



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani

PIANO NAZIONALE
DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC,
Big Data and Quantum Computing




High Performance Stingray

*Eleonora Veronica Lai, Matteo Bachetti, Daniela Huppenkothen
and Stingray developers*

DT-Lab Kick Off Meeting, Bologna, September 30, 2026


Scientific Rationale



Zdziarski (2000)

Spectra or variability?
Is it possible to use the full information?



Heil et al. (2015)



Huppenkothen et al. (2019)

Bachetti et al. (2024)

Lai et al. (2026)



Technical Objectives, Methodologies, and Solutions: what is able to do

•“Timing” analysis

- Pulsation searches and timing
- Aperiodic variability, periodogram modelling (ML, Bayesian)

•Spectral analysis -> connect to Xspec, Sherpa

- Continuum modeling
- Broad lines (e.g. Fe complex, cyclotron lines)



•Polarimetry

•Spectral-timing techniques:

- Time/Phase lags
- Coherence
- Covariance
- etc.

... all with instrument awareness

- Be aware of instrumental systematics: dead time, frame time, good time intervals, etc.
- Mission support



An open-source environment

- Python **open-source** software in a Github-based workflow:
 - Issue tracking
 - Assignments
 - Pull Requests
- Developers:
 - Astronomers
 - Google Summer of Code students
- Community outreach:
 - Public **Slack channel**
 - **Talks**
 - Hackatons/**Tutorials**



Google Summer of Code
2016, 2017, 2018, 2020,
2021, 2022, 2023, 2024,
2025, 2026

The screenshot shows the GitHub repository for 'Stingray'. The repository is described as 'Next-generation spectral-timing software for astronomy applications and beyond.' It has 38 followers and a website link 'https://stingray.science'. The page includes tabs for Overview, Repositories (12), Projects, Packages, and People (6). Under the 'Pinned' section, four items are listed: 'stingray' (Public, Python, 221 stars, 188 forks), 'HENDRICS' (Public, Jupyter Notebook, 26 stars, 18 forks), 'dave' (Public, JavaScript, 16 stars, 15 forks), and 'notebooks' (Public, Jupyter Notebook, 35 stars, 45 forks).

Main Results during the Spoke 3


LARGE data sets (> 32GB)

SMALL data sets (< 32GB)

1.Data loading

1. GPU porting

2.Parallelisation



Main Results during the Spoke 3

LARGE data sets (> 32GB)



1.Data loading

New class: “FITSTimeseriesReader”, able to load FITS files as memory-mapped arrays and read them in chunks.

2.Parallelisation

Two parallel implementations of the Bartlett periodogram using multiprocessing and MPI



Main Results: LARGE data sets



MPI is *faster*
than multiprocessing and
the serial version of the
code

Main Results during the Spoke 3

SMALL data sets (< 32GB)



Cluster OACa
GPU Nvidia A40

Porting in GPU?

Better performance
above a certain threshold


Conclusions and final results for Spoke 3

LARGE data sets (> 32GB)

1.Data loading



2.Parallelisation



SMALL data sets (< 32GB)

1. GPU porting



HENDRICS: command line interface for some libraries of Stingray

Latest improvements

- MPI implementation of the Bartlett periodogram and GPU histogram are included in HENDRICS now
- BM histogram functions (speedup GPU vs CPU - RTX 2070 SUPER vs Intel i7-10700)
 - 1D histogram: **1.5x** at 10^7 events (slower below)
 - 1D histogram + FFT: **~3x** at 10^7 , up to **~9x** at 10^5
 - 2D histogram: up to **1.7x** at 10^7
- BM histogram applications
 - Averaged power spectrum (64-256 segments $\times 2^{19}$ bins)
 - Fused (only final result copied to host): **8-14x**
 - Unfused (each lightcurve and FFT copied to host): **~3x**
 - Pulsation search (10^7 events): **~4.0x** end-to-end (**up to 11x** in the search loop)
- Porting in Julia of the main timing and spectral-timing functions (speedup Julia vs Python - CPU)
 - Better performance for a serial version of the code (e.g. *Averaged power spectrum* and *Averaged cross spectrum* **~5x**)

Thank you!





Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca




Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA




Additional slides

Main Results: SMALL data set (< 32GB)





Google Colab
GPU Nvidia T4



Porting in GPU?

Not convenient with this
architecture

Main Results: SMALL data set (< 32GB)



Cluster OACa
GPU Nvidia A40

Additional
execution time