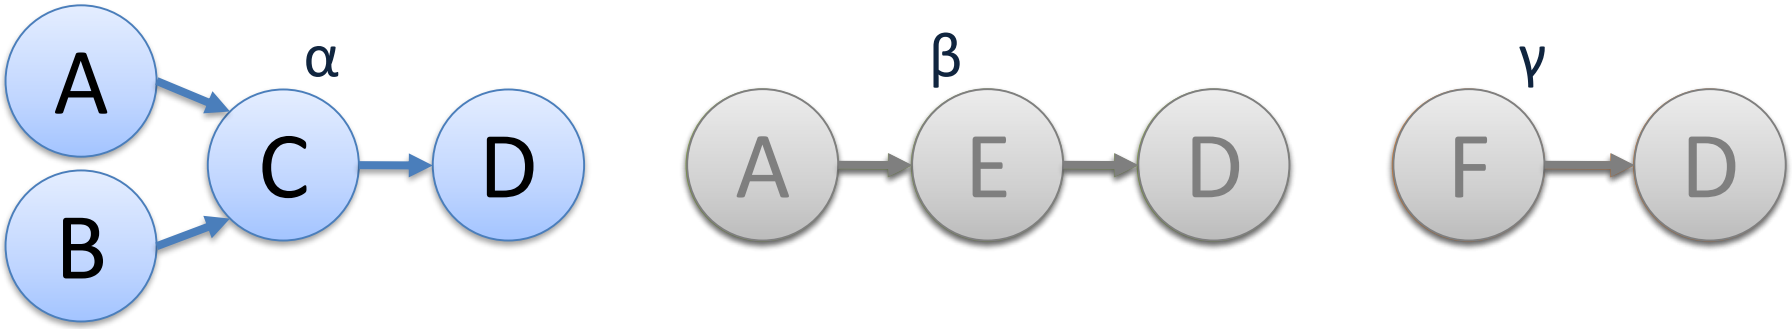

α execution: E and F are wasting power!

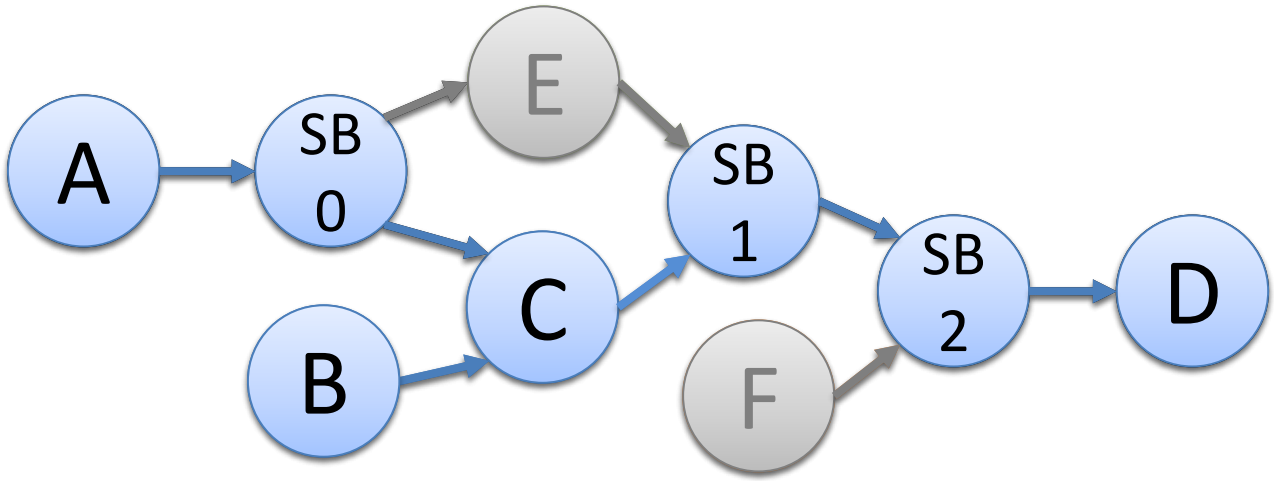

Dynamic Power Management

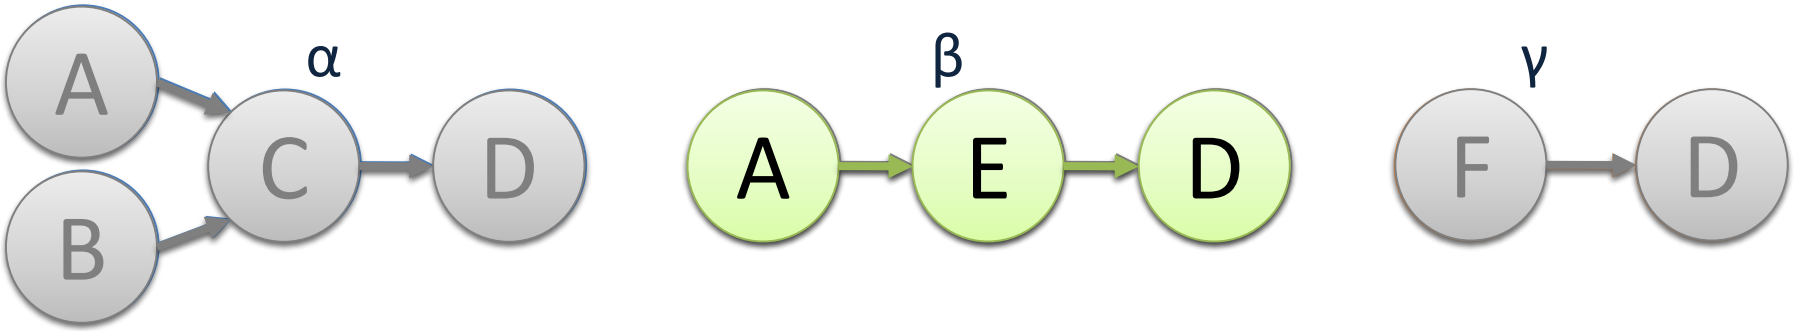

β execution: B, C and F are wasting power!

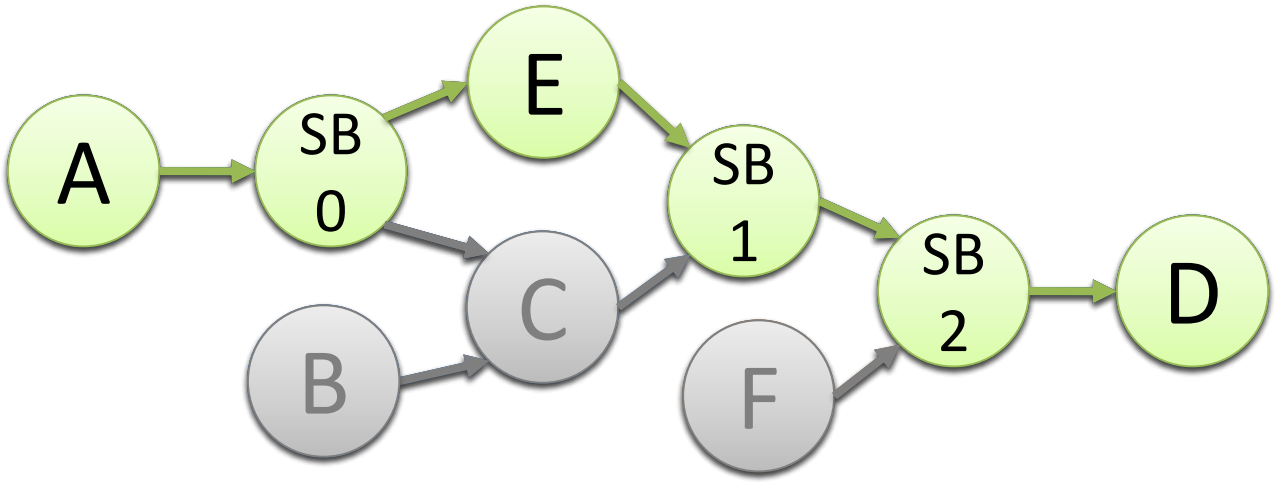

Dynamic Power Management

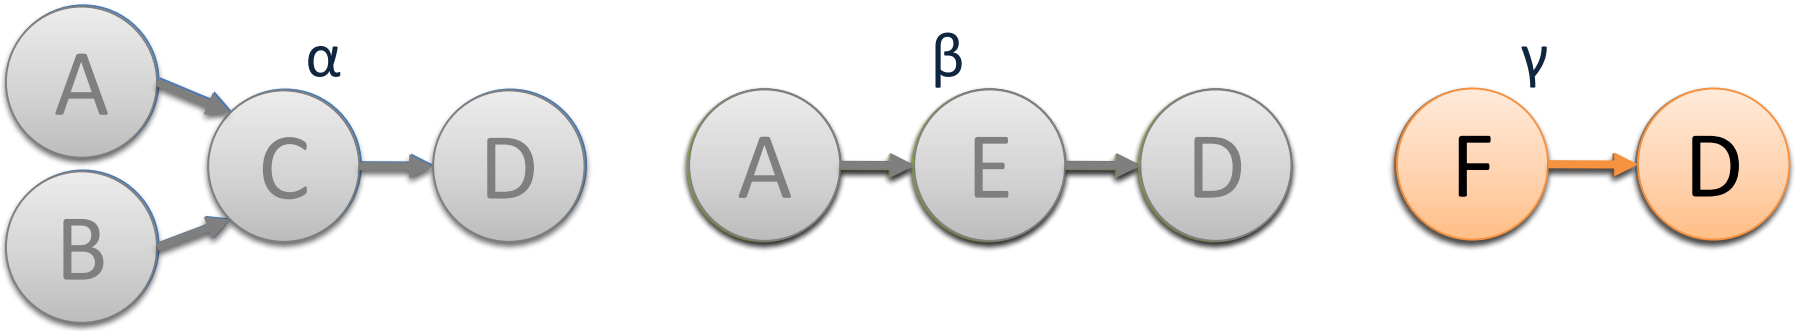

γ execution: A, B, C, E, SB0 and SB1 are wasting power!

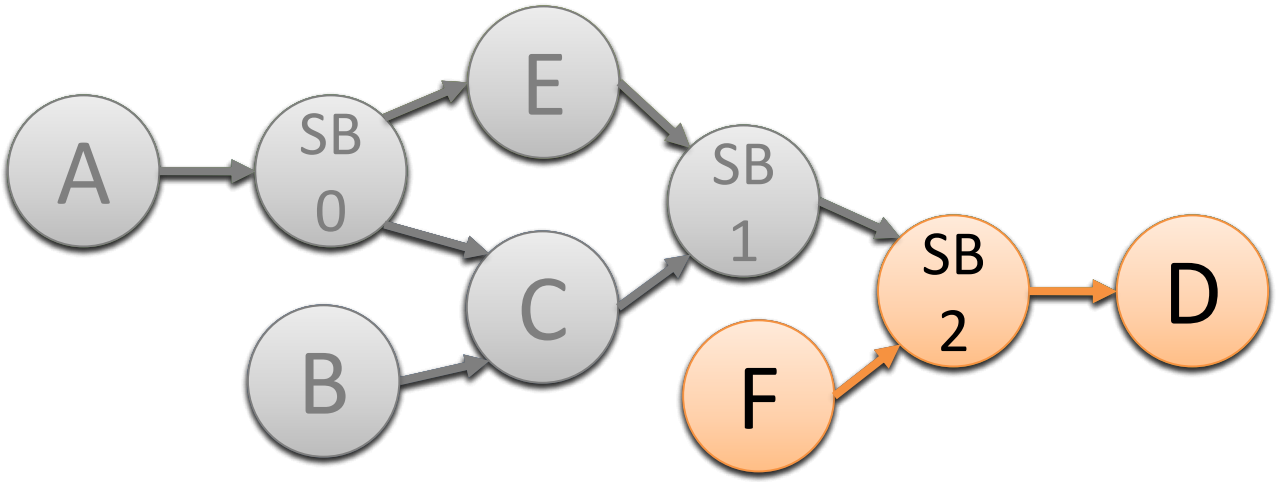

Dynamic Power Management

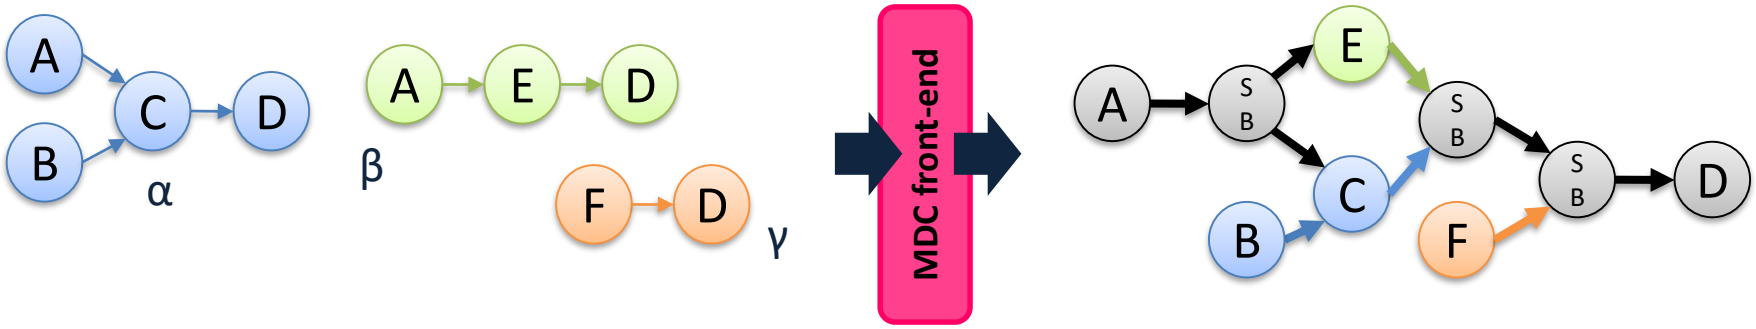

Dynamic Power Management

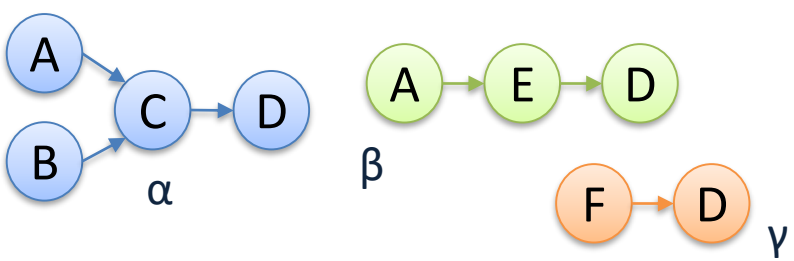

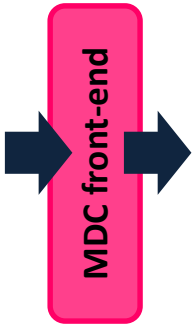

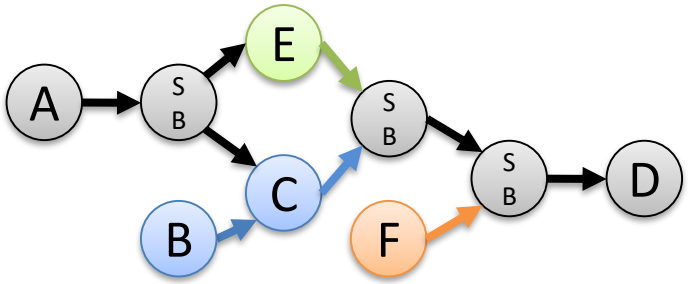

LR	1	2	3	4	5
actors	A	B,C	D	E	F
α	1	1	1	0	0
β	1	0	1	1	0
γ	0	0	1	0	1

Logic Regions (LRs) Identification

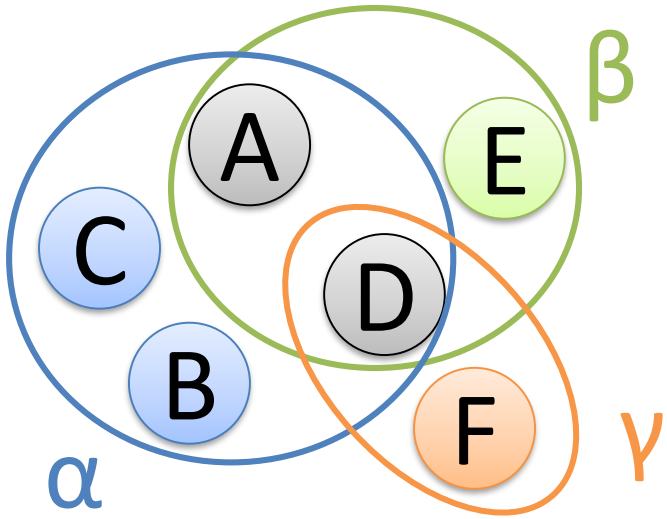

Dynamic Power Management

LR	actors	α	β	γ
1	A	1	1	0
2	B,C	1	0	0
3	D	1	1	1
4	E	0	1	0
5	F	1	0	1

MDC back-end

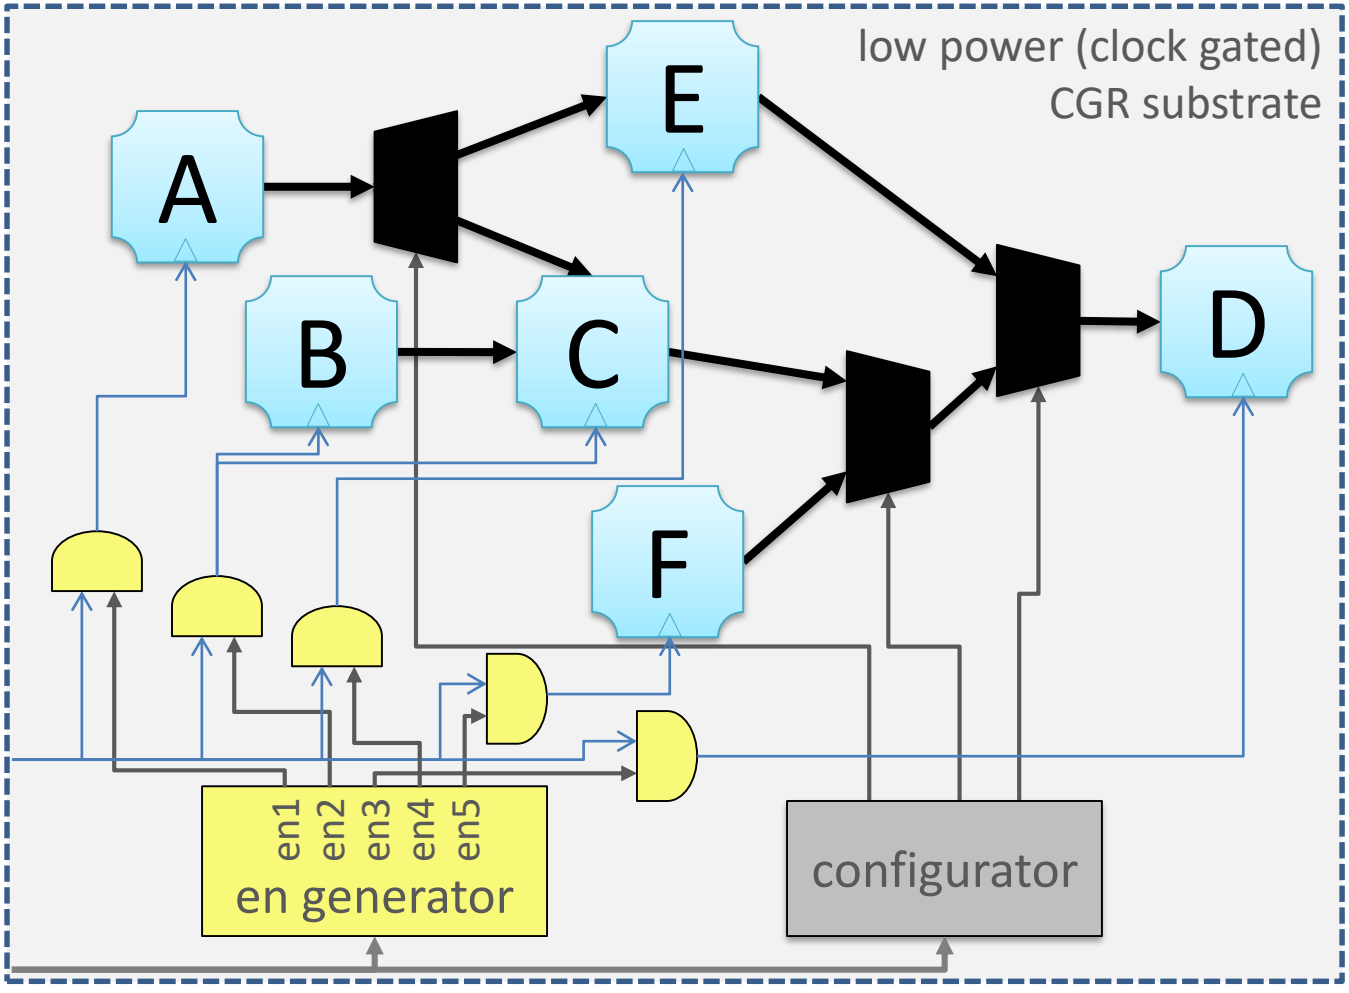

Co-Processor Generator

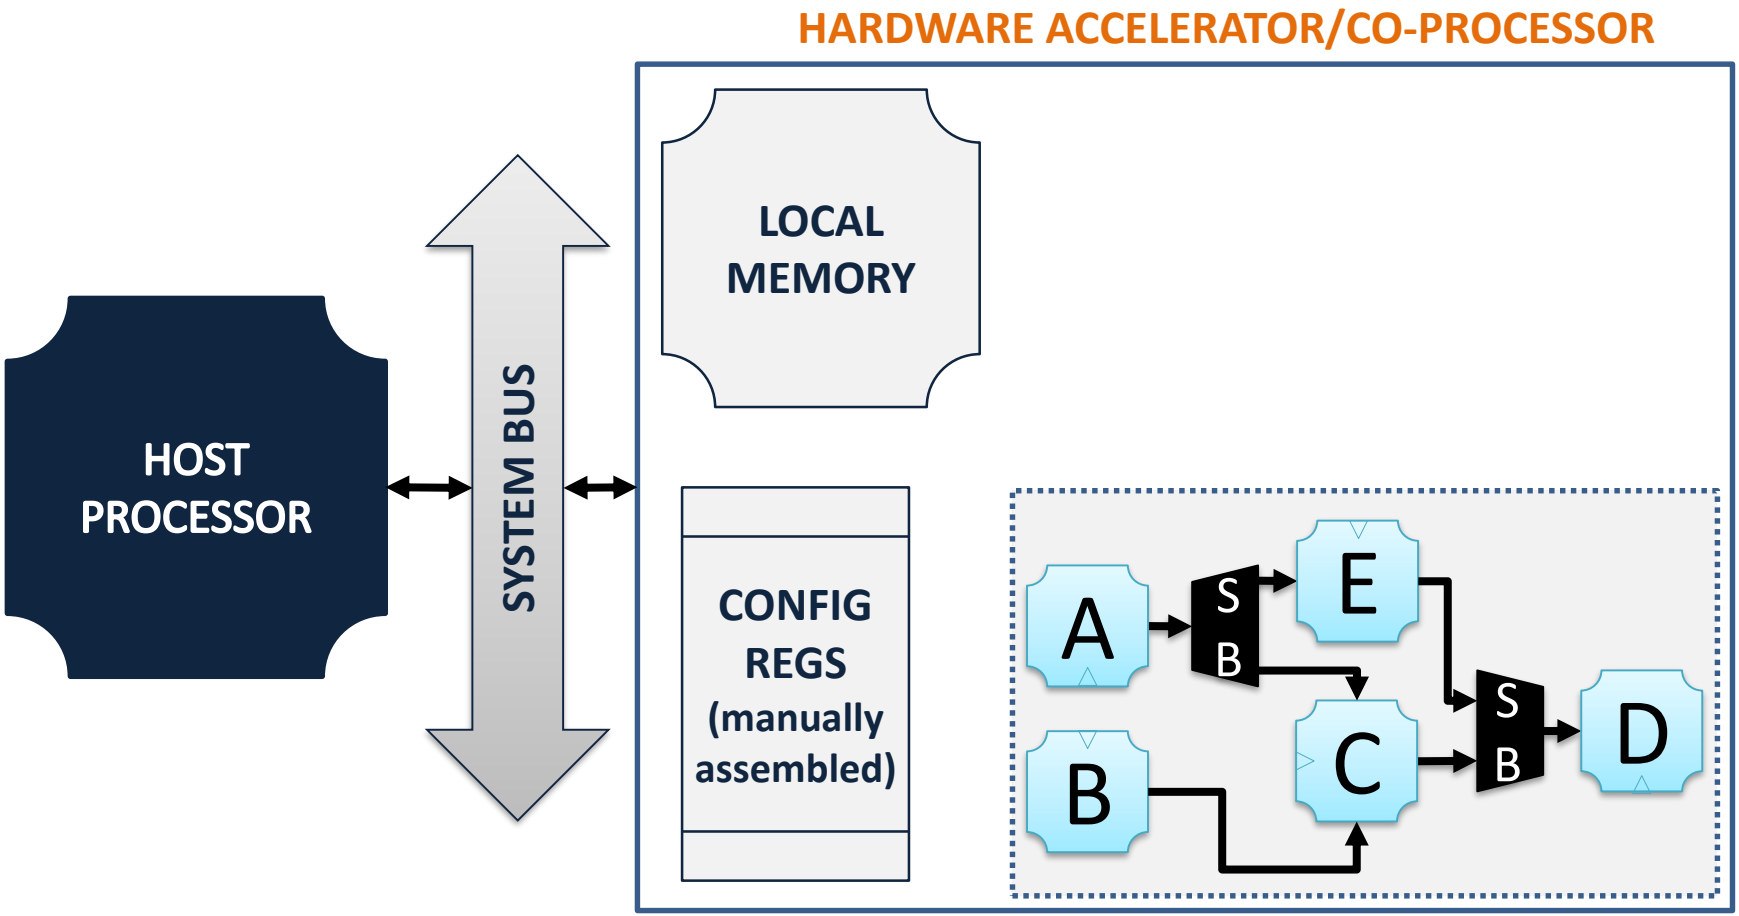

Co-Processor Generator

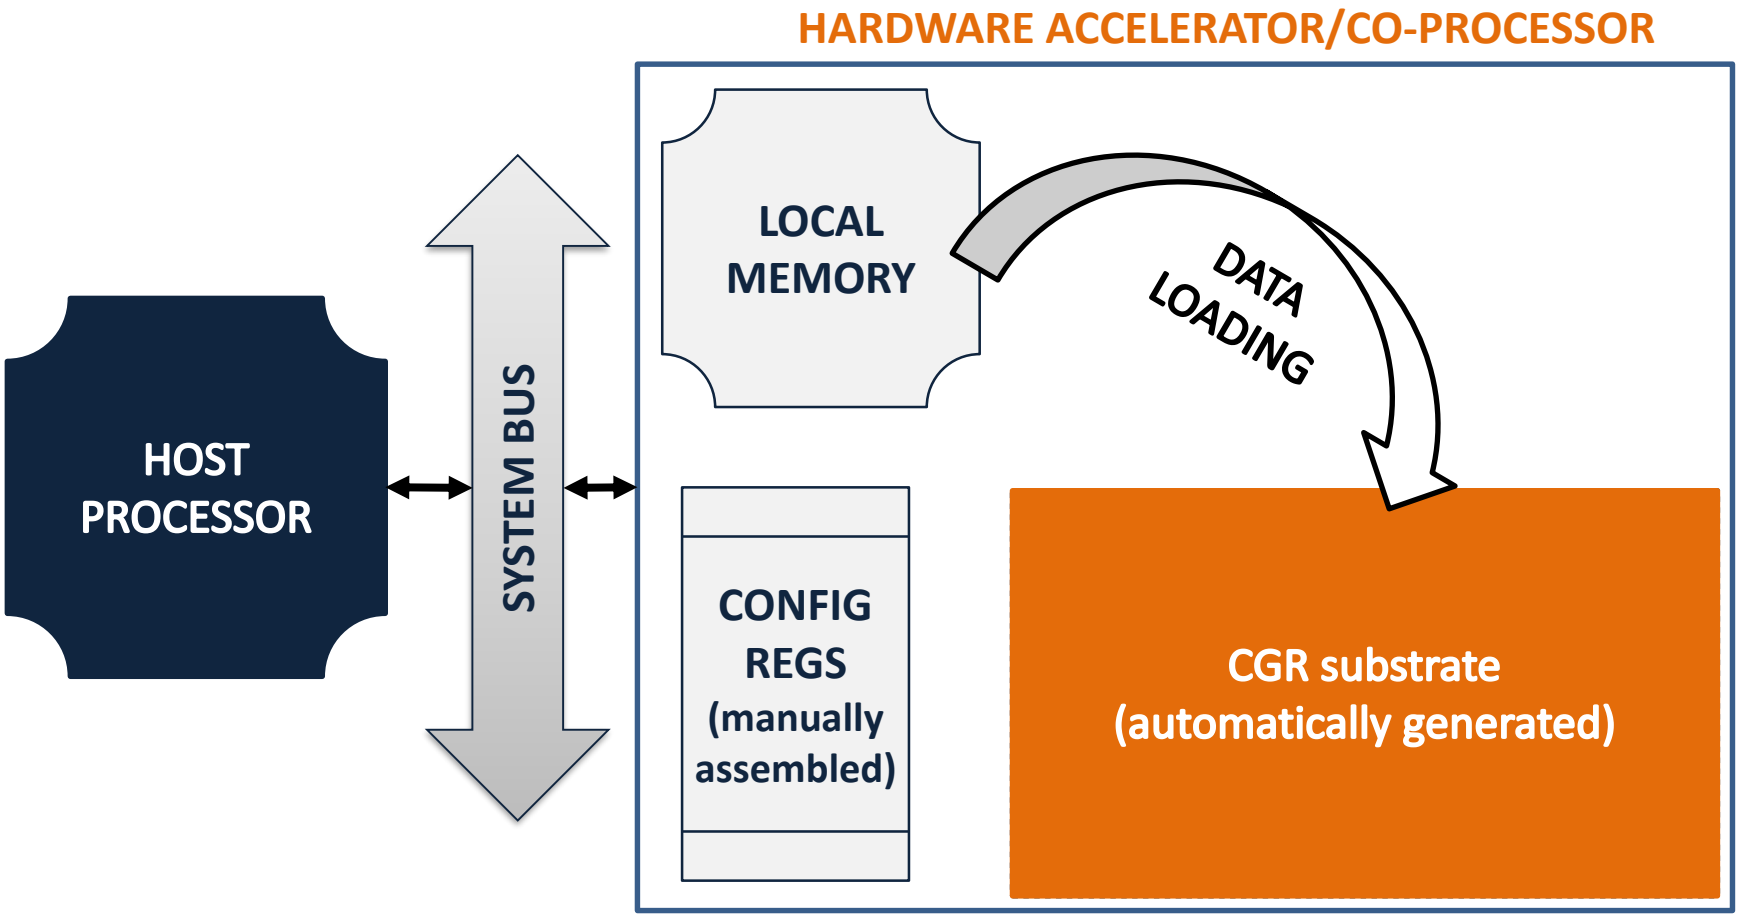

Co-Processor Generator

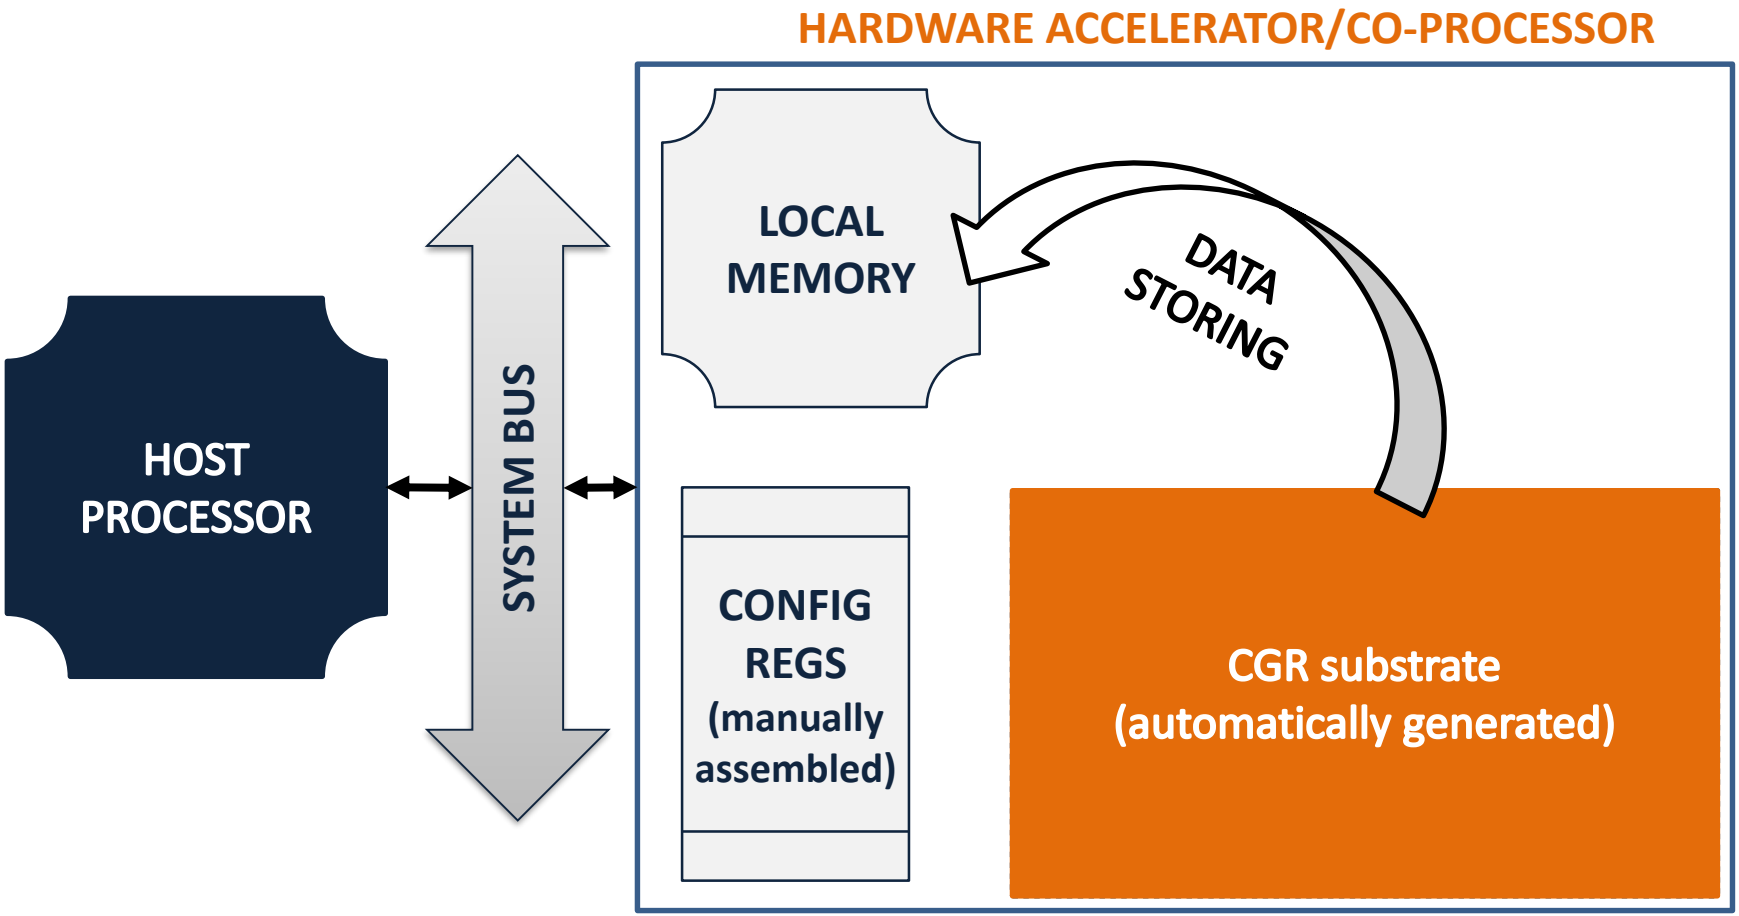

Co-Processor Generator

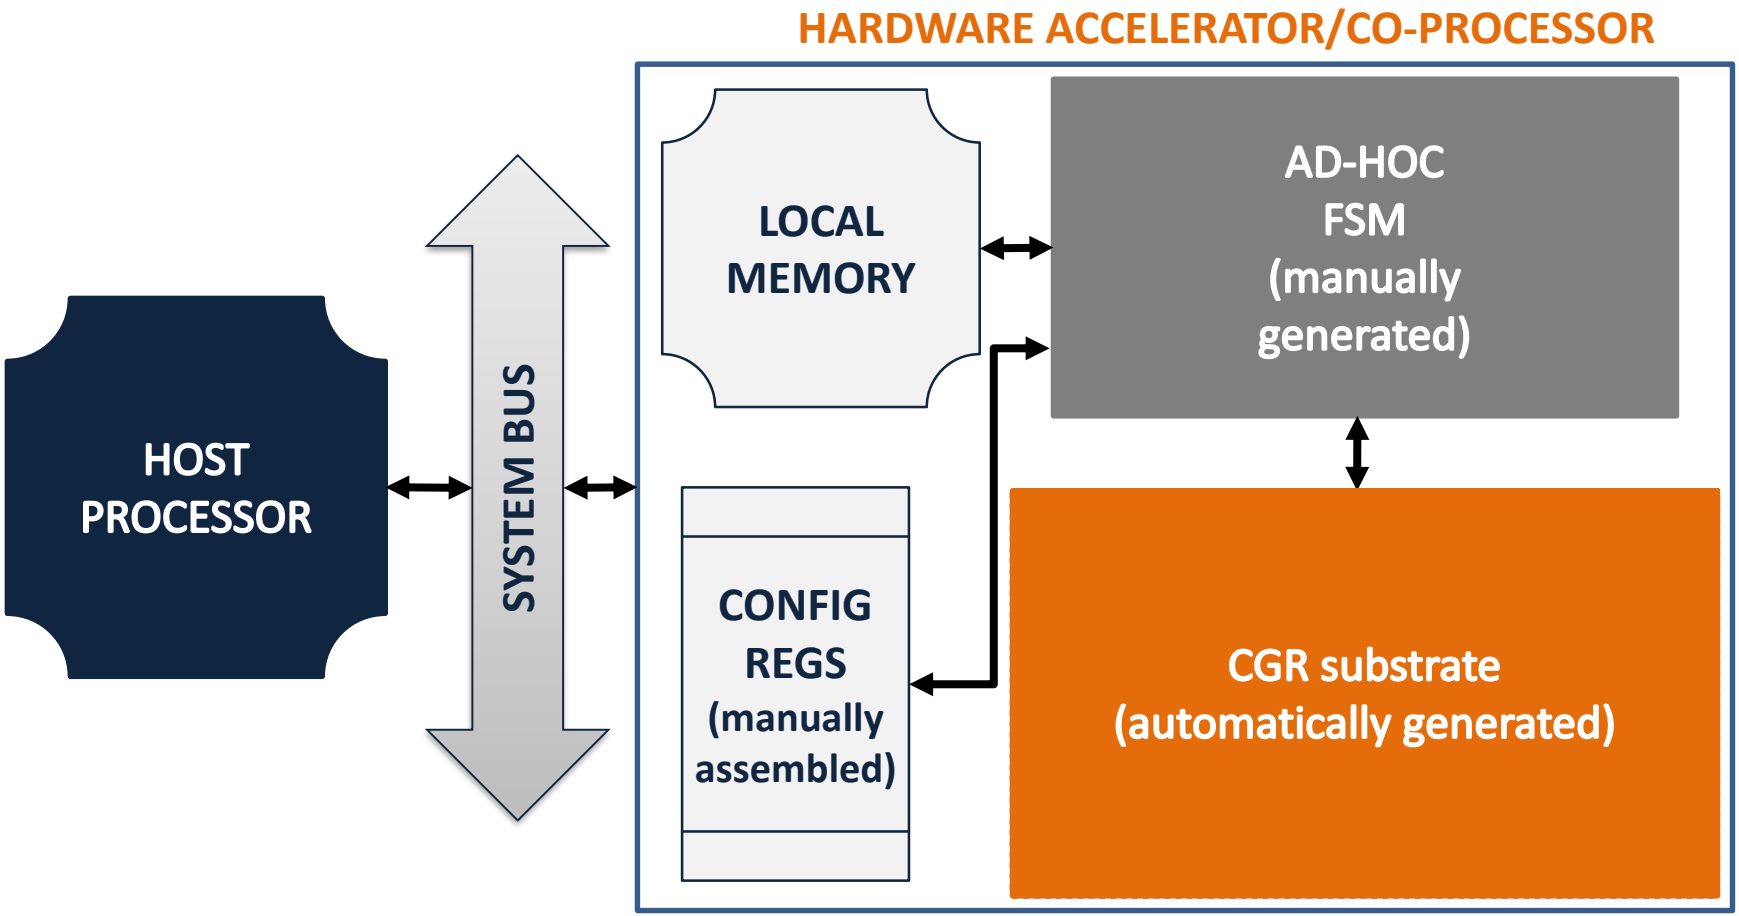

Co-Processor Generator

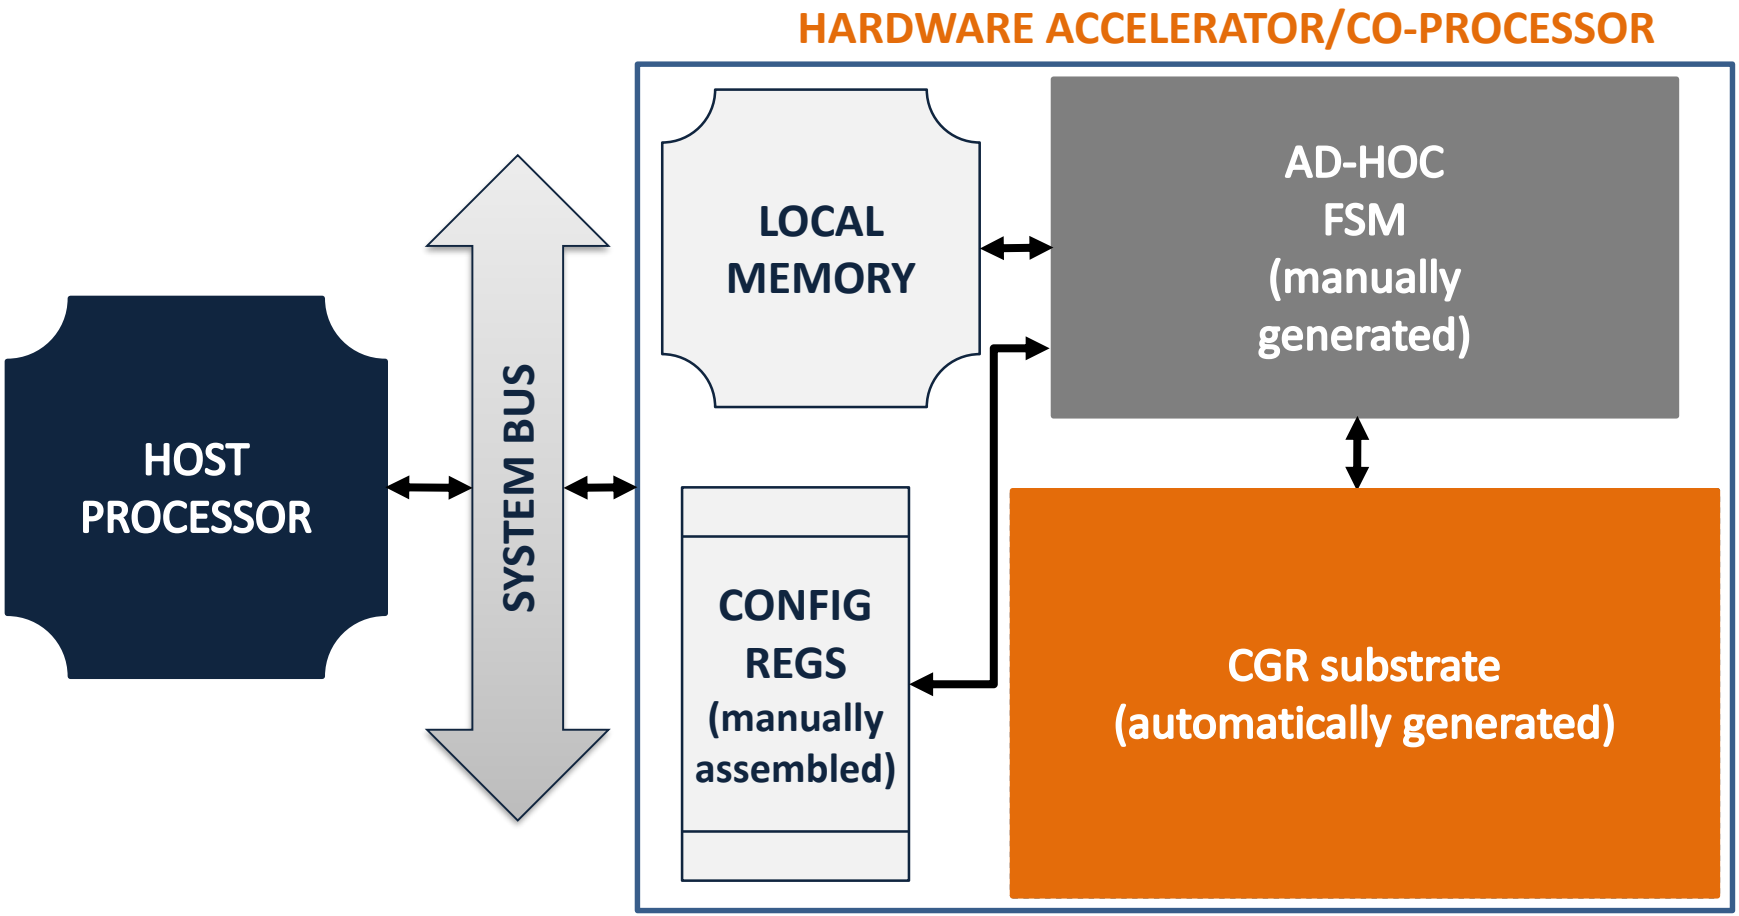

Co-Processor Generator

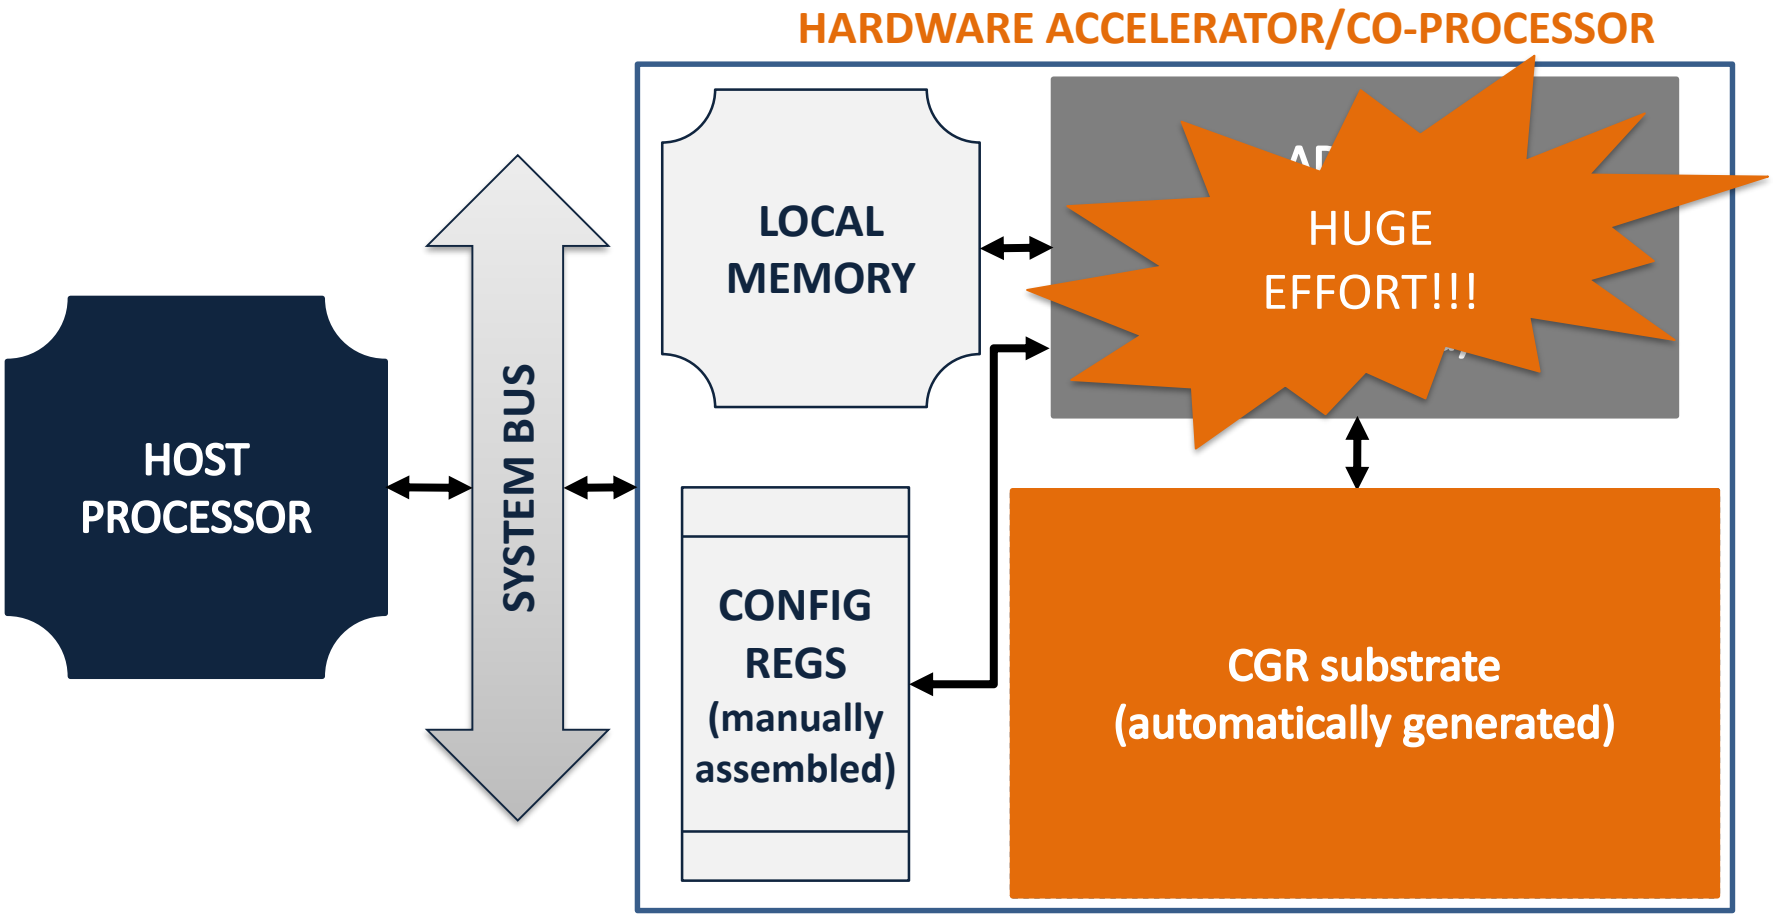

Co-Processor Generator

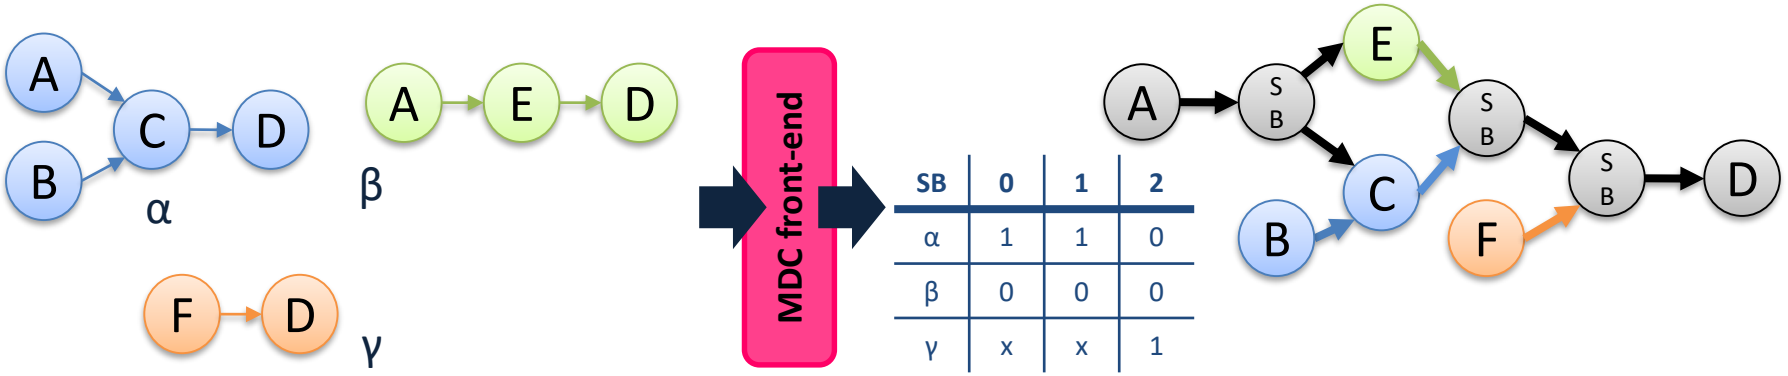

Co-Processor Generator

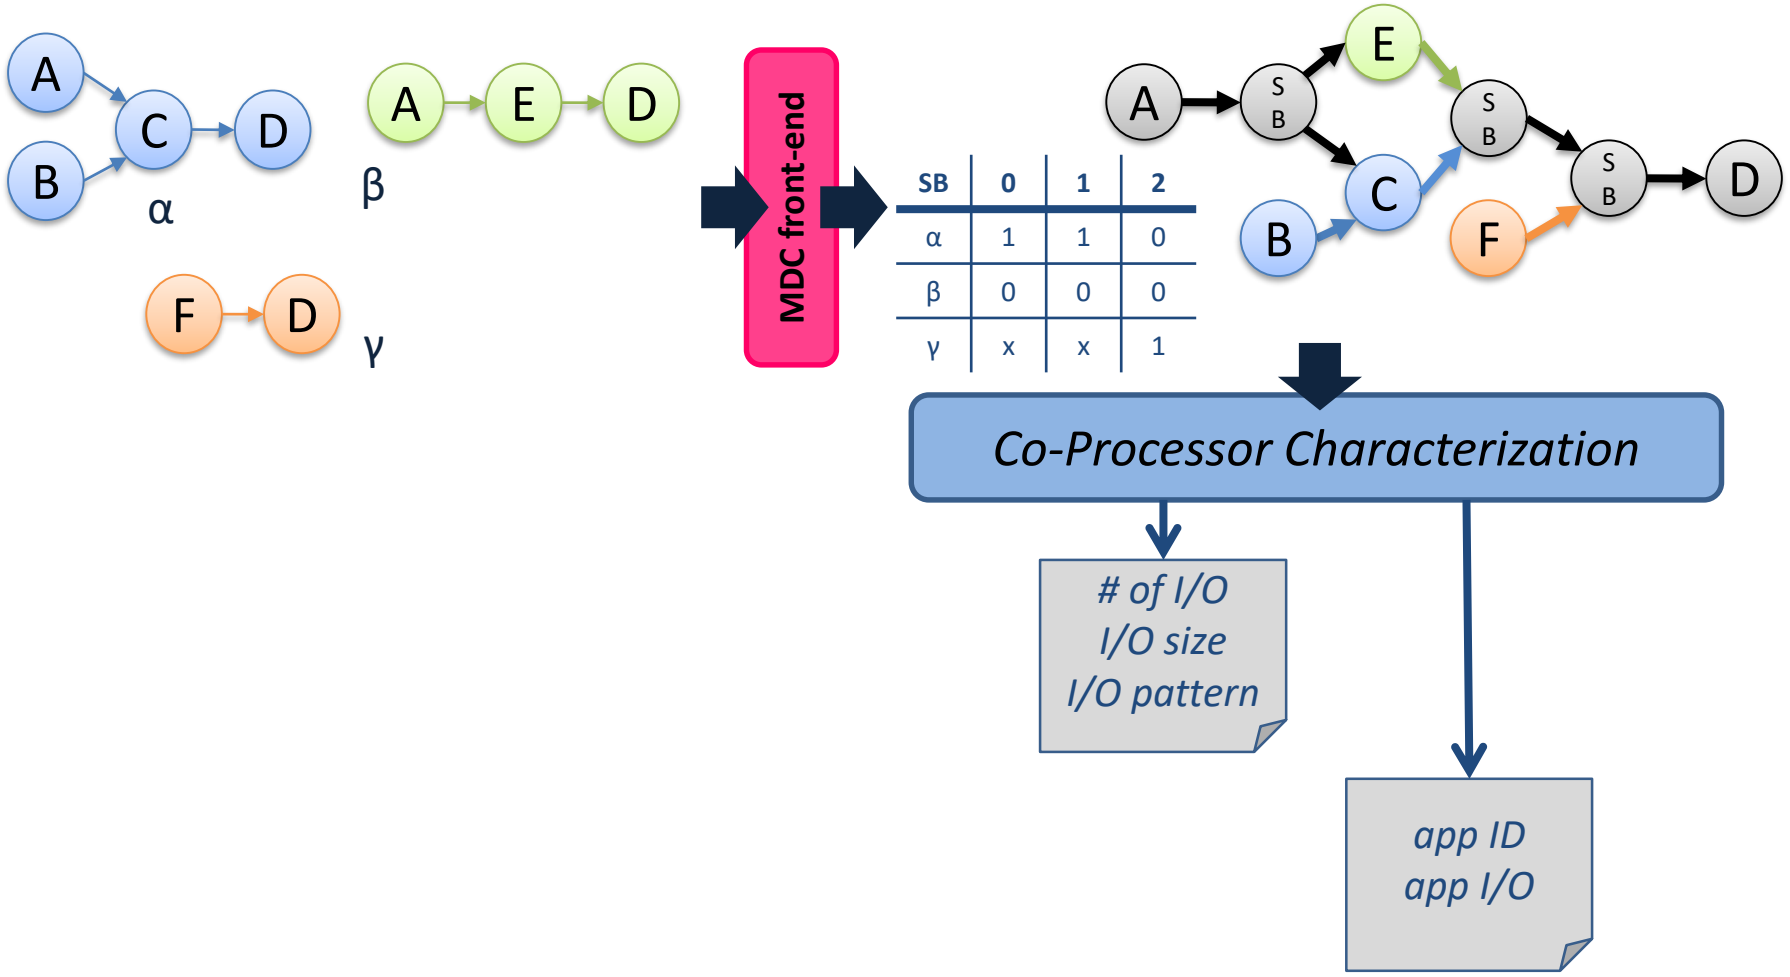

Co-Processor Generator

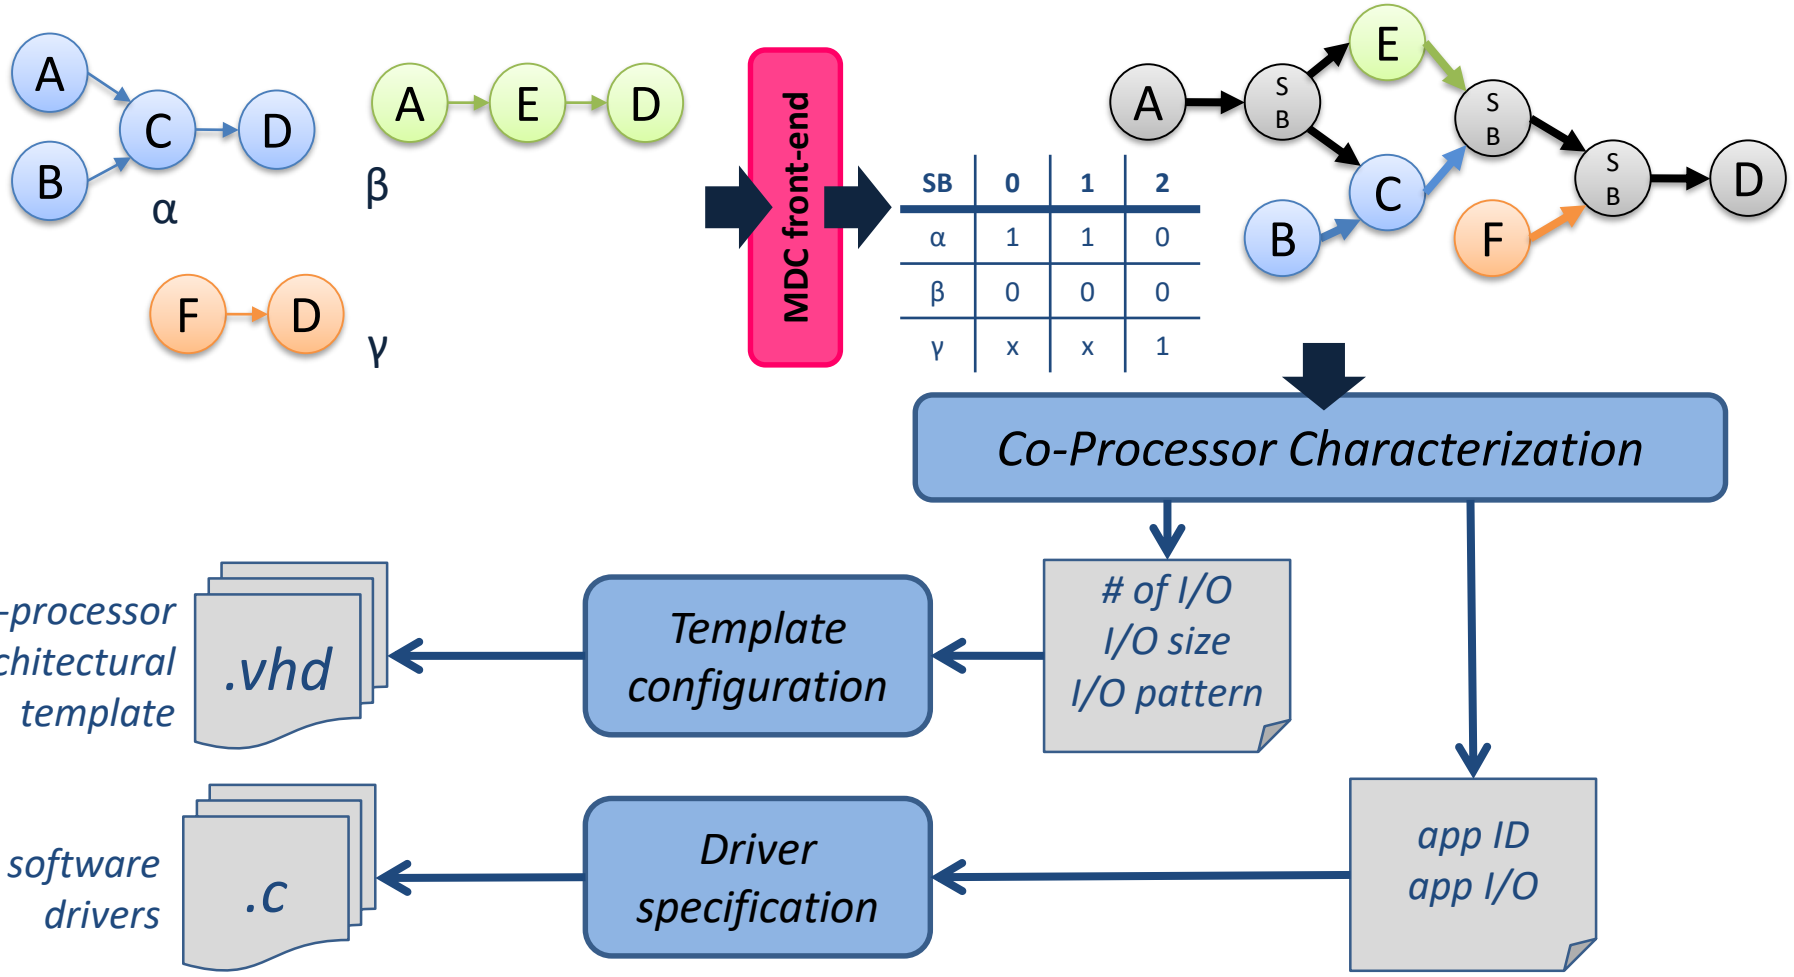

Co-Processor Generator

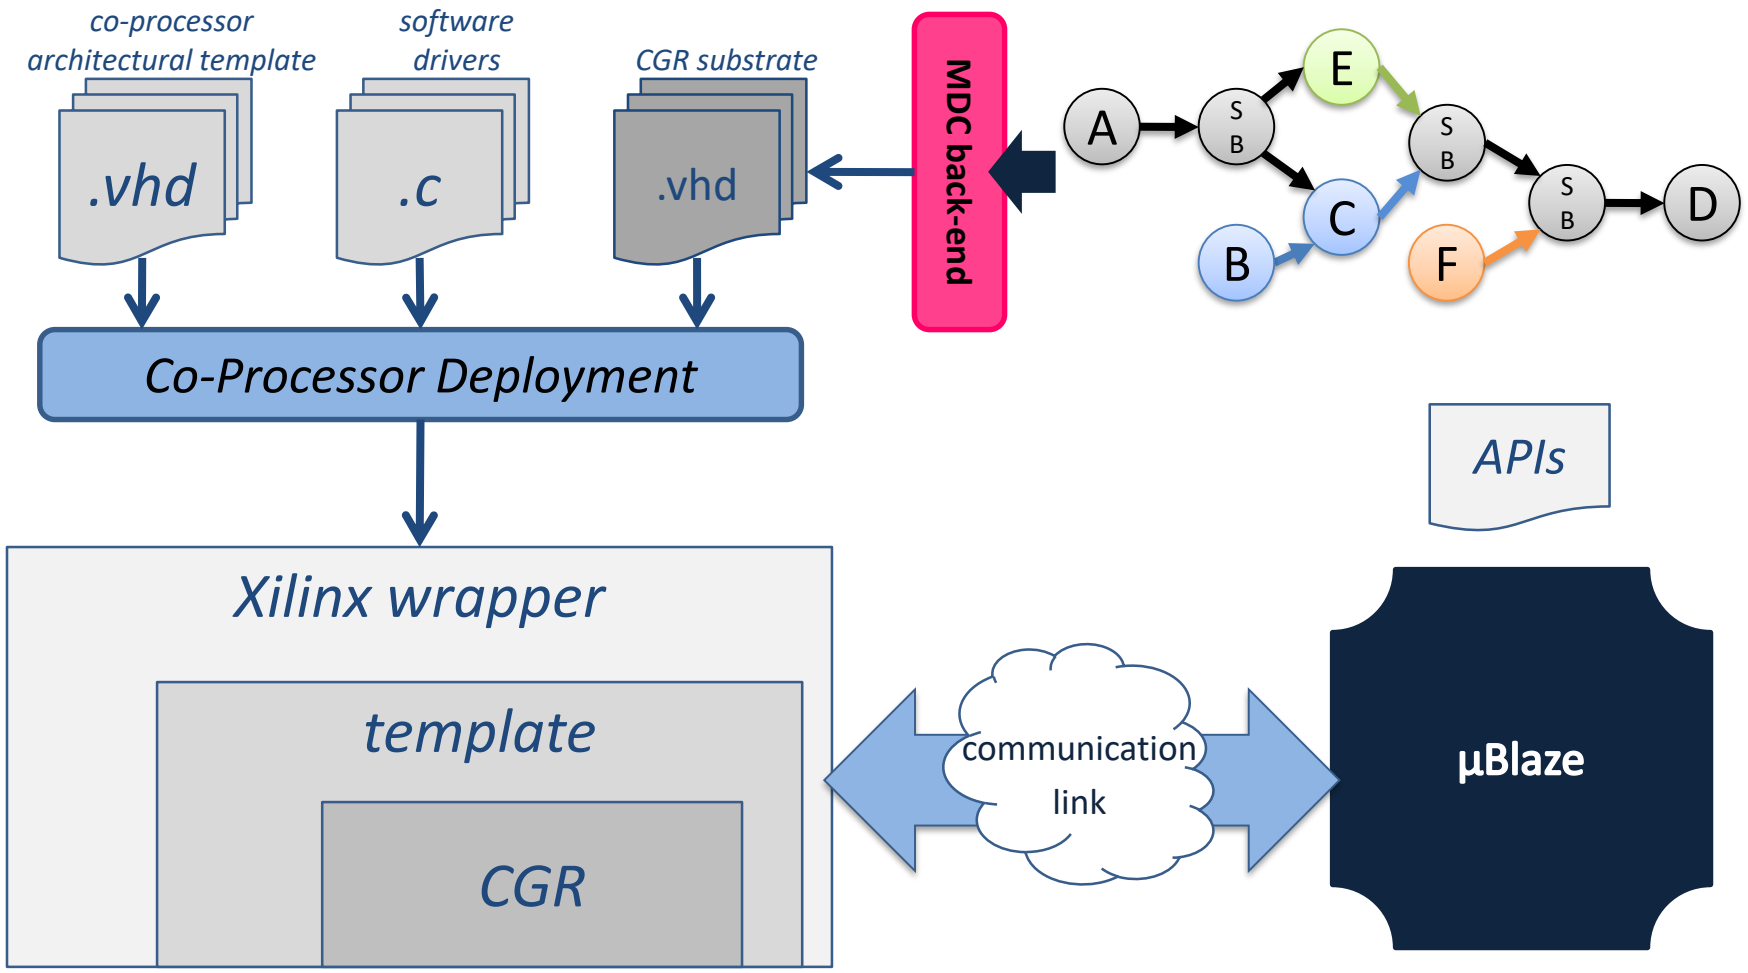

Co-Processor Generator

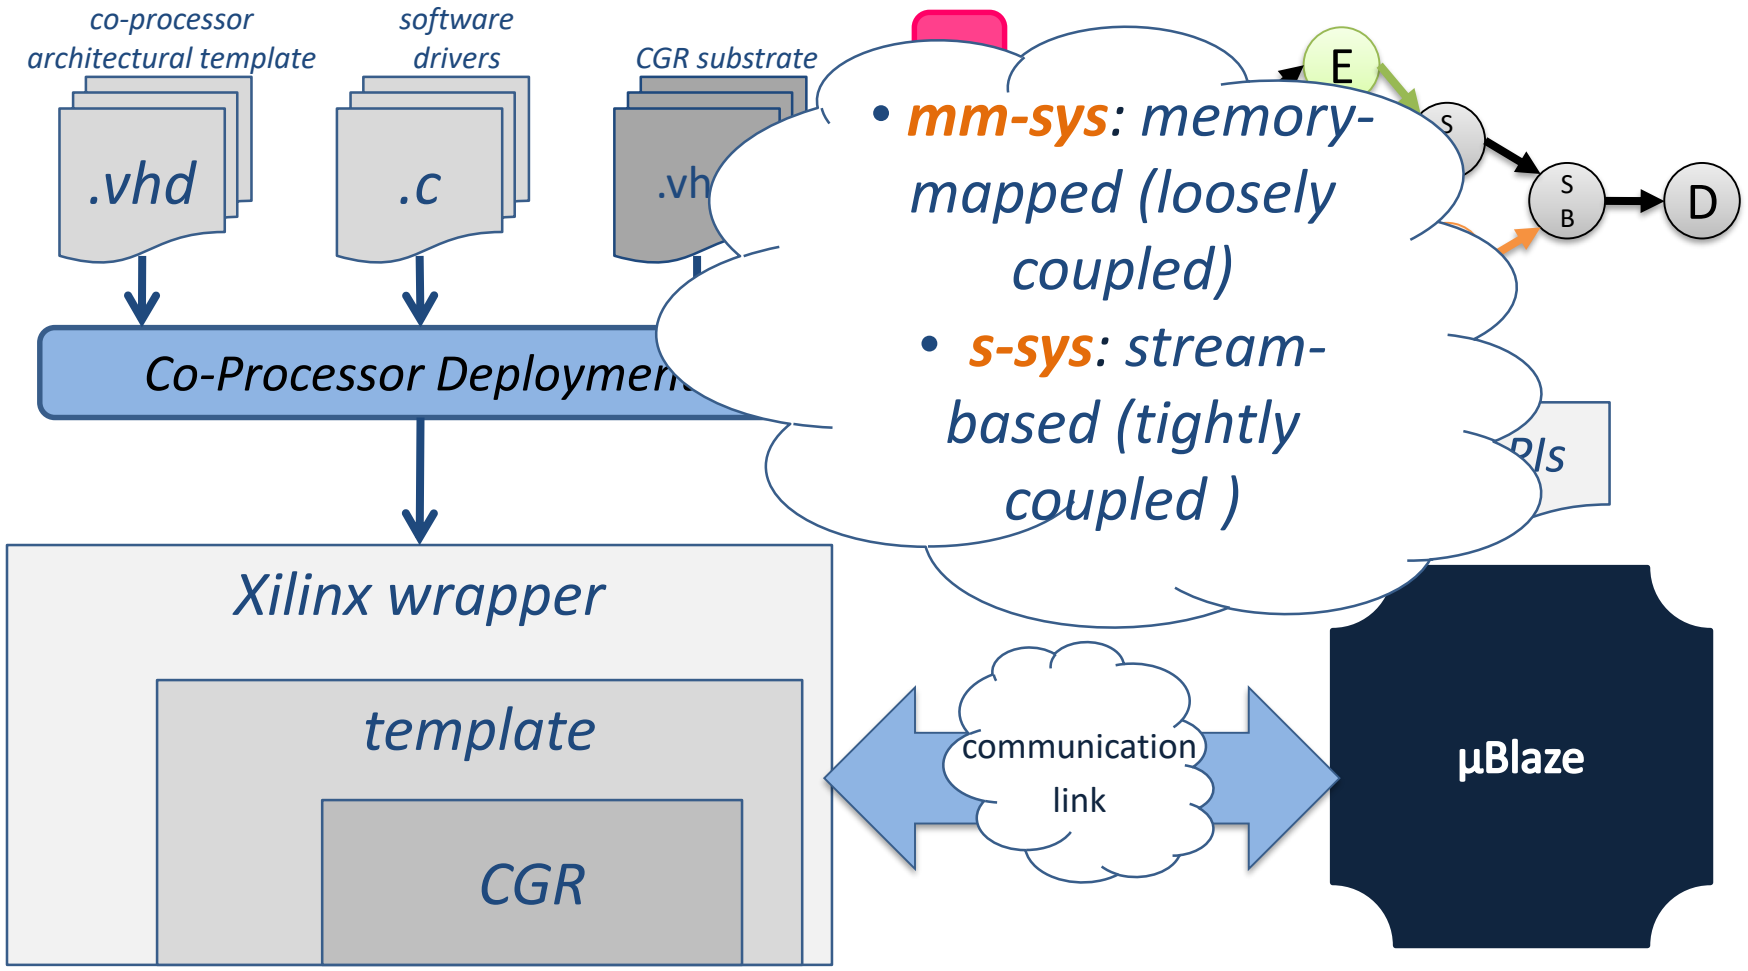

User Interface

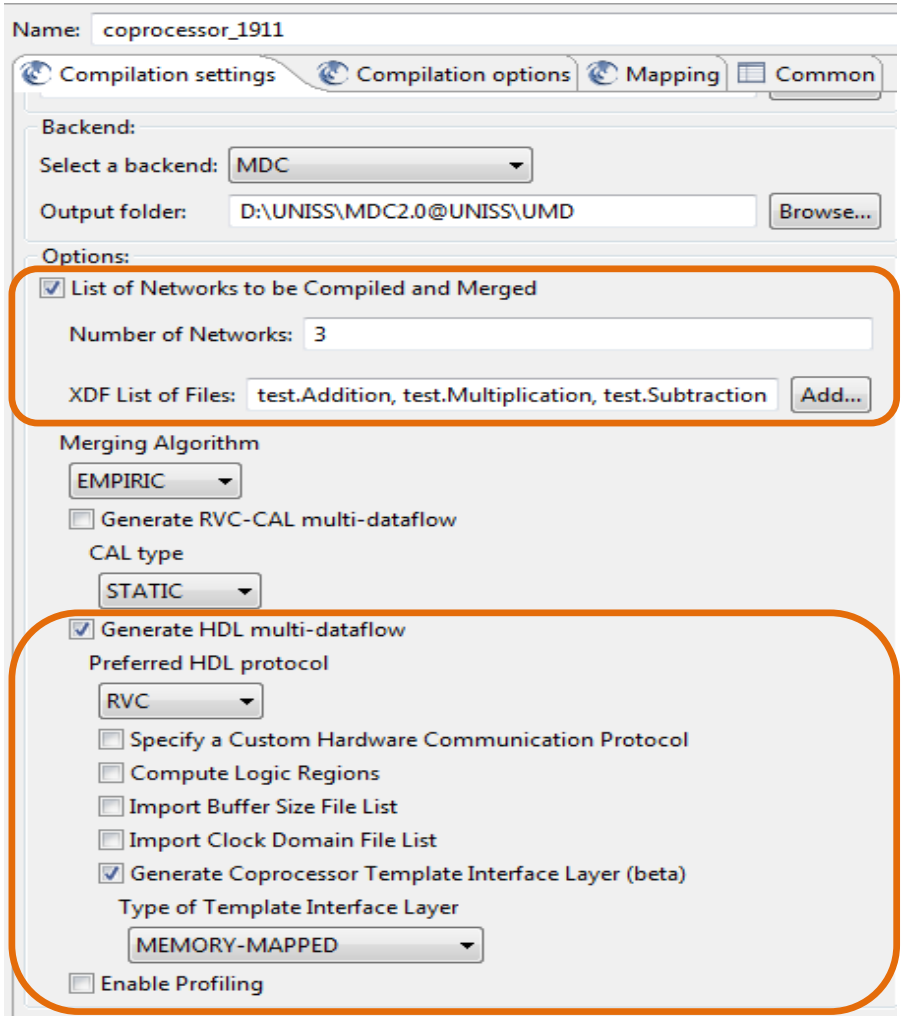

Input Dataflow Specifications

Specify the Extension to be used (if any).

Outline

- The origins of our dataflow to hardware studies: the RPCT Project
  - Context
  - Target Technologies
  - Project Development
- The MDC tool
  - Approach
  - Baseline Functionality and Extensions
- Contexts of application
  - Neural Signal Decoding
  - HEVC Interpolation Filters
- Final Remarks

Contexts of application

What kinds of applications can be combined with MDC?

Contexts of application

What kinds of applications can be combined with MDC?

1. **Different applications with common computational operations**: it is achieved by considering applications from the **same application field** or **small actor granularities**.

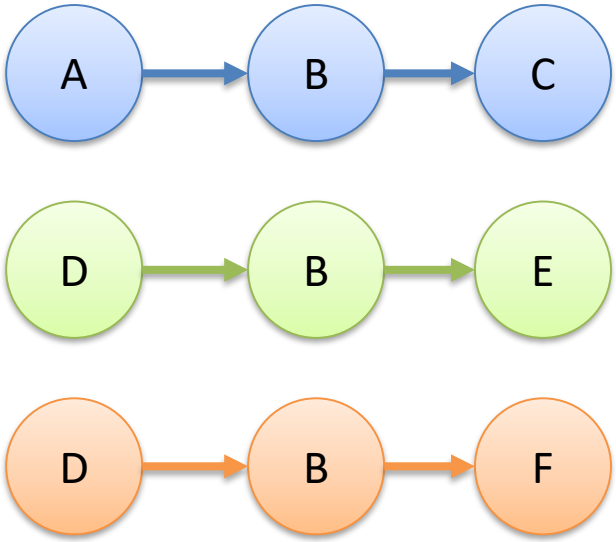

Contexts of application

What kinds of applications can be combined with MDC?

1. **Different applications with common computational operations**: it is achieved by considering applications from the **same application field** or **small actor granularities**.

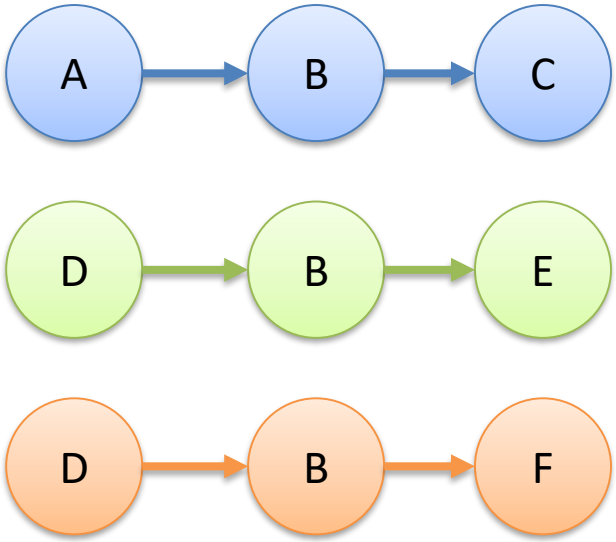

2. **Different working points of the same applications** obtained through several strategies (e.g. **actor parallelization**, actor variants, granularity modification, **approximate computing**, ...)

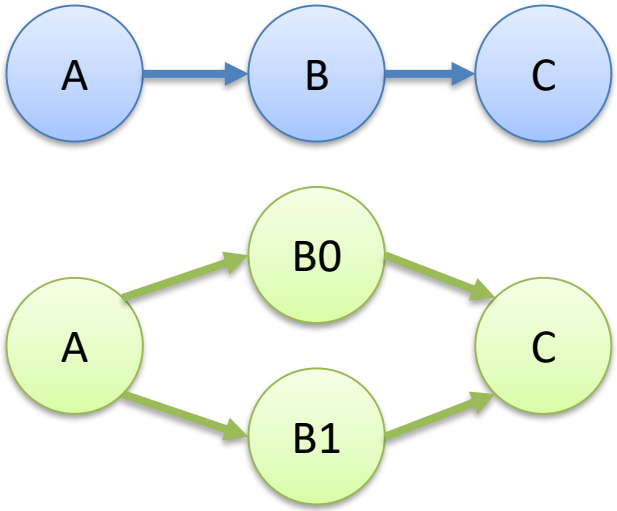

Contexts of application

What kinds of applications can be combined with MDC?

1. **Different applications with common computational operations**: it is achieved by considering applications from the **same application field** or **small actor granularities**.

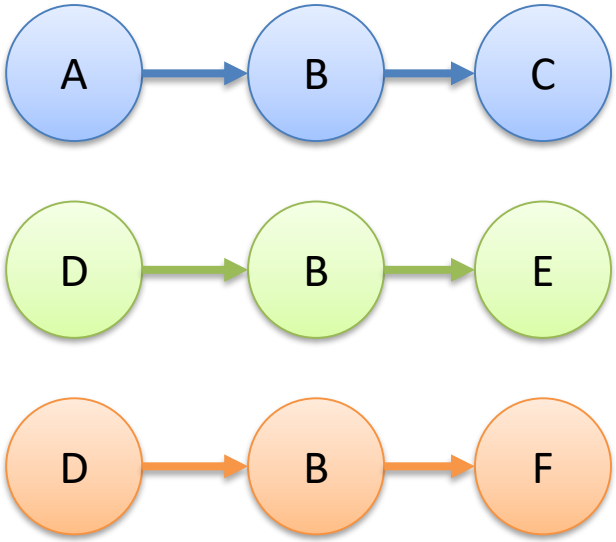

EXAMPLE: Neural Signal Decoding

2. **Different working points of the same applications** obtained through several strategies (e.g. **actor parallelization**, actor variants, granularity modification, **approximate computing**, ...)

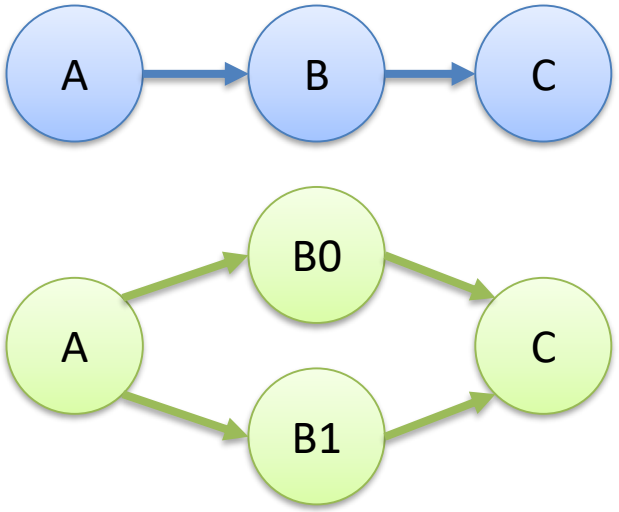

EXAMPLE: HEVC interpolation filters

Neural Signal Decoding

Resource Optimization

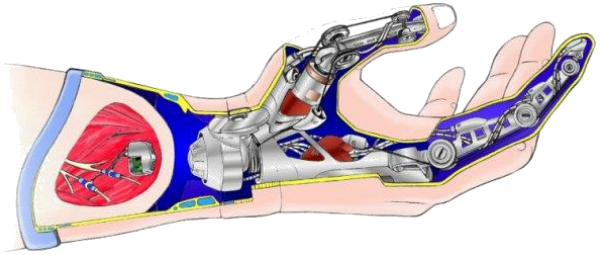

Implantable Devices: strict **area** & **power** requirements

Neural Signal Decoding
Resource Optimization

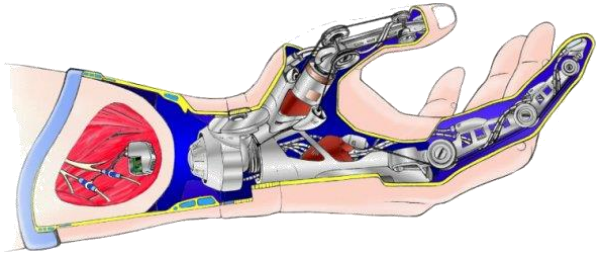

Implantable Devices: strict **area** & **power** requirements

Neural Signal Decoding:

- Fast
- Low Area
- Low Power

D. Pani, et al., «Real-time processing of tflife neural signals on embedded dsp platforms: A case study» *Neural Engineering*, 2011.

Neural Signal Decoding
Resource Optimization

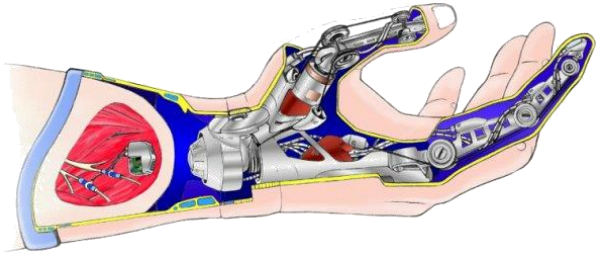

Implantable Devices: strict **area** & **power** requirements

Neural Signal Decoding:

- Fast
- Low Area
- Low Power

D. Pani, et al., «Real-time processing of tflife neural signals on embedded dsp platforms: A case study» *Neural Engineering*, 2011.

MDC can be used to build the accelerators compliant to those constraints.

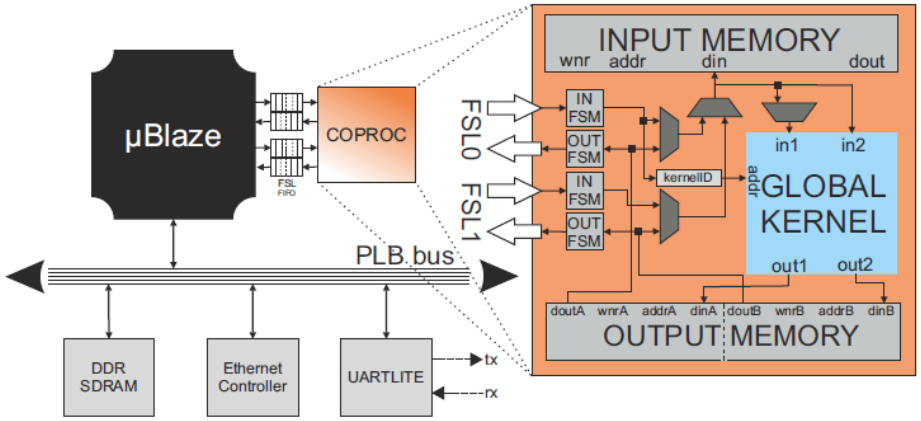

Neural Signal Decoding

Resource Optimization

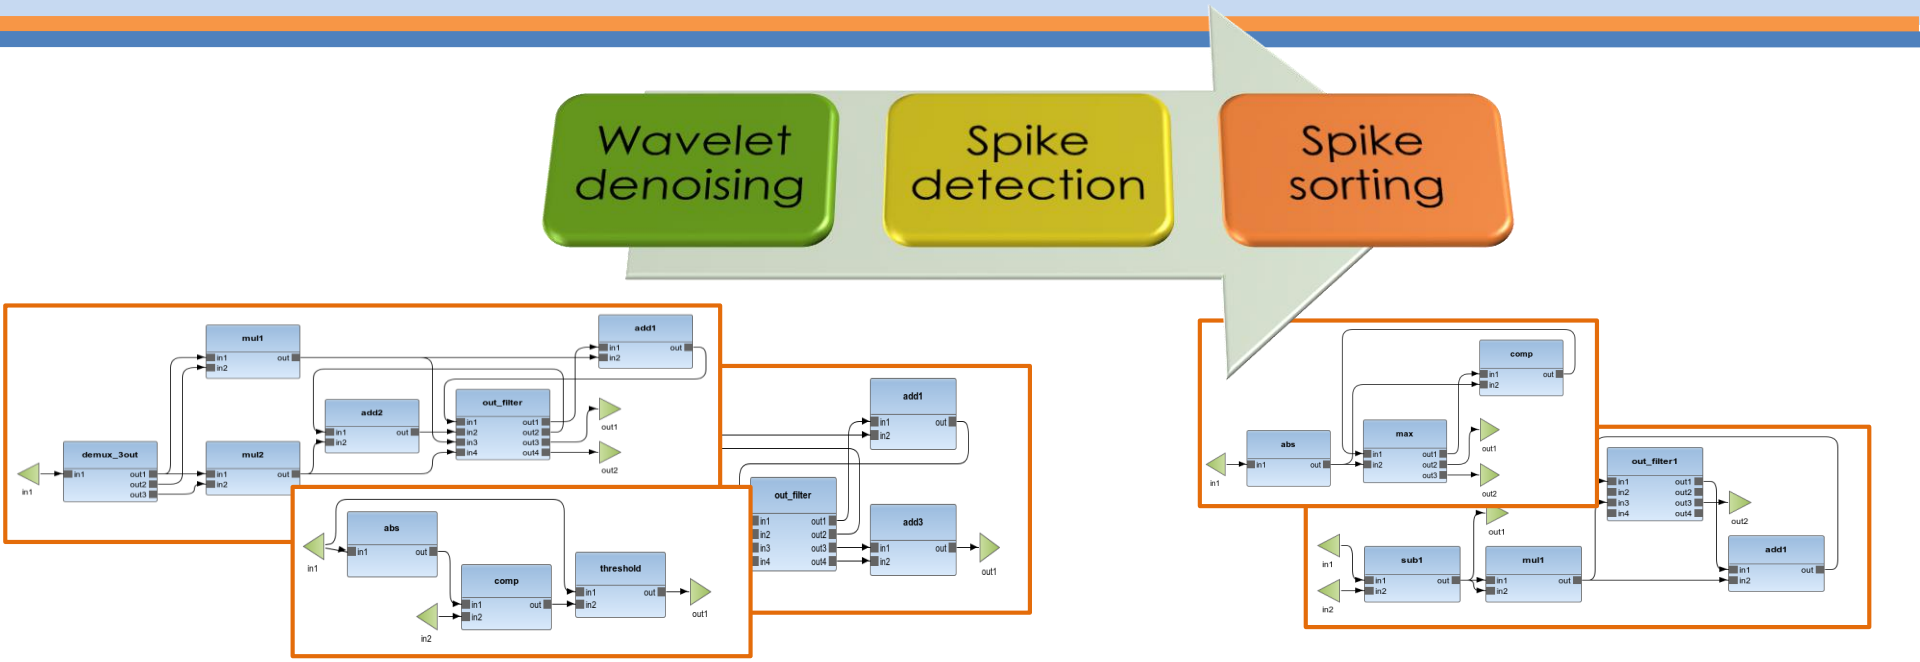

	# actors	#sbox
12 networks (dec_filter, Thr, rec_filter, NEO, idx_max_abs, Avg, sqr_sum, weight_mul, dot_prod, idx_max, sync_avg, sync_wavg)	46	0
MDC network	14	86

Neural Signal Decoding

Resource Optimization

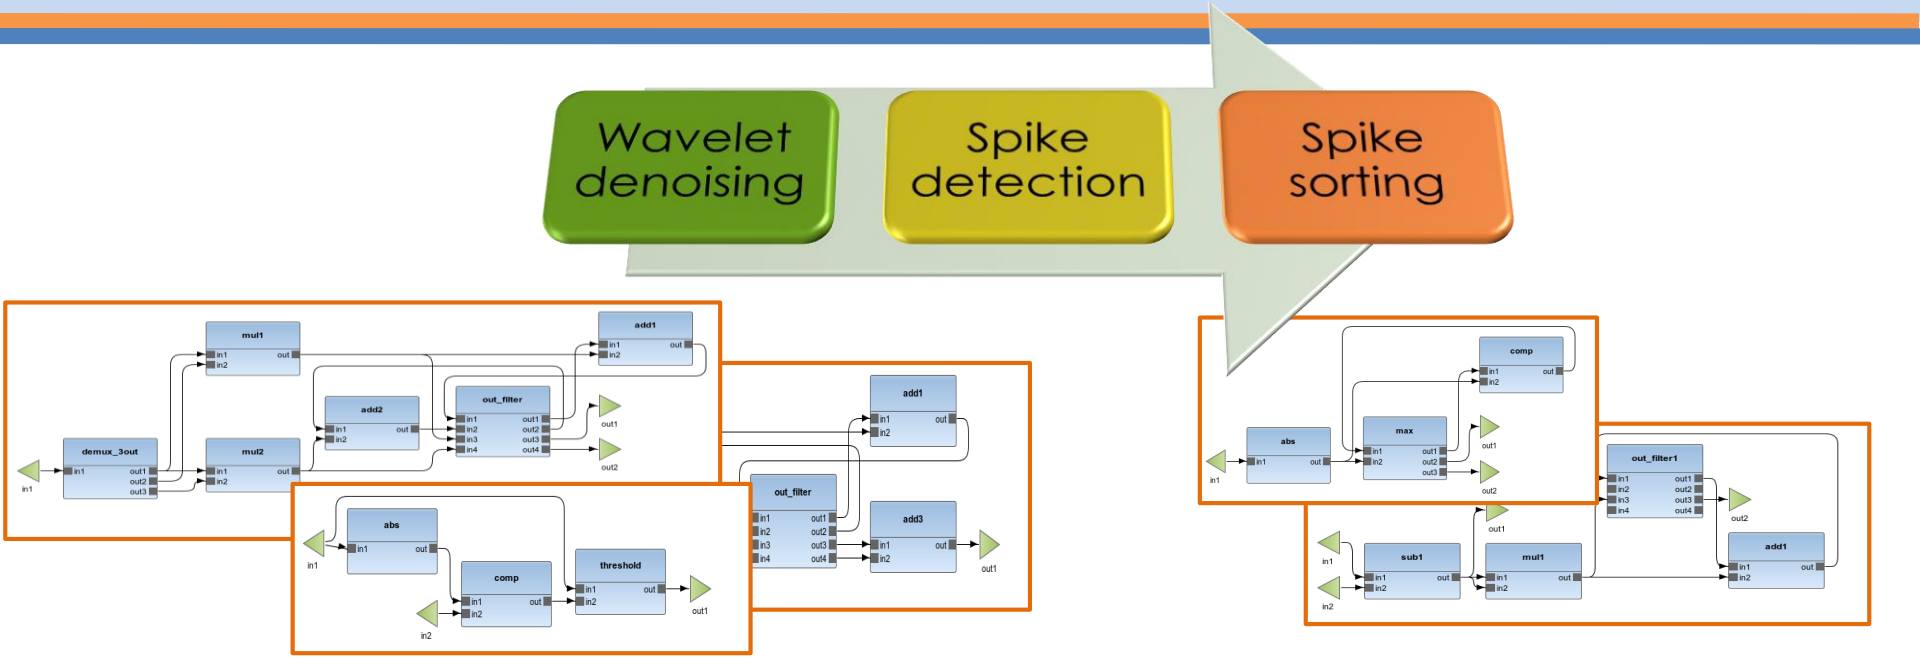

	# actors	#sbox
12 networks (dec_filter, Thr, rec_filter, NEO, idx_max_abs, Avg, sqr_sum, weight_mul, dot_prod, idx_max, sync_avg, sync_wavg)	46	0
MDC network	14	86

Neural Signal Decoding

Resource Optimization

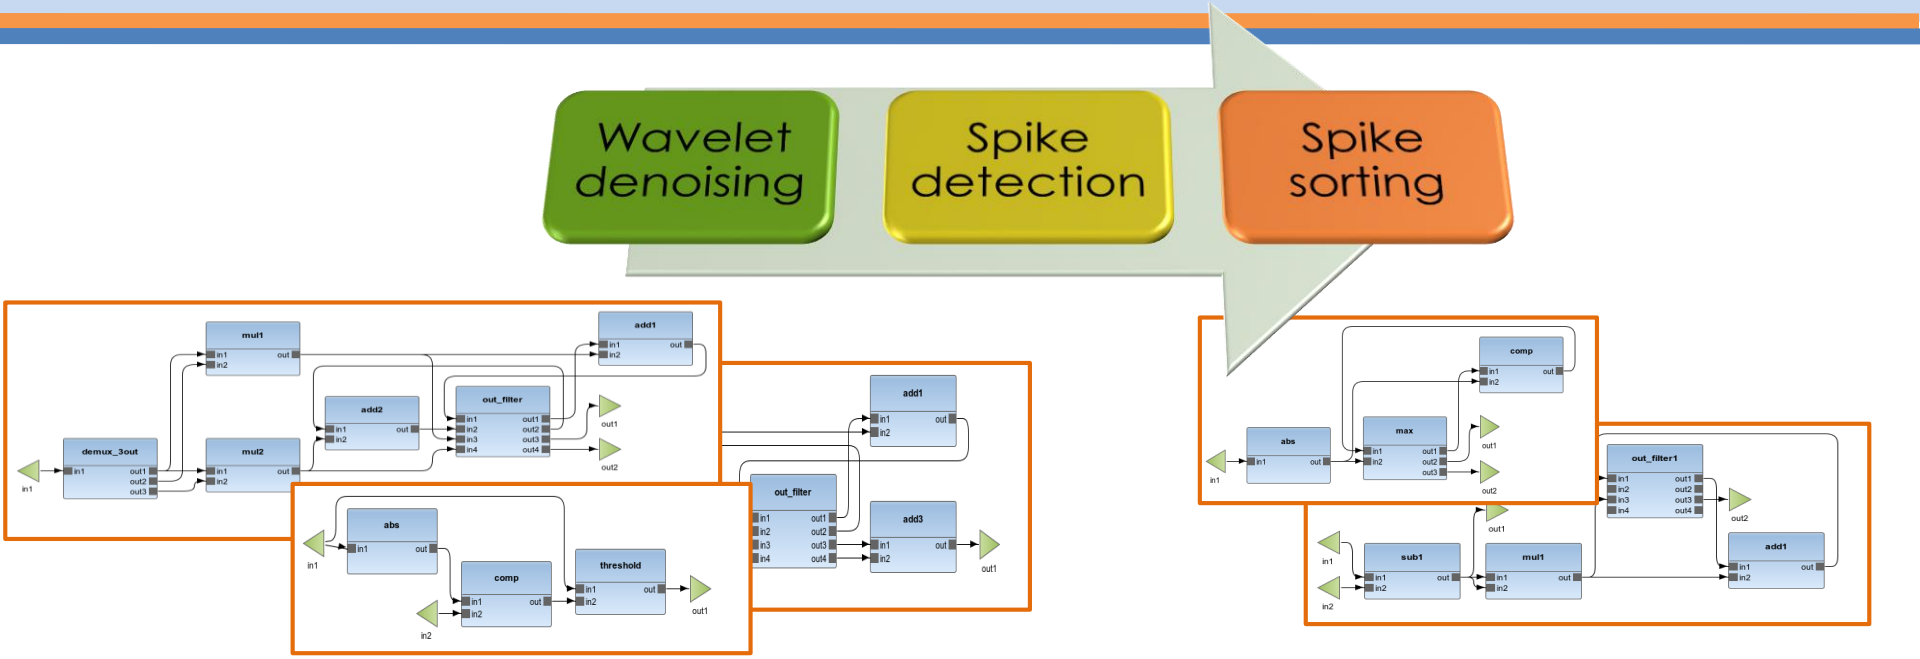

	# actors	#sbox
12 networks (dec_filter, Thr, rec_filter, NEO, idx_max_abs, Avg, sqr_sum, weight_mul, dot_prod, idx_max, sync_avg, sync_wavg)	46	0
MDC network	14	86

Neural Signal Decoding

Resource Optimization

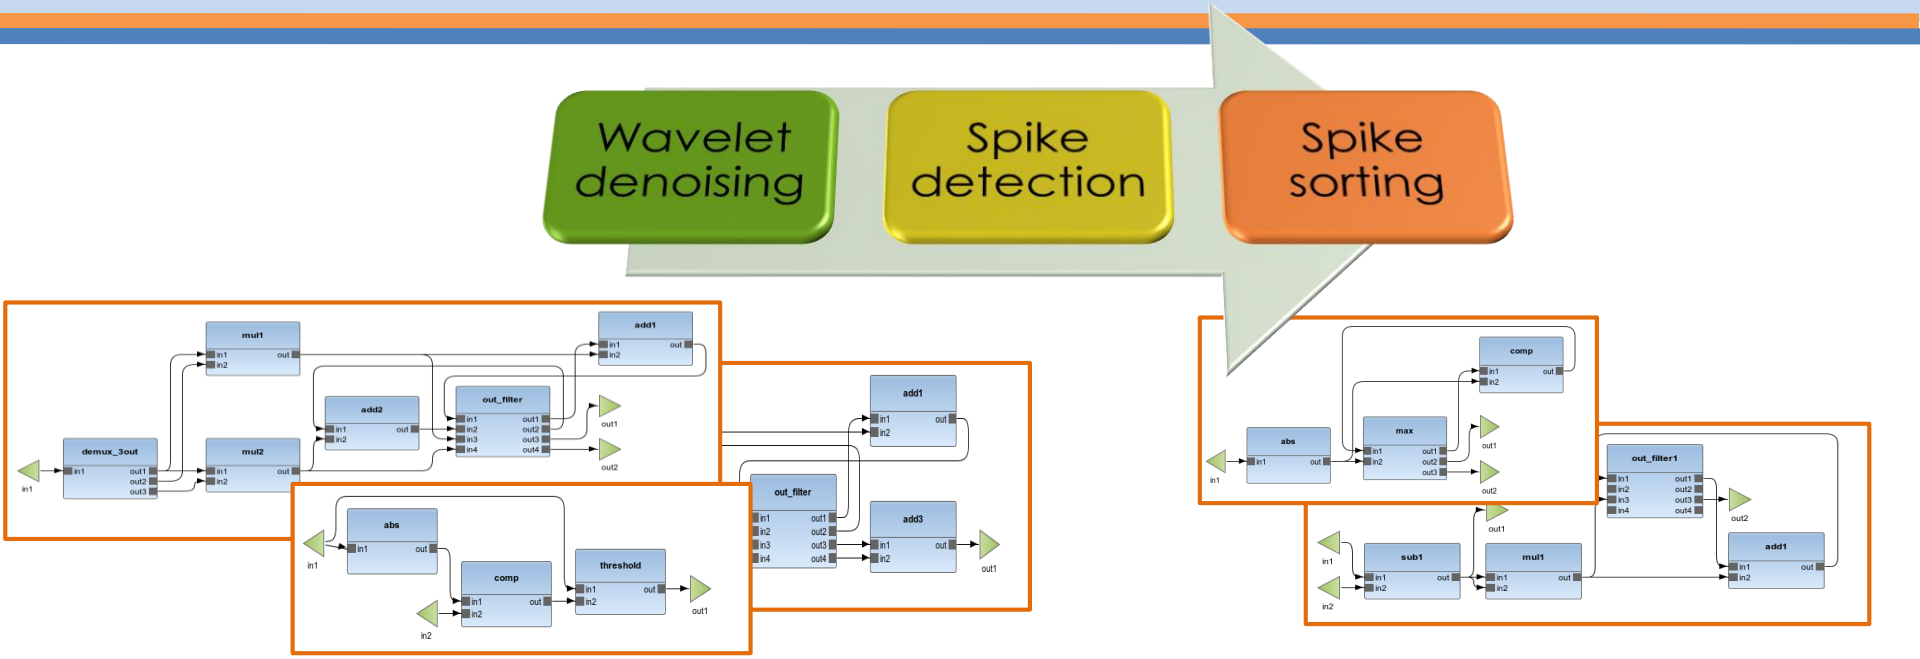

	# actors	#sbox
12 networks (dec_filter, Thr, rec_filter, NEO, idx_max_abs, Avg, sqr_sum, weight_mul, dot_prod, idx_max, sync_avg, sync_wavg)	46	0
MDC network	14	86

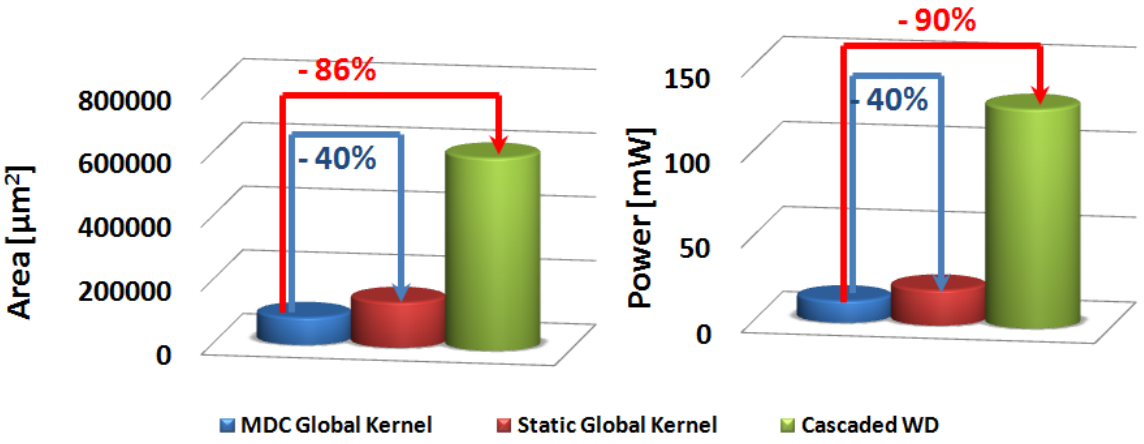

HEVC Interpolation Filters

Multiple Working Points

- **Approximate Computing:** trading a controlled quality degradation (# taps) for an increased energy efficiency
- **Software Implementation**: Erwan Raffin, et al., "*Low power HEVC software decoder for mobile devices*", JRTIP 12(2): 495-507 (2016)

HEVC Interpolation Filters

Multiple Working Points

- **Approximate Computing:** trading a controlled quality degradation (# taps) for an increased energy efficiency
- **Software Implementation**: Erwan Raffin, et al., "*Low power HEVC software decoder for mobile devices*", JRTIP 12(2): 495-507 (2016)

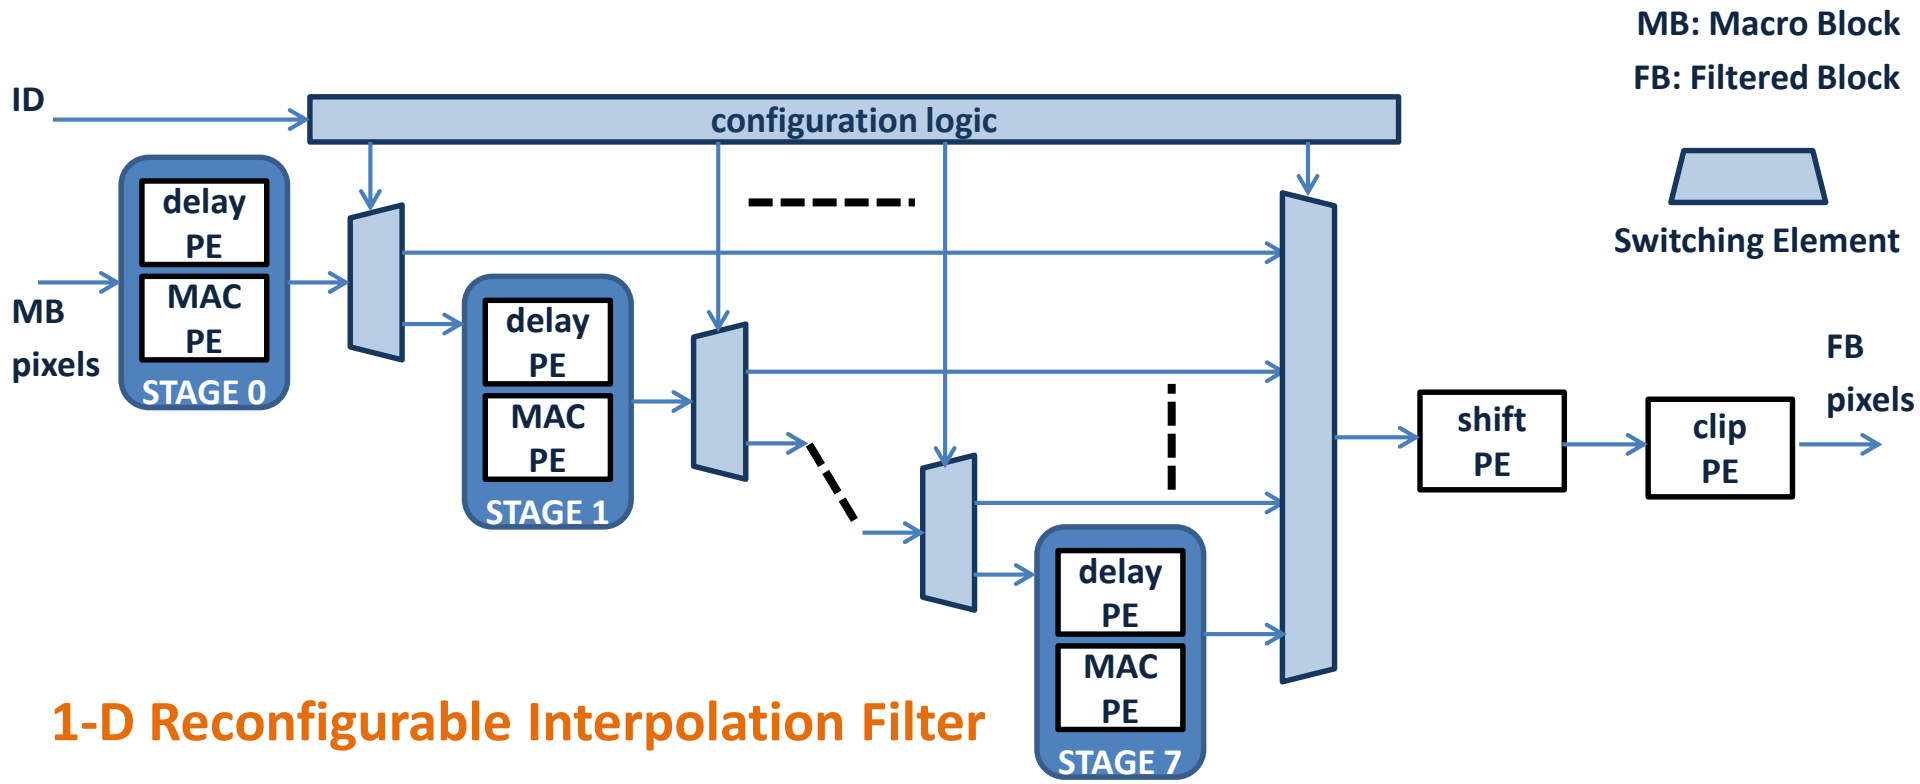

1-D Reconfigurable Interpolation Filter

HEVC Interpolation Filters

Multiple Working Points

design @200 MHz Xilinx XC7Z020	LUT	FF	BRAM	DSP	Fmax [MHz]	tap	dP (Vivado) [mW]	dE [µJ]	time per block [cycles]	# interpolated pixels in a fixed time
legacy_luma	212	37	4	16	213	8	11	0.248	460	57957
reconf_luma (vs legacy %)	582 **(+175%)**	85 (+130%)	4 (+0%)	16 (+0%)	200 **(-6%)**	8	12 (+9%)	0.270 (+9%)	460 (+0%)	57957 (+0%)
						7	11 (+0%)	0.245 (-1%)	395 (-14%)	59033 (+2%)
						5	10 (-9%)	0.217 (-12%)	265 (-42%)	61191 (+6%)
						3	10 **(-9%)**	0.211 **(-15%)**	135 **(-71%)**	63357 **(+9%)**
legacy_chroma	163	33	2	8	217	4	9	0.053	107	14753
reconf_chroma (vs legacy %)	383 **(+135%)**	65 (+97%)	2 (+0%)	8 (+0%)	200 **(-12%)**	4	9 (+0%)	0.053 (+0%)	107 (+0%)	14753 (+0%)
						3	8 (-11%)	0.045 (-13%)	73 (-32%)	15293 (+4%)
						2	6 **(-33%)**	0.033 **(-37%)**	39 **(-64%)**	15835 **(+7%)**

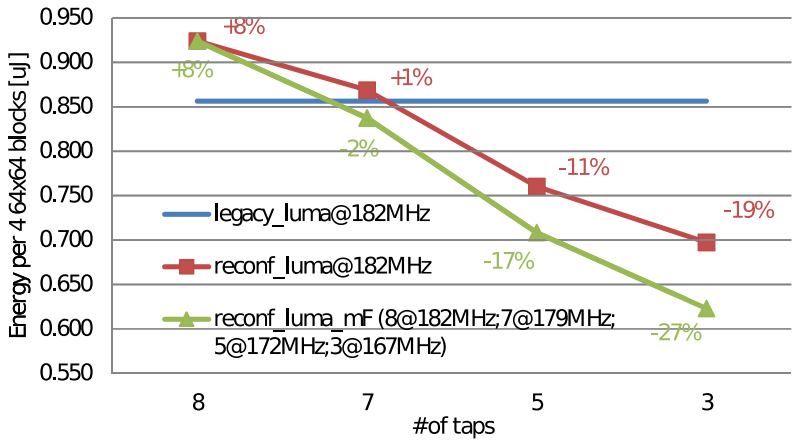

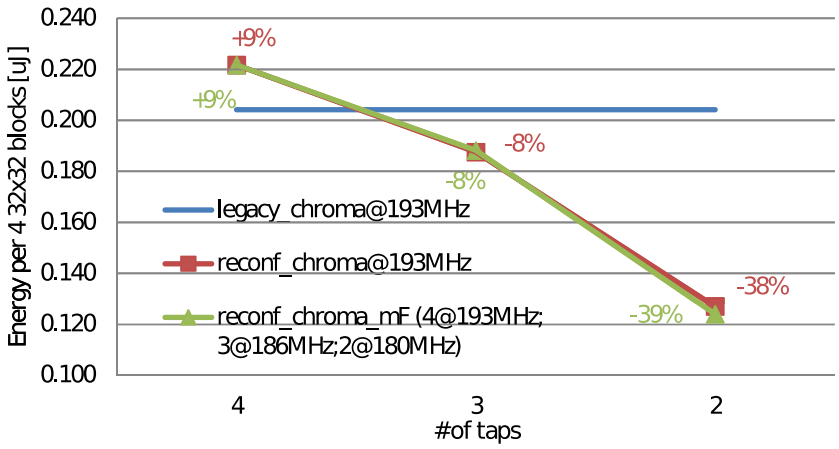

C. Sau et al. <<*Challenging the Best HEVC Fractional Pixel FPGA Interpolators with Reconfigurable and Multi-frequency Approximate Computing.*>> IEEE Embedded Systems Letters, 9 (3), pp. 65-68, 2017, ISSN: 1943-0663.

HEVC Interpolation Filters

Multiple Working Points

design @200 MHz Xilinx XC7Z020	LUT	FF	BRAM	DSP	Fmax [MHz]	tap	dP (Vivado) [mW]	dE [µJ]	time per block [cycles]	# interpolated pixels in a fixed time
legacy_luma	212	37	4	16	213	8	11	0.248	460	57957
reconf_luma (vs legacy %)	582 **(+175%)**	85 (+130%)	4 (+0%)	16 (+0%)	200 **(-6%)**	8	12 (+9%)	0.270 (+9%)	460 (+0%)	57957 (+0%)
						7	11 (+0%)	0.245 (-1%)	395 (-14%)	59033 (+2%)
						5	10 (-9%)	0.217 (-12%)	265 (-42%)	61191 (+6%)
						3	10 **(-9%)**	0.211 **(-15%)**	135 **(-71%)**	63357 **(+9%)**
legacy_chroma	163	33	2	8	217	4	9	0.053	107	14753
reconf_chroma (vs legacy %)	383 **(+135%)**	65 (+97%)	2 (+0%)	8 (+0%)	200 **(-12%)**	4	9 (+0%)	0.053 (+0%)	107 (+0%)	14753 (+0%)
						3	8 (-11%)	0.045 (-13%)	73 (-32%)	15293 (+4%)
						2	6 **(-33%)**	0.033 **(-37%)**	39 **(-64%)**	15835 **(+7%)**

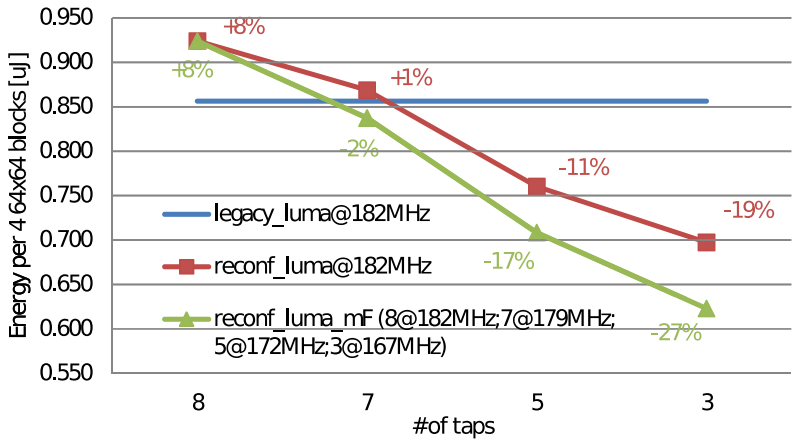

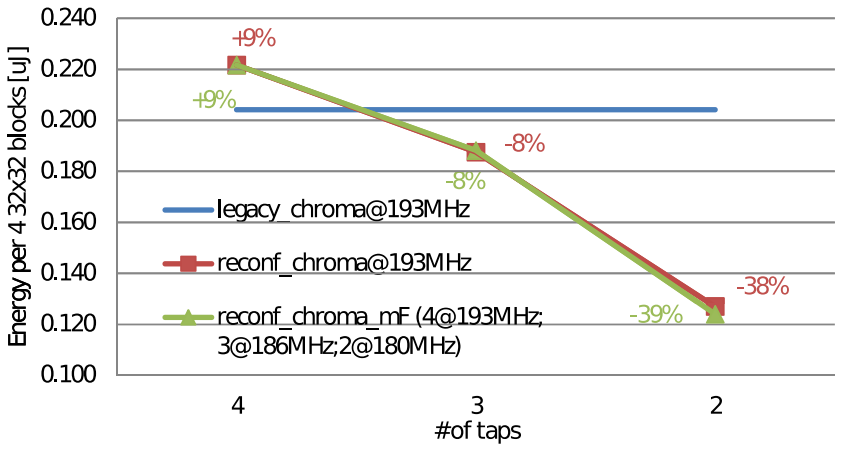

C. Sau et al. <<*Challenging the Best HEVC Fractional Pixel FPGA Interpolators with Reconfigurable and Multi-frequency Approximate Computing.*>> IEEE Embedded Systems Letters, 9 (3), pp. 65-68, 2017, ISSN: 1943-0663.

HEVC Interpolation Filters

Multiple Working Points

design @200 MHz Xilinx XC7Z020	LUT	FF	BRAM	DSP	Fmax [MHz]	tap	dP (Vivado) [mW]	dE [µJ]	time per block [cycles]	# interpolated pixels in a fixed time
legacy_luma	212	37	4	16	213	8	11	0.248	460	57957
						8	12 (+9%)	0.270 (+9%)	460 (+0%)	57957 (+0%)
reconf_luma (vs legacy %)	582 **(+175%)**	85 (+130%)	4 (+0%)	16 (+0%)	200 **(-6%)**	7	11 (+0%)	0.245 (-1%)	395 (-14%)	59033 (+2%)
						5	10 (-9%)	0.217 (-12%)	265 (-42%)	61191 (+6%)
						3	10 **(-9%)**	0.211 **(-15%)**	135 **(-71%)**	63357 **(+9%)**
legacy_chroma	163	33	2	8	217	4	9	0.053	107	14753
						4	9 (+0%)	0.053 (+0%)	107 (+0%)	14753 (+0%)
reconf_chroma (vs legacy %)	383 **(+135%)**	65 (+97%)	2 (+0%)	8 (+0%)	200 **(-12%)**	3	8 (-11%)	0.045 (-13%)	73 (-32%)	15293 (+4%)
						2	6 **(-33%)**	0.033 **(-37%)**	39 **(-64%)**	15835 **(+7%)**

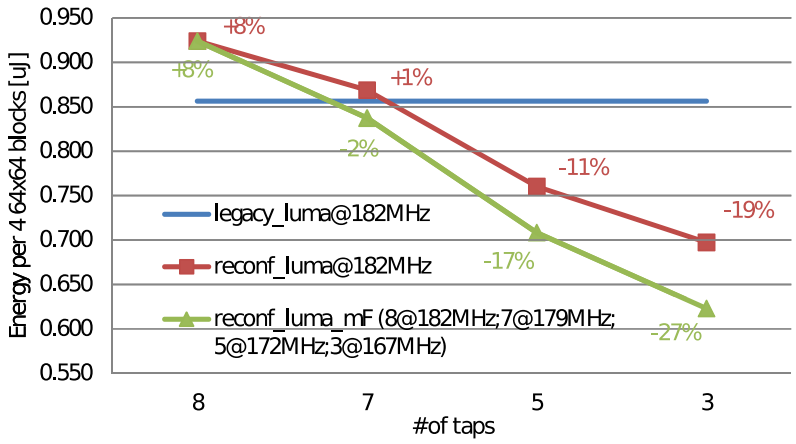

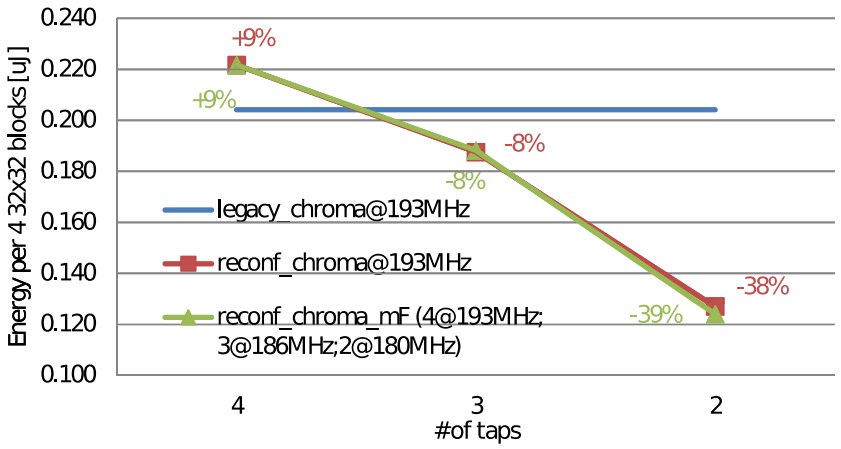

C. Sau et al. <<*Challenging the Best HEVC Fractional Pixel FPGA Interpolators with Reconfigurable and Multi-frequency Approximate Computing.*>> IEEE Embedded Systems Letters, 9 (3), pp. 65-68, 2017, ISSN: 1943-0663.

HEVC Interpolation Filters

Multiple Working Points

design @200 MHz Xilinx XC7Z020	LUT	FF	BRAM	DSP	Fmax [MHz]	tap	dP (Vivado) [mW]	dE [µJ]	time per block [cycles]	# interpolated pixels in a fixed time
legacy_luma	212	37	4	16	213	8	11	0.248	460	57957
						8	12 (+9%)	0.270 (+9%)	460 (+0%)	57957 (+0%)
reconf_luma (vs legacy %)	582 **(+175%)**	85 (+130%)	4 (+0%)	16 (+0%)	200 **(-6%)**	7	11 (+0%)	0.245 (-1%)	395 (-14%)	59033 (+2%)
						5	10 (-9%)	0.217 (-12%)	265 (-42%)	61191 (+6%)
						3	10 **(-9%)**	0.211 **(-15%)**	135 **(-71%)**	63357 **(+9%)**
legacy_chroma	163	33	2	8	217	4	9	0.053	107	14753
						4	9 (+0%)	0.053 (+0%)	107 (+0%)	14753 (+0%)
reconf_chroma (vs legacy %)	383 **(+135%)**	65 (+97%)	2 (+0%)	8 (+0%)	200 **(-12%)**	3	8 (-11%)	0.045 (-13%)	73 (-32%)	15293 (+4%)
						2	6 **(-33%)**	0.033 **(-37%)**	39 **(-64%)**	15835 **(+7%)**

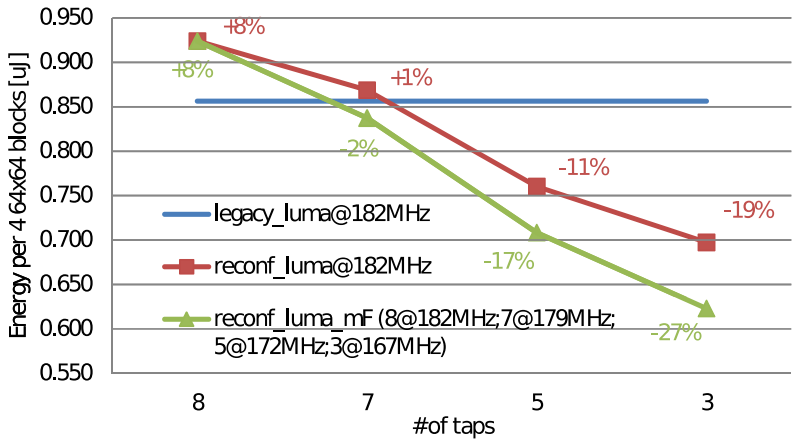

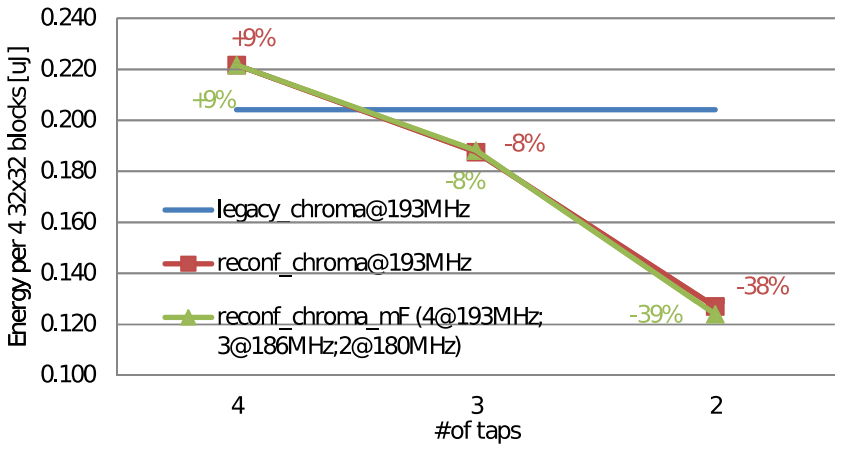

C. Sau et al. <<*Challenging the Best HEVC Fractional Pixel FPGA Interpolators with Reconfigurable and Multi-frequency Approximate Computing.*>> IEEE Embedded Systems Letters, 9 (3), pp. 65-68, 2017, ISSN: 1943-0663.

HEVC Interpolation Filters

Multiple Working Points

design @200 MHz Xilinx XC7Z020	LUT	FF	BRAM	DSP	Fmax [MHz]	tap	dP (Vivado) [mW]	dE [µJ]	time per block [cycles]	# interpolated pixels in a fixed time
legacy_luma	212	37	4	16	213	8	11	0.248	460	57957
						8	12 (+9%)	0.270 (+9%)	460 (+0%)	57957 (+0%)
reconf_luma (vs legacy %)	582 **(+175%)**	85 (+130%)	4 (+0%)	16 (+0%)	200 **(-6%)**	7	11 (+0%)	0.245 (-1%)	395 (-14%)	59033 (+2%)
						5	10 (-9%)	0.217 (-12%)	265 (-42%)	61191 (+6%)
						3	10 **(-9%)**	0.211 **(-15%)**	135 **(-71%)**	63357 **(+9%)**
legacy_chroma	163	33	2	8	217	4	9	0.053	107	14753
						4	9 (+0%)	0.053 (+0%)	107 (+0%)	14753 (+0%)
reconf_chroma (vs legacy %)	383 **(+135%)**	65 (+97%)	2 (+0%)	8 (+0%)	200 **(-12%)**	3	8 (-11%)	0.045 (-13%)	73 (-32%)	15293 (+4%)
						2	6 **(-33%)**	0.033 **(-37%)**	39 **(-64%)**	15835 **(+7%)**

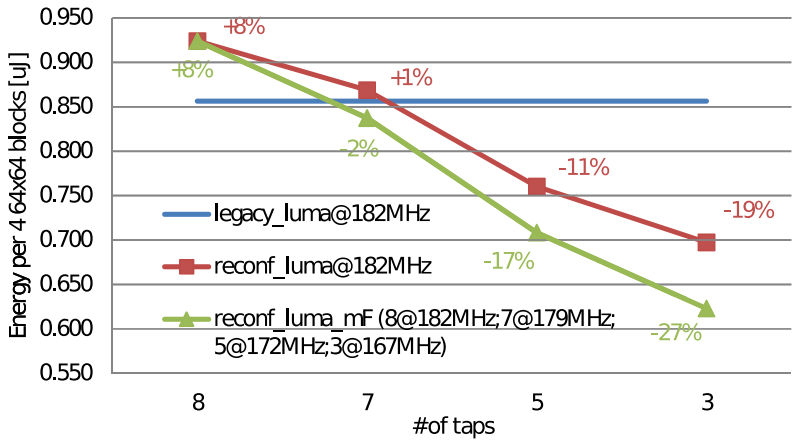

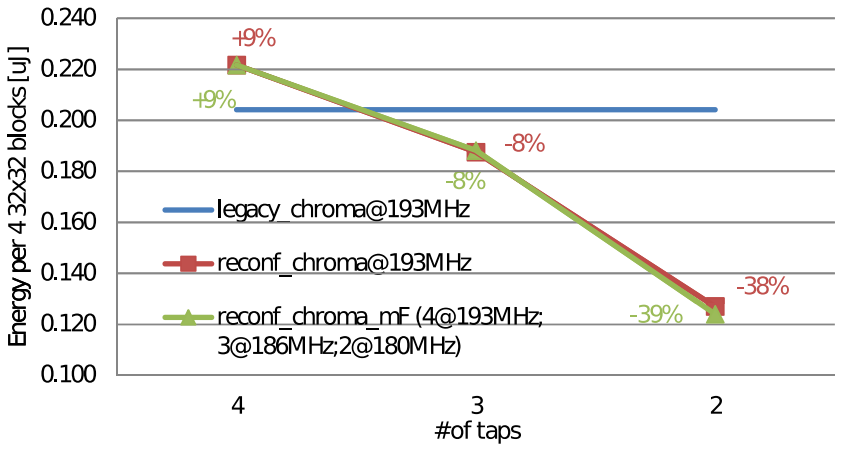

C. Sau et al. <<*Challenging the Best HEVC Fractional Pixel FPGA Interpolators with Reconfigurable and Multi-frequency Approximate Computing.*>> IEEE Embedded Systems Letters, 9 (3), pp. 65-68, 2017, ISSN: 1943-0663.

Outline

- The origins of our dataflow to hardware studies: the RPCT Project
  - Context
  - Target Technologies
  - Project Development
- The MDC tool
  - Approach
  - Baseline Functionality and Extensions
- Contexts of application
  - Neural Signal Decoding
  - HEVC Interpolation Filters
- Final Remarks

Conclusion and Future Plan

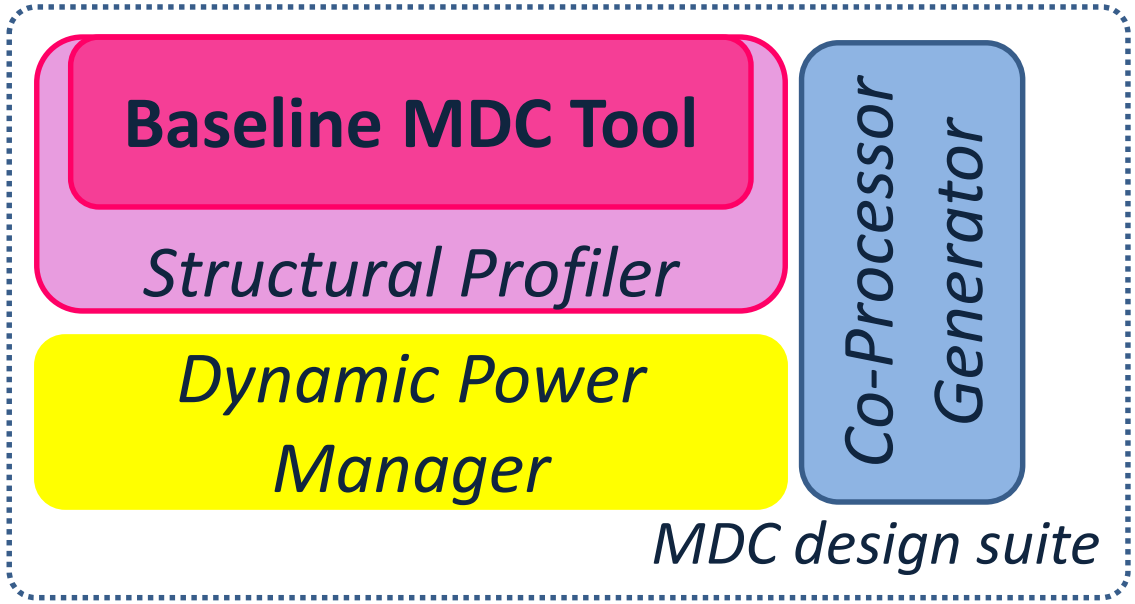

HW/SW Partitioning

Advanced HLS

.....& MORE

The RPCT project (2012-2015) has been funded by Sardinian Regional Government (L.R. 7/2007, CRP-18324).
http://sites.unica.it/rpct/

Thanks To …

EU Commission for funding the ***CERBERO*** (*Cross-layer modEl-based fRamework for multi-oBjective dEsign of Reconfigurable systems in unceRtain hybRid envirOnments*) project as part of the H2020 Programme under grant agreement No 732105.

Coordinator:

Michal Masin (IBM), michaelm@il.ibm.com

Scientific Coordinator:

Francesca Palumbo (UniSS), fpalumbo@uniss.it

Innovation Manager:

Katiuscia Zedda (Abinsula), katiuscia.zedda@abinsula.com

Dissemination-Communication Manager:

Francesco Regazzoni (USI), francesco.regazzoni@usi.ch

www.cerbero-h2020.eu

info@cerbero-h2020.eu

@CERBERO_h2020

CERBERO

Some References

1. Sau C, et al., "*Challenging the Best HEVC Fractional Pixel FPGA Interpolators With Reconfigurable and Multi-frequency Approximate Computing*", IEEE ESL 2017
2. Palumbo F., et al., "*Power-Awarness in Coarse-Grained Reconfigurable Multi-Functional Architectures: a Dataflow Based Strategy*", JSPS 2017
3. Sau C., et al., "*Automated Design Flow for Multi-Functional Dataflow-Based Platforms*", JSPS 2015
4. Palumbo F., et al., "*The multi-dataflow composer tool: generation of on-the-fly reconfigurable platforms*", JRTIP 2014

Exploiting Dataflows for Reconfigurable Hardware Accelerators

Francesca Palumbo, Claudio Rubattu,
Carlo Sau, Tiziana Fanni, Luigi Raffo

Rennes, 12-14 December 2017