Next-Generation Fast Radio Burst Telescopes in the Exascale Era – Extended Abstract*

Matteo Trudu[1]

[1]INAF-Osservatorio Astronomico di Cagliari, via della Scienza 5, 00947 Selargius, Italy

Abstract

Radio astronomy is one of the most data-intensive sciences: instruments run continuously, sampling wide bandwidths at fine time resolutions and, for interferometers, across every antenna pair. The result is high-dimensional data products that must be moved, buffered and processed at scale; recording everything at native resolution is rarely feasible. The field is entering the era of the "software telescope", where advances in HPC and algorithm design enable real-time discoveries of fast radio bursts (FRBs) and other transients.

FRBs are millisecond—sometimes sub-millisecond—radio flashes at cosmological distances that demand high time and frequency resolution to capture their brevity and to correct for propagation effects en route to Earth. In particular, they suffer dispersion as their waves traverse the ionised interstellar and intergalactic media: lower-frequency components arrive later, with a delay that follows a frequency^{-2} law and is proportional to the free-electron column along the line of sight—quantified by the dispersion measure (DM).

Historically, sensitivity was bought with hardware: large single dishes and fixed-function electronics carried most of the load. Today the trend is the reverse—sensitivity is aggregated from many small, inexpensive antennas—and correlating their signals shifts the cost from bespoke hardware to large-scale computing. Because the true DM is unknown a priori, data must be de-dispersed across dense grids of trial values—an inherently heavy computation—and, since native resolution makes continuous recording infeasible, this must run online in real time.

Software-defined pipelines therefore ingest high-rate time–frequency streams, mitigate interference, fan out into many coherent and incoherent beams, de-disperse across dense DM grids, and apply matched filtering plus compact machine-learning classification—all in-stream and under strict latency—capturing short baseband snippets around high-confidence candidates for offline vetting. As a result, FRB discovery is driven by algorithmic efficiency, memory-/I/O-aware kernel design, and the orchestration of heterogeneous, GPU-centric clusters. In this software-telescope paradigm, relatively simple front ends are coupled to powerful back-end systems that trade bulk storage for real-time computation.

This talk surveys the data challenges and software-centric architectures needed to scale FRB discovery to thousands per year. I will present updates on two next-generation FRB machines—the Northern Cross (Italy) and the Canadian Hydrogen Observatory and Radio-transient Detector (CHORD, Canada)—both supported under the PNRR "NG-Croce" programme. Northern Cross shows how a legacy transit array becomes an effective FRB instrument by coupling existing front ends to commodity GPU back ends and streaming, real-time pipelines. CHORD—a purpose-built software telescope with a compact core of small dishes and US outriggers—synthesises thousands of coherent beams; its performance is driven by large-scale beamforming, high-throughput dedispersion and low-latency triggering in a single real-time system.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

ITADATA-WS 2025: The 4th Italian Conference on Big Data and Data Science – Workshops, September 9–11, 2025, Turin, Italy
matteo.trudu@inaf.it (M. Trudu)
0000-0002-1530-0474 (M. Trudu)