

Monsef Makhlouf
Ecosystem Manager

Enabling Intelligent Hearables and Smart Sensors: GreenWaves' Ultra Low Power Processors

GreenWaves Technologies

- Fabless semiconductor startup founded in 2014
- We design and sell extreme performance processors for energy constrained devices
- 46 people, mostly in Grenoble, France
- Offices in Bologna, Italy and Shanghai, China. Global sales footprint.
- GAP9 is designed and produced in the EU



Best hardware product
Embedded World 2023



Embedded Technologies Award 2023
Les Assises de l'Embarqué




Cool Vendors in AI Semiconductors,
Alan Priestley, Saniye Alaybeyi, April
29, 2019.



GAP8

In production since 2020
one of the very first
commercially available
RISC-V processor and
AI microcontroller



GAP9

First design-wins,
Ramping up
volume
production




Where is embedded AI energy constrained?

Hearables & Wearables



Microphone
Speaker
IMU
Vibration

Energy constrained
Internet of Things



IR and visible light camera
Microphone
Radar
Vibration

Embedded AI/DSP

Adaptive ANC

Environmental Noise Reduction

Immersive spatial audio

Adaptive transparency mode

Hearing Enhancement

Source separation

Person / object detection

Face detection / identification


Speaker detection / identification

Abnormal sound detection

and more...

GAP9 in a nutshell

- + Real time, streamed autonomous time domain digital signal processing - Smart Filtering Unit - SFU
- + General digital signal processing - Cluster
- + Highly flexible neural network processing - Cluster / NE16
- + Ultra low power - Energy constrained devices
- + Highly Flexible - Easy to program versus alternatives



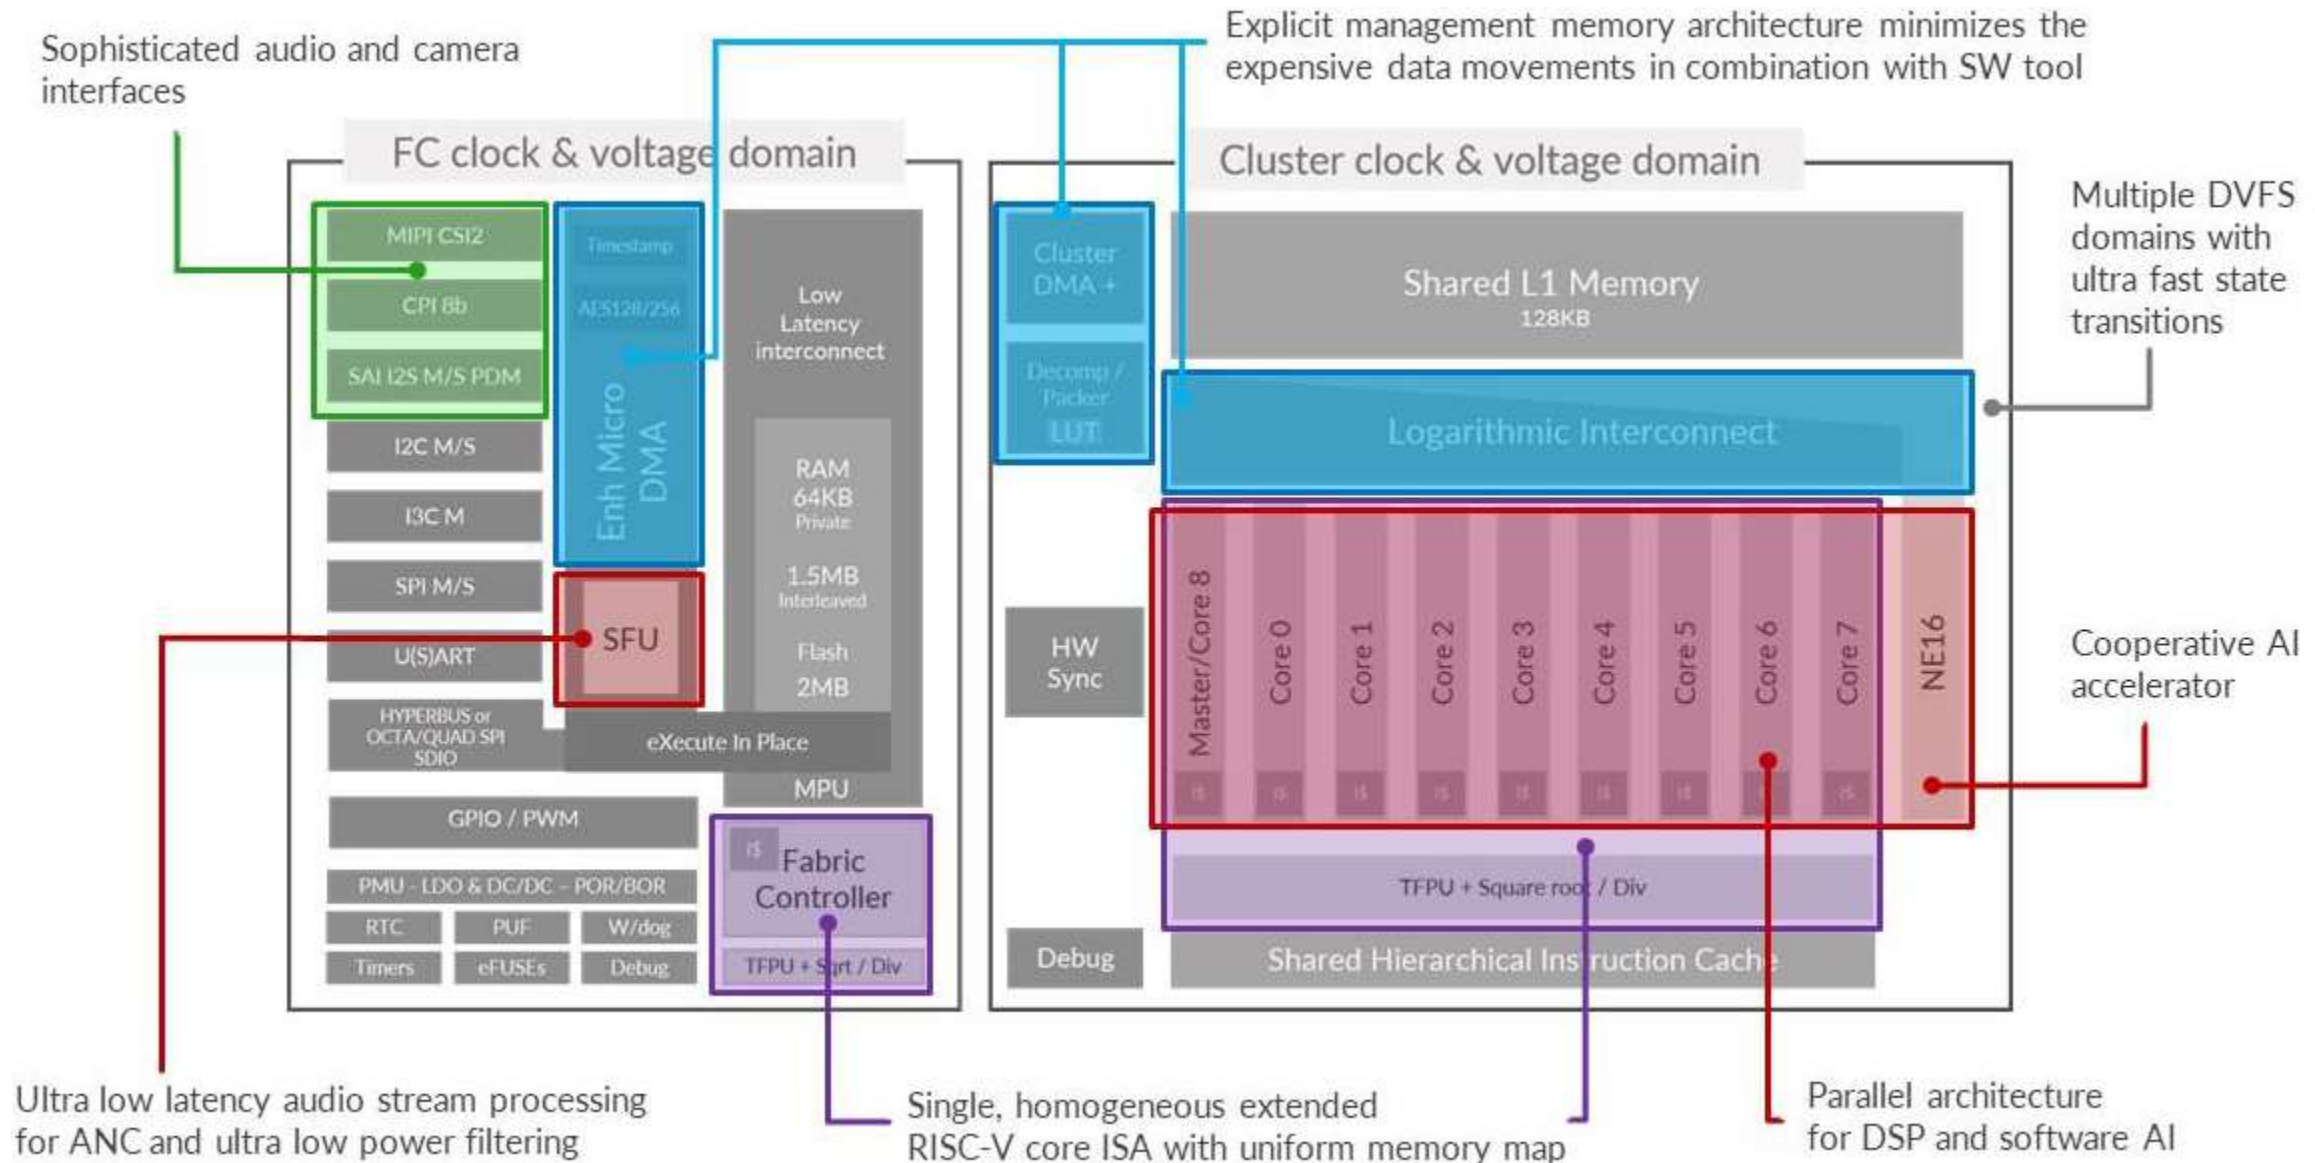
Production chip
available now

Hearables - Wearables - IoT Sensors

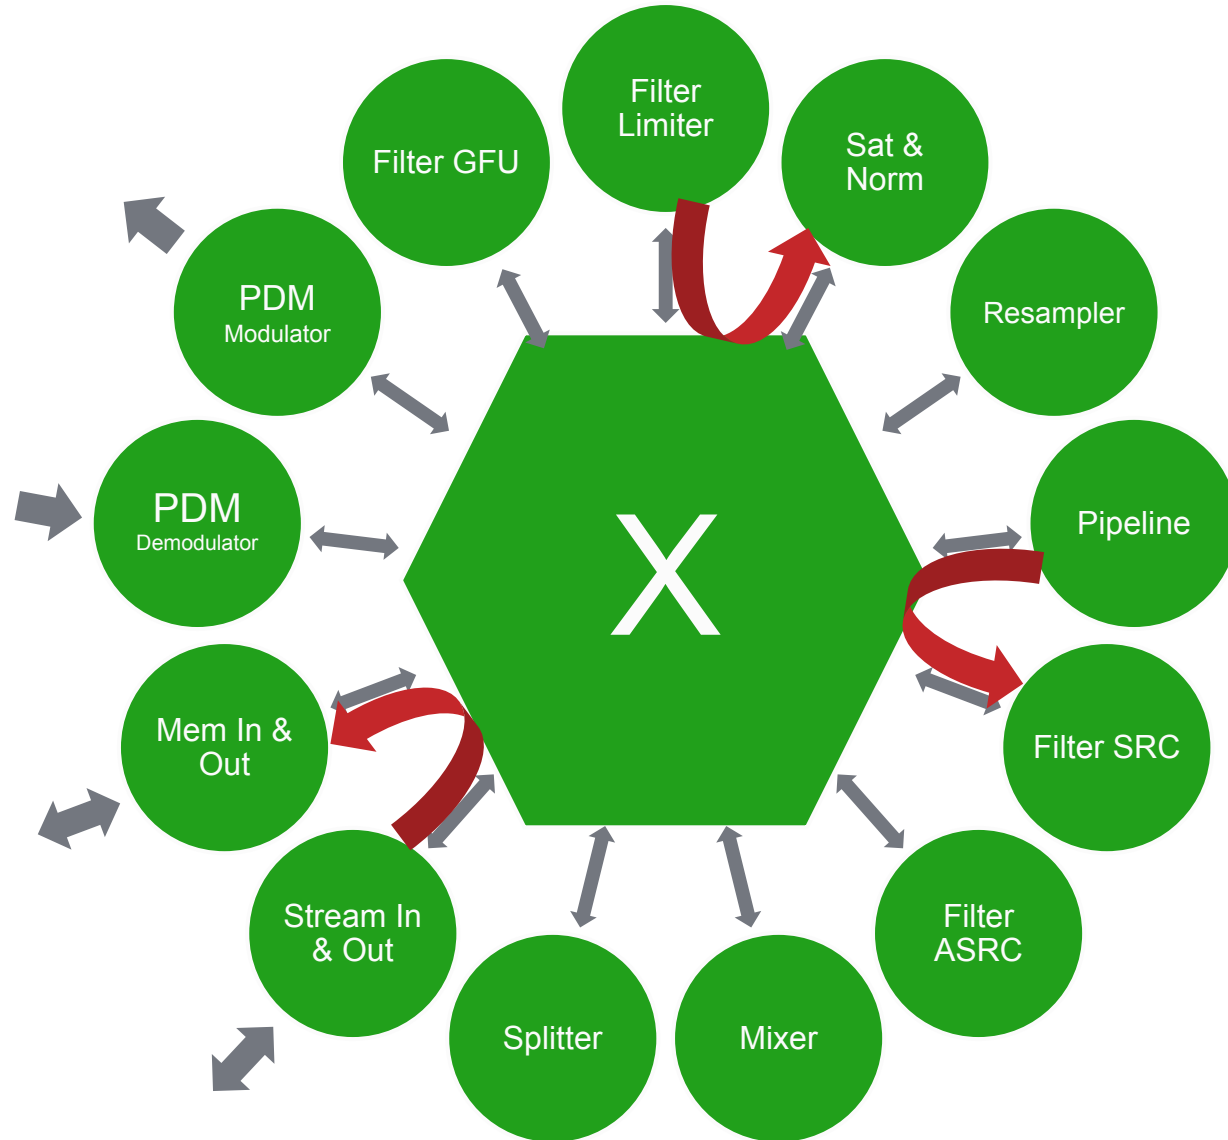


Designed and
produced in the EU

GAP9 Highlights



Configurable, Sample by Sample, Audio Filtering



What:

- Sample by sample, stream based autonomous time domain filtering
- Highly configurable hardware blocks, including 13 different filter patterns, that can be configured to form a data flow graph
- SFU can support multiple graphs and dynamic graph parameter update

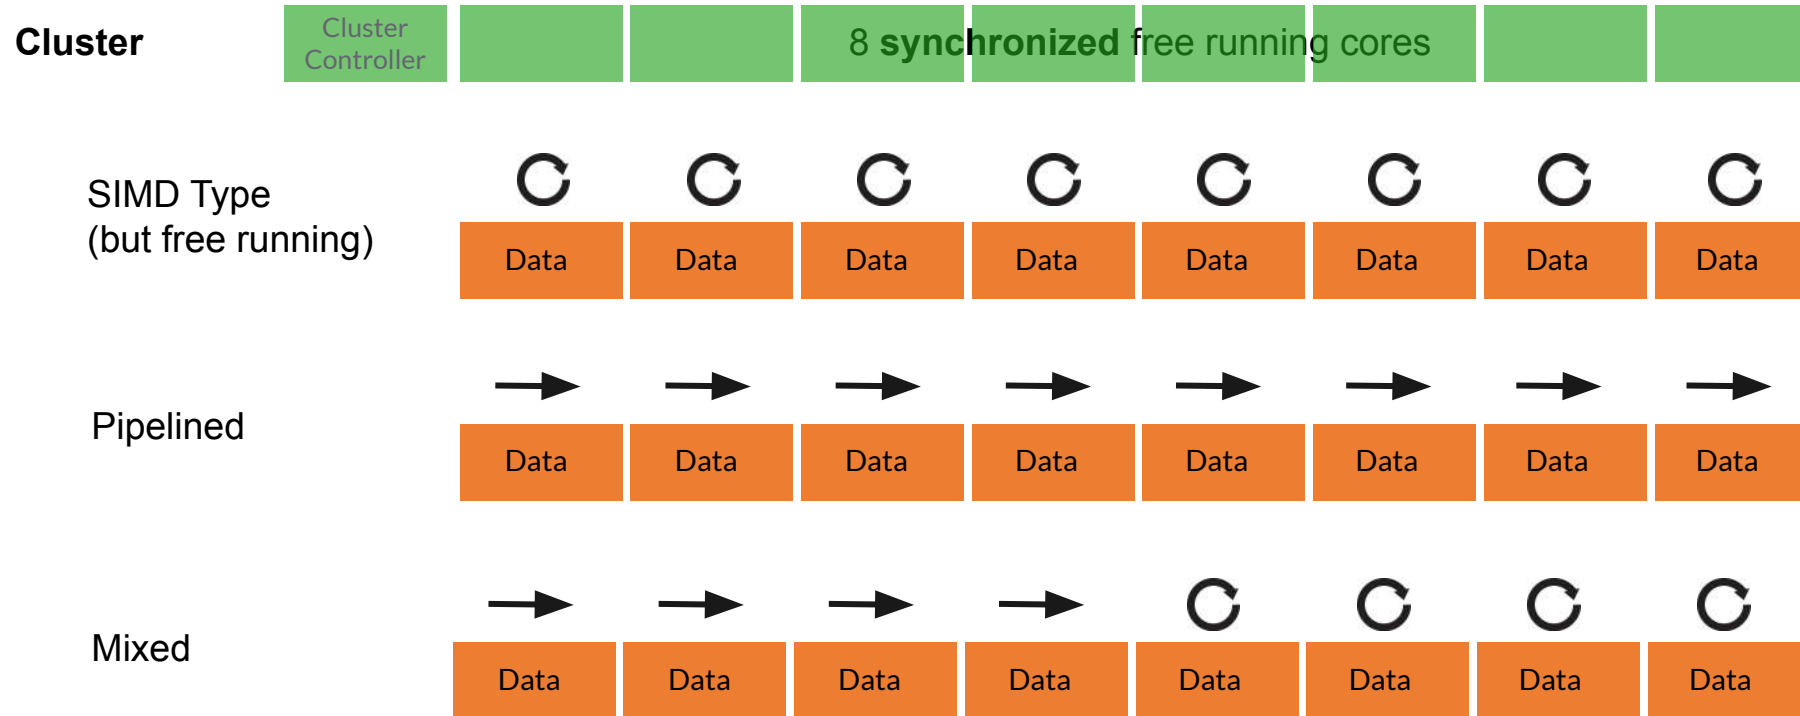
Why:

- Minimum latency from input to output
Result: 1.3 μ S ANC structural latency
- Eliminate overhead that comes with pure SW approach, (e.g. instruction load).
Result: 2.25mW / ANC Channel at 768KHz sample rate
- Fully deterministic execution time

Target:

- ANC, heavy duty filtering, sound spatialization, sound effects, ...

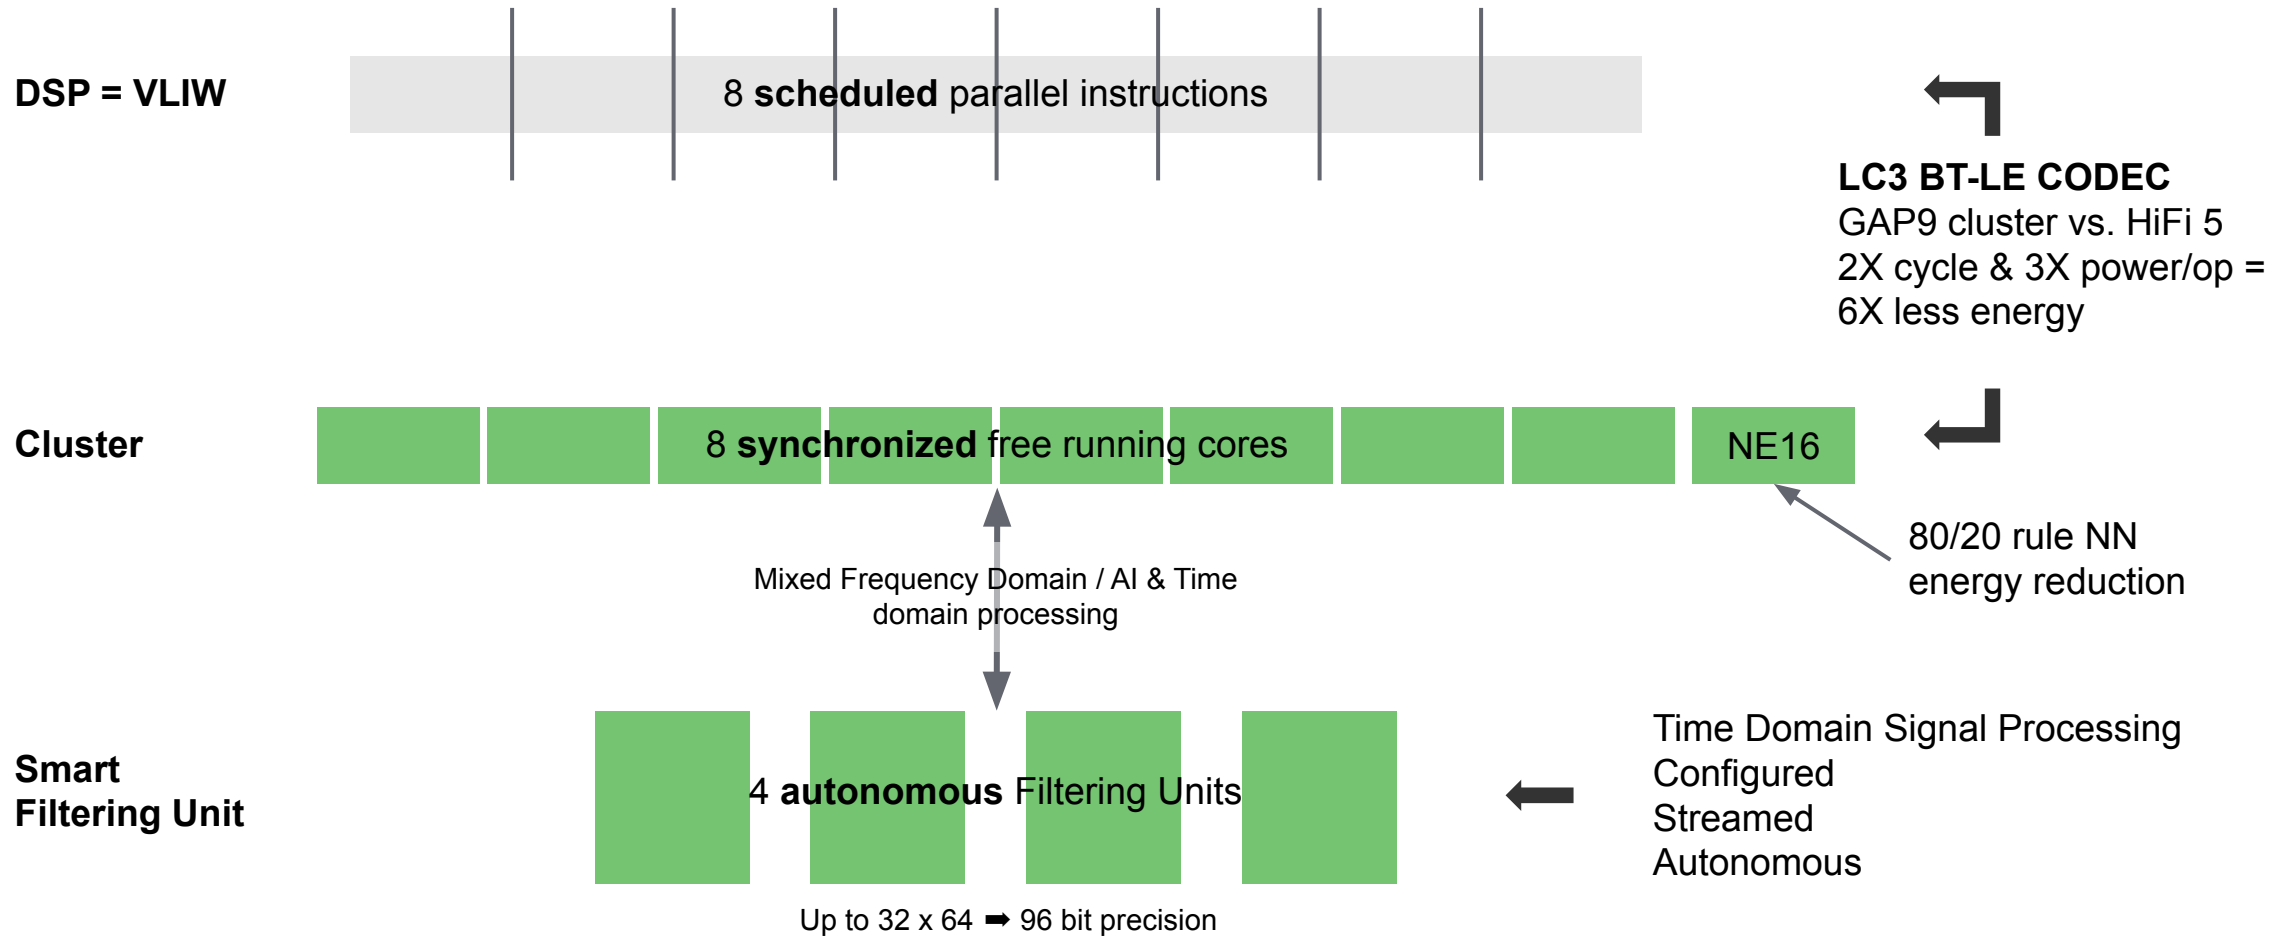
GAP Cluster - Parallelization



Quadratic energy benefit

Job Speedup => Less Cycles => Lower Frequency => Lower Core Voltage

How is Cluster + SFU different to a classic DSP?



Neural networks on GAP9

Powerful and flexible

- Different network architectures: perceptron, recurrent, convolutional, transformers, custom activations
- Powerful Python driven toolchain
- Flexible quantization
- NN compiler produces commented C code
No black boxes

```
res = G.execute_on_target(  
    directory="/tmp/testdir",  
    settings={  
        'l1_size': 128000,  
        'l2_size': 1300000,  
        'l3_flash_device': 'AT MEM L3 MRAMFLASH',  
        'cluster_stack_size': 1024,  
        'cluster_slave_stack_size': 512,  
        'graph_ll_promotion': 1  
    },  
    cmake=True,  
    at_loglevel=2,  
    memory=True  
)  
print(f"Run: {'ok' if res.returncode == 0 else 'error'}")  
print(res.pretty_performance())
```

Run ok

Layer	Cycles	Ops	Ops/Cycle	% ops
S26_Conv2d_128x1x32x1_Leaky	5,227	4,224	0.808	0.621
S29_Conv2d_64x1x1x128	5,527	8,192	1.482	1.211
S34_Conv2d_128x1x1x64_Leaky	6,085	8,320	1.367	1.231

From



to quite large ML...

Neural Network Tools

Lots of examples



TensorFlow Lite



NN Tool

Handwritten

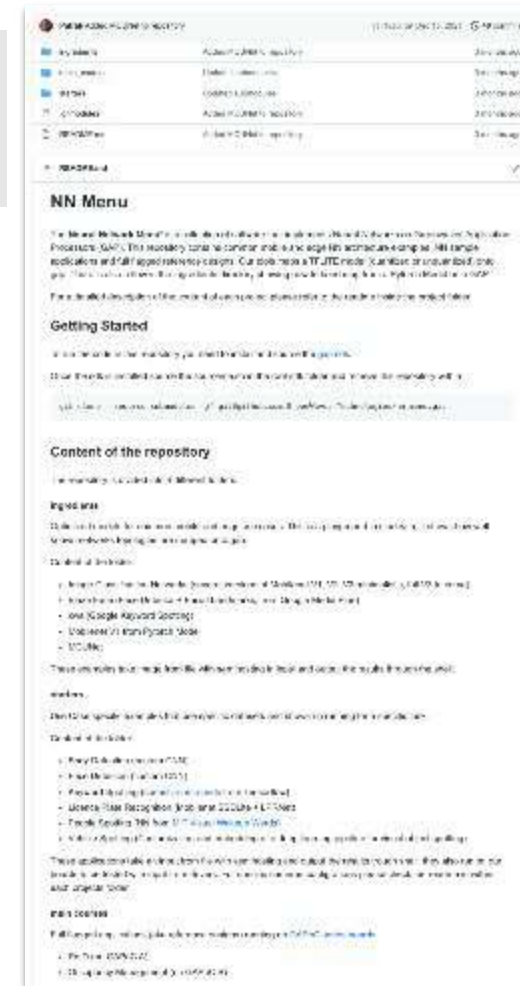
- Convolution
- Recurrent
- MatMul
- Pooling
- Linear
- Resizers
- SSD
- etc.

Compiled into C

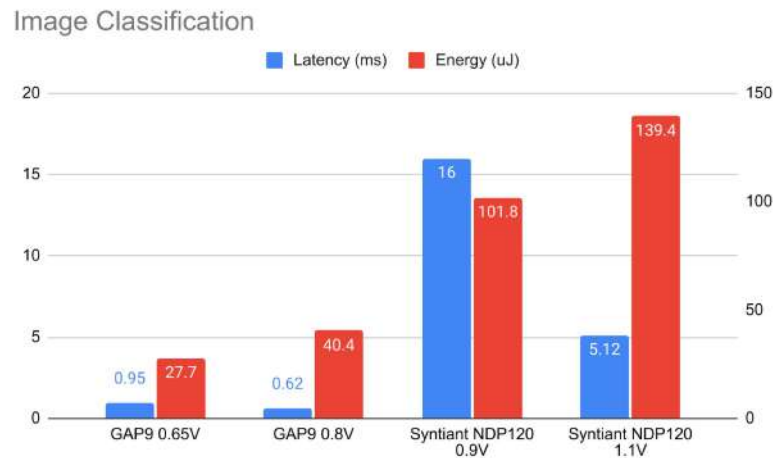
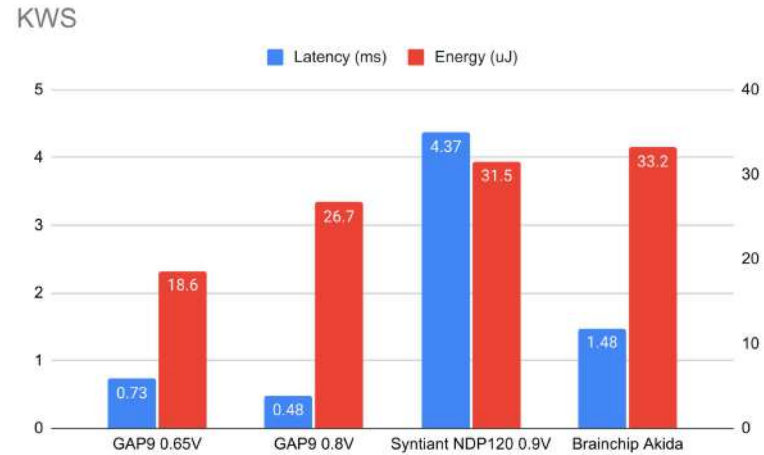
- Piecewise
- Fused Activations

AutoTiler

Optimised mixed IEEE16, BFloat16, 16 bit, 8 bit and sub byte quantization



GAP9: At remarkably high energy efficiency



GAP9 ~2-3X less energy than a highly specialized neuromorphic chip

GAP9 ~2-4X less energy than a specialized neural network accelerator chip

At the lowest latency in all tests

[1] All others: MLPerf™ v1.0 Tiny Closed KWS and VWW. Result verified by MLCommons Association.

[2] Brainchip: MLPerf™ v1.0 Tiny Closed KWS and VWW. Result not verified by MLCommons Association. Retrieved from Benchmarking AI Inference at the Edge: Measuring Performance and Efficiency for Real-World Deployments. Brainchip.


The MLPerf™ name and logo are trademarks of MLCommons Association in the United States and other countries. All rights reserved. Unauthorized use strictly prohibited. See www.mlcommons.org for more information."

GWT Software Tools – Customer’s Productivity is the ultimate KPI

PRODUCER


HIGH LEVEL SOFTWARE

Audio Designers




Graph Designer

Data Scientists



GAPflow

Application Developers



SDK

Unified C code

```
#include <stdio.h>
/* print Fahrenheit-Celsius table
 * for fahr = 0, 20, ..., 300; floating-point version w/
 * main()
 */
float fahr, celsius;
int lower, upper, step;

lower = 0; /* lower limit of temperature table w/
upper = 300; /* upper limit w/
step = 20; /* step size w/

fahr = lower;
while (fahr <= upper) {
    celsius = (5.0/9.0) * (fahr-32.0);
    printf("%i %i %i\n", fahr, celsius);
    fahr = fahr + step;
}
```



GAP



Evaluation boards available from GreenWaves

GAP9Mod

Core Module




22x30mm

GAP9 chip

Plug M.2 format Card into receptacle

GAP9_EVK

Generic Evaluation Kit



63x68mm

Mate GAP9_EVK with Audio Add-On

Audio Application Board

Audio-centric Expansion Board




92x82mm

What our partners are doing with GAP9



GAP9 enables dynamic spatial audio with head tracking consuming just 2.7mW with standard profile

- Stereo signals are reproduced through two virtual sound sources in a simulated listening room
- Performing the processing in the headphone (as opposed to a smartphone) avoids the detrimental latency introduced by Bluetooth
- Fast head tracking ensures that the sound image stays fixed and is stable in front of the listener



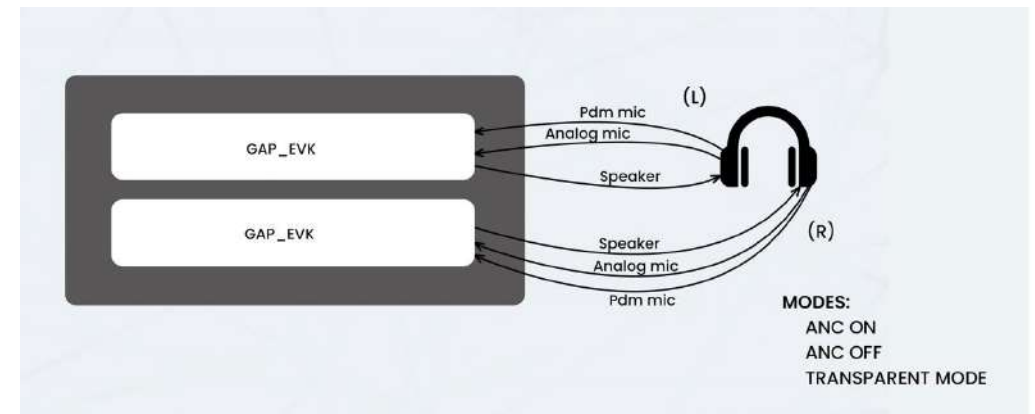
“IDUN audio’s dynamic spatial audio (DSA) platform allows for a vast amount of configuration possibilities. Having an application processor like GAP9 available really allows us to demonstrate the full extent of the DSA platform’s possibilities. At the same time, we are able to keep the power consumption low enough to run comfortably inside wireless headphones”.

What our partners are doing with GAP9



Selective and truly adaptive ANC running on GAP9 at 2,25 mW per channel 1

- ANC: Selective and truly adaptive ANC
- AI-NR: Voice pickup and speech separation
- Ultra low power consumption
- Breakthrough quality for audio algorithms



“GAP9's flexibility has allowed us to quickly port our market leading adaptive ANC solution onto it with great performance and power consumption results”.

¹ 768KHz sample rate

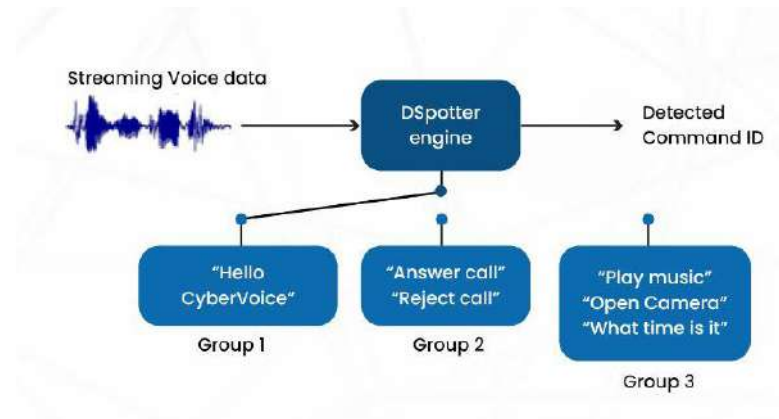
Eric BENHAIM, CTO of Orosound

What our partners are doing with GAP9



Always-listening voice trigger/ command recognition Dspotter operating on GAP9 at 10 MCPS and consuming 0.5mW

- High noise robustness
- Low resource requirement
- Global language support
- Easy command customization
- High portability



Commands understood by demonstration.
Wake up word: "Hey Gap"

Open camera Play music	Next song Volume louder
Stop music Previous song	Volume down

"We were able to port Cyberon's versatile, extendable voice command recognition engine, DSpotter to GAP9 in under a couple of weeks and are achieving the best energy efficiency that we have seen on any processor."

Alex Liou, Vice President of Cyberon Embedded Solution BU

Building a Complete Ecosystem around our Processors

	Partners
Algorithms and Tools	             
Sensors	       
BT Modem	 

Come and talk to us about how GAP9 could power your application

Thank You

GREENWAVES TECHNOLOGIES

28 Cours Jean Jaurès

38000 Grenoble

France

www.greenwaves-technologies.com