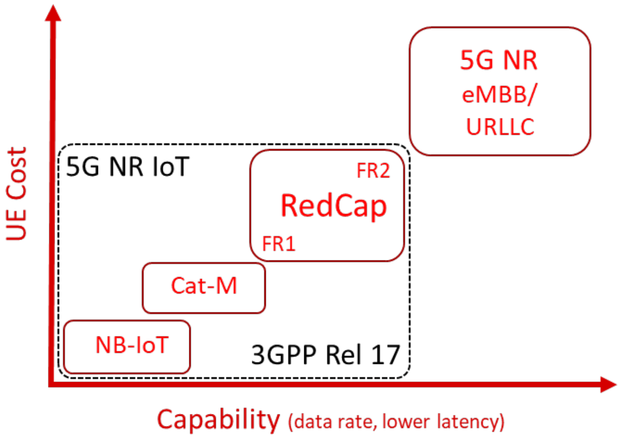

3GPP Release 17 for 5G IoT
(NB-IoT, Cat-M, RedCap)

Requirements and Workloads?

Algorithmic Approaches?

Suitable HW Platforms?

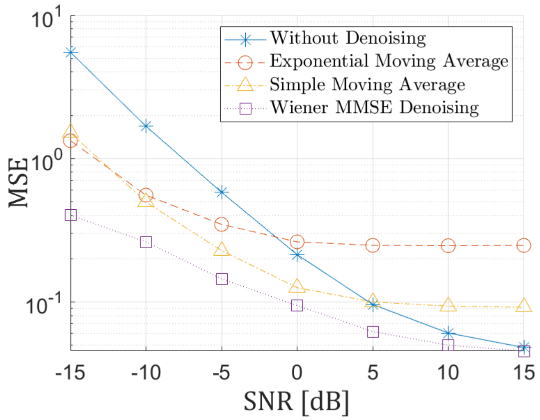

"Flexible Channel Estimation for 3GPP 5G IoT on a Vector Digital Signal Processor"

- *Generalisation of the moving average-based channel estimation and its mapping onto appropriately sized programmable platforms for 5G Release 17 IoT workloads, including the newly introduced RedCap.*
- *Identified workloads and requirements for channel estimation across all IoT specification.*
- *Devised an algorithmic approach that utilises flexibility of a programmable platform to provide robust approximate MMSE channel estimation quality using suboptimal moving average variants suitable for IoT, based on the best MSE performance criteria algorithm switching.*
- *Designed, implemented and measured denoising subkernels and the complete 5G IoT compatible channel estimation on 3 platforms, 32-bit RISC reference, vDSP128 and vDSP512 to cost effectively support IoT workloads.*
- *Processor SIMD architecture recommendation for each of the specified standards.*

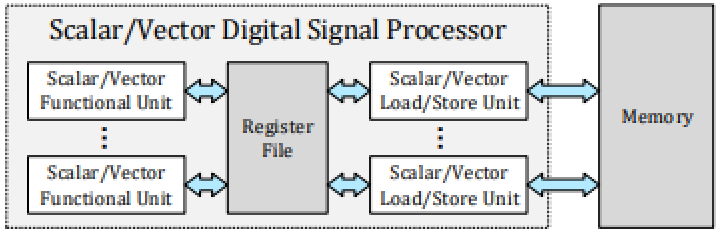

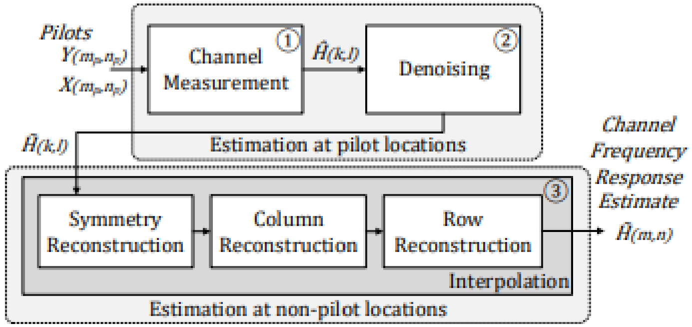

Stefan A. Damjancevic, Samuel Ajay Dasgupta, Emil Matus, Dmitry Utyanksy, Pieter van der Wolf, and Gerhard Fettweis,
Paper originally published in IEEE SiPS 2021, Oct 19-21, 2021

Flexible Channel Estimation for 3GPP 5G IoT on a Vector Digital Signal Processor

Stefan A. Damjancevic, Samuel Ajay Dasgupta, Emil Matus, Dmitry Utyanksy, Pieter van der Wolf, and Gerhard Fettweis

Synopsys ASIP University Day, November 17, 2021, Virtual Online Event

Research Interests:

Implementation aspects of 3GPP compliant physical layer algorithms, modem and processor architectures, and their codependency on SW, HW and system level

Funding:

In part by Synopsys Inc., under the Industry—University Cooperation Project “Efficient Implementation of 5G Baseband Kernels on a Vector Processor.”

Motivation: Introduction of RedCap for 3GPP 5G IoT

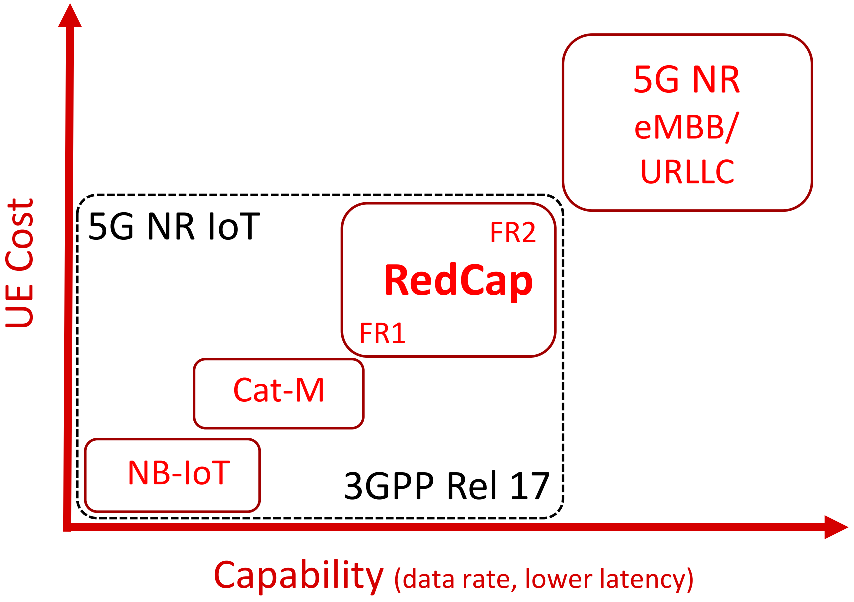

3GPP Release 17 for 5G IoT
(NB-IoT, Cat-M, RedCap)

Requirements and Workloads?

Algorithmic Approaches?

Suitable HW Platforms?

- ➢ 3GPP is introducing 5G NR Release 17 IoT:
  - NR Reduced Capability (RedCap/REDCAP/NR-Lite/NR-Light) to cover applications in the mid-range cost/capability devices of NR: [1] [2]
    - e.g. *wearables, industrial IoT, video surveillance*
  - Expected updates to 5G NR for Narrowband IoT and Cat-M [3]

- ➢ Goal: Generalisation of Channel Estimation for applications across all IoT standards on appropriately sized platforms

- ➢ Conditions and Requirements:
  - Weak radio signal – Reliable operation down to -10 dB SNR [4]
    => focus on denoising algorithms
  - High application variability
    => requires a parametrisable and scalable approach
  - Low-power operation – low computational complexity
    => requires low complexity algorithms

---

[1] A. Ghosh, R. Ratasuk, and A. M. Rao, "Industrial IoT Networks Powered by 5G New Radio." Microwave Journal, vol. 62, no. 12, 2019.

[2] N. Varsier, L.-A. Dufrene, M. Dumay, Q. Lampin, and J. Schwoerer, "A ` 5G New Radio for Balanced and Mixed IoT Use Cases: Challenges and Key Enablers in FR1 Band," IEEE Communications Magazine, vol. 59, no. 4, pp. 82–87, 2021.

[3] SierraWireless, "3GPP Enhancements Planned for 5G Release 17 Include RedCap, Coverage Improvements, Satellite Standards, and More ," SierraWireless Blog. [Online]. Available: https://bit.ly/2UetkZy

[4] 3GPP TSG RAN, TS 38.133 Requirements for support of radio resource management v17.1.0 , 3GPP Std., April, 2021. [Online]. Available: https://bit.ly/3hYo1qV

Workload, Platforms, Algorithms

Workload increase:

TABLE I
COMPARISSON OF RESOURCE BLOCK RATES OF 5G IOT DEVICES.

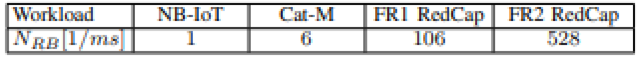

Workload	NB-IoT	Cat-M	FR1 RedCap	FR2 RedCap
$N_{RB}[1/ms]$	1	6	106	528

**Resource Block (RB) = 12 subcarriers x 14 OFDM symbols freq.-time data item grid*

- *With Rel 17 5G, RedCap IoT introduces higher data rates relative to prior IoT standards*
  - *18x increase for FR1 RedCap versus Cat-M*
  - *88x increase for FR2 RedCap versus Cat-M*
- *Need for a platform to provide the needed speed-up.*

Platform: programmable HW well suited

TABLE II
MOTIVATION FOR USE OF FLEXIBLE, PROGRAMMABLE PLATFORMS IN CHANNEL ESTIMATION FOR 5G IOT.

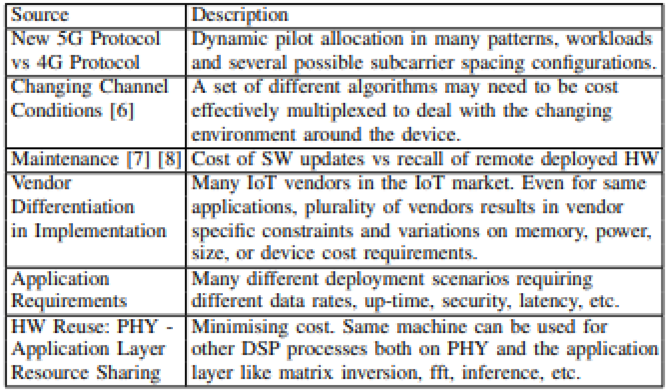

Source	Description
New 5G Protocol vs 4G Protocol	Dynamic pilot allocation in many patterns, workloads and several possible subcarrier spacing configurations.
Changing Channel Conditions [6]	A set of different algorithms may need to be cost effectively multiplexed to deal with the changing environment around the device.
Maintenance [7] [8]	Cost of SW updates vs recall of remote deployed HW
Vendor Differentiation in Implementation	Many IoT vendors in the IoT market. Even for same applications, plurality of vendors results in vendor specific constraints and variations on memory, power, size, or device cost requirements.
Application Requirements	Many different deployment scenarios requiring different data rates, up-time, security, latency, etc.
HW Reuse: PHY - Application Layer Resource Sharing	Minimising cost. Same machine can be used for other DSP processes both on PHY and the application layer like matrix inversion, fft, inference, etc.

[6] S. A. Damjancevic, E. Matus, D. Utyansky, P. van der Wolf, and G. P. Fettweis, "Channel Estimation for Advanced 5G/6G Use Cases on a Vector Digital Signal Processor," IEEE Open Journal of Circuits and Systems, vol. 2, pp. 265–277, 2021

[8] S. A. Damjancevic, E. Matus, D. Utyansky, P. van der Wolf, and G. P. Fettweis, "From challenges to hardware requirements for wireless communications reaching 6G," in Multi-Processor System-on-Chip 2, L. Andrade and F. Rousseau, Eds. London and New York: ISTE Ltd. and Wiley, 2021, ch. 1, pp. 3 – 33.

Channel Estimation Motivation and Proposed Approach

- *Optimal MMSE methods for denoising are too compute costly for IoT*
- *Suboptimal moving average filter variants are investigated as an approximate MMSE solution*

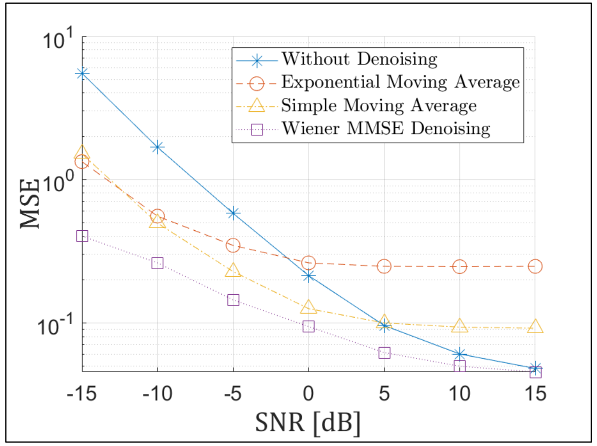

TABLE III
SIMULATION PARAMETERS.

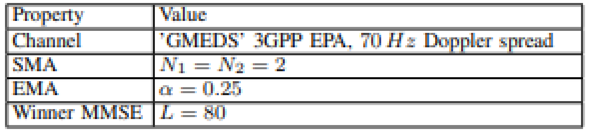

Property	Value
Channel	'GMEDS' 3GPP EPA, 70 Hz Doppler spread
SMA	$N_1 = N_2 = 2$
EMA	$\alpha = 0.25$
Winner MMSE	$L = 80$

- *A switching control approach is proposed that utilises the flexibility of the platform*
- *Switching: A priori knowledge SNR, ch. profile*

Simple Moving Average (SMA)

$$\tilde{H}(k,l) = \frac{1}{N_1 + N_2 + 1} \sum_{i=l-N_1}^{l+N_2} \hat{H}(k,i).$$

Exponential Moving Average (EMA)

$$\tilde{H}(k,l) = \alpha \hat{H}(k,l) + (1-\alpha)\,\tilde{H}(k,l-1),$$

Switching:

- *EMA good for -15 dB to -10 dB*
- *SMA good for -10 dB to 5 dB*
- *W/o denoising good for 5 dB to 15 dB*

Algorithmic Optimisation and 32-bit RISC Implementation

Further Algorithmic Optimisations:

- *SMA – i) mpy instead of div ii) discussion on less cycles of mac vs more cycles of add*
- *EMA – i) β instead of (1-α)*

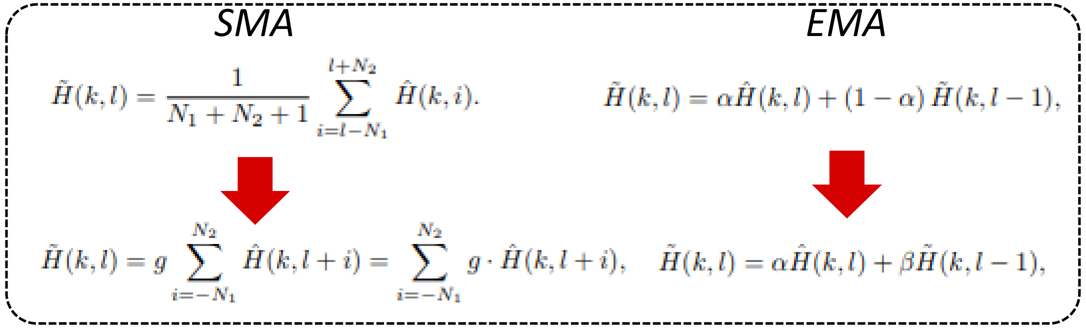

SMA

$$\tilde{H}(k,l) = \frac{1}{N_1 + N_2 + 1} \sum_{i=l-N_1}^{l+N_2} \hat{H}(k,i).$$

$$\tilde{H}(k,l) = g \sum_{i=-N_1}^{N_2} \hat{H}(k,l+i) = \sum_{i=-N_1}^{N_2} g \cdot \hat{H}(k,l+i),$$

EMA

$$\tilde{H}(k,l) = \alpha \hat{H}(k,l) + (1-\alpha)\,\tilde{H}(k,l-1),$$

$$\tilde{H}(k,l) = \alpha \hat{H}(k,l) + \beta \tilde{H}(k,l-1),$$

Operations [1/ms]:

- *See TABLE IV (352x more than prior art Cat-M)*

TABLE IV
NUMBER OF OPERATIONS.

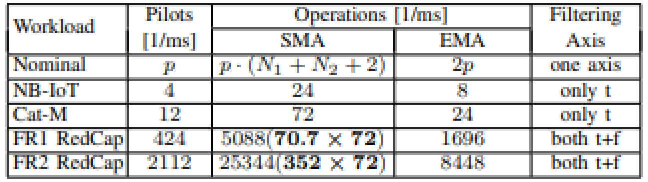

Workload	Pilots [1/ms]	Operations [1/ms]		Filtering Axis
		SMA	EMA	
Nominal	p	$p \cdot (N_1 + N_2 + 2)$	$2p$	one axis
NB-IoT	4	24	8	only t
Cat-M	12	72	24	only t
FR1 RedCap	424	5088(**70.7 × 72**)	1696	both t+f
FR2 RedCap	2112	25344(**352 × 72**)	8448	both t+f

Pseudo Code of Optimised Equations:

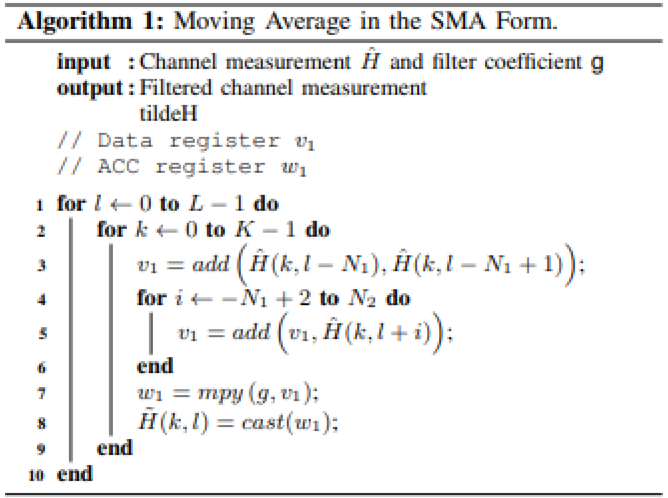

Algorithm 1: Moving Average in the SMA Form.

```
input  : Channel measurement $\hat{H}$ and filter coefficient g
output : Filtered channel measurement
         tildeH
// Data register $v_1$
// ACC register $w_1$
for $l \leftarrow 0$ to $L-1$ do
    for $k \leftarrow 0$ to $K-1$ do
        $v_1 = add\left(\hat{H}(k,l-N_1), \hat{H}(k,l-N_1+1)\right)$;
        for $i \leftarrow -N_1+2$ to $N_2$ do
            $v_1 = add\left(v_1, \hat{H}(k,l+i)\right)$;
        end
        $w_1 = mpy\,(g, v_1)$;
        $\tilde{H}(k,l) = cast(w_1)$;
    end
end
```

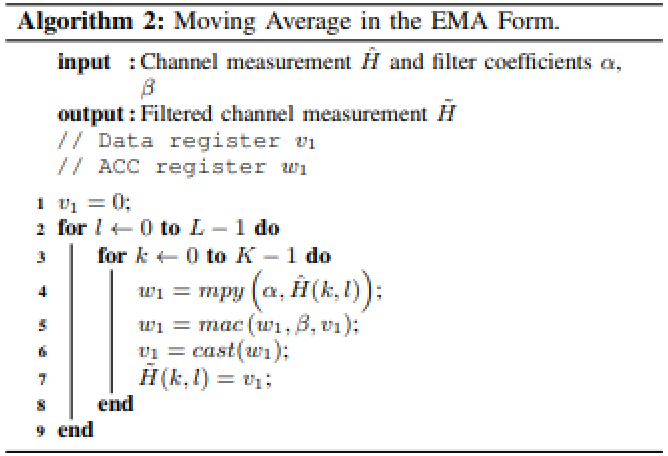

Algorithm 2: Moving Average in the EMA Form.

```
input  : Channel measurement $\hat{H}$ and filter coefficients $\alpha$,
         $\beta$
output : Filtered channel measurement $\tilde{H}$
// Data register $v_1$
// ACC register $w_1$
$v_1 = 0$;
for $l \leftarrow 0$ to $L-1$ do
    for $k \leftarrow 0$ to $K-1$ do
        $w_1 = mpy\left(\alpha, \hat{H}(k,l)\right)$;
        $w_1 = mac\,(w_1, \beta, v_1)$;
        $v_1 = cast(w_1)$;
        $\tilde{H}(k,l) = v_1$;
    end
end
```

32-bit reference Implementation:

- *Cycles [1/ms] -> Frequency [kHz]*
- *i) 25 MHz for just FR2 RedCap denoising is unacceptable, a SIMD solution is needed. ii) NB-IoT OK.*
- *Ops/cycles go up with higher workloads due to relatively smaller loop overhead and pipeline depth impact*

TABLE V
32-BIT SCALAR REFERENCE: REQUIRED CLOCK FREQUENCY AND MAC UNIT UTILISATION.

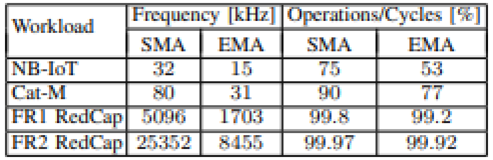

Workload	Frequency [kHz]		Operations/Cycles [%]	
	SMA	EMA	SMA	EMA
NB-IoT	32	15	75	53
Cat-M	80	31	90	77
FR1 RedCap	5096	1703	99.8	99.2
FR2 RedCap	25352	8455	99.97	99.92

Processor Architecture, 128-bit and 512-bit SIMD Implementation

Scalar/Vector VLIW Architecture:

- *Register file centric, computational FUs met with appropriate LD/ST units*

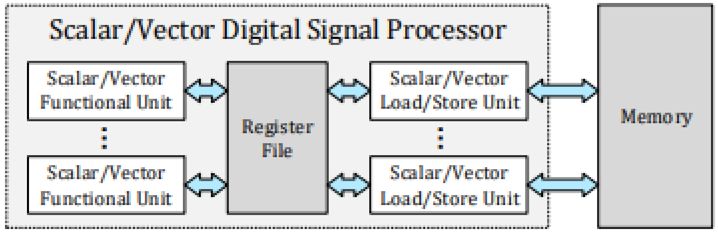

- *Easy to extend with SIMD FUs, developed with [17]*
- *Other RISC alternatives discussed for the 32-bit reference: single-issue processor with HW AGU and XY memory, however hard to extend to VLIW/SIMD [18]*

TABLE VI
PLATFORM CONFIGURATION: 32-BIT REFERENCE, 128-BIT SIMD, 512-BIT SIMD.

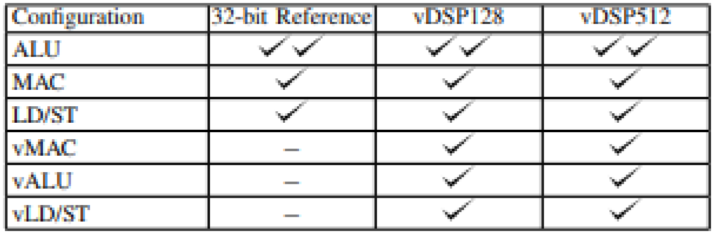

Configuration	32-bit Reference	vDSP128	vDSP512
ALU	✓✓	✓✓	✓✓
MAC	✓	✓	✓
LD/ST	✓	✓	✓
vMAC	–	✓	✓
vALU	–	✓	✓
vLD/ST	–	✓	✓

SIMD vDSP Implementation:

TABLE VII
128-BIT VLIW-SIMD VECTOR PROCESSOR: REQUIRED CLOCK FREQUENCY, VECTORISATION GAIN AND EFFICIENCY.

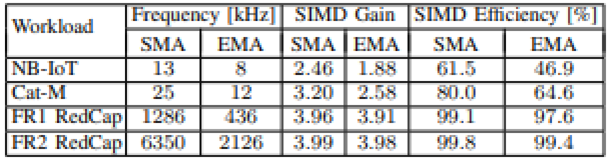

Workload	Frequency [kHz]		SIMD Gain		SIMD Efficiency [%]	
	SMA	EMA	SMA	EMA	SMA	EMA
NB-IoT	13	8	2.46	1.88	61.5	46.9
Cat-M	25	12	3.20	2.58	80.0	64.6
FR1 RedCap	1286	436	3.96	3.91	99.1	97.6
FR2 RedCap	6350	2126	3.99	3.98	99.8	99.4

TABLE VIII
512-BIT VLIW-SIMD VECTOR PROCESSOR: REQUIRED CLOCK FREQUENCY, VECTORISATION GAIN AND EFFICIENCY.

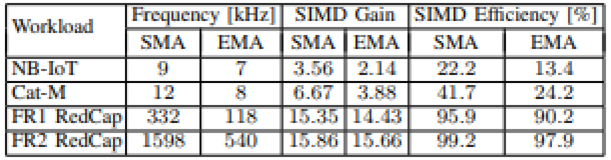

Workload	Frequency [kHz]		SIMD Gain		SIMD Efficiency [%]	
	SMA	EMA	SMA	EMA	SMA	EMA
NB-IoT	9	7	3.56	2.14	22.2	13.4
Cat-M	12	8	6.67	3.88	41.7	24.2
FR1 RedCap	332	118	15.35	14.43	95.9	90.2
FR2 RedCap	1598	540	15.86	15.66	99.2	97.9

- *Denoising subkernel Required Clock Frequency, SIMD Gain and SIMD Efficiency*
- *SIMD Gain = Cycle count reduction against the 32-bit reference*
- *SIMD Efficiency = SIMD Gain versus ideal SIMD Gain*

- *vDSP512 achieves best gains at FR2 RedCap workloads*
- *vDSP128 achieves best gains already at FR1 RedCap workloads*

- *Both vDSP128 and vDSP512 have $\approx 1MHz$ denoising option in their high gain operation area, making the range of 128-bit to 512-bit platforms quite well suited for RedCap workloads.*

---

[17] Synopsys Inc, "ASIP Designer Website," accessed 30-06-2021. [Online]. Available: https://www.synopsys.com/asip

[18] P. van der Wolf and Y. Tanurhan, "Processors for the internet of things," in Multi-Processor System-on-Chip 1, L. Andrade and F. Rousseau, Eds. London and New York: ISTE Ltd. and Wiley, 2021, ch. 10, pp. 159 – 181.

Implementation of the Complete Channel Estimation Procedure Including Channel Measurement and Interpolation

NR Compliant *Channel Estimation* Procedure [6] adapted to IoT:

- *Three major subkernels: i) Channel Measurement; ii) Denoising; iii) Interpolation*

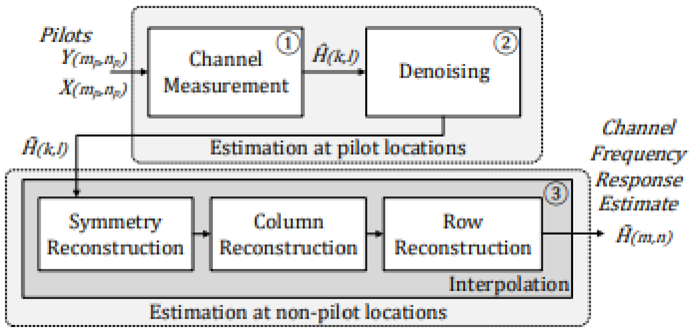

Open Questions:

- *What is the frequency requirement for the whole procedure, not just denoising?*
- *What is the share of denoising in the entire procedure?*

[6] S. A. Damjancevic, E. Matus, D. Utyansky, P. van der Wolf, and G. P. Fettweis, “Channel Estimation for Advanced 5G/6G Use Cases on a Vector Digital Signal Processor,” IEEE Open Journal of Circuits and Systems, vol. 2, pp. 265–277, 2021

SIMD vDSP Implementation:

TABLE IX
COMPLETE CHANNEL ESTIMATION PROCEDURE: CLOCK FREQUENCY REQUIREMENTS AND VECTORISATION GAIN.

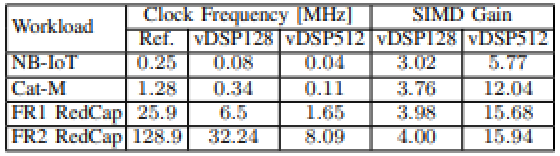

Workload	Clock Frequency [MHz]			SIMD Gain	
	Ref.	vDSP128	vDSP512	vDSP128	vDSP512
NB-IoT	0.25	0.08	0.04	3.02	5.77
Cat-M	1.28	0.34	0.11	3.76	12.04
FR1 RedCap	25.9	6.5	1.65	3.98	15.68
FR2 RedCap	128.9	32.24	8.09	4.00	15.94

TABLE X
COMPLETE CHANNEL ESTIMATION PROCEDURE: COMPUTATIONAL SHARE BETWEEN STEPS, FIG. 1 NOTATION.

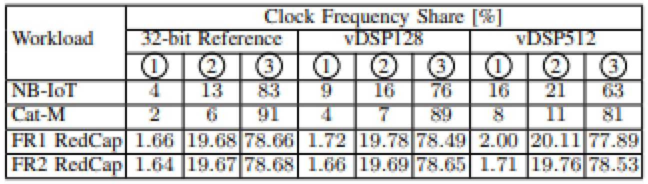

Workload	Clock Frequency Share [%]								
	32-bit Reference			vDSP128			vDSP512		
	①	②	③	①	②	③	①	②	③
NB-IoT	4	13	83	9	16	76	16	21	63
Cat-M	2	6	91	4	7	89	8	11	81
FR1 RedCap	1.66	19.68	78.66	1.72	19.78	78.49	2.00	20.11	77.89
FR2 RedCap	1.64	19.67	78.68	1.66	19.69	78.65	1.71	19.76	78.53

- *TABLE IX: SIMD gain remains high and even improves when the whole channel estimation is considered.*
- *With the whole channel estimation taken into account, the high gain area extends for the vDSP128 also towards Cat-M workloads, signalling that the range of 32-bit to 128-bit machines would also be suitable low-rate RedCap or Cat-M*

** Assumes SMA as the denoising subkernel*

- *TABLE X: i) Interpolation dominates the frequency budget, ii) denoising takes ≈ 20% share in RedCap independent of the platform*
- *For FR2 RedCap CE, 512-bit platforms are well suited*
- *For FR1 RedCap CE, platforms in the range of 128-bit to 512-bit are well suited*
- *For Cat-M CE, platforms in the range 32-bit to 128-bit are well suited*
- *For NB-IoT CE, 32-bit platforms are well suited*

Results: Key Contributions

➢ *Workloads:*

- *Programmable HW Platforms like a Vector Processor can provide the required flexibility and speedup to support RedCap IoT*

➢ *Algorithm:*

- *Moving Average filter variations offer good MSE performance and low computational cost – good for IoT*
- *MSE performance near optimal (MMSE) with SW alg. switching*
- *SW Switching is easy on a flexible, programmable platforms*

Switching:

- *EMA good for -15 dB to -10 dB*
- *SMA good for -10 dB to 5 dB*
- *W/o denoising good for 5 dB to 15 dB*

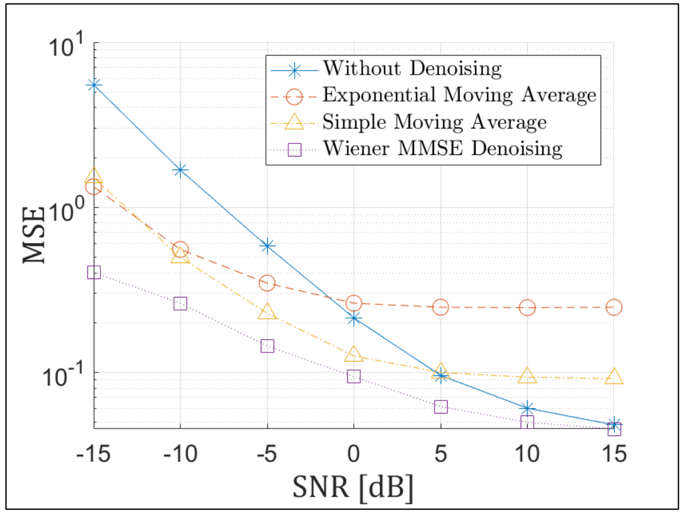

➢ *Implementation:*

- *Kernels achieving good cycle efficiency across IoT workloads*
  - *VLIW effectively masks LD/ST instruction overhead*
  - *SIMD Efficiency close to ideal*

Algorithm	Moving Average based Channel Estimation					
Architecture	**32-bit scalar**		**vDSP128**		**vDSP512**	
Metric	Req. Clock [MHz]	Arithmetic Ops/Cycles [%]	Req. Clock [MHz]	SIMD Efficiency [%]	Req. Clock [MHz]	SIMD Efficiency [%]
NB-IoT	**0.25**	**89**	0.08	67	0.04	33
Cat-M	**1.28**	**97.8**	**0.34**	**92**	0.11	74
FR1 RedCap	25.9	99.89	**6.5**	**99.48**	**1.65**	**97.9**
FR2 RedCap	128.9	99.98	32.4	99.89	**8.09**	**99.58**

➢ *Architecture:*

- *32-bit RISC reference, vDSP128 and vDSP512 can cost-effectively support IoT workloads*
- *Identified good architecture – standard pairs*

Standard	NB-IoT	Cat-M	FR1 RedCap	FR2 RedCap
Suitable Architecture	**32-bit**	**32-bit to 128-bit**	**128-bit to 512-bit**	**512-bit**

*SIMD Gain = Cycle count reduction against the 32-bit reference
**SIMD Efficiency = SIMD Gain versus ideal SIMD Gain