ITER MONACO POSTDOCTORAL FELLOWSHIP 2026 – RESEARCH TOPICS

Topic 1) Fuelling, density transport and control with the PDS with benchmarks and validation on current-day devices

Description

This work falls within the activities of the ITER Pulse Design Simulator (PDS) currently under development using the standard of the ITER Integrated Modelling and Analysis Suite (IMAS). The goal of the PDS is to configure an optimal pulse schedule (PS) for all the scenarios of the ITER Research Plan. It is the primary tool for scenario designers to prepare pulses that meet performances requirements within physics, machine and plant operational limits.

The initial version of the PDS includes self-consistent coupling between a free-boundary equilibrium code (NICE) and a transport solver (TORAX or METIS), including Heating and Current Drive (H&CD) sources provided by the H&CD workflow. However, actuator models for fuelling, to be used together with density transport and control algorithms are still missing in the PDS: this is the topic of this postdoctoral proposal.

TORAX currently uses the QuaLiKiz-Neural Network (QLKNN) for turbulent transport, allowing discharges from current-day devices to be simulated, including the simultaneous transport of heat and particles.

The goals of this postdoctoral position are:

- To study existing surrogate models for plasma fuelling, choose the most appropriate and integrate them into the PDS to simulate the plasma response to gas puffing and pellet injection. Open-Source models are favoured.
- To implement density controllers into the PDS to simulate the outcome of a density request provided through specific density waveforms. To the largest extent possible these controllers should be taken from the Plasma Control System Simulation Platform (PCSSP); more basic controllers can also be developed for test purposes.
- To activate density transport in the transport models currently implemented in the PDS and to check the evolution of the density and its deviation from the pulse schedule, accounting for actuator response times.
- To apply the PDS to specific scenarios of the ITER Research Plan to both thoroughly test the evolving PDS and assist Science Division in the general area of scenario development.

To actively propose and participate in validation activities against experiments past and present on current research tokamaks or against simulations performed with alternative simulation frameworks.

Qualifications/ Expertise Required

- Expertise in transport modelling of fusion plasmas, with experience in any of the following: fuelling models (gas puffing, pellet injection), plasma response, density transport and

density control. Experience in Integrated Modelling, code development and programming (familiarity with IMAS an asset).
- Knowledge of Python (Fortran is a plus).
- Proximity to plasma operation on current-day devices an asset.

Topic 2) ITER scenario modelling with the PDS including tungsten transport with benchmarks and validation on current-day devices

Description

This work falls within the activities of the ITER Pulse Design Simulator (PDS) currently under development using the standard of the ITER Integrated Modelling and Analysis Suite (IMAS). The goal of the PDS is to configure an optimal pulse schedule (PS) for all the scenarios of the ITER Research Plan. It is the primary tool for scenario designers to prepare pulses that meet performances requirements within physics, machine and plant operational limits.

The initial version of the PDS includes self-consistent coupling between a free-boundary equilibrium code (NICE) and a transport solver (TORAX or METIS), including Heating and Current Drive (H&CD) sources provided by the H&CD workflow. TORAX can currently use QuaLiKiz-Neural Network (QLKNN) or TGLF for turbulent transport. In parallel, the short-term implementation of FACIT in TORAX will enable prediction of impurity transport, particularly adapted for tungsten (W). The ITER PDS thus allows the simulation of discharges from present devices, including the simultaneous transport of heat and particles, for the purposes of validation and benchmarking its application to ITER.

The goal of this postdoctoral position is twofold:

1) Develop and apply a methodology for the estimation of the W wall and divertor source as an actor in PDS which then provides a time-dependent, scenario consistent boundary condition (e.g. the W concentration at the magnetic separatrix) to the impurity transport solver (e.g. FACIT and TORAX). Given the requirement for the PDS to run on relatively rapid timescales, this W source model will need to be fast and thus very likely to be a surrogate model trained on existing high fidelity (time-independent) simulations of impurity release and transport performed with workflows combining simulated boundary plasma backgrounds and separate erosion/migration calculations.
2) Once the impurity source actor has been developed and coupled into the PDS, the latter should be applied to selected scenarios of the ITER Research Plan with emphasis on situations with sensitivity to W transport. A particularly important example is the use of Electron Cyclotron Heating (ECH) to avoid/mitigate W accumulation through combined temperature screening and increased core turbulence (e.g. during current ramp-up in L-mode, or in Start of Research Operation H-mode plasmas). Benchmarking against existing or new higher fidelity plasma simulations (using DINA-JINTRAC integrated modelling) should be performed with an option also to validate against real experimental data, for example from the WEST tokamak, one of the operating devices on which PDS is being tested.

Qualifications/ Expertise Required

▪ Expertise in transport modelling, H&CD physics, plasma turbulence.

▪ Experience in Integrated Modelling, code development and programming (familiarity with IMAS an asset).

▪ Knowledge of Python language (Fortran is a plus).

▪ Proximity to plasma operation on current-day devices an asset.

Topic 3) Developing a Digital Twin for Fusion Energy: A Foundational World Model Approach

Description

This project proposes the development of a Fusion World Model (FWM), a generative AI system capable of learning the complex, non-linear dynamics of Tokamak plasmas to support the operation of future devices like ITER. While our ultimate goal is a foundational FWM for ITER, we will leverage the available open-source repositories of tokamak experimental data (such as that from the MAST tokamak) for inital training. Such datasets are uniquely suited for this task since they contain thousands of plasma pulses spanning many years of operation and comprising hundreds of terabytes of multi-modal data, including high-resolution wide-angle camera footage synchronized with scalar diagnostics. By training on this type of resource, we aim to build a system that can "dream" valid plasma scenarios, enabling robust pulse planning, real-time control, diagnostic health monitoring, and the discovery of new physics.

The operation of current Fusion devices relies heavily on analytical models and pre-programmed sequences. We envision a paradigm shift where a FWM learns the "physics of the machine" directly from experimental data and simulations. Just as a human world model allows you to navigate a dark room by predicting the outcome of your steps before you take them, a FWM allows the control system to predict the future state of the plasma given current actuator inputs. This capability unlocks "Dreaming": the ability to simulate millions of potential pulse trajectories in a latent space. This is not just about control; it is about robust pulse planning. Much like planning your route through a busy subway station, the FWM allows us to foresee potential collisions or delays (instabilities) and chart a safe path through the high-dimensional operational space.

This system would provide proactive operational boundaries, mitigating the risk of damage to critical components. As the machine learns what scenarios induce disruptions or bring the machine close to its operational limits, such as approaching force limits on coils or density limits that trigger unstable MHD activity in the plasma, it can adapt so that future cases do not get themselves into these situations. Furthermore, it allows us to optimize machine utilization by knowing where to "shine the light" next, illuminating areas of the operational space where new physics may be hiding. We can rapidly iterate on natural language pulse objectives, formulating derisked, robust plans that may be passed to agentic AI for validation via integrated simulation before a single experiment is run.

Qualifications/ Expertise Required

▪ Expertise in autonomous agents, multi-step planning, and LLM

▪ Advanced Python proficiency, specifically in introspection, meta-programming, or library design

▪ Experience with knowledge representation or semantic data structures

▪ Experience with Retrieval-Augmented Generation (RAG) systems

▪ Experience designing automated workflows for HPC systems

▪ Background in interfacing with or wrapping legacy C++ and Fortran code

▪ Experience developing multi-modal systems that combine computer vision with natural language processing.

Topic 4) Experimental Exploitation of SPIDER and MITICA

Description

This post-doctoral project supports the scientific exploitation of the ITER Neutral Beam Test Facility, with a focus on negative-ion source physics, RF plasma production and beam extraction. The postdoctoral researcher will work as an integrated member of the NBTF team, spending extensive periods at the facility, located in Padova, Italy, working on the two main experiments: SPIDER and MITICA.

For SPIDER, the researcher is expected to contribute to experimental campaigns, diagnostic interpretation and data analysis, with the objective of characterizing and optimizing RF plasma generation, negative-ion extraction, and beam performance in terms of divergence and accelerated current. The work will also help consolidate operating procedures and transfer the experimental knowledge gained on SPIDER to MITICA.

For MITICA, the researcher will support the integrated commissioning of the test bed and its transition to experimental operation. This will include preparation and analysis of commissioning tests, support to high-voltage conditioning, participation in the first beam extraction experiments and characterization.

The overall scientific objective of this postdoc is to connect SPIDER results with MITICA operation, strengthening the understanding of giant RF-driven negative-ion sources for fusion application.

Qualifications/ Expertise Required

▪ Qualifications/ Expertise Required

▪ Expertise in basic plasma physics, plasma diagnostics or particle beam physics.

▪ Knowledge of Python or MATLAB for data analysis and visualisation.

▪ Hands-on experience on experimental test beds an asset.

▪ Knowledge of Neutral Beam Injection (NBI) physics, including negative-ion sources and beam extraction, an asset.

▪ Ability to work in a team at an experimental facility, including extended stays in Padova, Italy.

Topic 5) Agentic AI for the Exploitation of Fusion Simulations and Experiments

Description

Plasma modelling serves as the critical link between theoretical prediction and experimental validation, but the drive for high-fidelity simulations has created a significant operational bottleneck. While the Integrated Modelling and Analysis Suite (IMAS) provides the standard data model necessary to couple diverse solvers, the orchestration of these codes remains complex. Interfaces are often fragile, making the preparation of consistent input decks error prone. This challenge is foundational, affecting even single-code simulations which must be mastered before multi-code workflows can be attempted. Consequently, physicists are forced to become HPC specialists, diverting effort from physics research to concentrate instead on framework development. Here, simple parameter errors in the inputs to a high-fidelity analysis chain can translate into substantial delays, ineffective resource utilization, and ultimately opportunity costs as key insights are missed.

To address this, this project proposes an Agentic AI system to orchestrate physics codes via natural language. Moving beyond simple text generation, this autonomous operator will formulate strategies, submit HPC jobs, and monitor simulation states. By grounding the AI in the IMAS Data Dictionary and utilizing runtime introspection of physics codes, the system enforces deep semantic consistency. It proactively resolves compatibility issues, such as verifying solver validity ranges, before committing computational resources. Furthermore, by connecting to experimental databases, the agent can autonomously formulate validation cases to test hypotheses against real-world data.

The work plan proceeds in three development phases. First, the project will develop standardized Python wrappers for core solvers to convert opaque inputs into query-able objects, allowing an AI agent to inspect valid parameters and data exchange interfaces at runtime. Second, a Retrieval-Augmented Generation (RAG) system will be constructed to ingest documentation and validation databases, capturing the "soft" knowledge of code usage and limitations. Finally, interfaces to a multi-scale Coupling Library (MUSCLE3) will be defined to enable agents to connect diverse physics solvers. Here, successful workflows and usage patterns persist in a central knowledge graph to create a reusable foundation for autonomous workflow construction and execution.

Qualifications/ Expertise Required

▪ Expertise in autonomous agents, multi-step planning, and LLM

▪ Advanced Python proficiency in introspection, meta-programming, or library design

▪ Experience with knowledge representation or semantic data structures