Pattern-matching Unit for Medical Applications (PUMA)

22nd IEEE Real Time Conference

Speaker: Orlando Leombruni

An interdisciplinary approach to the problem

Istituto Nazionale di Fisica Nucleare

Expertise on **highly innovative hardware solutions**

IMAGO7

IRCCS FONDAZIONE STELLA MARIS

Ground-breaking research on modern medical imaging techniques

Introduction: (qualitative) MRI

- Current state-of-the-art
- MR signal acquisition done when system is in steady-state after applying magnetization
- Generate high-resolution contrast image
  - Voxels do not have a quantitative meaning
  - Image must be interpreted by a radiologist
- Quantitative reconstruction is possible by repeated acquisitions (long time)

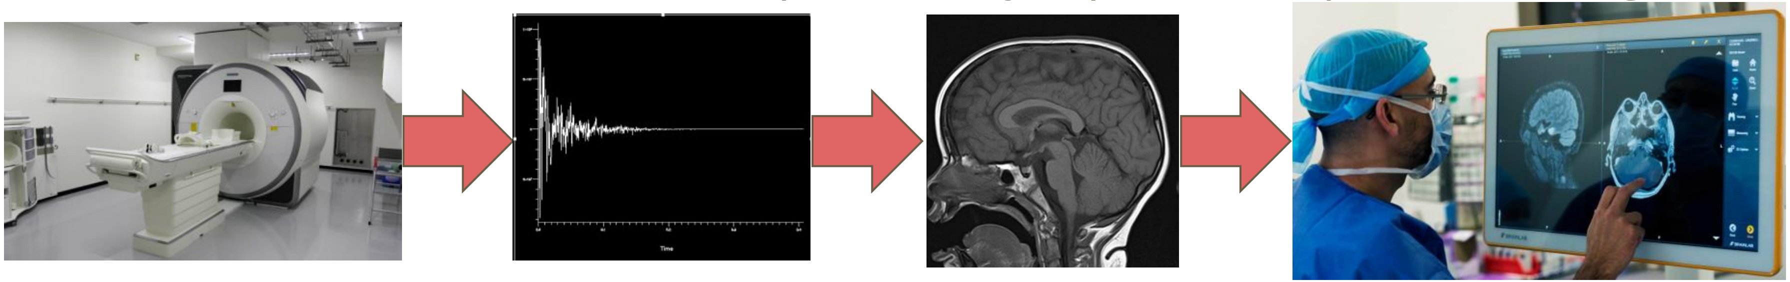

Intro: Magnetic Resonance Fingerprinting (MRF)

- Modern approach
- Exploit the *transient state* in-between frequent repeated acquisitions
- Extract *quantitative* information from an MRI signal in a short time
- Each point of the map is associated to specific values for the parameters
- Voxels now have a quantitative measure associated to them!

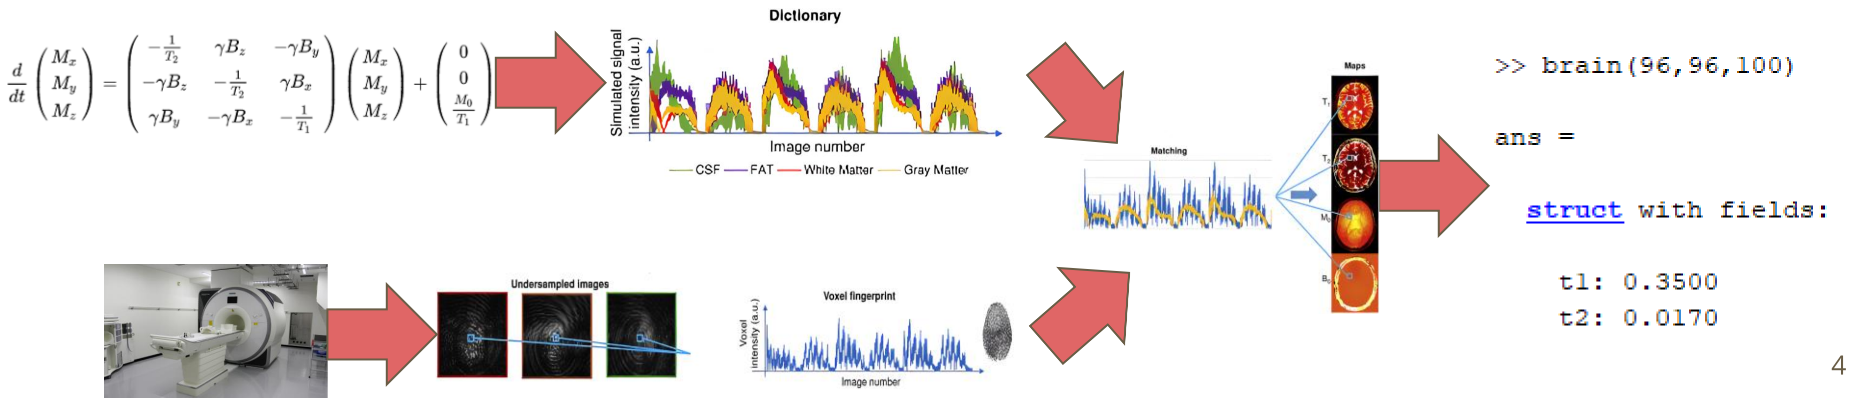

Fingerprinting state-of-the-art: the dictionary

- All current fingerprinting techniques use a *dictionary of simulated signals*
- For each possible combination of values of the chosen tissue parameters, simulate a signal
    - as close as possible to the “true” signal that would be acquired from the MR machine when it encounters a tissue with those specific values of the parameters
    - using Bloch’s equations

$$\frac{d}{dt}\begin{pmatrix} M_x \\ M_y \\ M_z \end{pmatrix} = \begin{pmatrix} -\frac{1}{T_2} & \gamma B_z & -\gamma B_y \\ -\gamma B_z & -\frac{1}{T_2} & \gamma B_x \\ \gamma B_y & -\gamma B_x & -\frac{1}{T_1} \end{pmatrix} \begin{pmatrix} M_x \\ M_y \\ M_z \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ \frac{M_0}{T_1} \end{pmatrix}$$

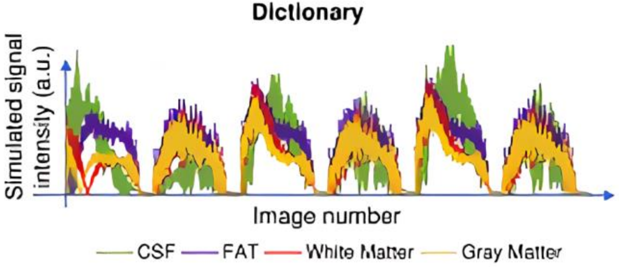

Fingerprinting state-of-the-art: computation

- Both signal voxels and dictionary entries are vectors of complex numbers
- for *each* voxel *i*, the standard method calculates the **dot product** between the *i* vector and the vector of ***each entry*** of the dictionary, to choose the dictionary entry that produces the **maximum** *dot product*
  - If entry *j* of the dictionary, generated from the parameters $p^j{}_1, \dots, p^j{}_n$, is such that $v_i \times v_j$ has the highest dot product for pixel *i*, then $p^j{}_1, \dots, p^j{}_n$ are assigned as the tissue parameters of *i*
- Dot product is a computationally easy operation (few μs even on general purpose processors)
- The problem is the sheer number of dot products to be computed even in simple cases!

Some numbers

- A dictionary built from two parameters at a medium granularity has ~200k entries, a brain MR image is made by 200x200x200 = $8 \cdot 10^6$ voxels
    - → $1600 \cdot 10^6$ dot product operations!
    - On a fast consumer PC (~5μs per dot product), this amounts to **8000 seconds** (>2hrs)!
    - On more powerful specialized hardware (GPGPUs): ~350 seconds
- Adding more parameters (or increasing their granularity) to the dictionary greatly increases its size
    - 4 parameters: 3M entries
    - 6 parameters: 300M entries
    - 10 parameters: 15B entries!
    - With 10 parameters, we'd take **1.5 years** for a single brain image, even with the fastest GPU!

The FTK Associative Memory

- Our idea: adapt the Associative Memory developed for High Energy Physics projects to this problem
- Capable of millions of comparisons in a single clock cycle!
- Estimated performance:
    - Using a 200k dictionary: ~0.5s
    - Using a 15B dictionary: ~1h*

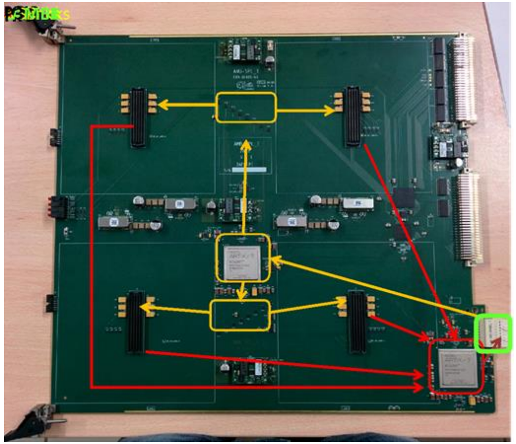

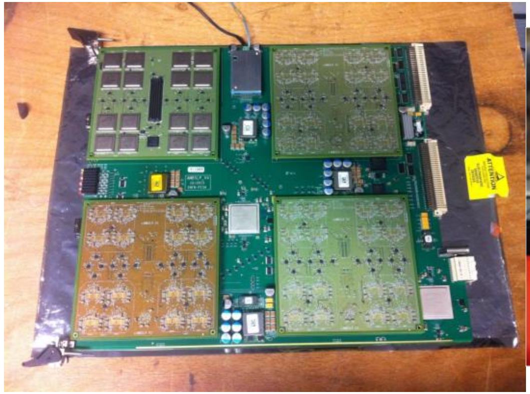

The Associative Memory

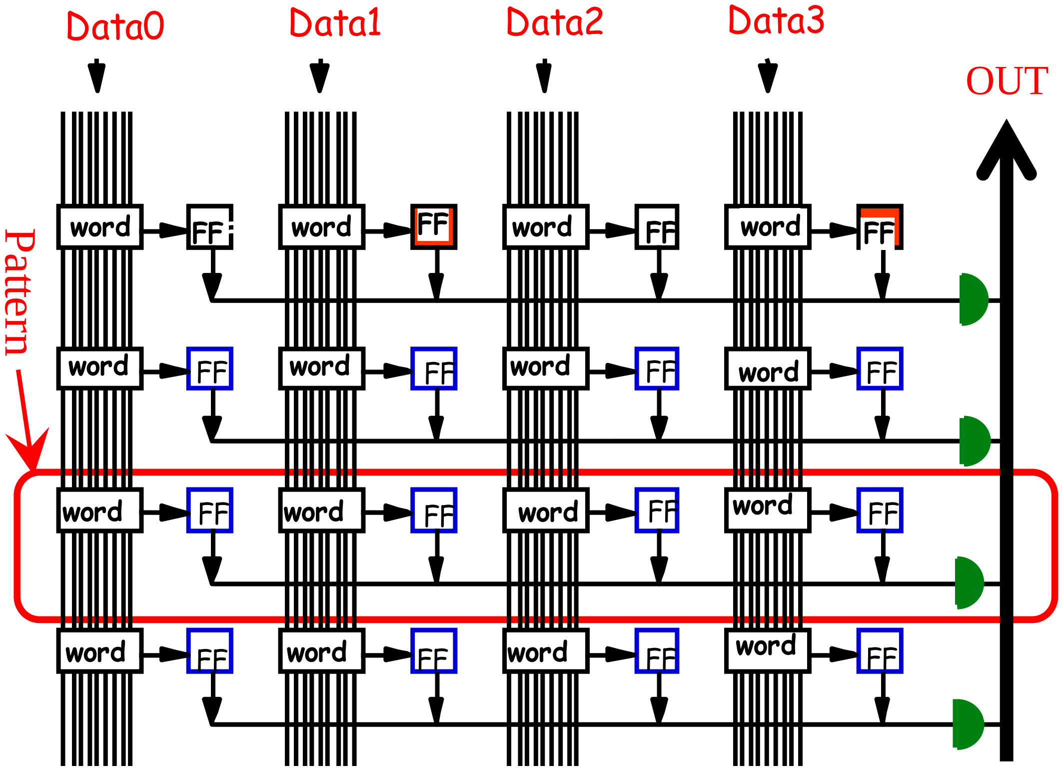

Using the Associative Memory for MRF

- If the simulated signals that compose the dictionary were “perfectly equal” to the corresponding acquired MRI signals, we could use the AM “as-is”: load the dictionary into the AM banks, then use the MRI acquisition as input
- Unfortunately, real MRI signals show acquisition defects
  - noise from the environment
  - “aliasing” effects from nearby tissues
- Therefore, we have to “pre-process” both dictionary and acquisitions

Using the AM in practice

- Both voxels of the acquired MR image and entries of the dictionary are represented by a vector of m complex numbers (often: $m = 8$)
  - each complex number is called a *component*
- In order to be loaded into the AM banks, dictionary entries must be *converted* into vectors of integers
- We employed a **binning** technique
  - Fact: components are always in the interval [-1, 1]
  - Divide the interval into n bins of size 2/n
  - If the i-th component falls into the k-th bin → assign the number k to that component

Binning: considerations

- Binning is effectively a form of *discretization*
- Multiple distinct dictionary entries could be represented by the same bin string (called *pattern* in our terminology)
- Nevertheless, if done correctly, all entries that are represented by the same pattern are “close” to each other in the dictionary space
  - i.e. they have similar values for all the parameters
- Of course, we have to perform the same binning transformation to each voxel of the acquired MR image
- The AM is then used to perform pattern matching

How to use the AM output

- The AM will output the address of the matched pattern
- Next step: associate each pattern with the list of the dictionary entries that are represented by it, then perform the dot product between the voxel and these (few) entries - **filtering!**
- Filtering requires a fine tuning of the procedure parameters
  - bin number and size
  - gaussian noise applied to dictionary entries (to simulate environmental noise)
  - ...
- «Best» parameters chosen via heuristics
  - Minimize the number of voxels with no matched pattern, while keeping the error w.r.t. the gold standard reconstruction low

Using the AM to filter “good” entries

- Each voxel of the acquired MR image is associated to a (hopefully short) list of dictionary entries
- If the AM is used correctly, the list contains the dictionary entry with the highest dot product among all entries
- Basically, the AM outputs for each voxel the list of “closest” dictionary entries
- For each voxel, must perform dot products on a subset of the dictionary after using the AM (postprocessing)

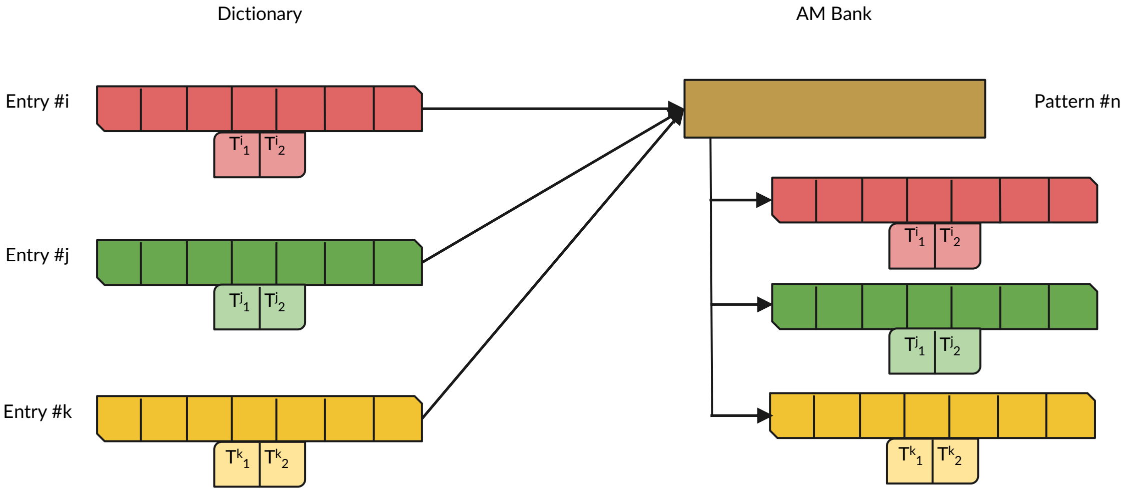

Dot-products step

- If done correctly, after filtering (step 1, see image) a few (~5000) dot products per voxel are required (step 2)
- Can be done in hardware, e.g. using FPGAs

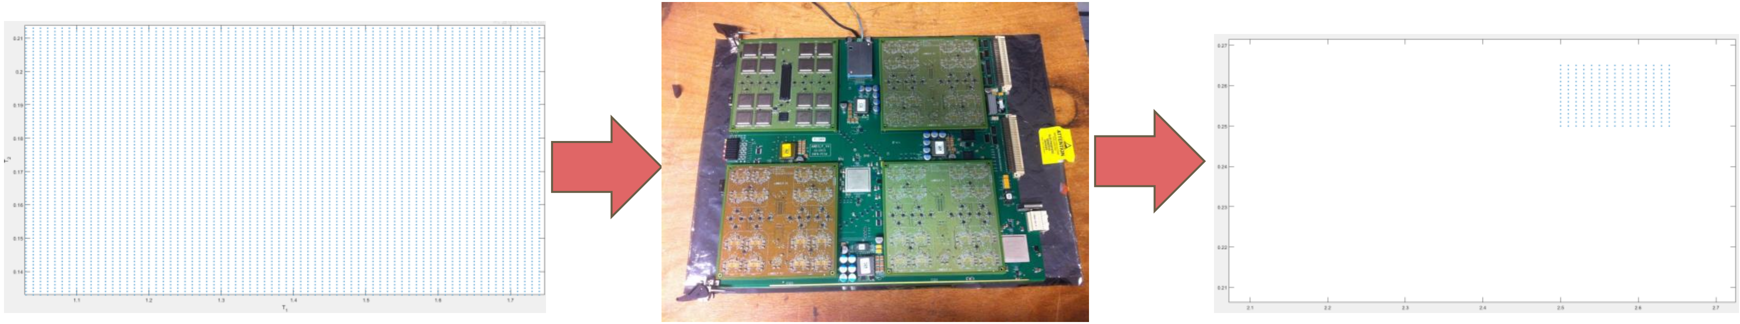

Some results

- Extremely good results when reconstructing a brain using a dictionary with two parameters: longitudinal (T_1) and transverse (T_2) relaxation times
  - 99.97% reconstruction efficiency
  - Less than 3% error
  - **36 times** less dot products than the standard method
- Actual timing performance to be validated on real hardware (next step of the work)

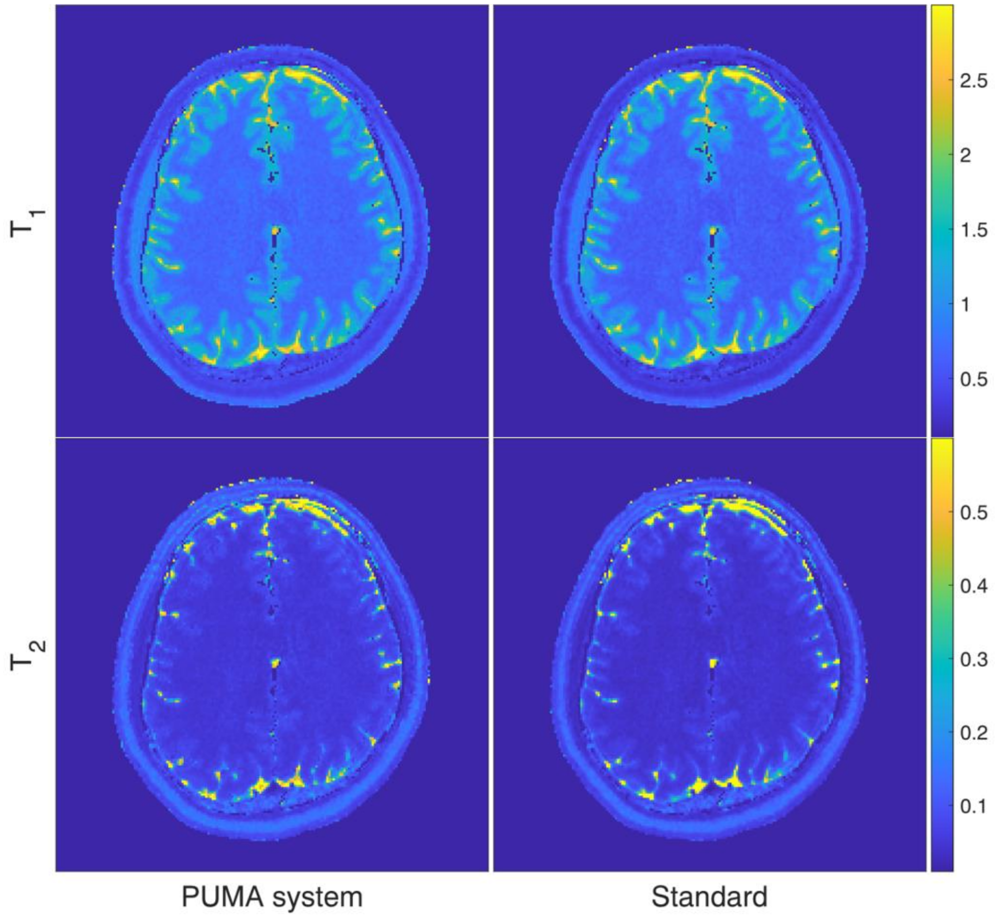

Conclusions

- Highly innovative approach to a well-known problem
- «Transfer of knowledge» between High Energy Physics and Medical Imaging scenarios
- Promising preliminary results to be confirmed with tests on actual hardware

Thanks for your attention!

- Questions?
- Feel free to drop me an email: orlando.leombruni@pi.infn.it

References:

- Ma, Dan et al. “Magnetic Resonance Fingerprinting.” Nature 495.7440 (2013): 187–192. PMC. Web. 28 June 2018.
- C. L. Sotiropoulou et al., ”The Associative Memory System Infrastructures for the ATLAS Fast Tracker,” in IEEE Transactions on Nuclear Science, vol. 64, no.6, pp. 1248-1254, June 2017.
- A. Annovi et al., ”Characterization of an Associative Memory Chip in 28 nm CMOS Technology,” 2018 IEEE International Symposium on Circuits and Systems (ISCAS), Florence, Italy, 2018, pp. 1-5.
- P.A. Gómez, M. Cencini et al, ”Rapid three-dimensional multiparametric MRI with quantitative transient-state imaging”. Sci Rep 10, 13769 (2020).
- G. Buonincontri et al, “MR fingerprinting with simultaneous B1 estimation”, MRM 2015.
- M. Cencini, M. Tosetti et al., ”Magnetic resonance fingerprinting with dictionary-based fat and water separation (DBFW MRF): A multicomponent approach”, Magn Reson Med. 2019.
- M. Cencini, “Magnetic Resonance Fingerprinting for multi-component estimations”, Ph.D. Thesis, 2020.