ERI ELECTRONICS RESURGENCE INITIATIVE SUMMIT

& MTO Symposium

MTO MICROSYSTEMS TECHNOLOGY OFFICE

DARPA

GOVERNMENT PASSIVE SENSING NEEDS

THOMAS DALRYMPLE

TECHNICAL ADVISOR

AFRL SENSORS DIRECTORATE

AFRL SENSORS DIRECTORATE

AFRL/RY Mission

- *Lead the discovery and development of future capabilities providing integrated Intelligence, Surveillance, and Reconnaissance (ISR), combat identification, and spectrum warfare effects*

AFRL/RY Vision

- *Enable ubiquitous Situation Awareness and Spectrum Dominance for Global Vigilance, Reach, and Power*

ENDURING CHALLENGES

Protect Aircrews and Aircraft

"forces that can deploy, survive, operate, maneuver, and regenerate in all domains while under attack"

National Defense Strategy

Defeat Integrated Air Defense Systems

"capabilities to penetrate the highly contested environment [and] deliver effects from stand-off ranges … to counter the A2/AD strategy"

Air Force Air Superiority 2030 Flight Plan

Dominate Airborne Threats

"new operational concepts and capabilities to win without assured dominance in air, space, and cyberspace domains"

National Security Strategy

Hold Critical Mobile Targets at Risk

"be able to strike diverse targets inside adversary air and missile defense networks"

National Defense Strategy

STRATEGIC CAPABILITY AREAS FOR USAF

Global Persistent Awareness

"Support continuous and timely knowledge of adversaries throughout the operating environment via distributed sensing across all domains"

Resilient Information Sharing

"Coordinate across all Joint Force assets through assured communications and precise position, navigation and timing resilient to any denial methods"

Rapid Effective Decision-Making

"Increase the speed of battlespace understanding and decision-making through automation to act faster than any adversary"

Complexity, Unpredictability, and Mass

"Overwhelm adversaries with complexity, unpredictability and numbers through a collaborative and autonomous network of systems and effects"

Speed and Reach of Disruption and Lethality

"Rapidly disrupt and neutralize dynamic and mobile targets using new methods to attack with speed and global reach"

GLOBAL PERSISTENT AWARENESS

Pervasive Sensing Grid

<u>Vision:</u> ***Penetrating ISR*** **platforms detect, track, and share time critical intelligence to enable battle management**

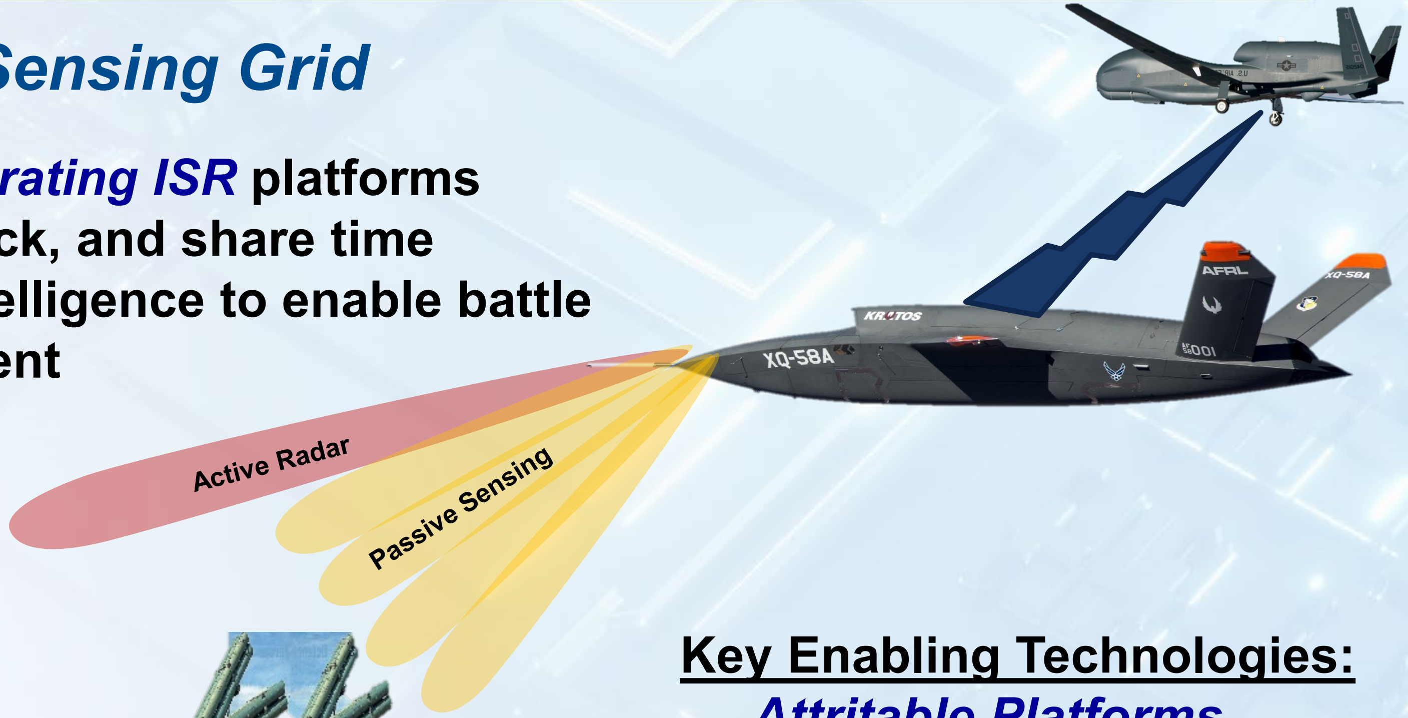

<u>Key Enabling Technologies:</u>

Attritable Platforms

Software Defined Payloads

GLOBAL PERSISTENT AWARENESS

Expanding Spectrum Coverage

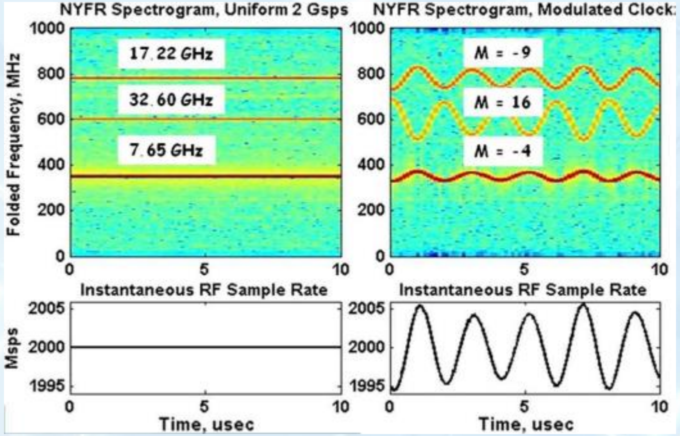

<u>Vision:</u> ***Staring receivers*** **capable of detecting hard to detect signals over wide frequency ranges**

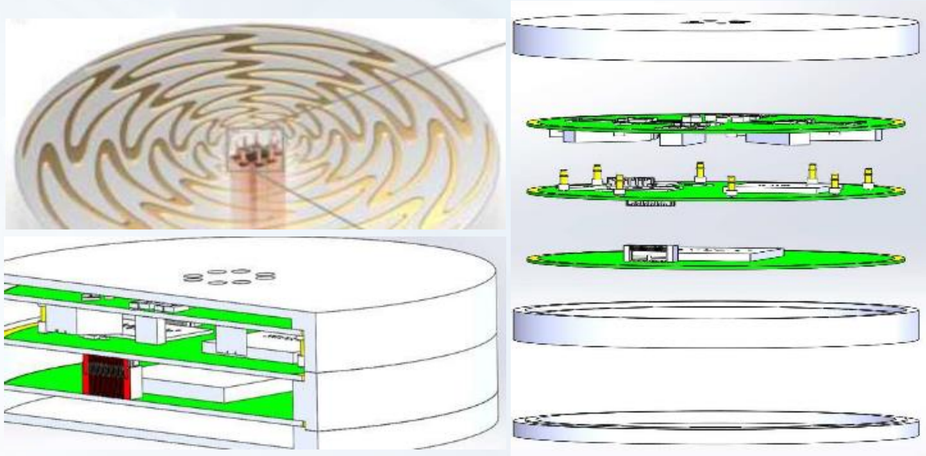

<u>Key Enabling Technologies:</u>

Wideband Components

Compressed Sensing

GLOBAL PERSISTENT AWARENESS

Passive Sensing and Wideband Digital Beamforming

Vision: ***Dynamically*** **utilize environmental illuminators for passive surveillance**

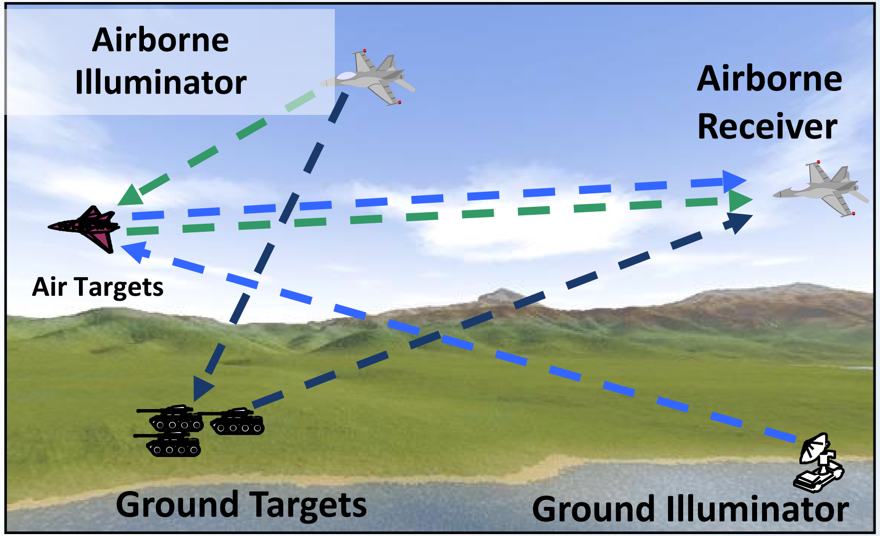

"ACT" Digital Common Module

Wideband Digital Array Technology

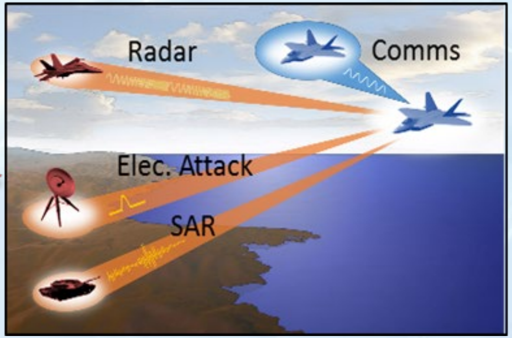

Key Enabling Technologies:

Digital-at-the-Element RF

Illumination Selection Manager

SUMMARY

- Significant operational challenges in contested operations
- Pervasive spectrum access critical to future military success
- Enduring challenges and strategic capabilities provide frameworks to guide advanced sensor developments
- Seeking advanced technologies that can evolve with agile threats
  - Performance tailored for operational impact
  - Techniques to consume spectrum and drive decisions
  - Enabling components and devices to improve size and performance

Ubiquitous Situation Awareness and Spectrum Dominance

ERI ELECTRONICS RESURGENCE INITIATIVE SUMMIT & MTO Symposium

GOVERNMENT PASSIVE SENSING NEEDS

ERI ELECTRONICS RESURGENCE INITIATIVE

SUMMIT

& MTO Symposium

MTO MICROSYSTEMS TECHNOLOGY OFFICE

DARPA

ANALOG SIGNAL PROCESSING

BORIS MURMANN

PROF. OF ELECTRICAL ENGINEERING
STANFORD UNIVERSITY

SIGNAL PROCESSING ENERGY

Energy per cycle per pole in a filter

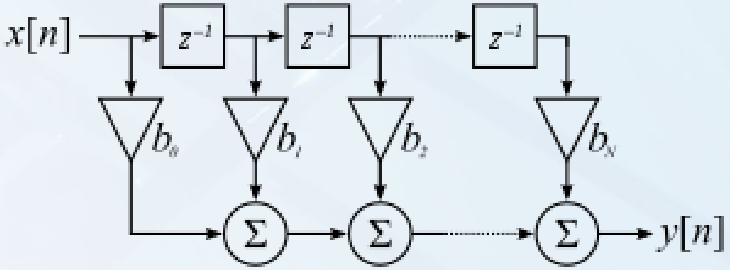

Digital $E \propto \log^2(\mathrm{DR})$

Analog $E \geq 8\mathrm{kT} \cdot \mathrm{DR}$

E. A. Vittoz, “Future of analog in the VLSI environment,” ISCAS 1990

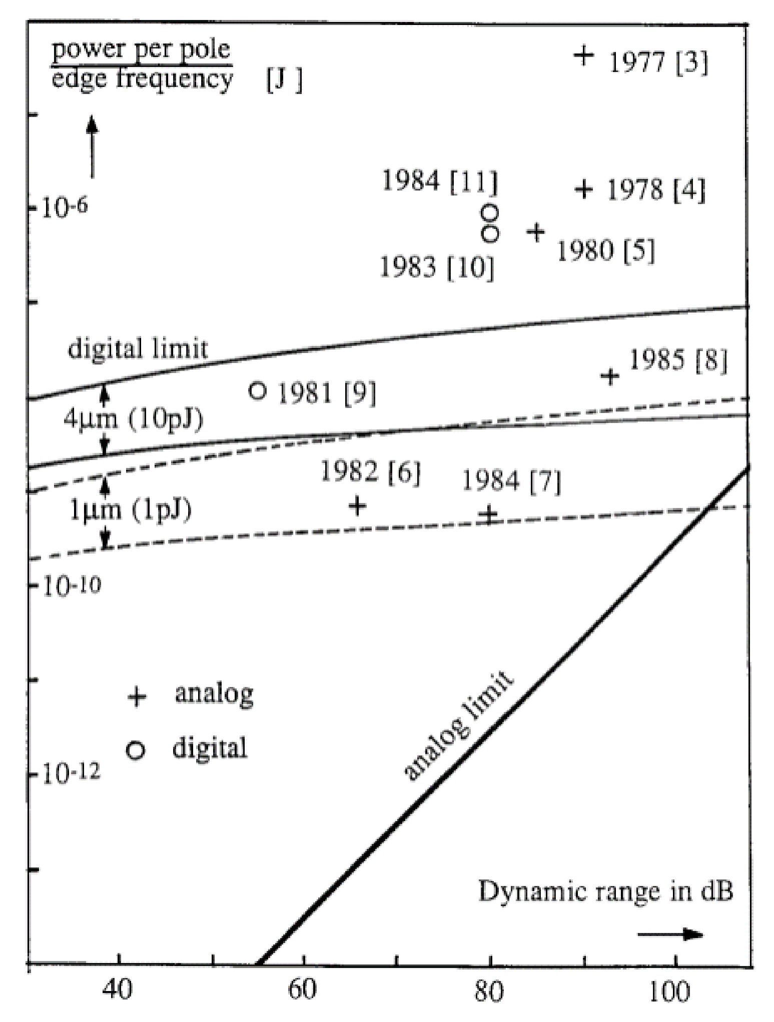

TODAY'S PERSPECTIVE

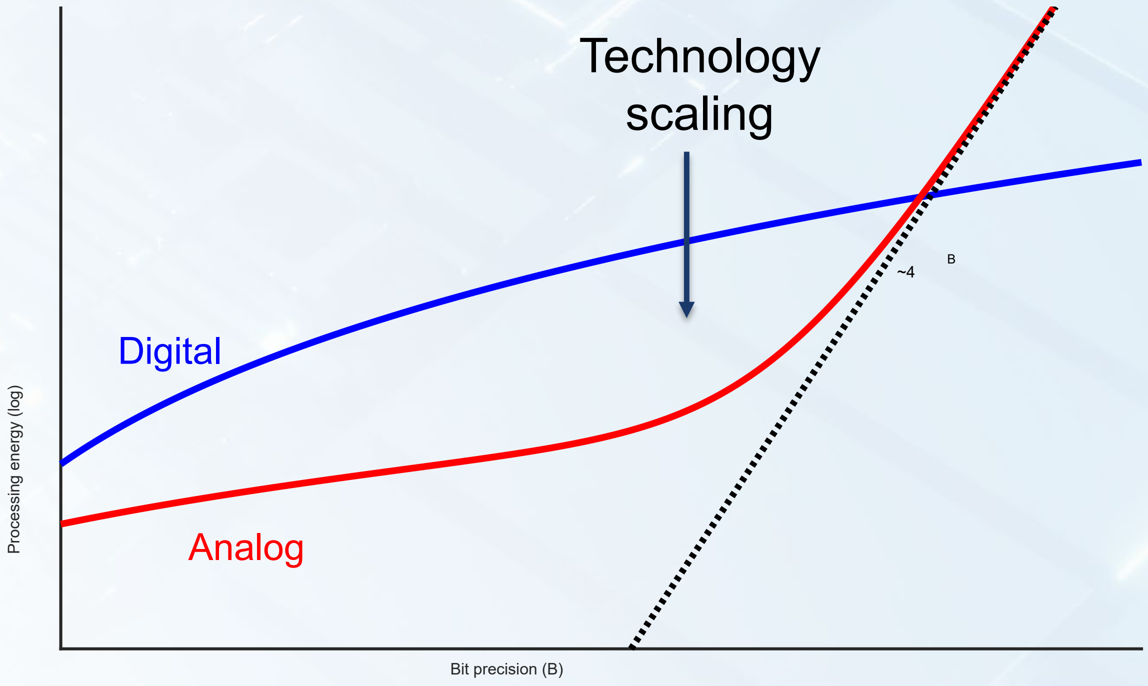

- Low bit precision
  - Analog ultimately approaches digital (digital ≈ 1-bit analog)
- Moderate bit precision
  - Analog can be about one order of magnitude more efficient than digital
- High bit precision
  - Analog tends to be inferior beyond 8 bits

ENERGY CHEAT SHEET

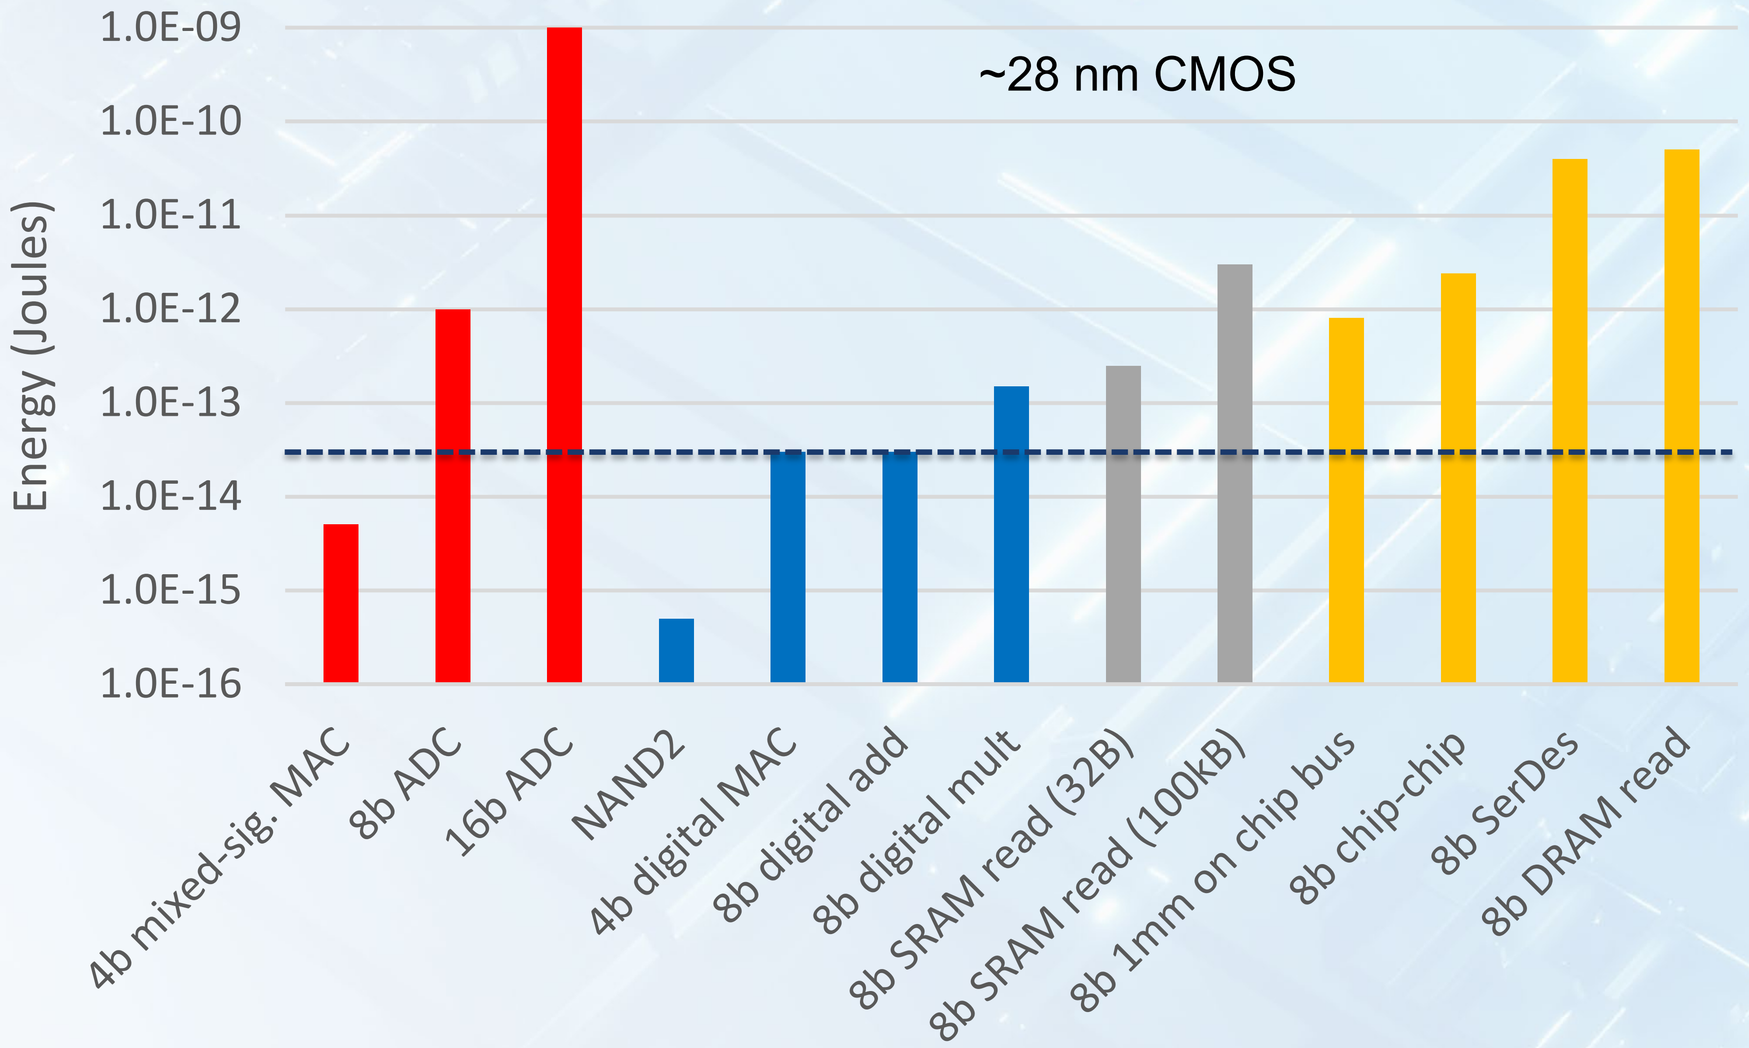

OBSERVATIONS

- Data movement and memory access are more expensive than compute and often limit the efficiency of modern systems
- Why bother with analog signal processing?
- Two opportunities
  - Reduction of data at the analog interface
  - Reduction of data movement for massively parallel compute

INPUT DATA DELUGE

Large arrays
Complex waveforms

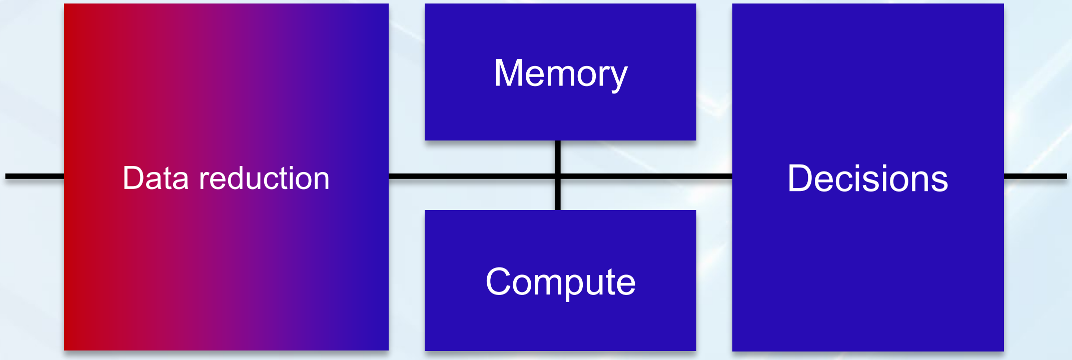

- Many modern systems are fed with complex or high-dimensional inputs
- Most data is “destroyed” via computation and data movement
- Up-front analog data reduction can reduce backend energy consumption

EXAMPLE: AUDIO FEATURE EXTRACTION

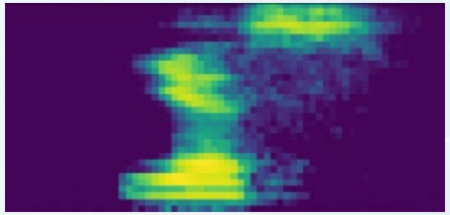

Spectrogram of “Yes”
Speech ≈39 bit/sec

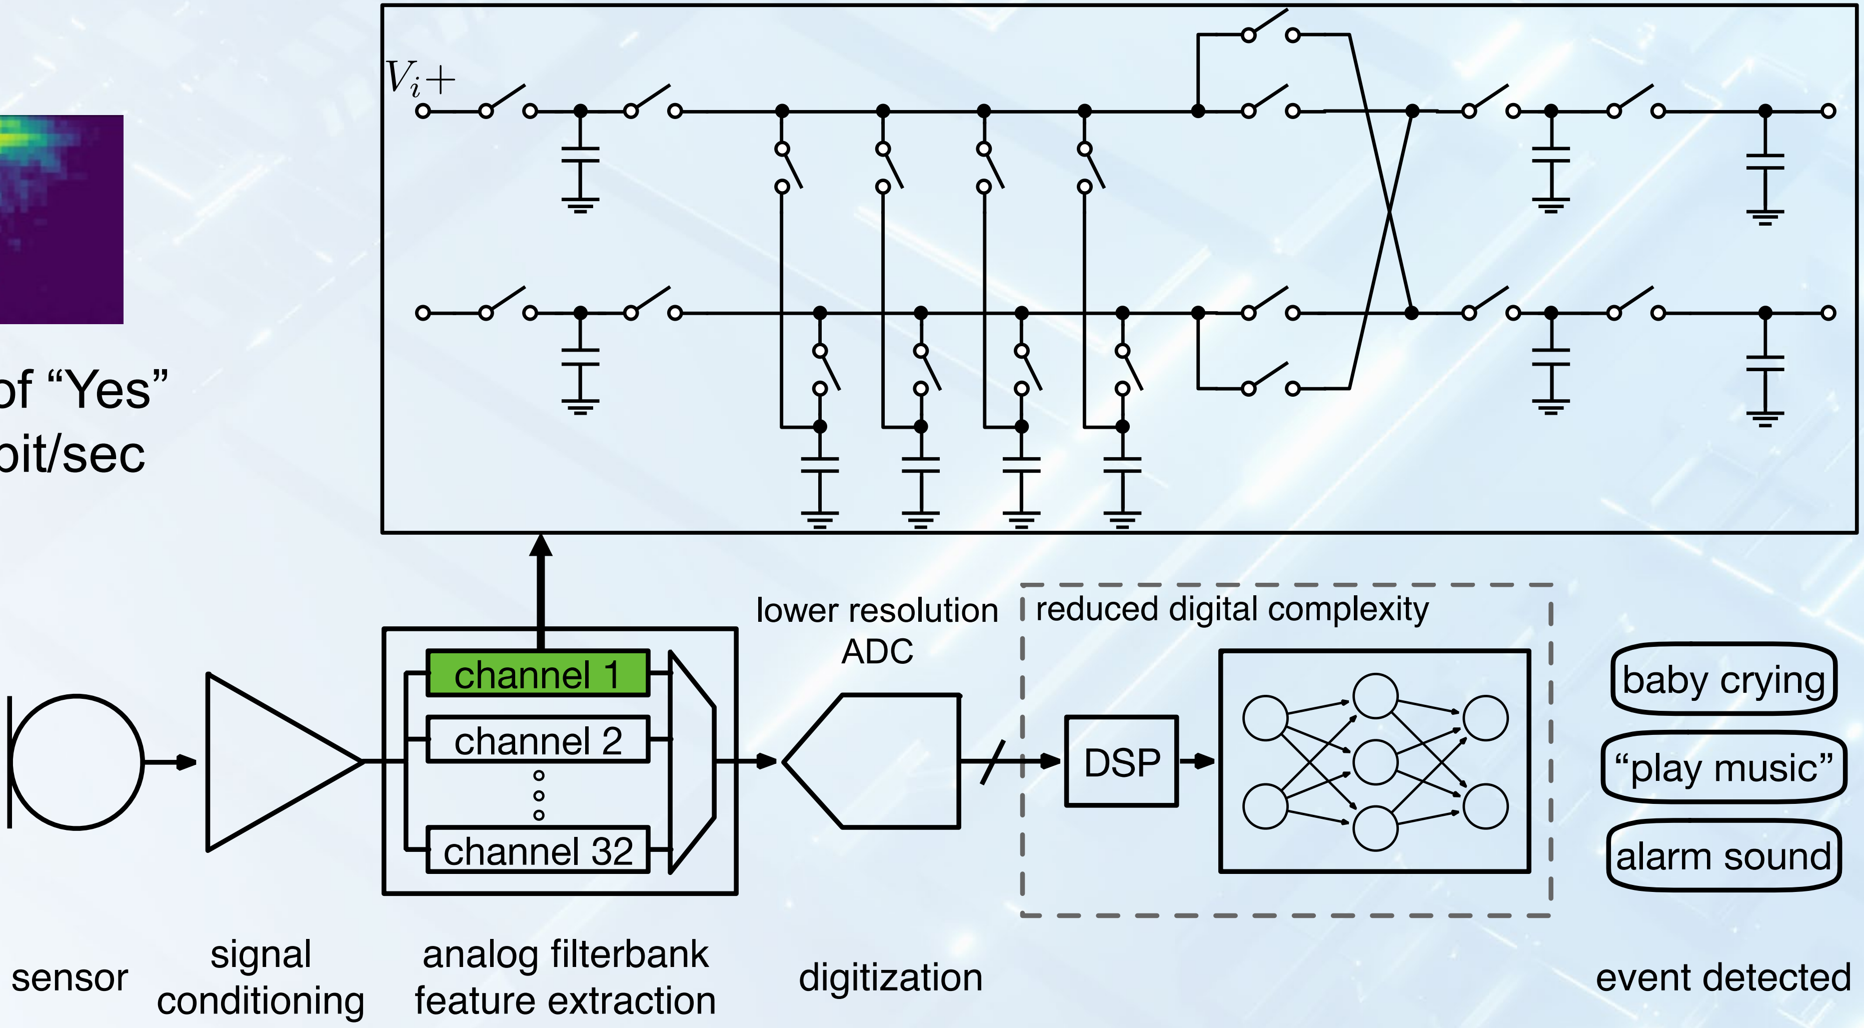

Villamizar, TCAS1, 2021

OPPORTUNITY FOR MANY APPLICATIONS

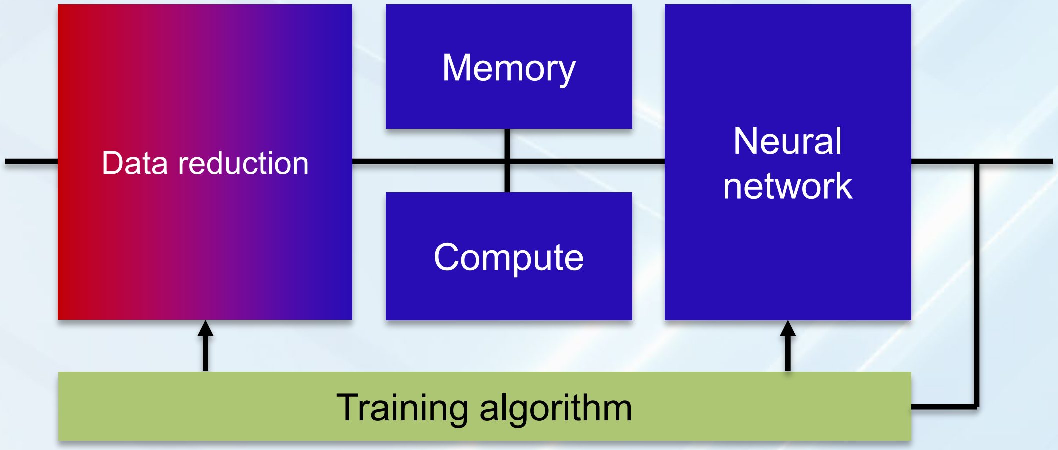

- End-to-end training with analog in the loop
- Neural network loss must be differentiable with respect to analog parameters
- “Differentiable signal processing” (Engel, ICLR 2020)

ANALOG COMPUTE?

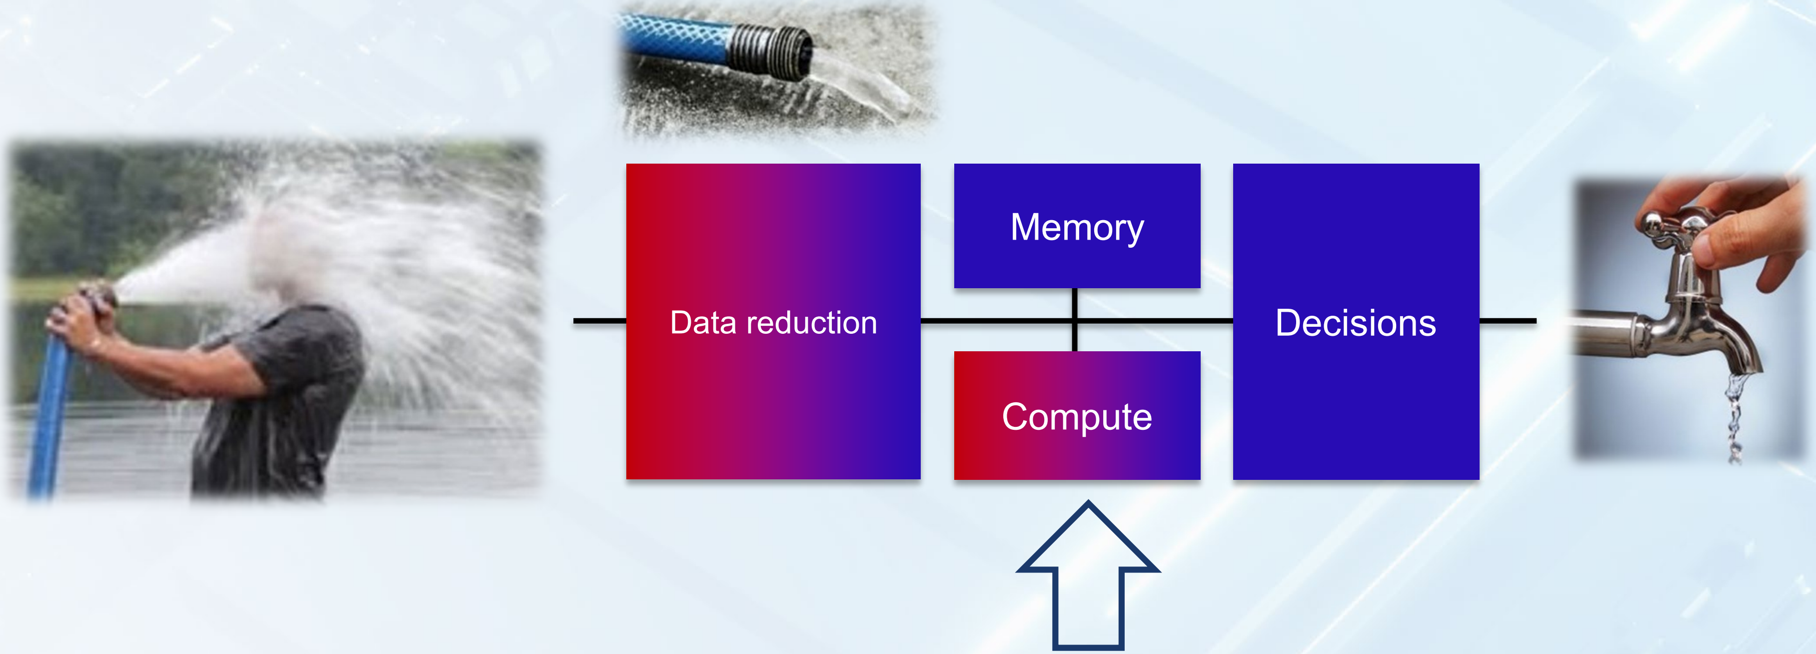

- Does it ever make sense to go back to analog for compute?
- Forgotten fact: We routinely go back to analog during memory readout
  - For example, SRAM sense amplification

MIXED-SIGNAL COMPUTE FOR DNNS

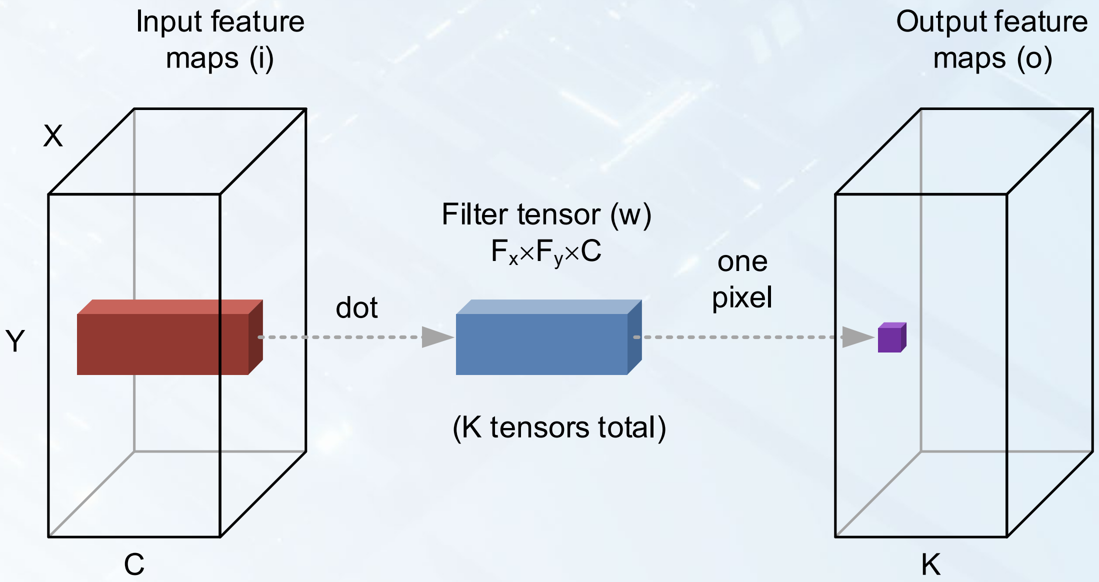

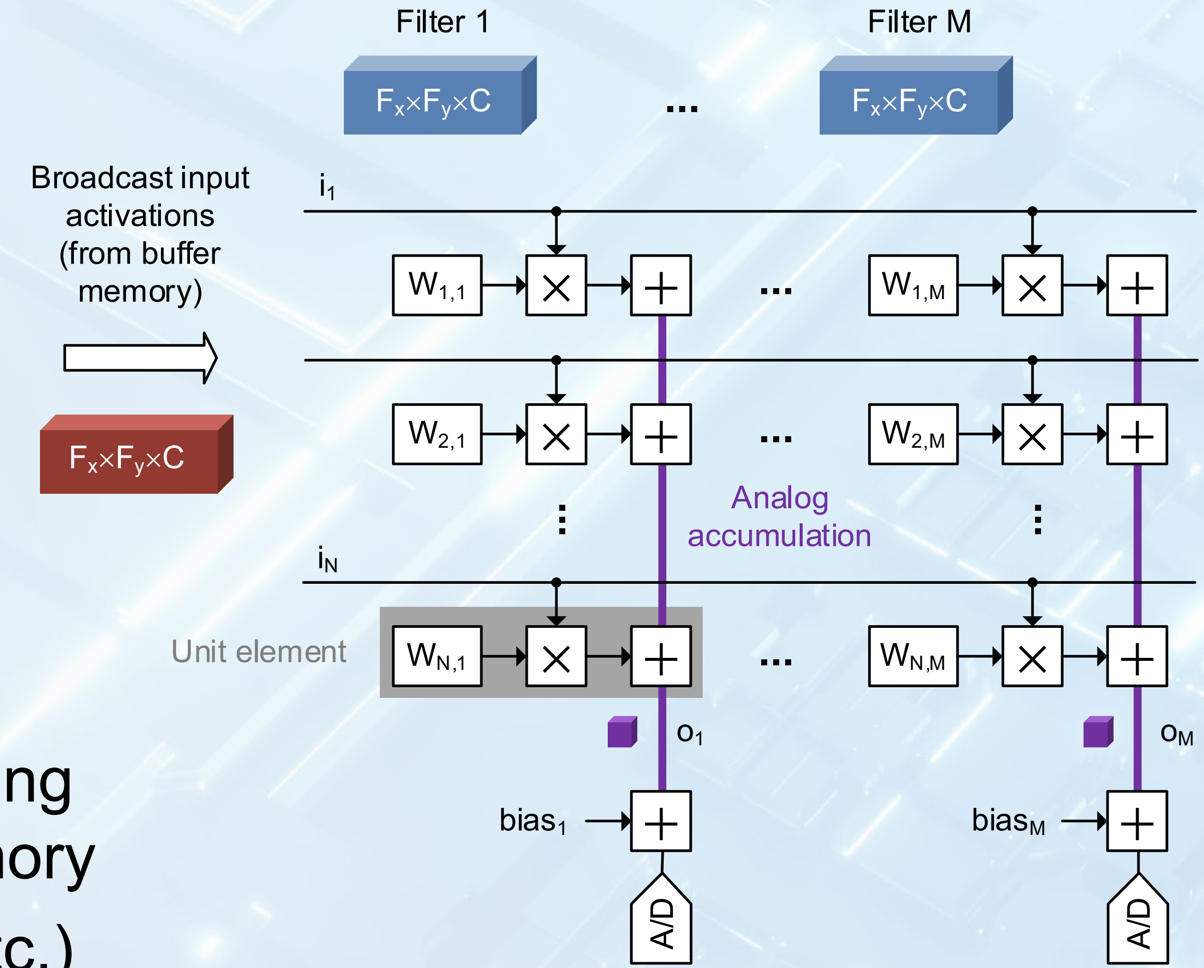

- Reduced data movement by co-locating massively parallel compute with memory
- Many incarnations (SRAM, RRAM, etc.)
- Key issues include flexibility, array utilization

THE BATTLE OF ANALOG VERSUS DIGITAL

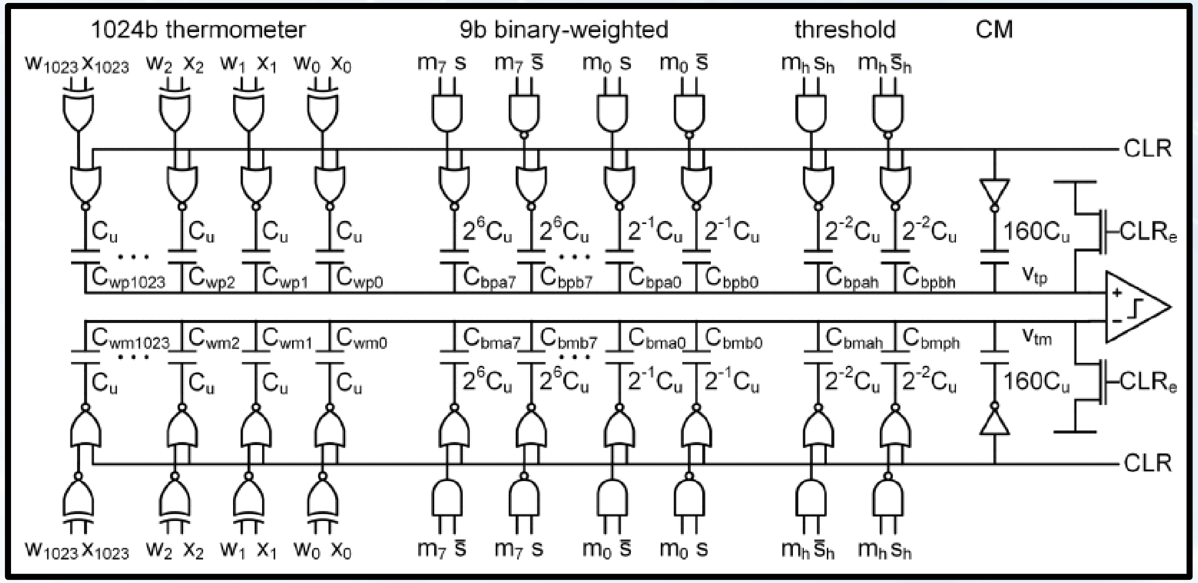

[Bankman, ISSCC 2018]
28 nm CMOS
478 1b-GOPS
532 1b-TOPS/W

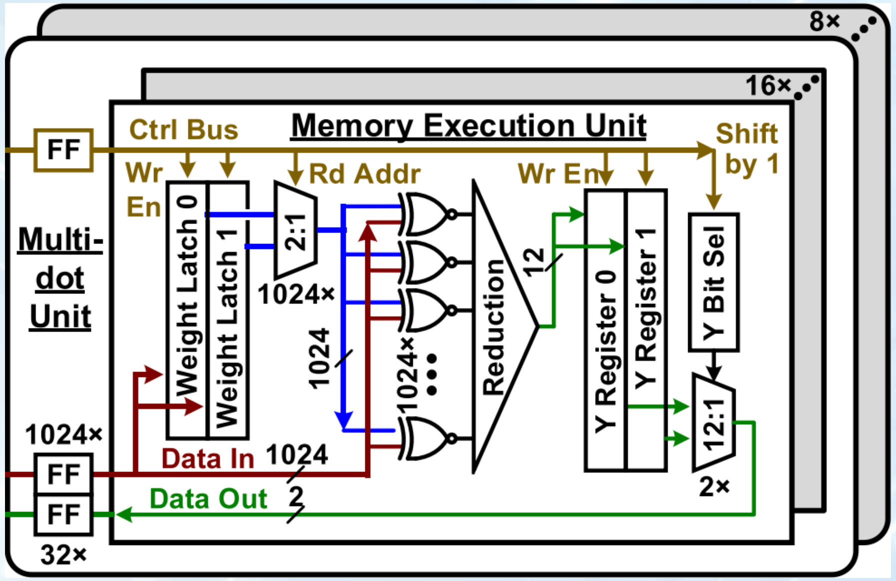

[Knag, VLSI 2020]
10 nm CMOS
163 1b-GOPS
269 1b-TOPS/W

SUMMARY

In the context of modern systems, analog processing can be intriguing for:

1. Mitigating data deluge at the system input
    - Many known solutions (some as “simple” as analog beamforming)
    - Many unexplored opportunities in coupling front-end with ML backend
    - Challenge is that most concepts are one-offs (not generalizable)

2. Reducing data movement in “embarrassingly parallel” compute loads
    - Literature shows promising results for small demonstrators
    - Challenge lies in integration with (somewhat) flexible compute platform
    - And the digital folks will fight back…

End-to-End Autoencoder Communications with Interference Suppression

Signal Processing in Neural Networks (SPiNN)

Kemal Davaslioglu, Intelligent Automation Inc., Senior Research Scientist

Tugba Erpek, Intelligent Automation Inc., Lead Scientist

Yalin Sagduyu, Intelligent Automation Inc., Director

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA).

End-to-End Autoencoder Communications with Interference Suppression

- **Motivation:** Next-generation communication systems need to be optimized beyond conventional communication designs that are typically based on simple analytical models or expert knowledge.
  - Support high reliability and high rate subject to complex channel and interference effects.
  - Jointly optimize the transmitter and receiver operations w.r.t channel and interference conditions.

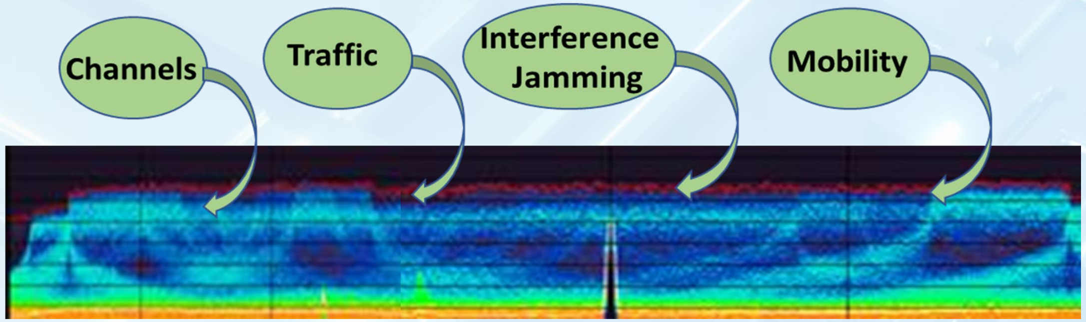

- **Data-driven Solution: Autoencoder (AE) based communication** system for which the transmitter and receiver are represented as **deep neural networks**.
  - Adapts to channel and interference dynamics.

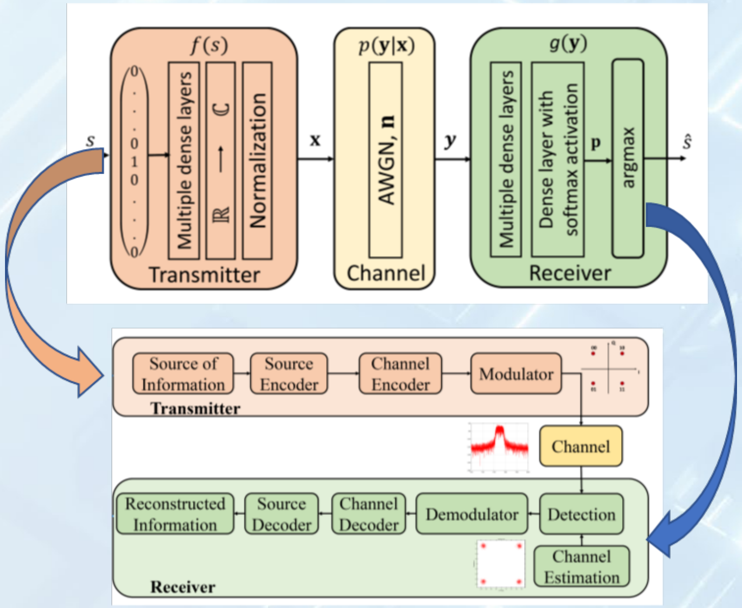

- **Conclusion:** Major improvement demonstrated in terms of reliability and interference suppression compared to conventional communication schemes.
  - **Suppress >30 dB multi-symbol interference.**

System Model & Channel Adaptation

- The **encoder** at the transmitter and the **decoder** at the receiver are represented as deep neural networks to adapt to spectrum (channel and interference) dynamics.
- AE system is **trained end-to-end** using **interference training** and **randomized smoothing** to mitigate the effects of interference (jamming).

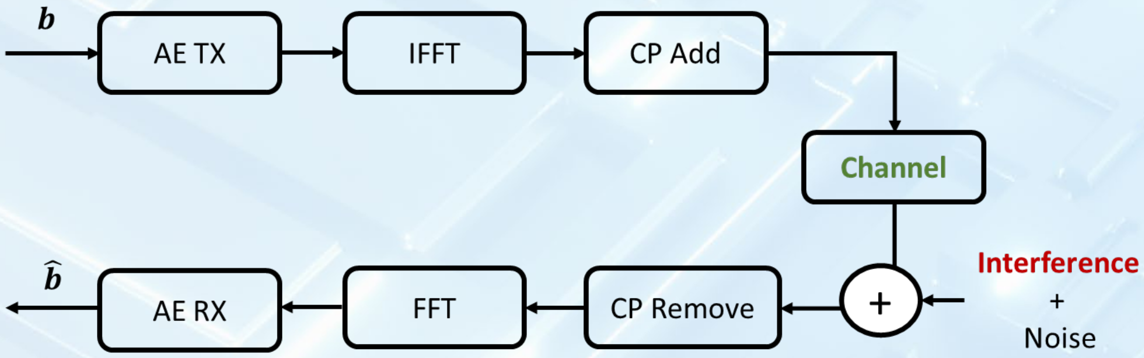

AWGN

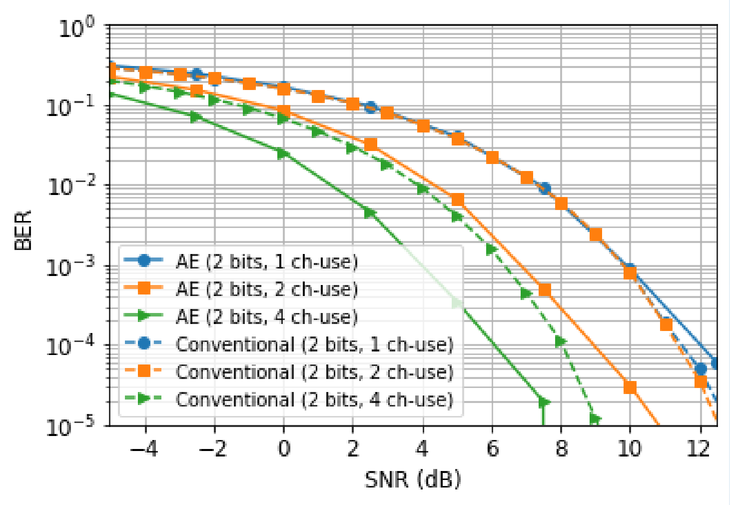

AWGN + Channel impairments
(phase offset = 10 deg,
frequency offset = 30 Hz)

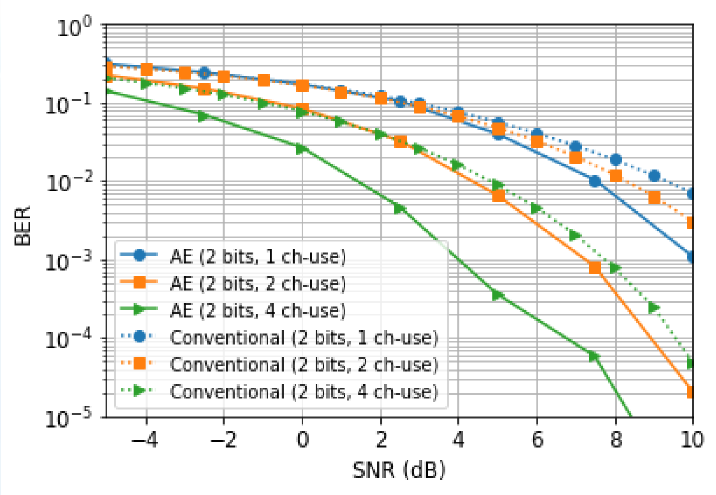

- AE system outperforms the conventional methods under different channel conditions.
- The n ch-use determines the redundancy for channel coding purposes.

Hardware and Data Limitations

- **Quantization** is needed for embedded hardware implementation on FPGA or embedded GPU.
- Compared to floating point implementation, model size and memory usage are reduced while maintaining the BER and inference time.

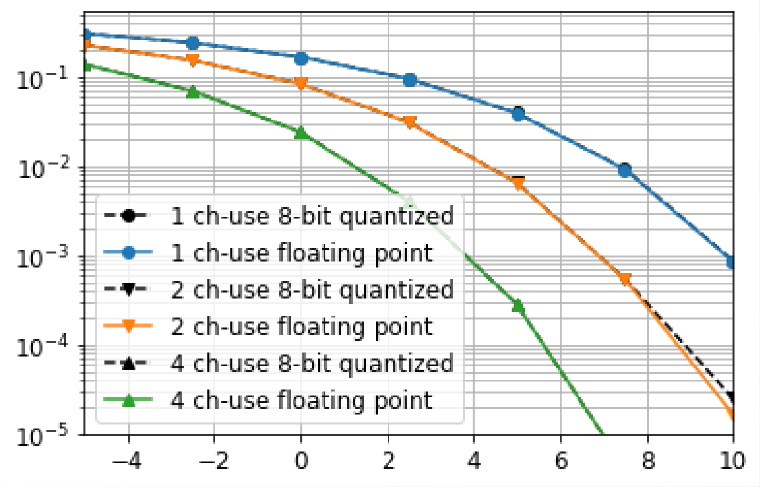

- Model size

Channel Use	Floating point model	8-bit quantized model
1 ch-use	2.88 Mbyte	0.73 Mbyte
2 ch-use	3.03 Mbyte	0.77 Mbyte
4 ch-use	3.34 Mbyte	0.85 Mbyte

- Inference time

Channel Use	Floating point model	8-bit quantized model
1 ch-use	0.000109 sec	0.000114 sec
2 ch-use	0.000333 sec	0.000335 sec
4 ch-use	0.000354 sec	0.000358 sec

- What if training data is limited?
  - Augment the training with **Generative Adversarial Network (GAN)** to add synthetic data and improve the training accuracy.

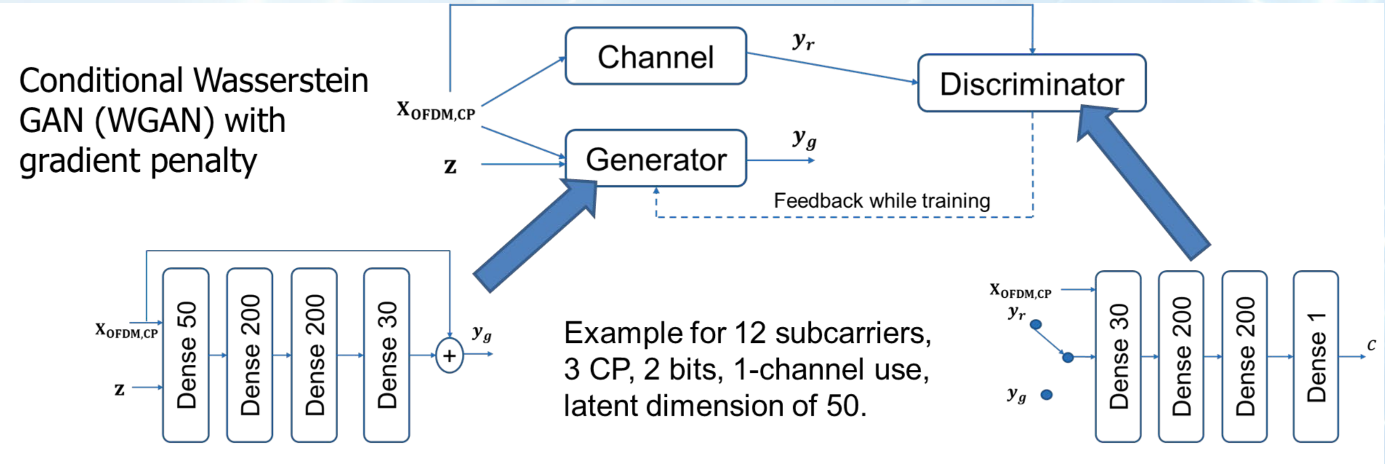

Training Data Augmentation with GAN

- GAN learns the distribution of the real data over time and generates the synthetic data.

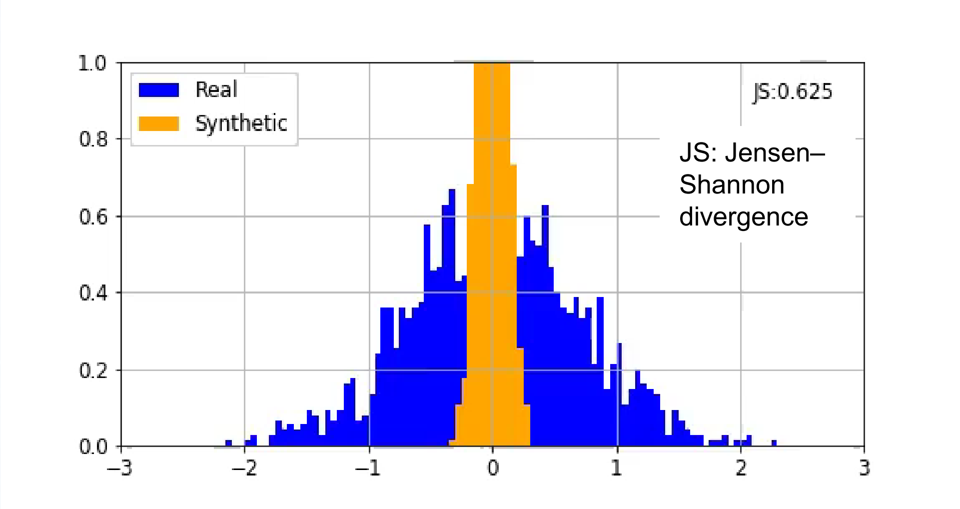

Nr. Synthetic Samples	BER
No Augmentation	0.0260
2x samples	0.0150
4x samples	0.0054
6x samples	0.0077
8x samples	0.0123
10x samples	0.0075

- Optimize the BER over the # of synthetic data samples (2 bits/channel use).

100 real + synthetic samples

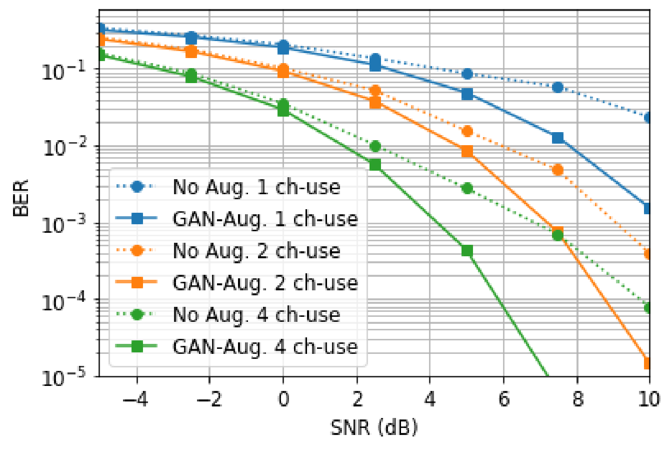

1000 real + synthetic samples

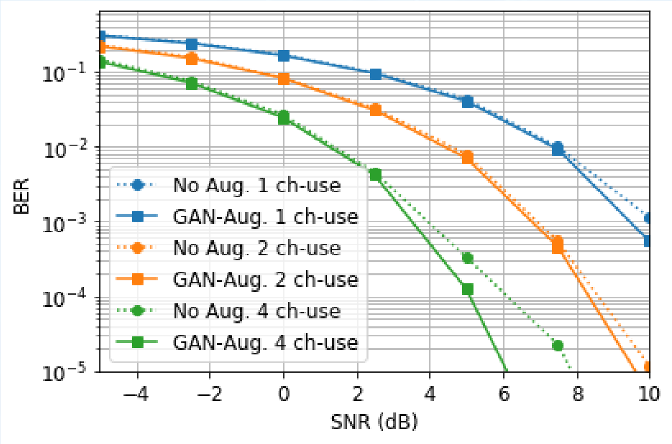

Interference (Jamming) Suppression

- ❑ Mitigate both channel and interference (jamming) effects
- ❑ Autoencoder with interference mitigation using Randomized Smoothing
  - Uses Gaussian data augmentation during training and increases robustness against perturbations introduced by noise and interference.

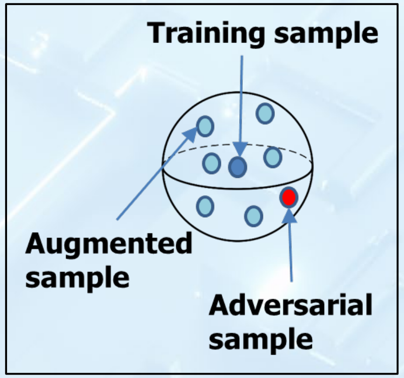

Randomized Smoothing

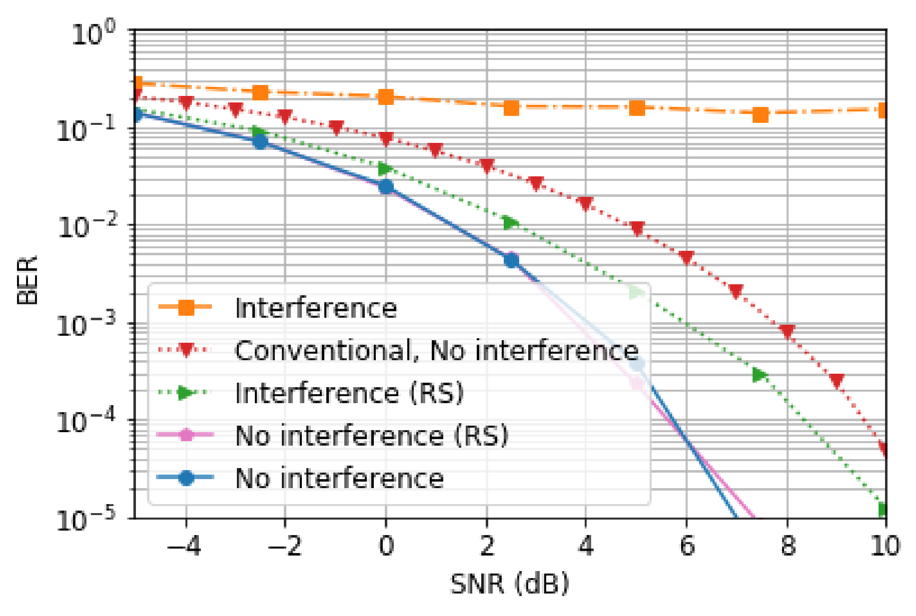

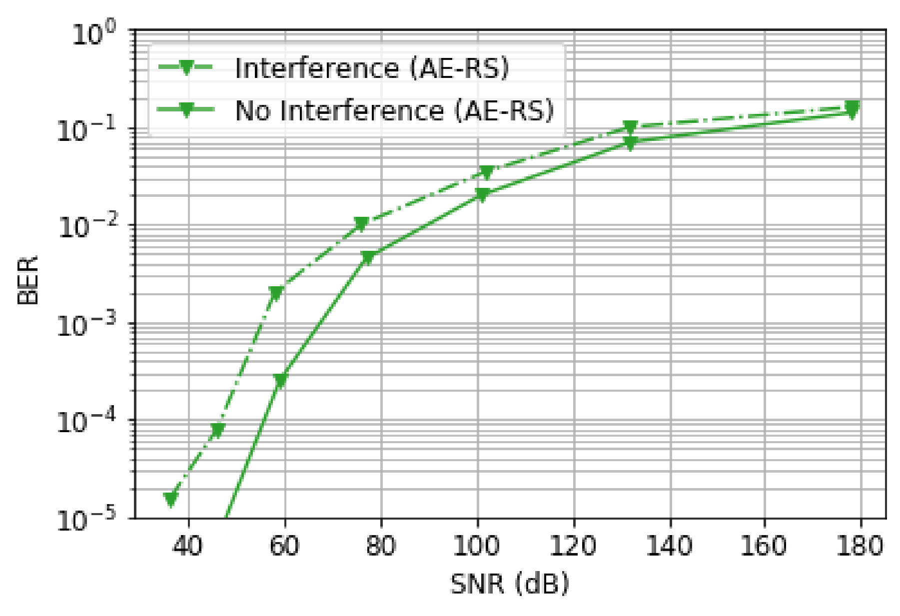

BER and EVM results are for 4 ch-use.

Error Vector Magnitude (EVM) measures how far the points are from the ideal locations at the RX as a result of impairments such as phase noise.

Training for Interference Suppression

Interference training

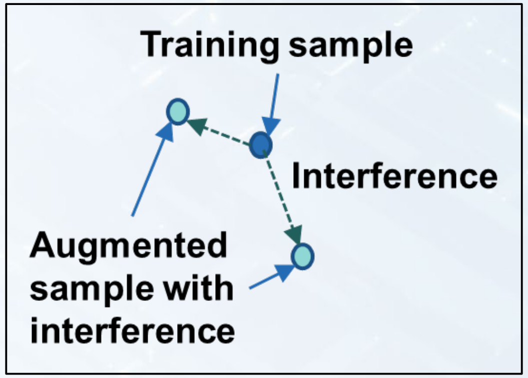

+

Randomized Smoothing

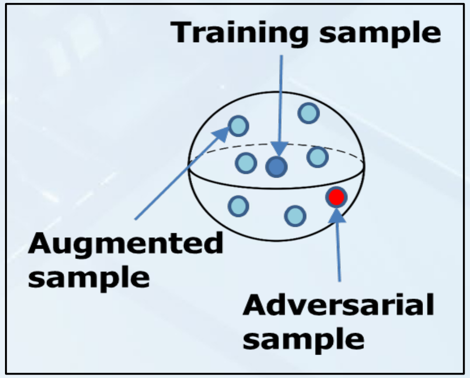

1-symbol **12-symbol interference**

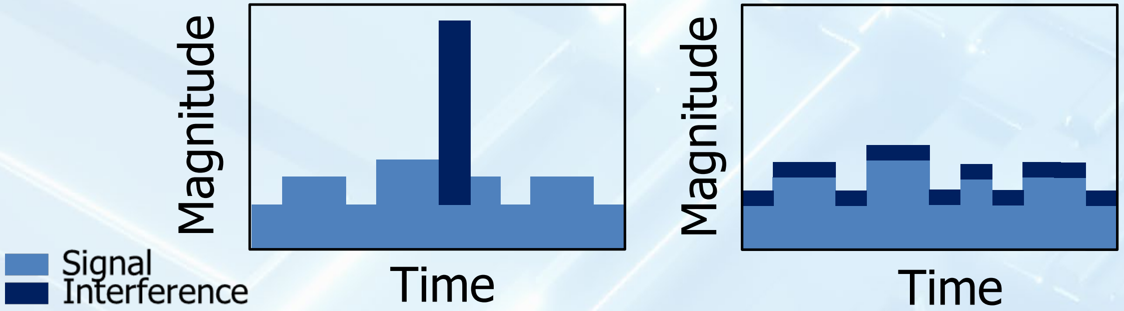

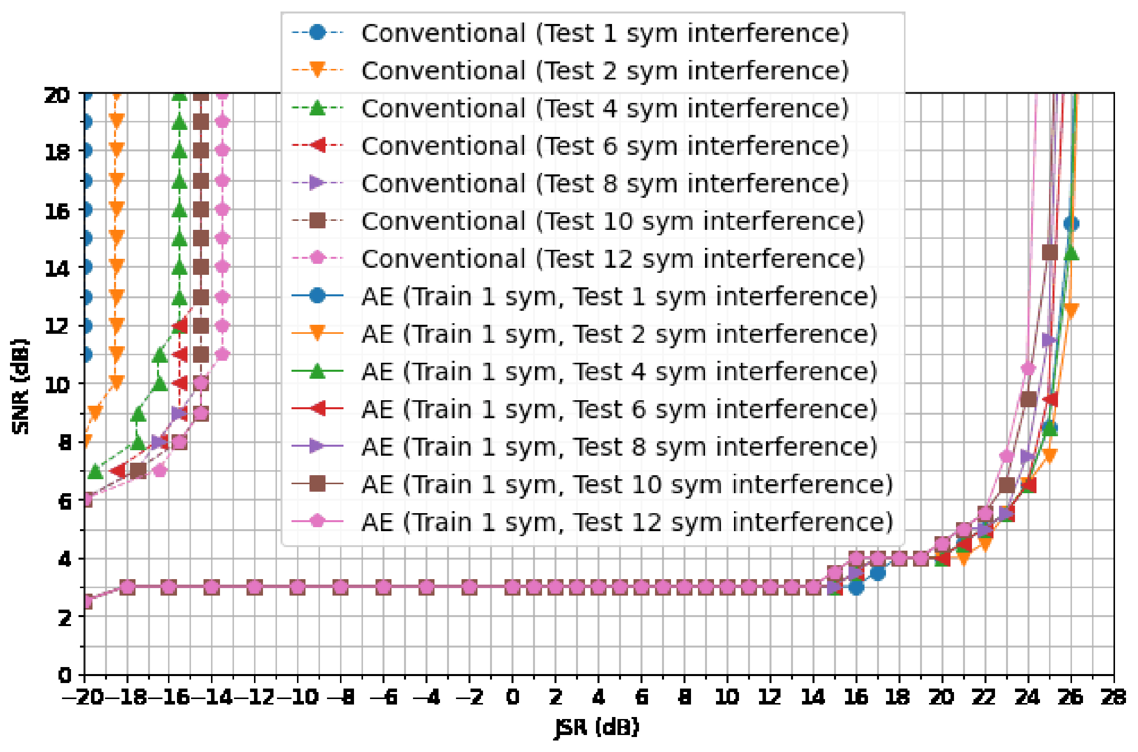

- ❑ The maximum tolerated jamming-to-signal-ratio (JSR) to achieve BER ≤ 1e-2 at a given SNR is compared for 4 ch-use.
- ❑ **AE suppresses >30 dB multi-symbol interference with respect to conventional communications.**

ERI ELECTRONICS RESURGENCE INITIATIVE
SUMMIT
& MTO Symposium

MTO MICROSYSTEMS TECHNOLOGY OFFICE

DARPA

ERI ELECTRONICS RESURGENCE INITIATIVE SUMMIT & MTO Symposium

EMBEDDED AI AT THE EDGE FOR COMMUNICATIONS IN NONSTATIONARY CHANNELS AND BUSY SPECTRUM

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA).

GIL RAZ

SYSTEMS & TECHNOLOGY RESEARCH
RAZ@STR.US

WORK PERFORMED UNDER DARPA SPINN PROGRAM Y.K. CHEN (PM)

G. RAZ, M. CHILENSKI, N. DESTEFANO, C. DOHRMAN, B. DONOVAN, G. GUBASH

I. MATEI, J. DE KLEER

A CHALLENGE IN COMMUNICATIONS AND SENSING

- **Communications and sensing in nonstationary & busy environments require fast / smart decisions at the edge**
- **HF Ionospheric channels provide beyond line of site capabilities, but bandwidth and reliability have historically limited its usage**

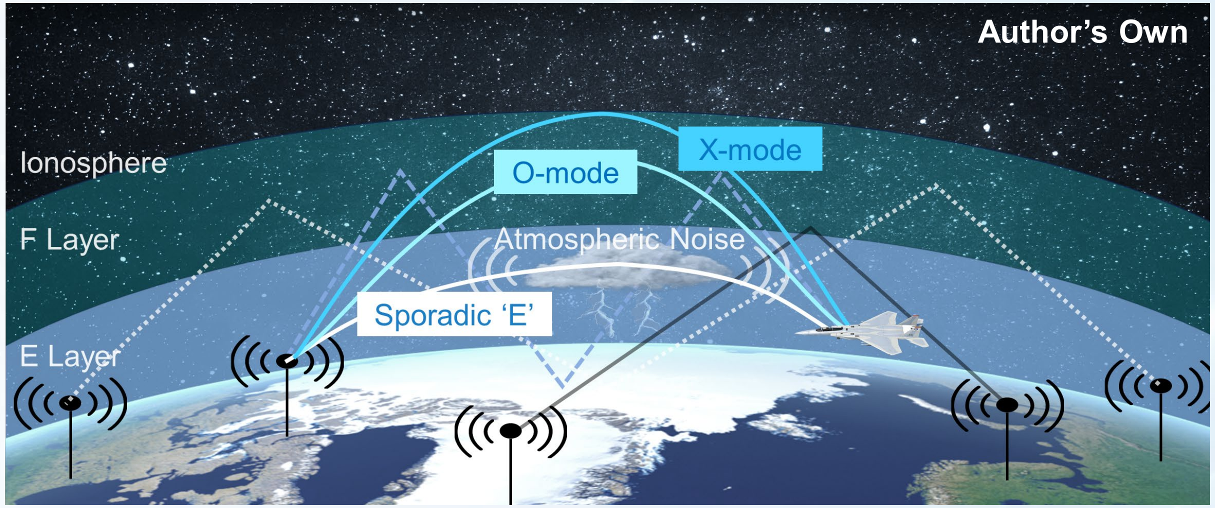

Author's Own

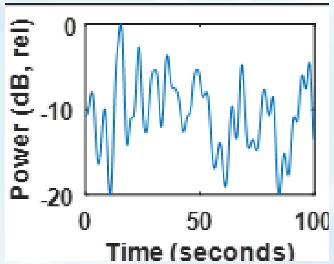

Challenging channel: Refraction from Multiple layers and polarization splitting in ionosphere creates multiple time-varying modes that recombine at the receiver, causing signal fading

Busy spectrum: Ability to sense manmade and natural interference from 1000s of km away created dense spectrum with nonstationary characteristics

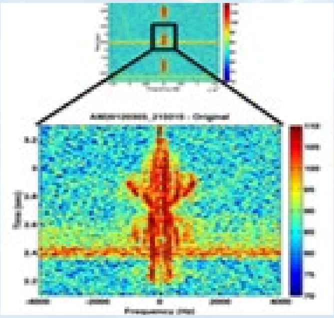

Author's Own

The advent of 3rd wave AI methods allows exploitation of prior insights (e.g., physics, signal structure). Combined with efficient ML architectures and training are well-suited for edge processing with relatively little training to address these challenges.

HIGH-NOTE ARCHITECTURE

HIGH-FREQUENCY COMMUNICATIONS NEURAL NET OPTIMIZATION AND TRAINING ENGINE

- **Use NNs to demodulate symbols in time-varying frequency-selective HF channels more accurately than achievable by state-of-the-art adaptive equalizers**
- **Exploit signal structure informed learning methods to identify and cancel interferers**

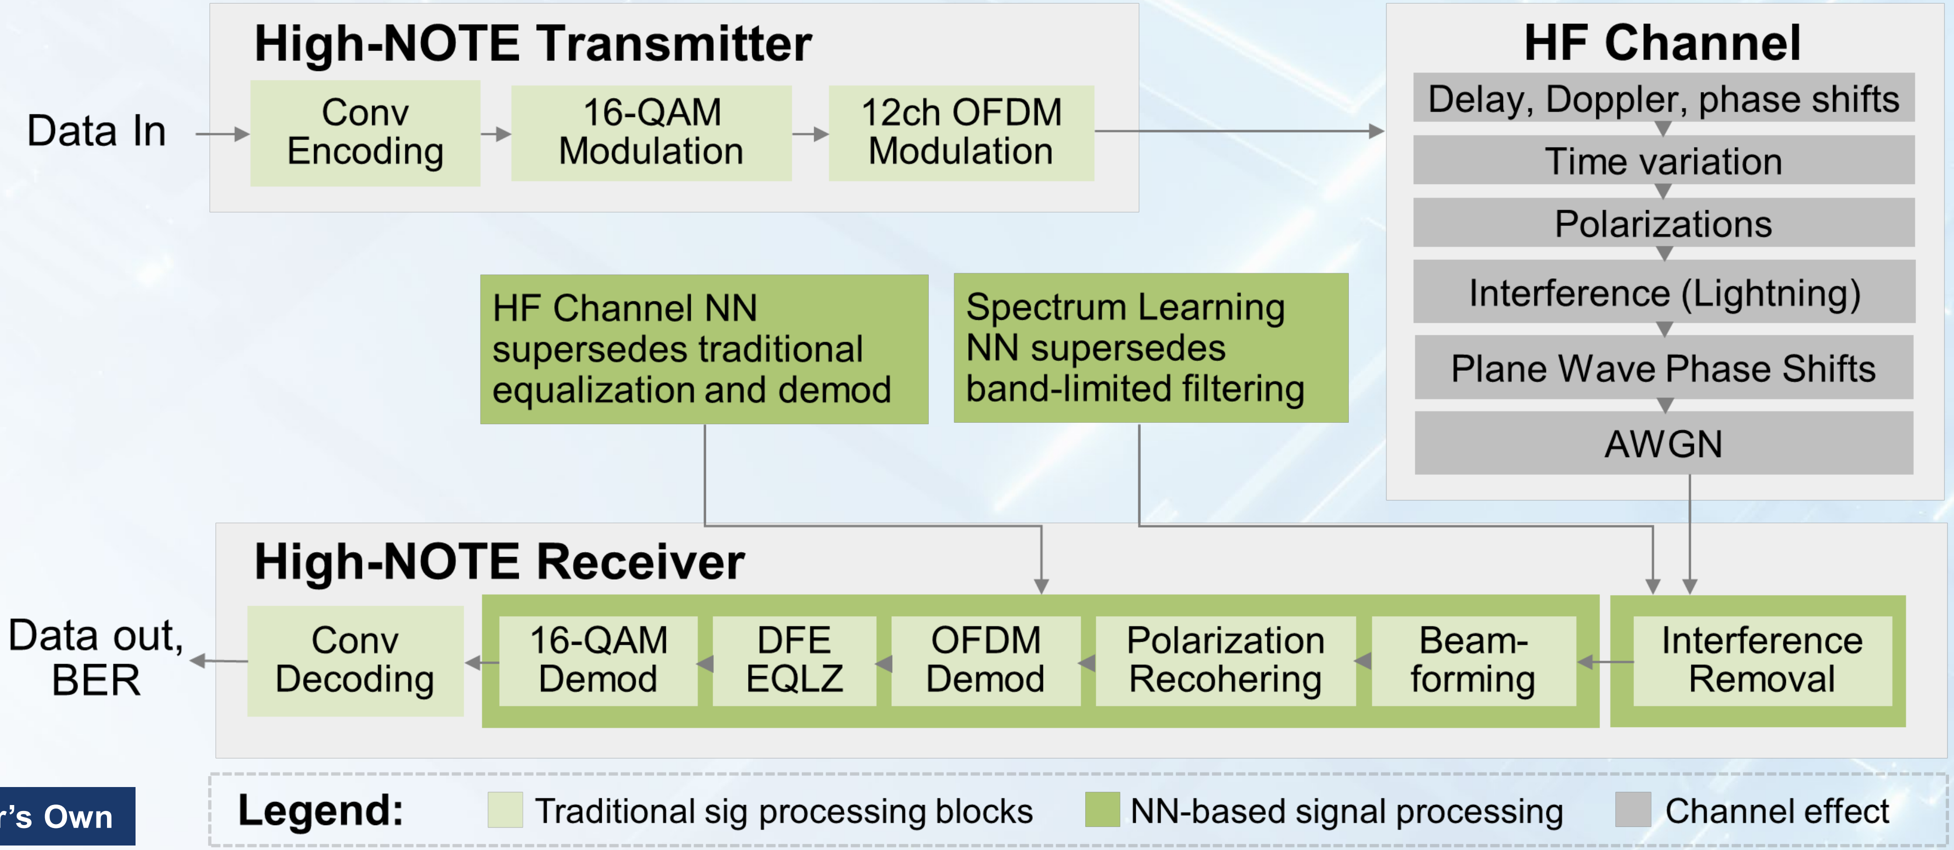

Author's Own

SPECTRUM LEARNING ARCHITECTURES:
COMPARING PURE NN AND HYBRID APPROACHES

Author's Own

- **Hybrid approach: Neural Annotated Signature Pursuit**
  - Combines matching pursuit with a lightweight neural network to label the selected atoms
  - Domain knowledge is encoded in signature dictionary to minimize training data requirements
- **Pure neural network approach: fully convolutional U-net**
  - Single neural network maps received signal directly to estimated SOI
  - Replaces iterative matching pursuit operation with cascaded convolution operations
  - Enabled by novel loss function which imposes domain knowledge during training process

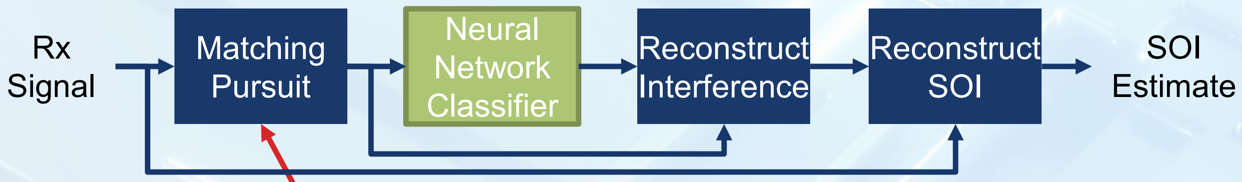

Matching pursuit step requires iterative selection of >100 atoms: limits ability to parallelize processing

Green = neural network-based component

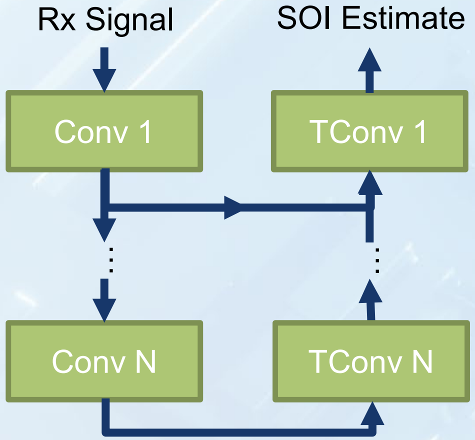

Only iterative operation is sequential application of ~10 layers: convolutions within each layer can be efficiently parallelized

SPECTRUM LEARNING ARCHITECTURES:
RESULTS: INTERFERENCE (LIGHTNING) MITIGATION

Author's Own

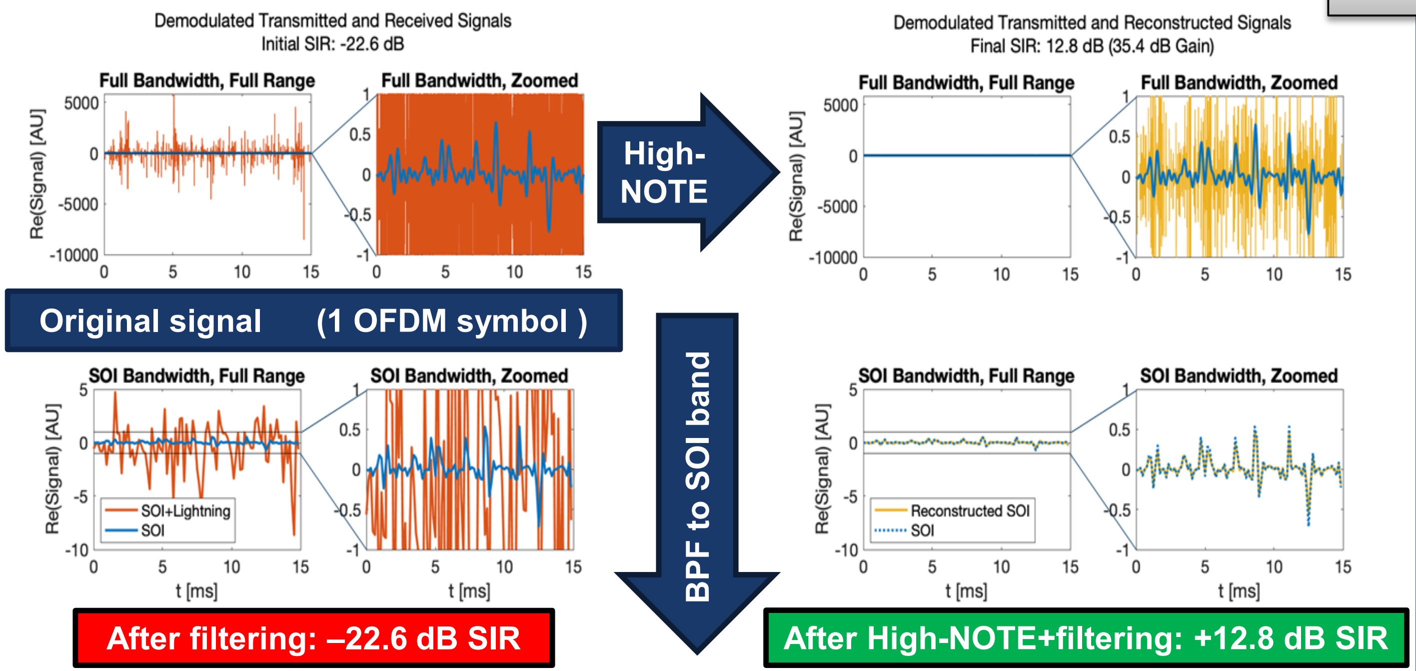

SIR gains of ~35 dB demonstrated for signal with structured interferer (lightning)

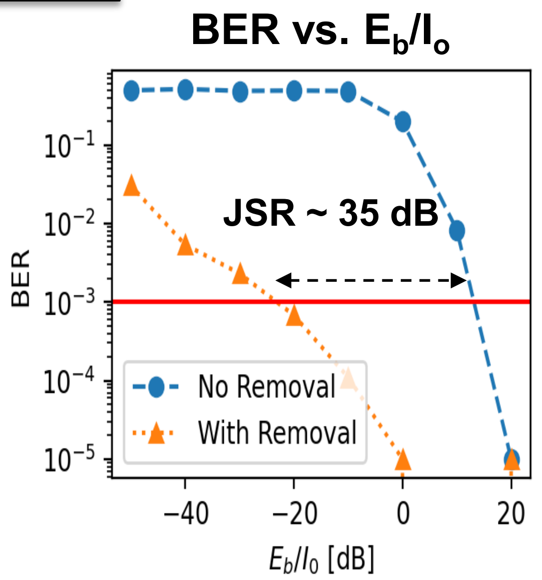

~35 dB improvement validated in BER Monte Carlo simulation results

- **Interference removal successfully recovers low BER even when starting from negative SIR, despite overlap with SOI**
- **Achieved using single-channel Rx signal (no spatial / polarimetric DOFs)**

HF CHANNEL LEARNING – ARCHITECTURE

- **High-NOTE uses NNs to demodulate symbols in time-varying frequency-selective HF channels**
  - Goal: improve on performance of state-of-the-art adaptive equalizers.
- **We trained NN-based symbol decoders for a wideband HF channel**
  - Two NN architectures explored: memoryless (dense NN) and with memory (LSTM-based RNN)
- **Channel estimation achieved using a hybrid transceiver model**
  - Learn the multipath delays, gains, and Doppler shifts using gradient descent algorithms
  - Use sparsity constraints (e.g., L1 regularization) on gains to reduce the model complexity
- **A NN-based predictive model is updated online and used for estimating future channel behavior**

LSTM-based RNN architecture for symbol decoder

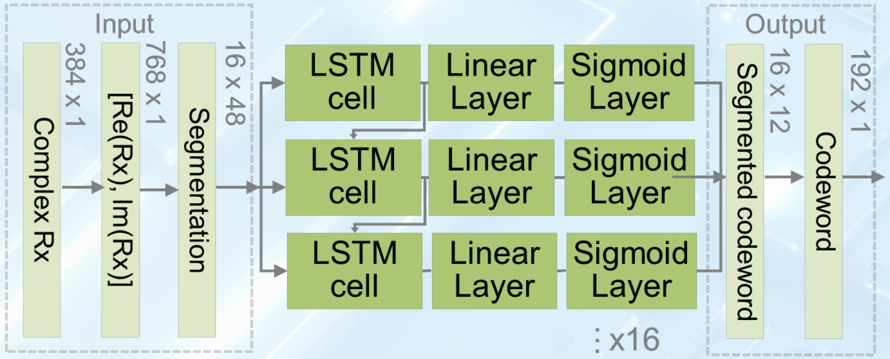

Author's Own

Channel Type	Multipath Spread	Fading Bandwidth (Doppler Spread)
Stationary AWGN Channel	0	0
HF Moderate Conditions**	1 msec	0.5 Hz

*** - Source: Recommendation ITU-R F.520, "Use of high frequency ionospheric channel simulators."*

HF CHANNEL LEARNING – RESULTS

- For stationary channels High-NOTE meets the fundamental theoretic performance bounds
- In narrowband time varying channels with single receiver, High-NOTE improves BER equivalent to 6dB improvement in SNR
- In wideband time varying conditions High-NOTE improves performance dramatically over standard equalization

Stationary HF Multipath

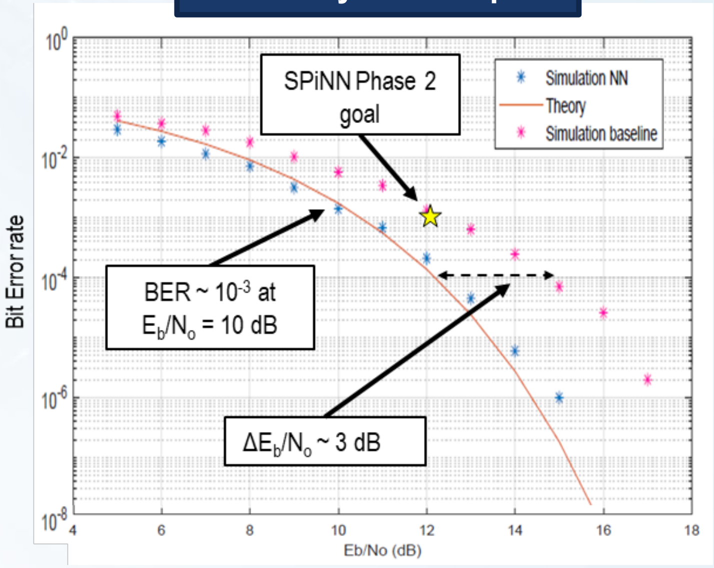

Time-varying HF Multipath

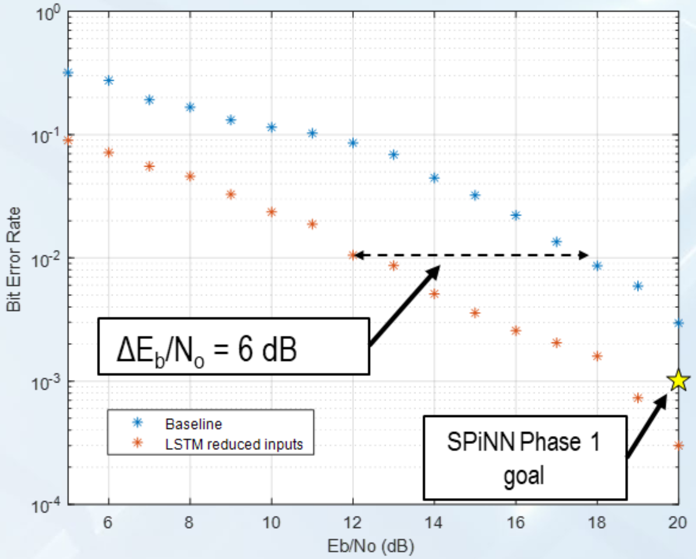

Time-varying HF Multipath with wideband HF channel model

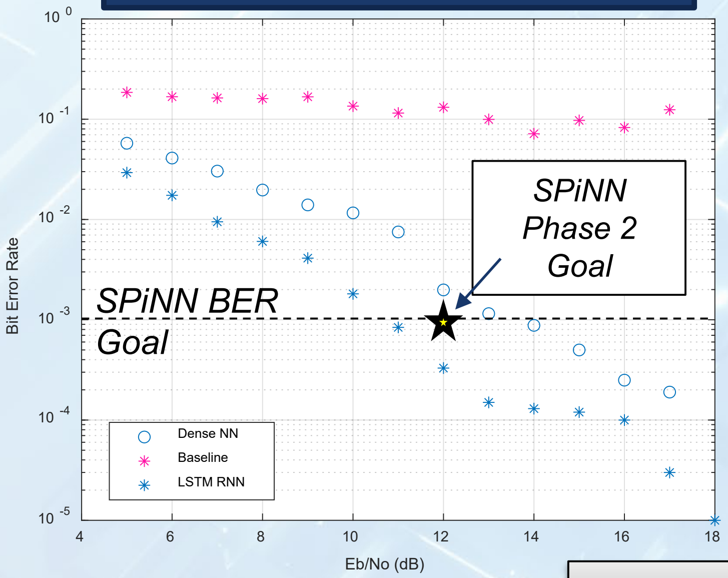

Author's Own

NOTES: 1) Baseline Comparison: distributed feedback adaptive equalizer, tested against identical channel conditions
2) State-of-Art (SoA) Comparison: MIL-STD-188-110C-compliant HF radios, 4800 bps mode, 8-PSK waveform

HIGH-NOTE ARCHITECTURE IMPLEMENTED ON CLOUD-BASED FPGA

Parsimonious architecture!

1476 trainable parameters in channel learning NN - Estimated size in memory 47.2 kB

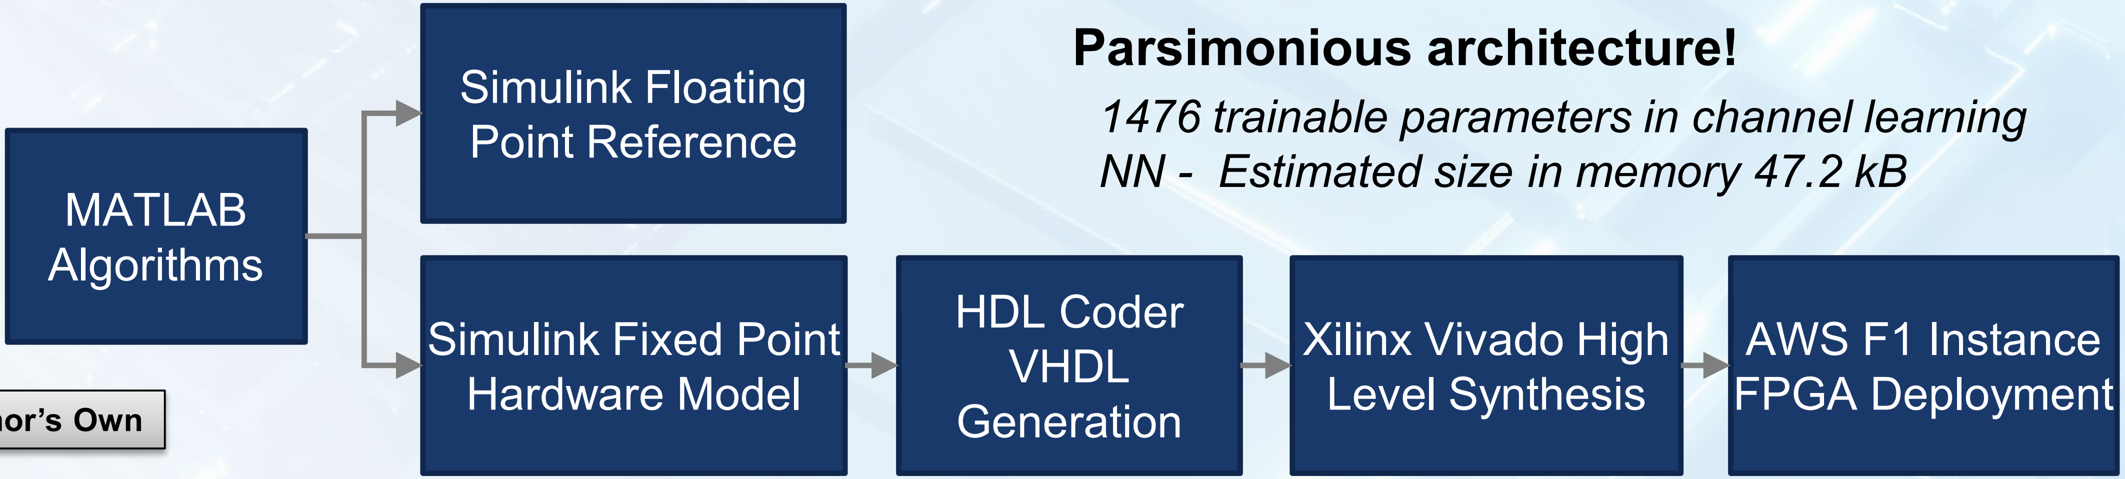

Author's Own

- **Identified Light NN candidate for FPGA implementation on AWS F1 instance**
  - Cloud based Xilinx XCVU9P FPGA - midsize FPGA
- **Utilized HDLCoder code generation tools**
  - Generates VHDL code for target hardware from a Simulink model description. Greatly increases IP prototype design pace by allowing us to skip the RTL development stage
- **FPGA implementation uses fixed-point datatypes with well defined bit widths, leads to quantization error**
  - We balanced resource utilization while maintaining algorithm performance

SUMMARY

- **High-NOTE demonstrates the ability to exploit domain insights along with advances in AI to overcome sensing and communications challenges**
  - Demonstrating these methods in a communications application allows a comparison to both current systems as well as to fundamental theoretical limits
  - Improved structured interference mitigation (>35 dB)
  - Enhanced equalization and demodulation in presence of nonstationary channel
  - Results achievable in single channel processing and arrays / vector-sensors
    - Alleviating the need for larger arrays
- **Framework is not specific to HF or communications and can be applied to other domains:**
  - Mobile communications; Sensing: (e.g., radar, sonar); LPI/LPD signal processing

NGMS HD Matched Filter and Edge Supercompute (ESC)

DARPA HyDDENN

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA).

NORTHROP GRUMMAN

Nishant Zachariah

Jamal Molin

Isidoros Doxas

June 24th, 2021

Overview and Matched Filter Formalism

GOAL : Matched Filtering seeks to uncover an unknown signal buried in white Gaussian noise by careful design of the "optimum" filter.

This can be expressed as

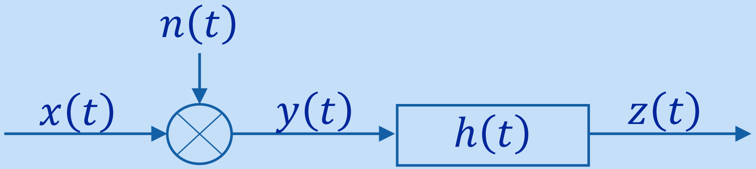

where:

$x(t)$: Uncorrupted Desired Signal

$y(t)$: Corrupted Received Signal

$n(t)$: Additive White Gaussian Noise

$h(t)$: Uncorrupted Desired Signal

$z(t)$: Filtered Signal

It can be shown that the optimal $h(t)$ in the frequency domain is

$$h(\omega) = \mathrm{c}x^*(\omega)e^{-jwt}$$

for some complex scaling factor c. Thus in the time domain, we have that the optimal filter

$$h(t) = x(T - t)$$

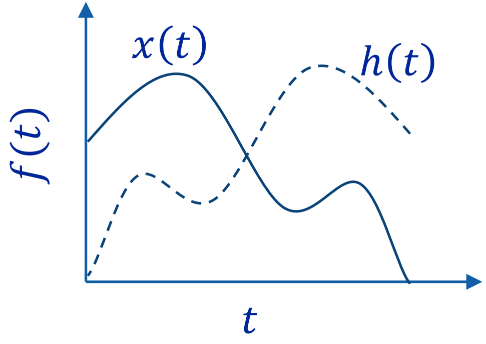

Computation Complexity: HD Matched Filtering vs Traditional Matched Filtering

NORTHROP GRUMMAN

Matched Filtering Pipeline

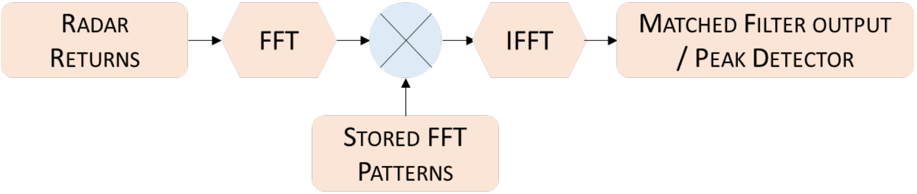

Metric	FFT based MF	Time Based MF
Number of Samples	$N = Tf_s$	
Zero Padded Vector Length	$Z = N + f_s(\tau_{max} - T)$	
Computation Cost (FFT)	$NlogN$	N/A
Computation Cost (IFFT)	$NlogN$	N/A
Precision	$64\ bits$	
Search Grid Size	MP	
Total Memory Overhead	$64MPZ$	$64MPZ$
Total Compute Overhead (Training)	$MPZ(logZ)$	N/A
Total Compute Overhead (Inference)	$MPZ + 2ZlogZ$	MPZ

HD Matched Filtering Pipeline

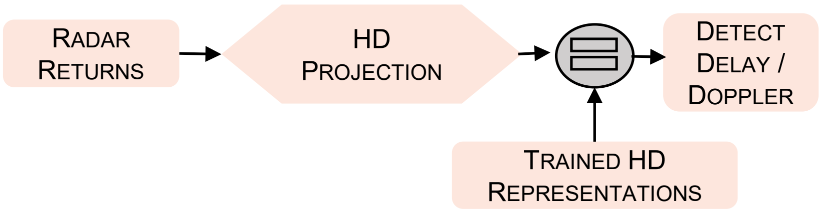

Metric	HD based MF
Number of Samples	$N = Tf_s$
Zero Padded Vector Length	$Z = N + f_s(\tau_{max} - T)$
HD Projection	DZ
HD Decoding	N/A
Number of Classes	MP
Precision	$64\ bits$
Total Memory Overhead	$64MPD$
Total Compute Overhead (Training)	$MPZ(ED)$
Total Compute Overhead (Inference)	$MPD\varepsilon$

T: Duration of Received Pulse f_s: Sampling Rate τ_{max}: Max Time Delay M: Number of Delays P: Number of Dopplers D: HD Dimensionality
E: Number of HD Training Exemplars per Class ε: Computational Complexity Factor (<<1) of Vector Comparison Relative to MACC

HD MF Outperforms Physics-Based MF

Noise Resilience

HD reduces error by much as 50% for Delay / Doppler estimation relative to physics-based MF

Grid Jitter

HD reduces error by much as 80% for Delay / Doppler estimation relative to physics-based MF

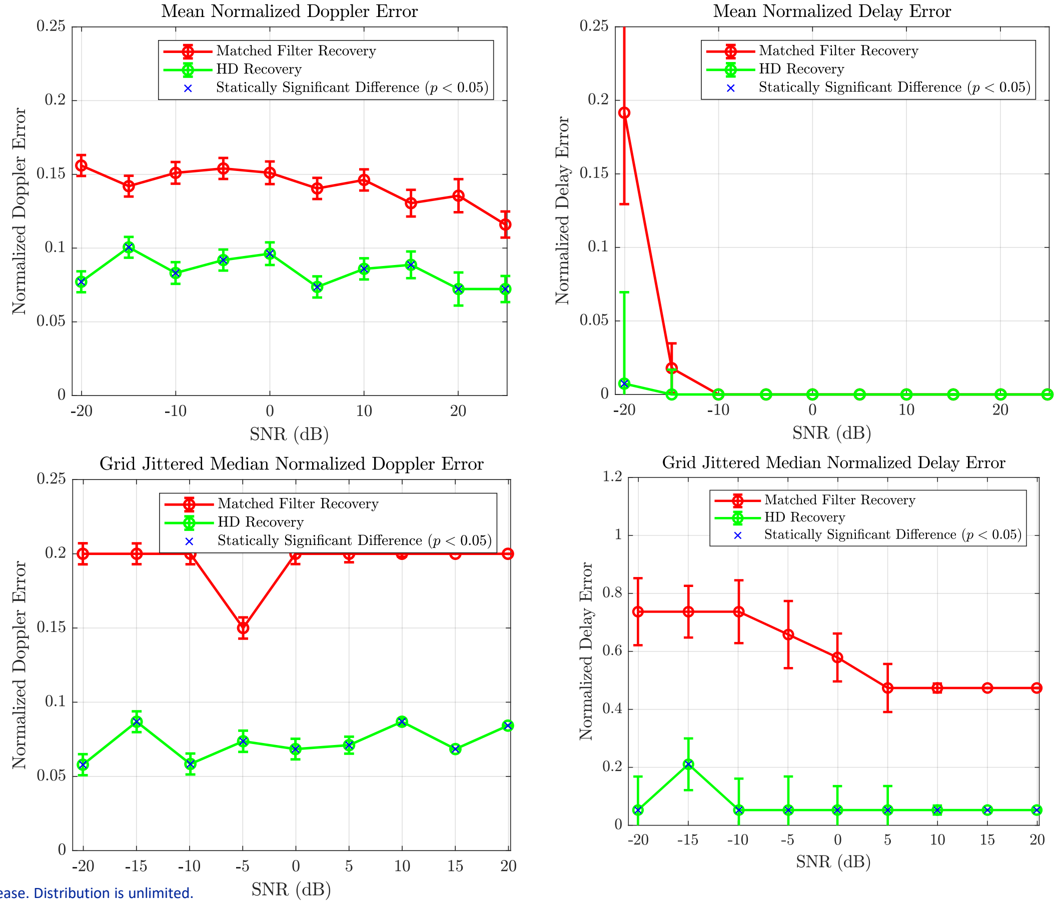

NORTHROP GRUMMAN

NG Compute-In-Memory Edge Supercompute (ESC)

The following computation is used for illustration:

$y = a1*x1 + a2*x2$

$z = y*w$

Step number in the sequence

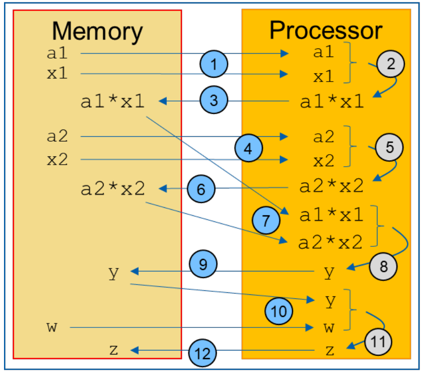

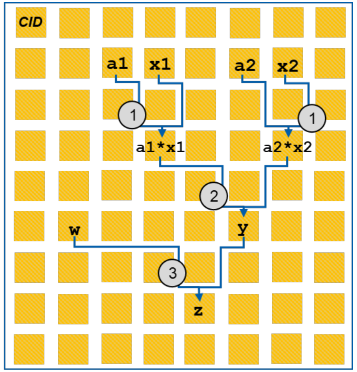

Von Neumann architecture requires constant movement of data to/from memory

Solution: We take advantage of compute-in-memory technology in the analog, charged-based domain to overcome the memory bottleneck

Key Features

- 6.4M CIM Elements
- 50 Gbps High Speed Interface
- x2 NoCs (Network on Chip)
- x2 SiFive CPUs
- 0.4 TOps/sec (1-bit MAC)
- 22.8 fJ per 1-bit MAC

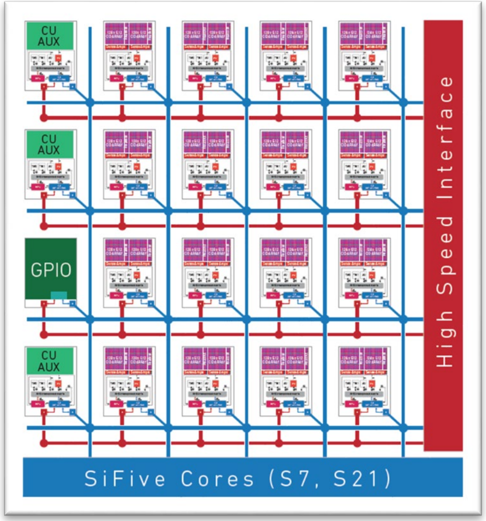

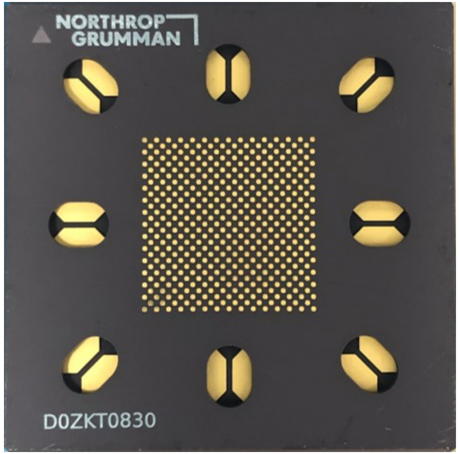

ESC Outperforms Standard Architectures

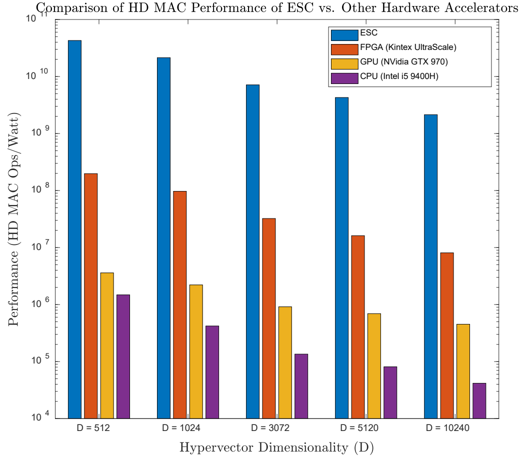

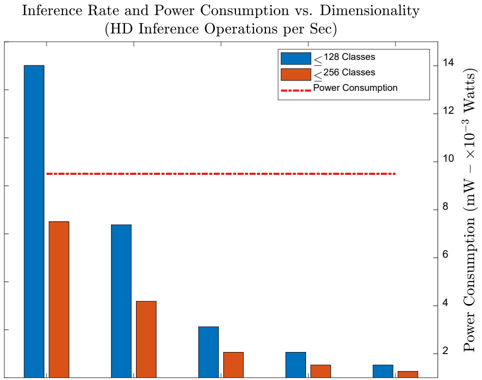

ESC performs > 1000x better than GPU and > 250x better than FPGA in the HD Based Matched Filtering task

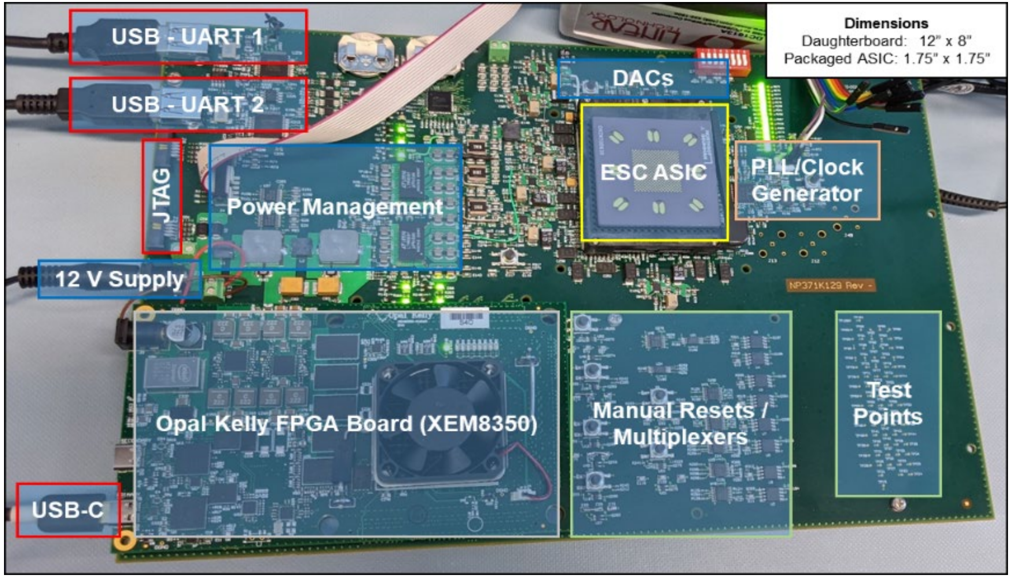

Metric	Measurement
Energy / Op	22.8 fJ
Throughput	0.42 TOps

NORTHROP GRUMMAN