EMBRAPII
Empresa Brasileira de Pesquisa e Inovação Industrial
UNIT

SENAI CIMATEC

Technology, Innovation and Education for Industry

SENAI CIMATEC
SISTEMA FIEB
Federação das Indústrias do Estado da Bahia

The SENAI CIMATEC Campus

SENAI CIMATEC

Highlights

- 4 buildings
- More than 35,000 m²
- Over US$200 million of investment
- 42 competence areas
- More than 800 employee

SENAI CIMATEC Supercomputing Center

Supercomputing Center Timeline

SENAI CIMATEC

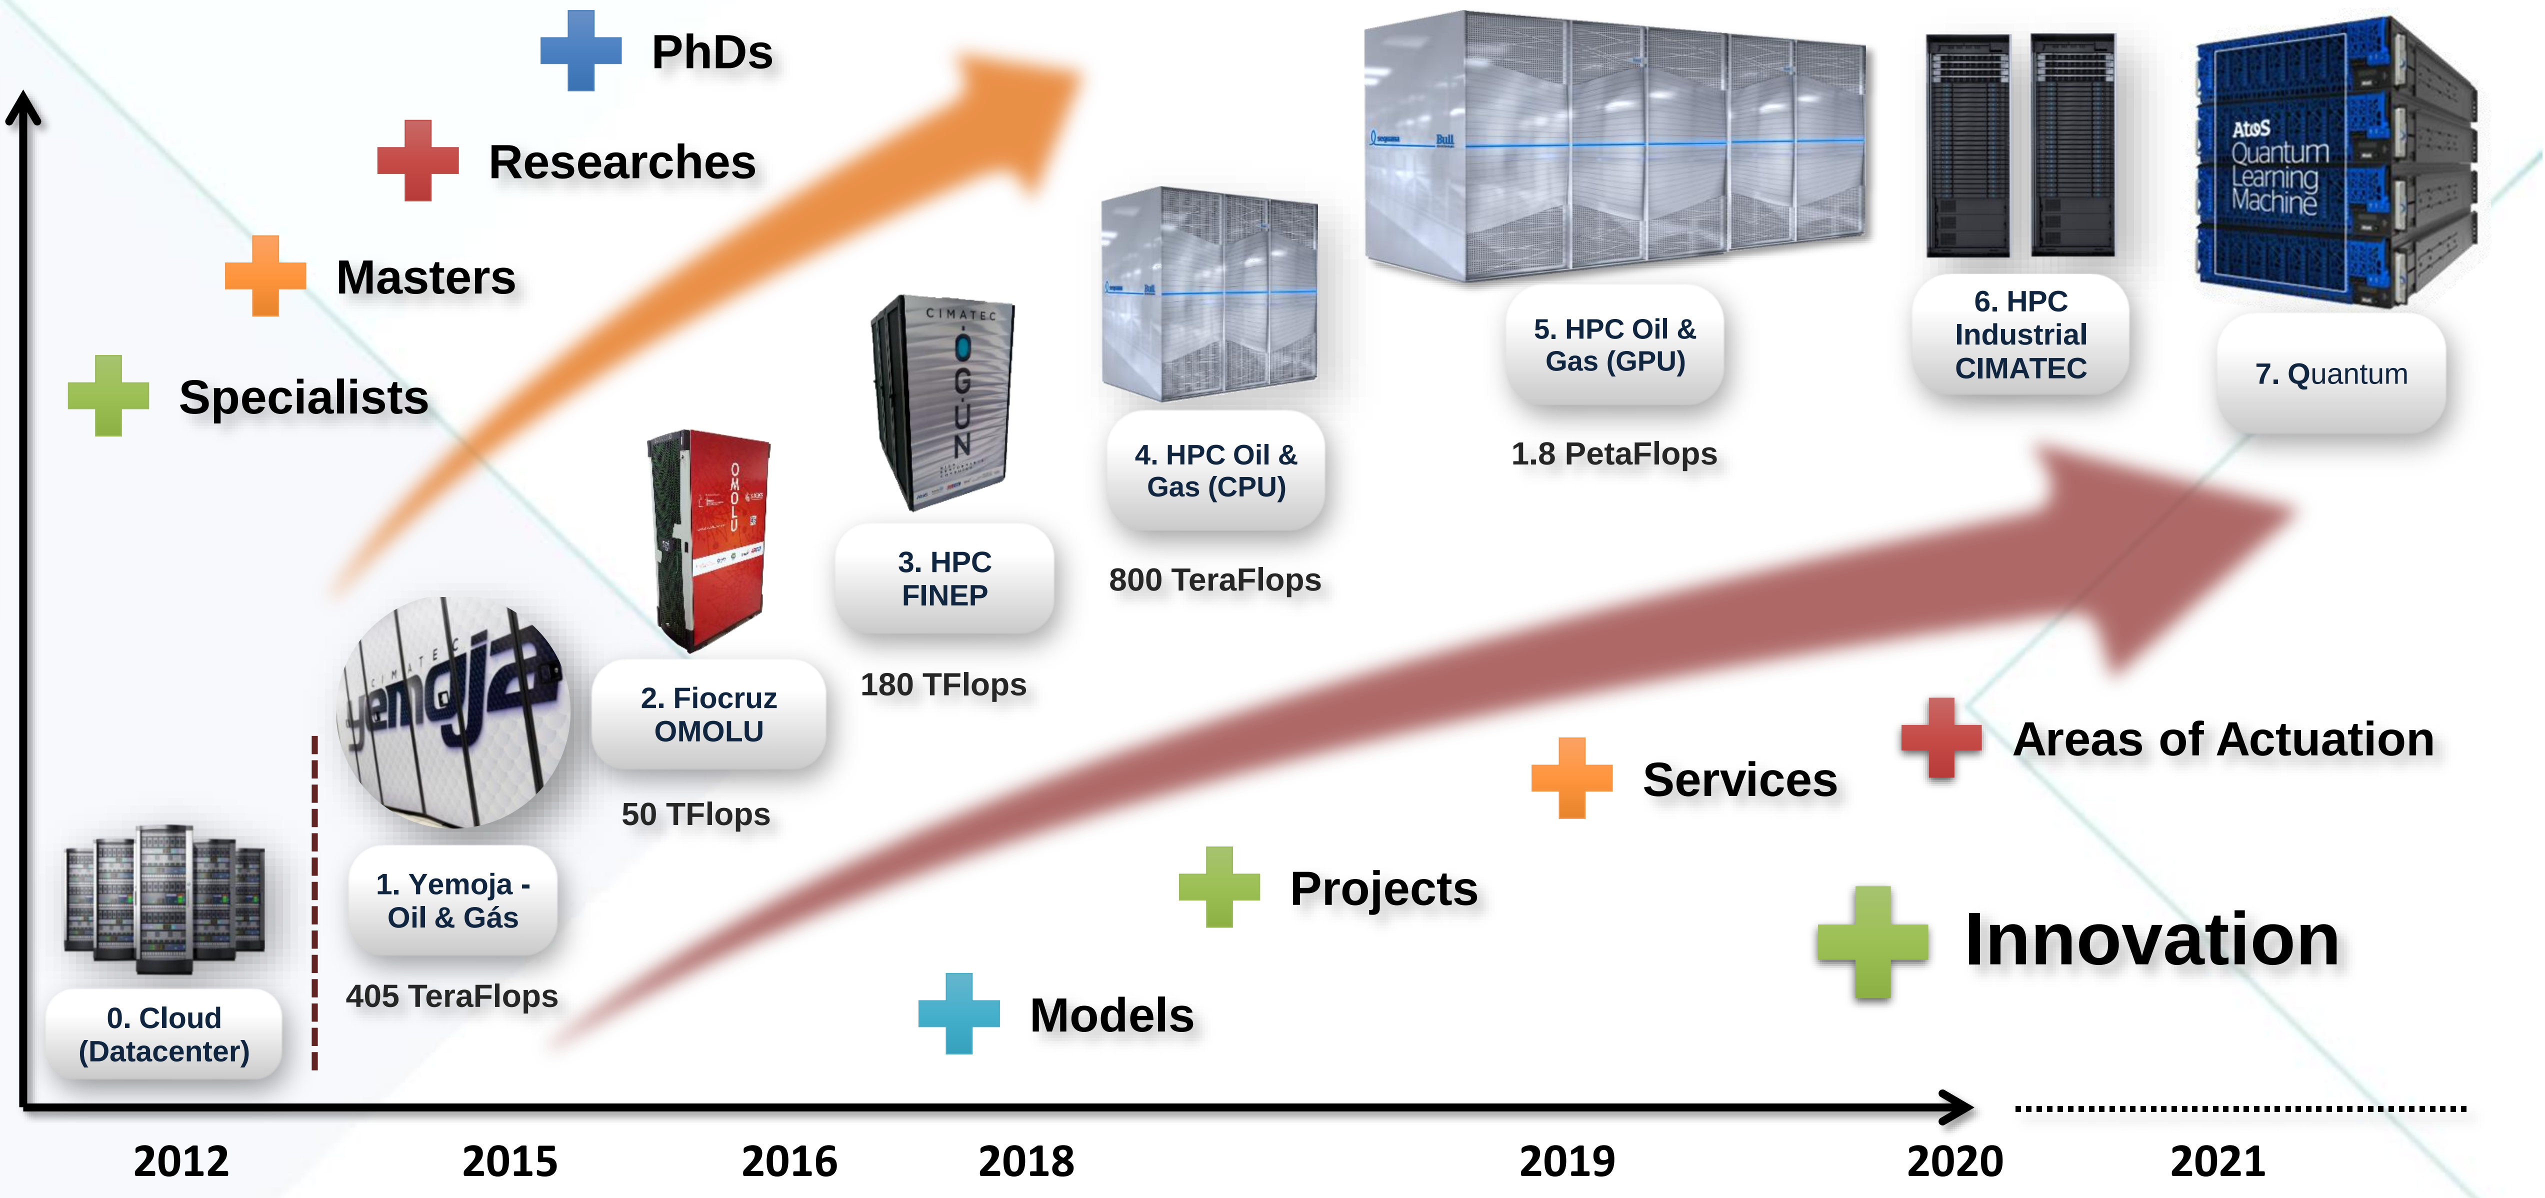

HPC Ògún: Heterogeneous Computing

SENAI CIMATEC

Bull
atos technologies

SENAI CIMATEC SISTEMA FIEB

CS2I

SINAPAD

SINAPAD

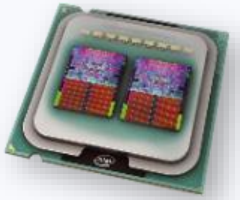

Intel Xeon

- 59 Nós Xeon 6148
- Total: 127 TF

Xeon Phi

- 4 Nós
- 8 TFlops

NVidia GPU

- 2 Nós x 2 P100 Nvlink
- Total: 13 TF

FPGA

- 2 x Arria10
- Total: 2 TFlops

+ARM

+GPU

Performance and Energy Efficiency Analysis of a Reverse Time Migration Design on FPGA

SENAI CIMATEC

Summary

- RTM Brief Review
- Main Computational Challenges
- Reducing Memory Requirements
- Hardware-based Acceleration
- RTMCore's Architecture
- Performance Tests
- Conclusions

João Carlos Bittencourt, Joaquim Oliveira, Anderson Nascimento, Rodrigo Tutu, Lauê Jesus, Georgina Rojas, Deusdete Matos, Leonardo Fialho, André Lima, Erick Nascimento, **João Marcelo Souza**, Adhvan Furtado, and Wagner Oliveira

Overview on Reverse Time Migration (RTM)

SENAI CIMATEC

- RTM is a Seismic Migration technique for accurate imaging of subsurfaces with great structural and velocity complexities
- Largely used in Seismic Imaging Flow for refining boundaries in velocity model building processes (FWI, PSO, Tomography, etc.)

Impulse Source

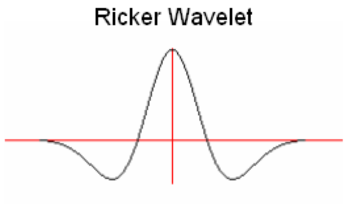

Input Velocity Model

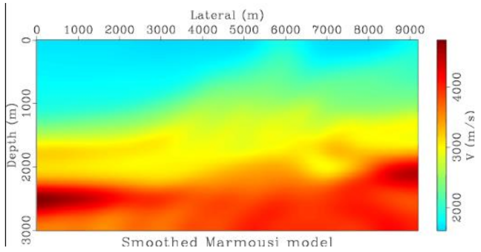

Seismogram Data

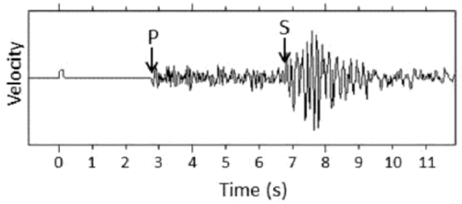

RTM

Enhanced Subsurface Image

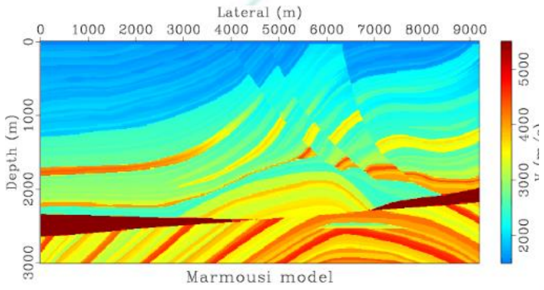

Overview on Reverse Time Migration (RTM)

- Project's specific:
  - 2D RTM
  - Point Source and Receiver
  - Second-order acoustic wave

$$P_{i,j}^{n-1} = 2P_{i,j}^{n} - P_{i,j}^{n+1} + c^2 \Delta t^2 D_k(P_{i,j}^{n})$$

  - Finite-difference based solution
  - P-waves only

Wave Propagation Geometric Layout

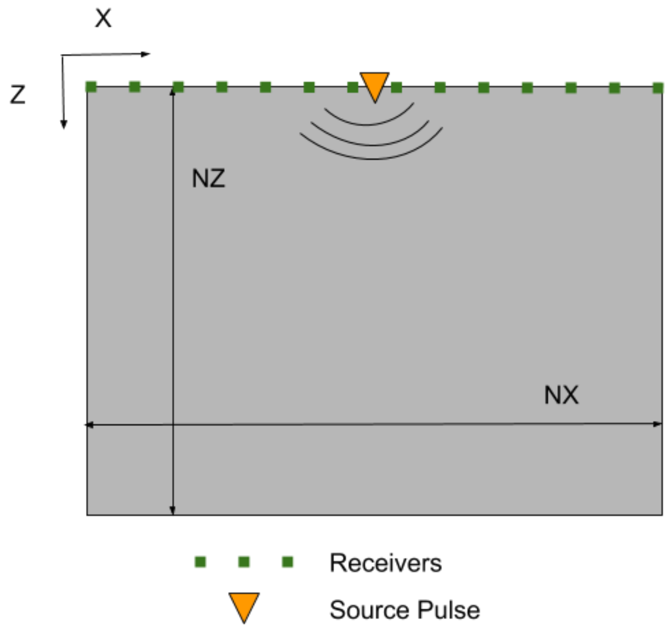

Imaging Condition

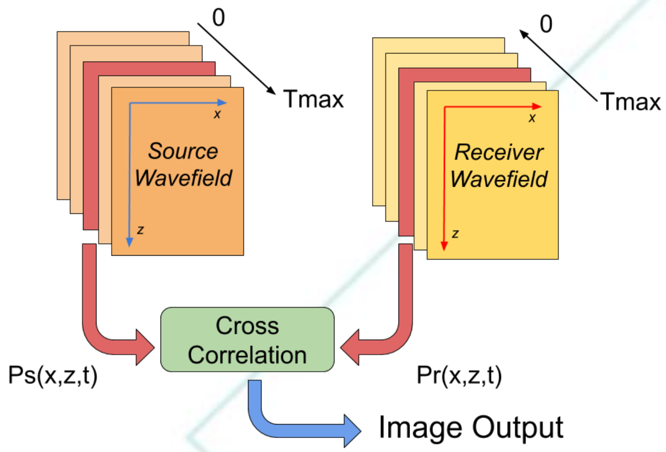

Main Computational Challenges

- RTM requires a massive computation power, memory and storage to migrate even small fields
- Finite-difference (Stencil) operators require several memory accesses
- Migration time and associated energy costs may be prohibitive on production scale

Main Computational Challenges

- Optimization Goals:
  - Reducing memory requirements
  - Reducing migration time and energy consumption
- Design Strategy:
  - Choosing memory efficient algorithms
  - Optimizing memory access
  - Efficient design of heterogeneous computing accelerators on FPGA and GPU

Reducing Memory Requirements

SENAI CIMATEC

- Focus on boundary treatment strategies:
  - Traditional Check Point strategy [1]
  - Random Boundary Condition (RBC) [2]
  - Hybrid Boundary Condition (HBC) [3]
    - During forward propagation, a slice of the pressure field upper boarder is saved, for each time step
    - On backward propagation, the border slices are used for source wave reconstruction
- Test specification:
  - Pluto 2D model (6,960 x 1,201)
  - Number of Shots: 1
  - Time Steps: 12,860

Boundary Condition Mem. Requirements

Strategy	Required Memory (GB)	Image Quality
Checkpoint	311.4	High
RBC	0.25	Low
HBC	**1.04**	**High**

2D Pluto Velocity Model

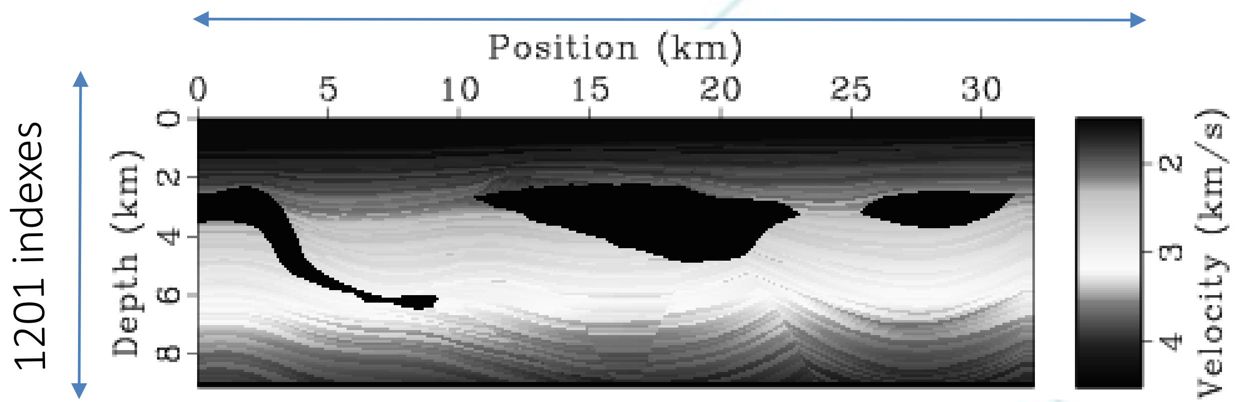

Reducing Memory Requirements

- Fixed-point representation
  - Fixed-point operations generally require less clock cycles
- Word length fixed in 24 bits
- Memory efficiency is increased
- HW/SW Validation
  - A fixed-point reference software model was developed and its outputs were verified

24-bit Fixed-point Numeric Representation*

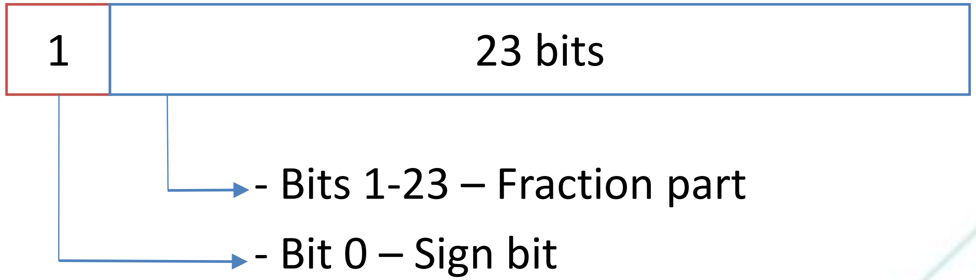

*No Integer part, all values between –1 and 1

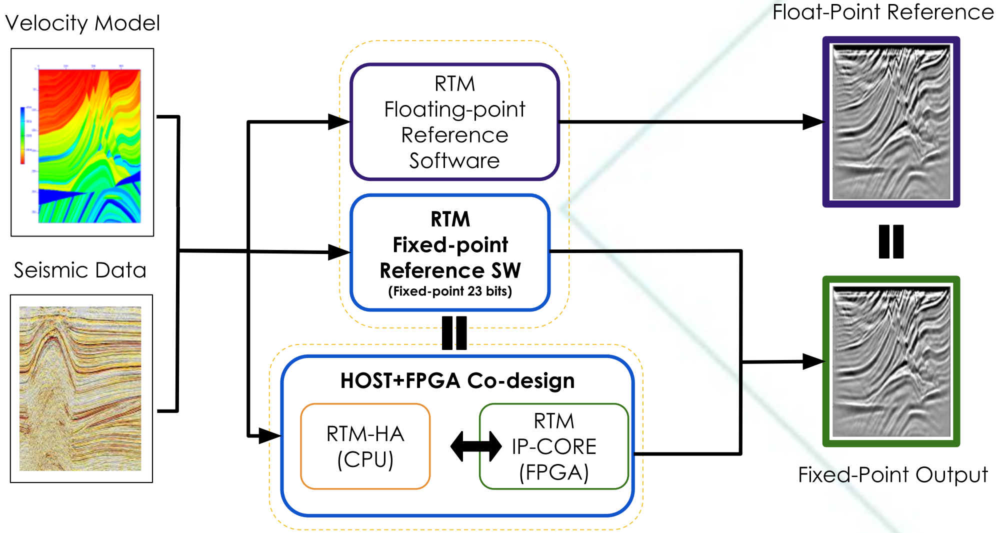

Hardware-based Acceleration

- Complete solution is a hw/sw co-design:
  - RTM CPU-based host application
  - RTM FPGA-based acceleration kernel
- The Host application is responsible for:
  - Configuring kernel parameters
  - Processing input and output data
  - Distributing shots among multiple FPGA
  - Stacking output images
- Each kernel performs an full image migration

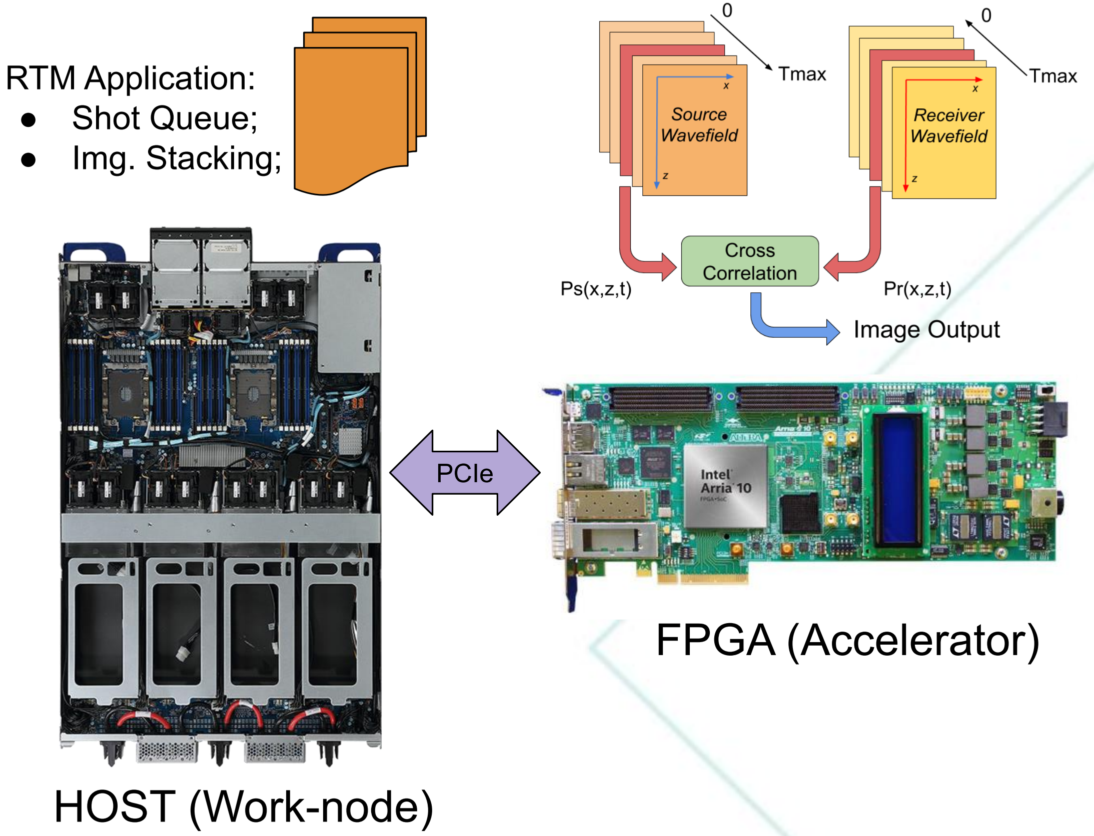

Co-design Architecture

SENAI CIMATEC

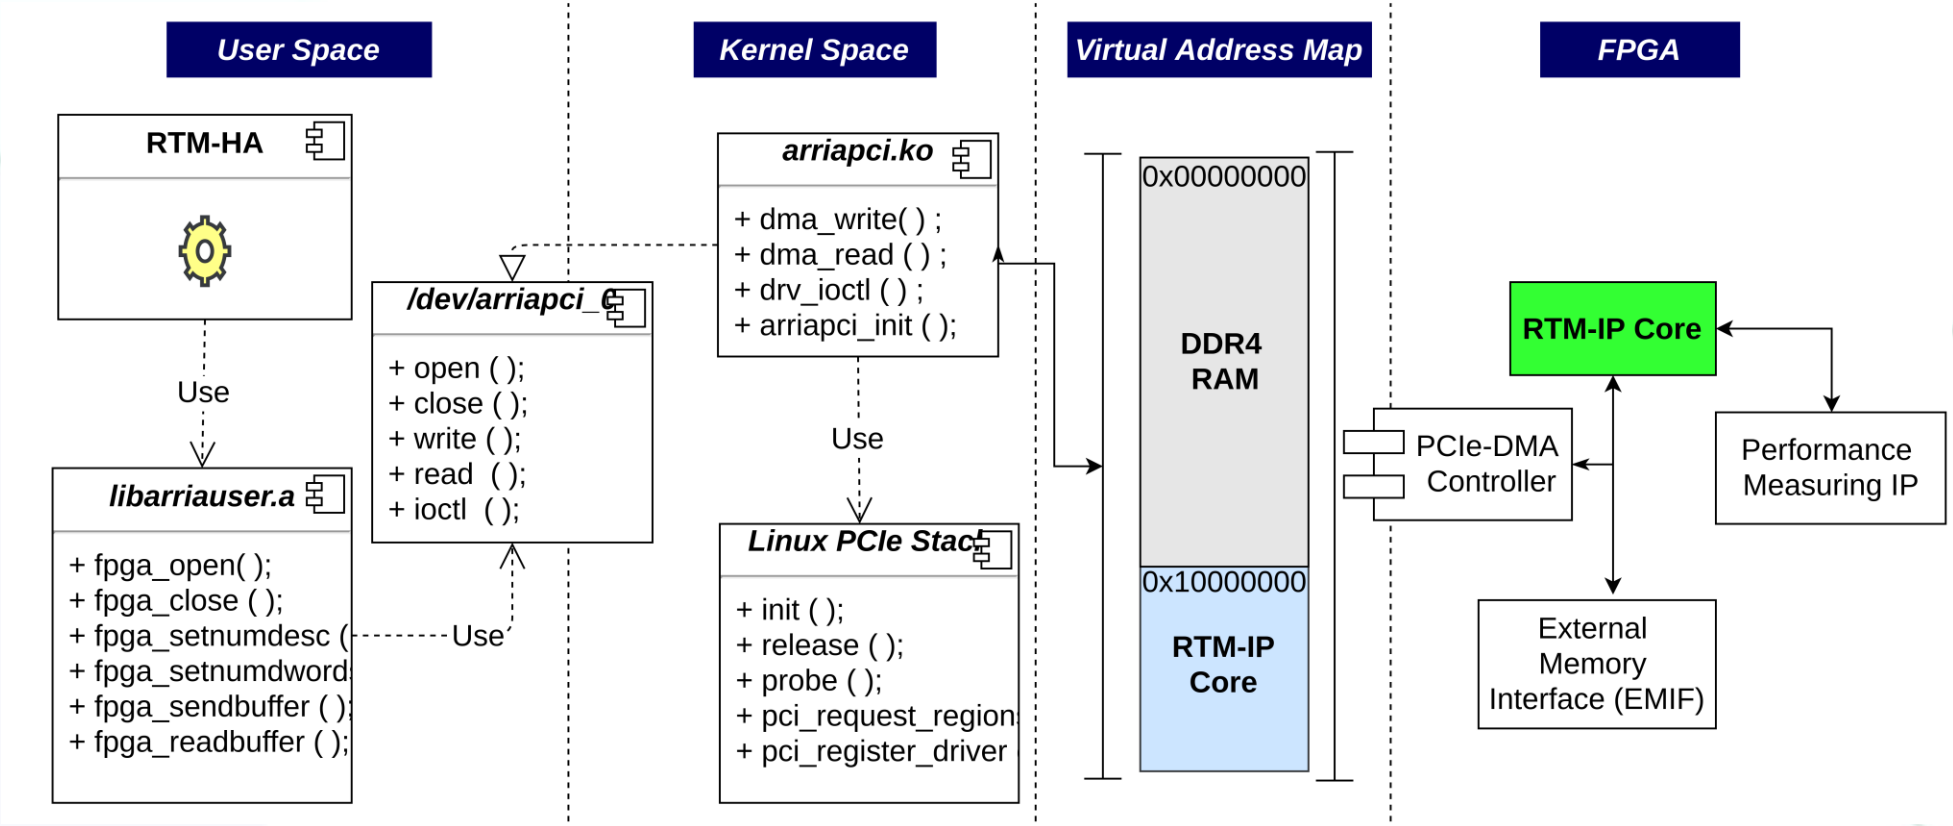

RTMCore's Architecture

- Space Parallelism:
  - All pressure fields of the same time step can be updated simultaneously
  - Multiple Processing Elements update up to 21 pressure points per iteration
- Time Parallelism:
  - Consecutive time steps can be computed in pipeline
  - A total of 24 cascading Pipelined Staged Modules (PSM) stream time iterations

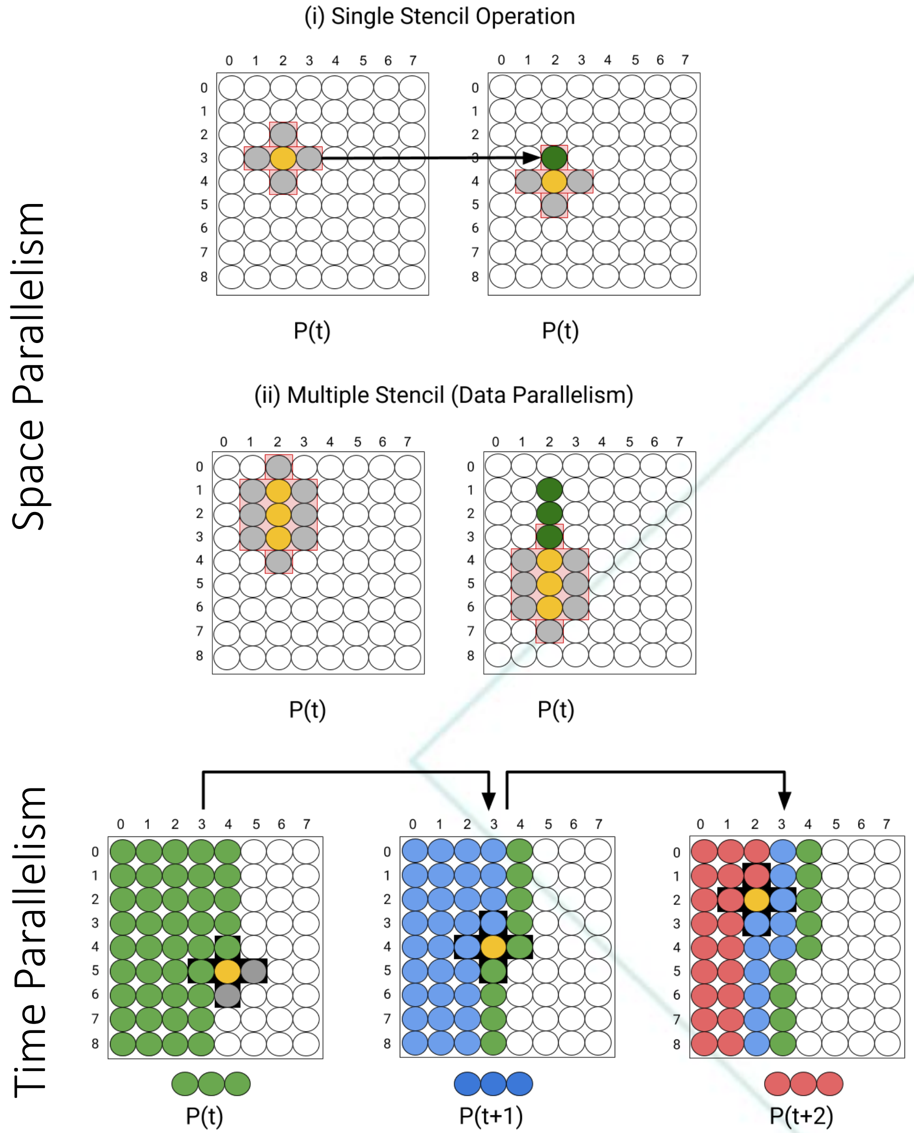

RTMCore's Architecture

SENAI CIMATEC

- The design model is based on research presented in [4]

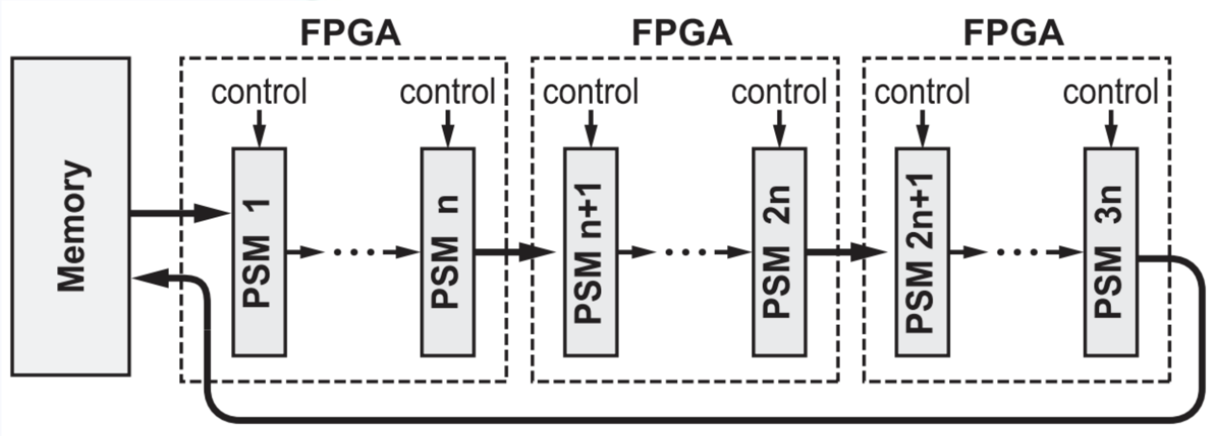

Proposed Kernel Architecture

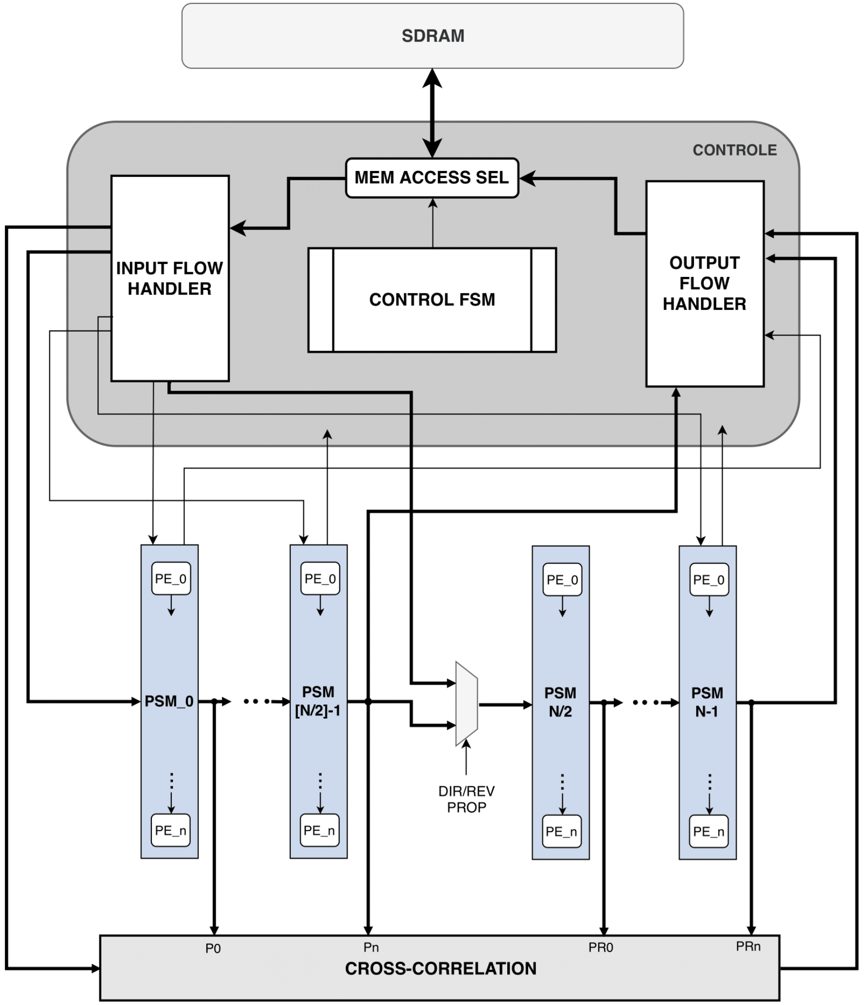

Performance Evaluation

- Evaluation of the FPGA performance against traditional acceleration alternatives, such as GPU and Multithreading
- Two aspects were considered
  - Migration Time: how fast is a seismic shot migrated?
  - Energy efficiency: which accelerator delivers more performance, while requiring less energy?

Migration Time

$$T_{mig} = T_{cpu} + T_{write} + T_{read} + T_{kernel}$$

Consumed Energy

$$P = T/N \sum_{i=1}^{N} (P(i))$$

T = $T_{mig\ (Hour)}$

N = Number of Power Samples

P(I) = Instantaneous Power (W)

Performance Evaluation

SENAI CIMATEC

- RTM implementations for performance comparison:
  A. Serial CPU: used as target reference for speed up analysis
  B. Multithread CPU: 40 CPUs computing pressure fields in parallel for each time step (space parallelism)
  C. GPU CUDA: NVidia's Titan X (11 TFLOPs) exploring massive space parallelism
  D. FPGA: RTM kernel exploring both space and time parallelism

Multithread CPU

NVidia's Titan X

Intel's Arria 10 Dev. Kit

Performance Evaluation

SENAI CIMATEC

- Power Measuring Methodology
  - A power meter device was placed between power supply and host
  - Both host and device power were measured during RTM executions
  - Power meter device was configured to collect samples at 10Hz
  - Only GPU and FPGA power were measured

Energy Measuring Setup

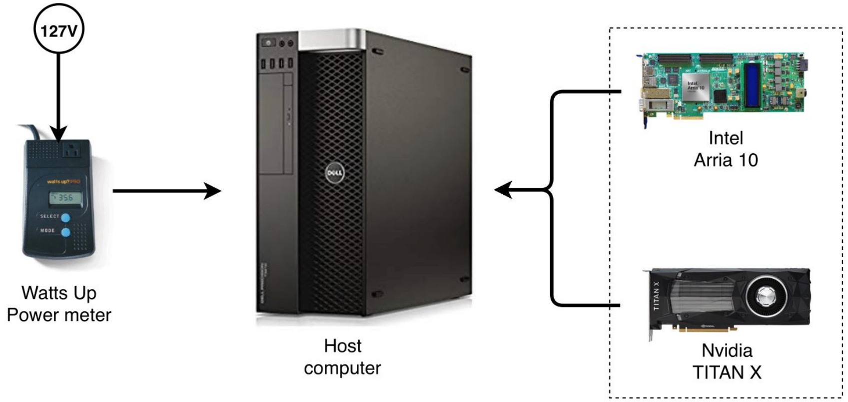

Example: 1 Min. Migration Samples

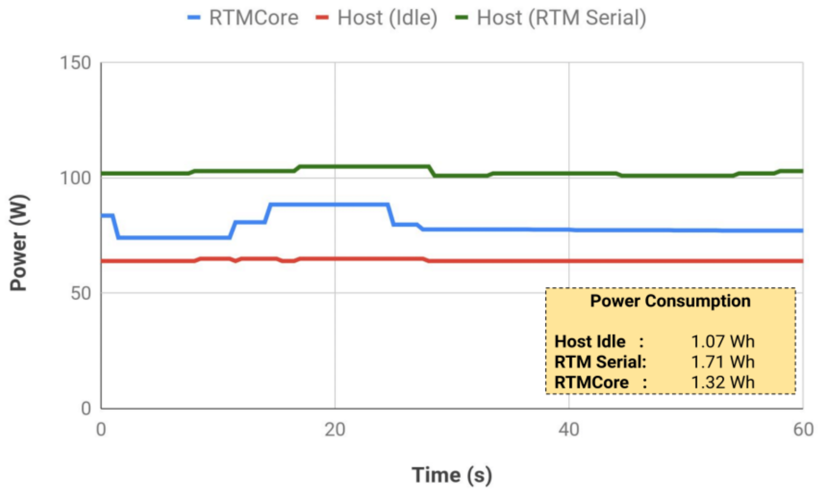

Performance Results

- Input Parameters
  - Pluto 2D (6,960 x 1,201)
  - 12,860 Time steps
  - Shot Position: 3,480 x 0
  - Number of Shots: 1
  - Overall Workload: 1.4 GB

- Efficiency measured in Speedup/Wh

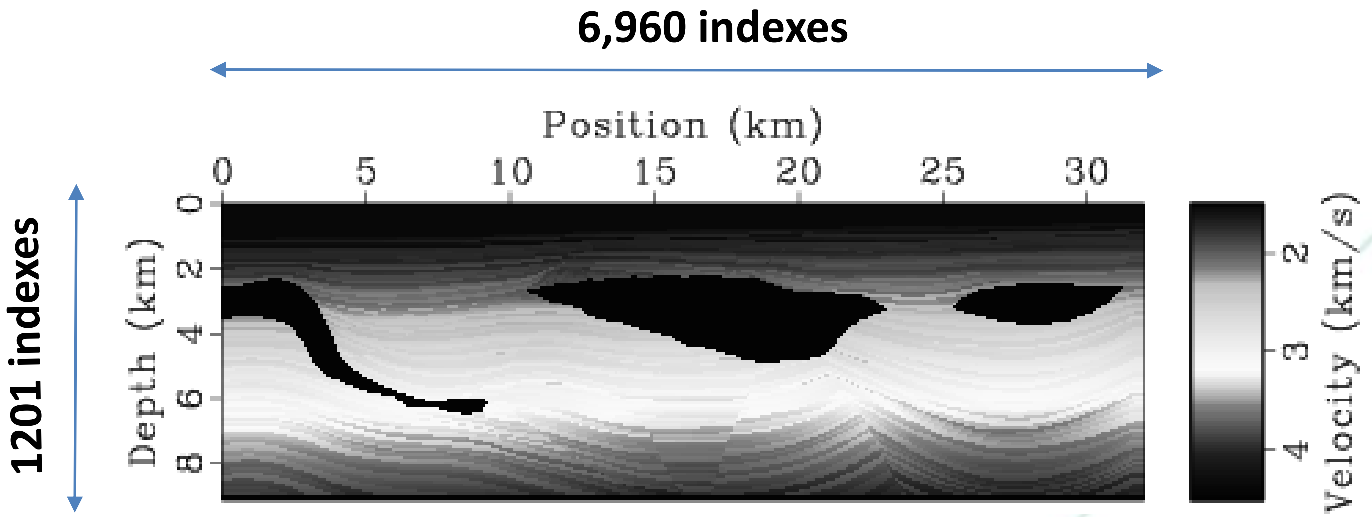

Performance Results

Implementation	Runtime (s)	Speed up	Energy (Wh)	Efficiency
Serial CPU	**21,873.85**	1	-	-
Multithread	2,429.5	9	-	-
GPU Titan X	**182.7**	124	36	3.44
FPGA Arria 10	194	112	20	5.60

Concluding Remarks

SENAI CIMATEC

- Scalability of the solution lies in the parallelization of shots
    - Multiple FPGA boards in one or more compute nodes
- Higher scalability can be achieved by exploring temporal parallelism
    - Increasing the number of Pipeline Stage Modules
    - More iterations could be computed in parallel
- Exploration of fixed-point computation
    - Possibility to explore such method in 3D stencil operators

Concluding Remarks

SENAI CIMATEC

- Speedups of 112x can be achieved, when compared to a Sequential CPU implementation
  - GPU is only 9% faster
  - Consideration: FPGA achieved such a performance with 8 times lower frequency
- Although the design present lower speed up compared to GPU, our FPGA accelerator achieved better energy efficiency
  - The power consumption when compared to a GPU has been reduced up to 55% with an efficiency 60% greater

Acknowledgments

References

SENAI CIMATEC

[1] Symes, William W. "Reverse time migration with optimal checkpointing." *Geophysics* 72.5 (2007): SM213-SM221.

[2] Clapp, Robert G. "Reverse time migration with random boundaries." *Seg technical program expanded abstracts 2009*. Society of Exploration Geophysicists, 2009. 2809-2813.

[3] Liu, Hongwei, et al. "Wavefield reconstruction methods for reverse time migration." *Journal of Geophysics and Engineering* 10.1 (2012): 015004.

[4] Sano, Kentaro, Yoshiaki Hatsuda, and Satoru Yamamoto. "Multi-FPGA accelerator for scalable stencil computation with constant memory bandwidth." *IEEE Transactions on Parallel and Distributed Systems* 25.3 (2013): 695-705.

Thank You

SENAI CIMATEC

João Marcelo Silva Souza

joao.marcelo@fieb.org.br

SENAI CIMATEC

FIEB | www.fieb.org.br

+55 (071) 3462-8449

SENAI CIMATEC SISTEMA FIEB

Federação das Indústrias do Estado da Bahia