TRANSMUTEX

R&D for High Power Proton Cyclotrons

FFA'25 WORKSHOP, IMPERIAL COLLEGE LONDON
17 September 2025

Malek HAJ TAHAR

Table of contents

Subcritical Transmuting Accelerated Regenerative Technology (START)

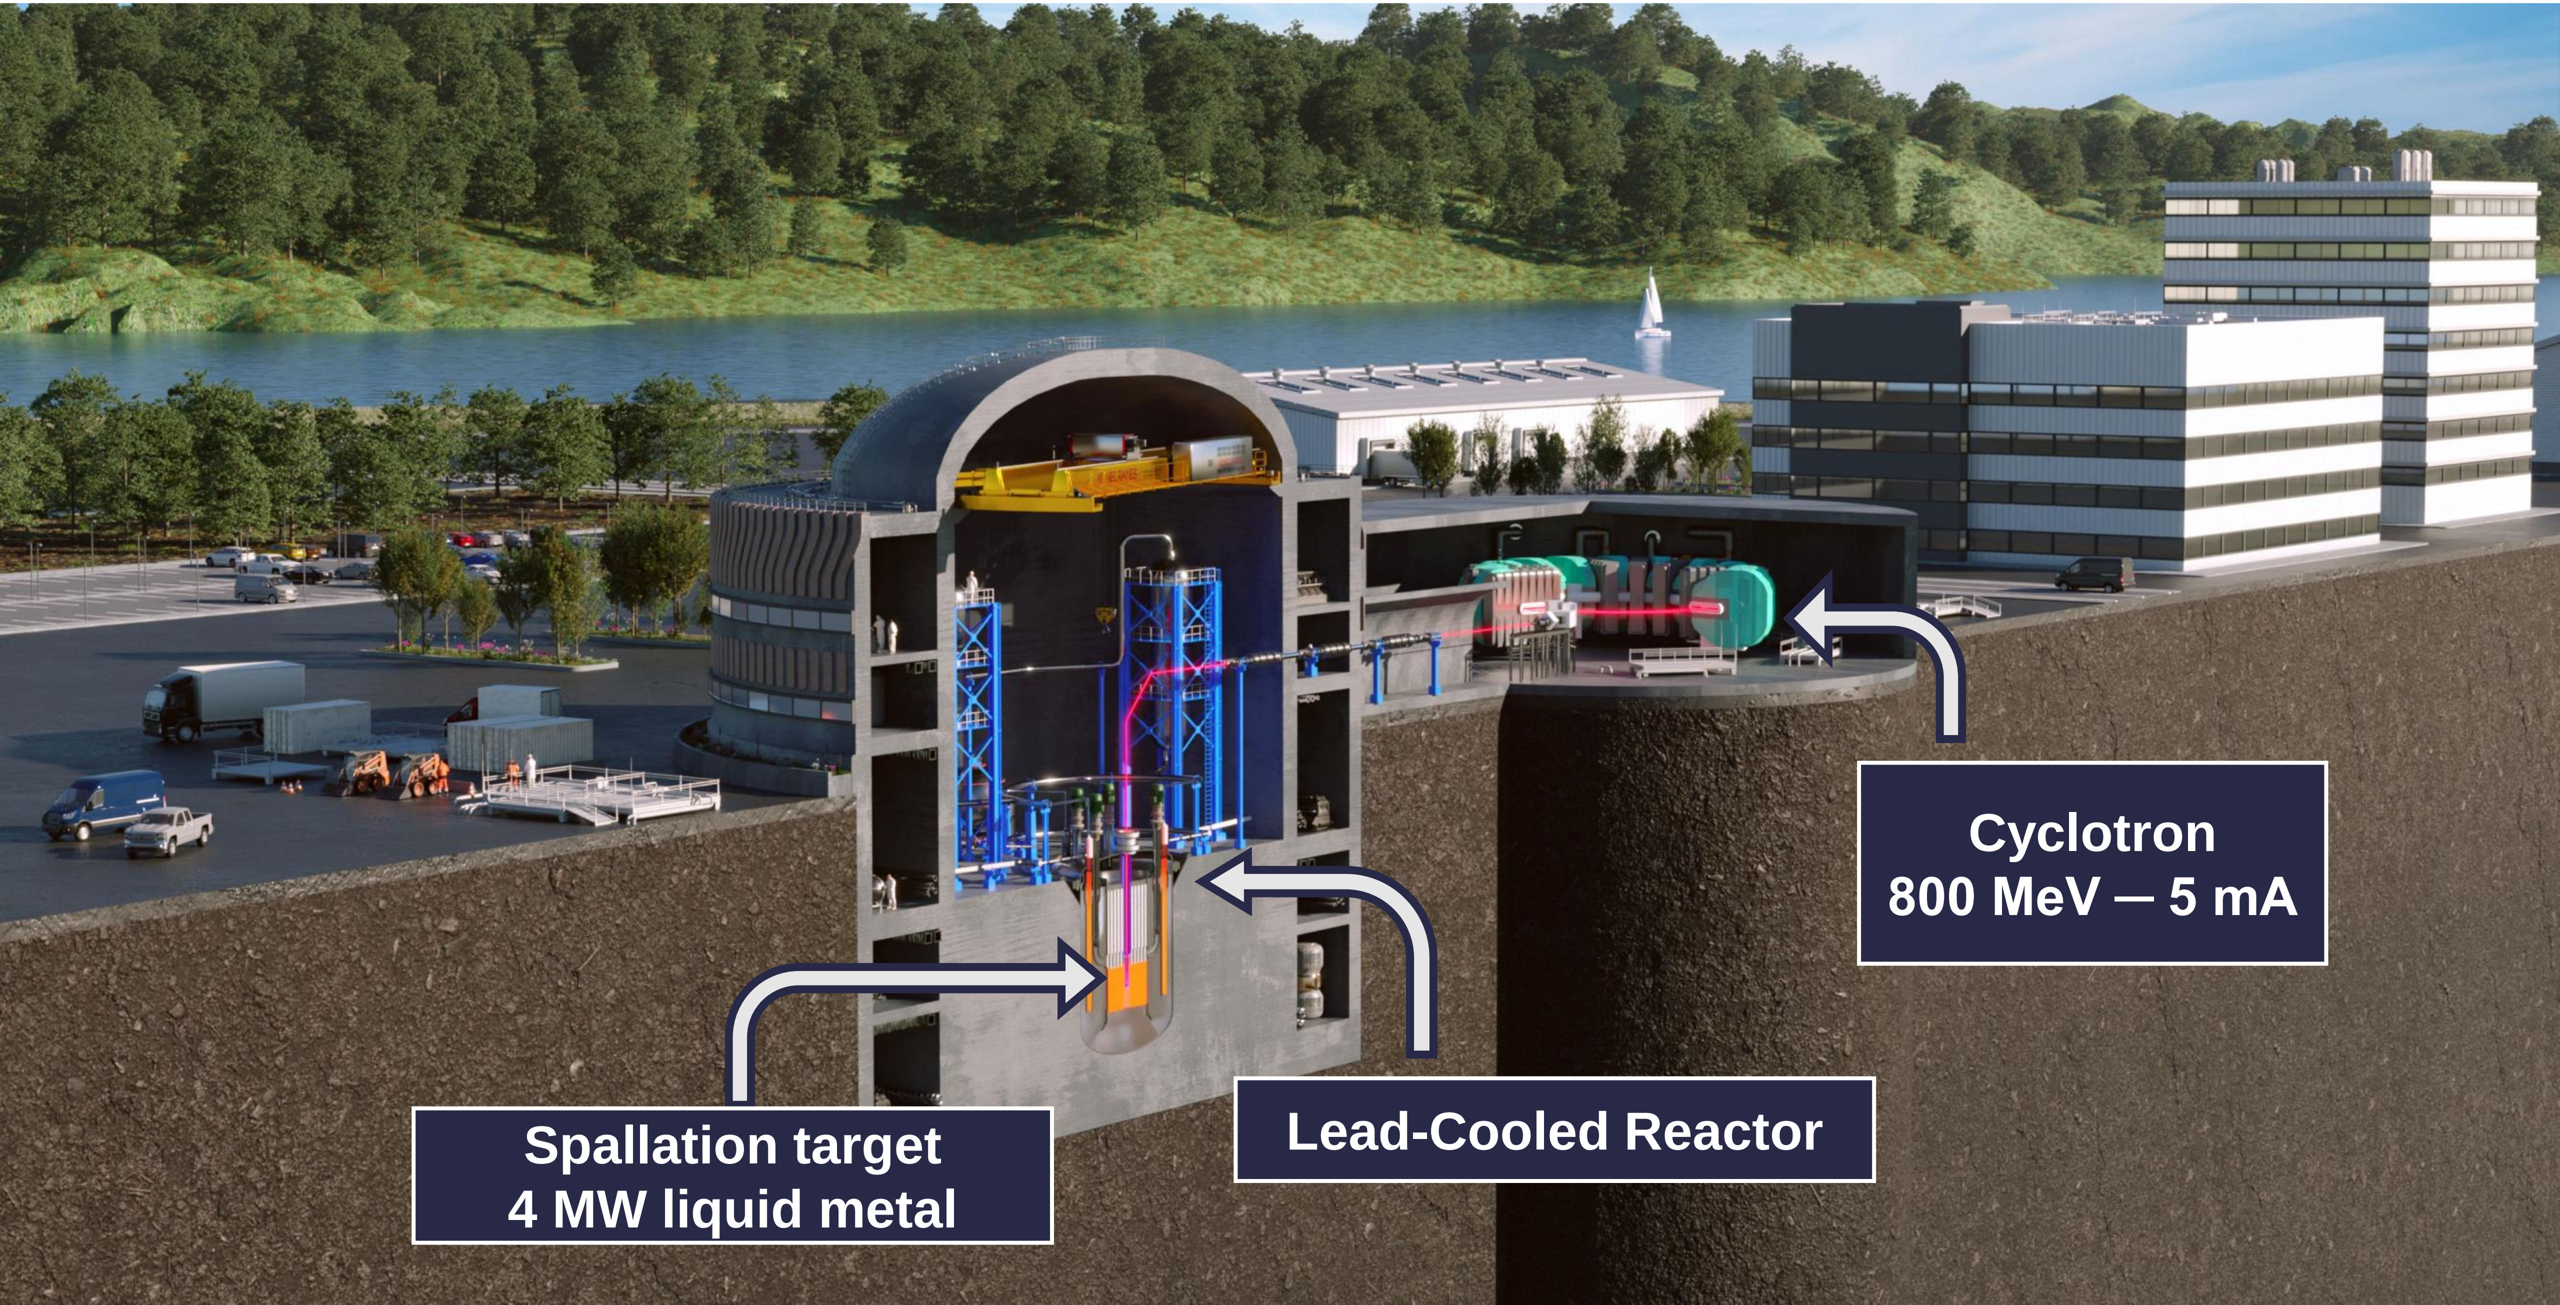

Challenges for High Reliability

- **High Reliability Requirements:**
  - ➢ Ensure high plant availability
  - ➢ Limit thermal stress on critical components (target, subcritical blanket, fuel assemblies)
- **Current Data**:
  - ➢ Reliability is 2-3 orders of magnitude below ADSR requirements
- **Next steps**:
  - ➢ Most R&D efforts focused on how to ensure the most reliable high power cyclotron design & operation

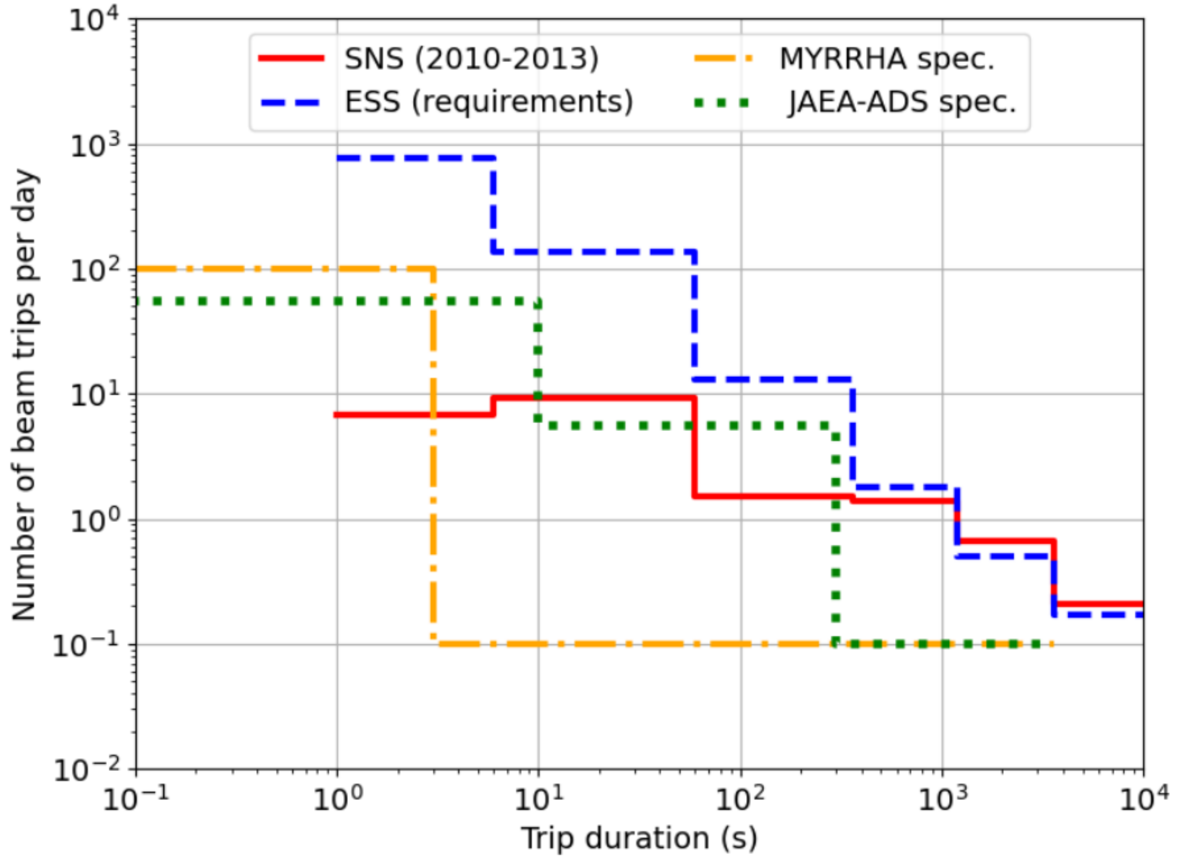

Beam trips frequencies

Overview of ADS Projects in the World (cern.ch)

M. HAJ TAHAR - 17/09/2025

Key R&D Themes in High Power Proton Cyclotrons

Hardware Design

- Ion Source & Injection
- Magnet Systems
- RF Systems
- Collimation
- Beam dumps

Sets physical performance limits

Design Optimization

- Beam Dynamics & Loss Mitigation
- Fault Compensation Strategies
- Redundancy & Robustness in Layout

Improve performance through modelling and optimization even in the event of failure

Operations & Control

- Beam Diagnostics & Monitoring
- Automation & ML-Based Tuning
- Predictive Maintenance

Meet reliability & availability goals with intelligent operation

Hardware Design

- Ion Source & Injection:
  - Emittance control, beam current stability → crucial to achieve a stable vortex effect and good matching
  - Reliability Challenges to be addressed in ARPA-E Project (**Advanced Research Projects Agency**)

TRANSMUTEX SA – Los Alamos, NM

Highly Reliable Ion Source and Injection Beamline to Maximize Proton Beam Availability – $4,293,007

Transmutex is developing a highly reliable ion source by enhancing existing commercial technology through innovative engineering. Transmutex will modify a commercial ion source to reach the reliability required for efficient operation of the accelerator and the overall system. The project includes broad engineering improvements, followed by an extensive testing campaign to identify potential sources of failures. Using advanced data analysis, the team will continuously monitor and optimize the system's performance. The project's overarching goal is to enable nuclear waste transmutation to transform long-lived radioactive elements into shorter-lived ones, reducing their hazardous lifetime from 1 million years to a few hundred years.

ARPA-E NEWTON_Project Descriptions_FINAL.pdf

Hardware Design

- Ion Source & Injection:
  - Emittance control, beam current stability → crucial to achieve a stable vortex effect and good matching
  - Reliability Challenges to be addressed in ARPA-E Project

- Magnet Systems:
  - Field quality
  - Iron-dominated design is robust, energy efficiency aspects

- RF Systems:
  - High-gradient cavities, rapid detuning/recovery in case of failure
  - PSI experience shows long term reliability goals achievable
  - Development and testing of new RF for cyclotron: 3 to 5 years

Hardware Design

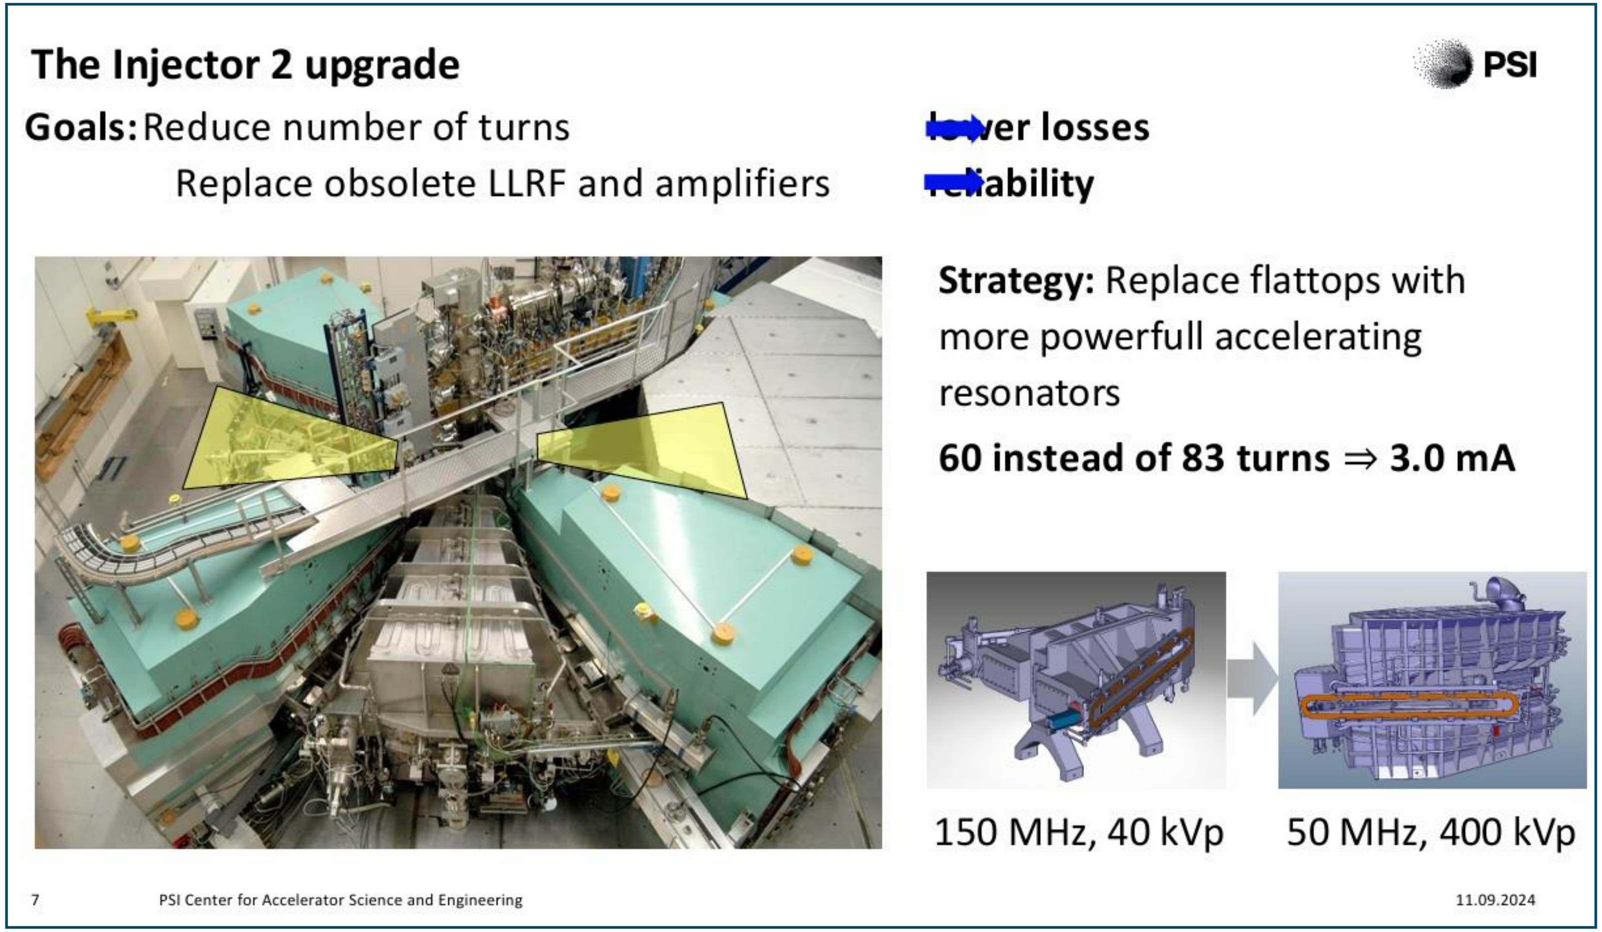

Courtesy Markus Schneider

20240912_CWRF2024_Status_Injector_2_upgrade_Schneider.pdf

Design Optimization

- Beam Dynamics & Loss Mitigation
  - Halo control, collimation, space charge effects
  - Still an active area for predictive modelling & mitigation
  - Crucial to define the right matching conditions from the ion source and the feasibility of high current

<table>
<tr><th>1mA</th><th>5mA</th><th>10mA</th></tr>
<tr>
<td>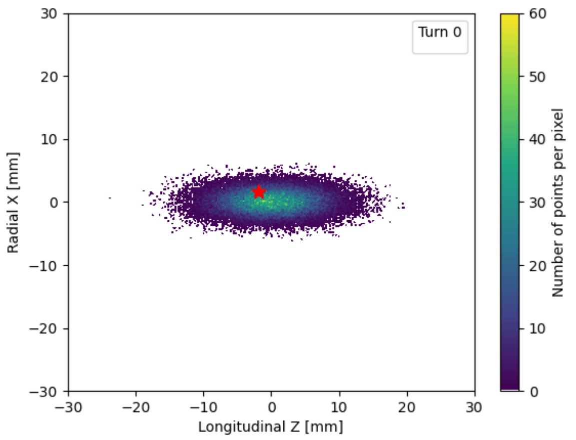
</td>
<td>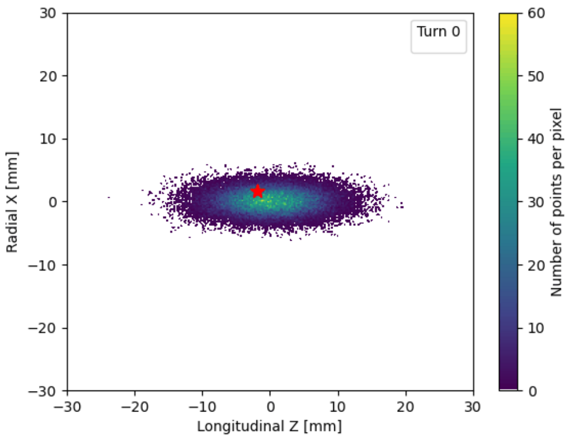
</td>
<td>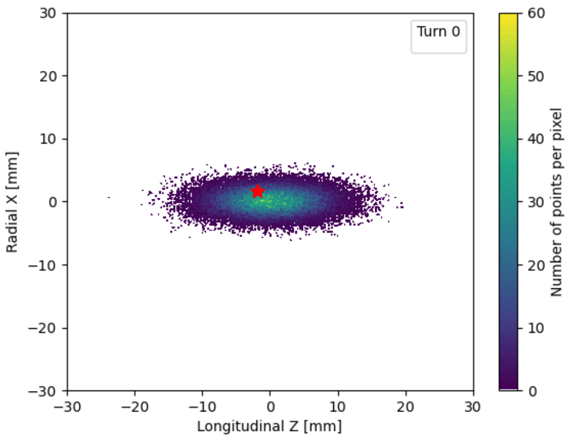
</td>
</tr>
</table>

Stable formation of round bunches due to vortex effect in injector 2, at different currents

→ Amount of halo increases substantially with the current, requiring an optimized matching and a proper collimation, ideally @ injection

Design Optimization

- Beam Dynamics & Loss Mitigation
  - Halo control, collimation, space charge effects
  - Still an active area for predictive modelling & mitigation
  - Crucial to define the right matching conditions from the ion source and the feasibility of high current

- Redundancy & Robustness in Layout:
  - Flexible beamlines, modular injector layout
  - Design choices to maximize reliability

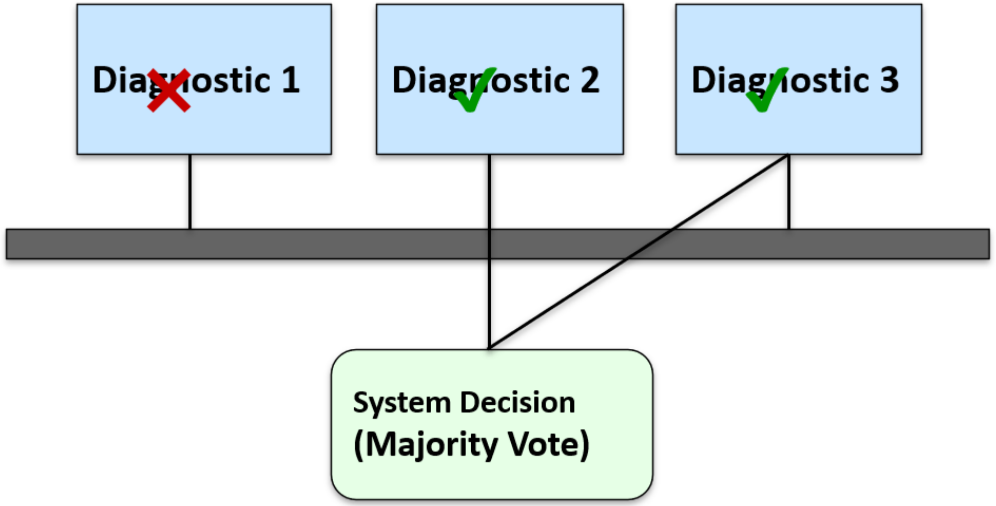

Design Optimization

- Beam Dynamics & Loss Mitigation
  - Halo control, collimation, space charge effects
  - Still an active area for predictive modelling & mitigation
  - Crucial to define the right matching conditions from the ion source and the feasibility of high current

- Redundancy & Robustness in Layout:
  - Flexible beamlines, modular injector layout
  - Design choices to maximize reliability

- Fault Compensation Strategies:
  - Retuning in case of cavity failure
  - Demonstrated in 590 MeV main ring at PSI, yet speed and automation need improvement → hardware design changes

    More details in FFA'24 Workshop, Fault Compensation scheme in cyclotrons
  - Compensation Strategies: Global vs Local

Fault Compensation

- **Normal Operation:**
  - **Beam Current**: 2 mA
  - **Cavity Settings**: All 4 cavities @ 850 kV
- **Experiment @ PSI 590 MeV ring:**
  - **Objective**: Determine the maximum extractable current with only 3 cavities in operation
  - **Approach**:
    - **Cavity Settings**:
      - Cavity 1 & 2 @ 900 kV
      - Cavity 3 @ 860 kV (scan)
      - **Cavity 4 off**
  - **Question**: How much current can be extracted under these conditions?

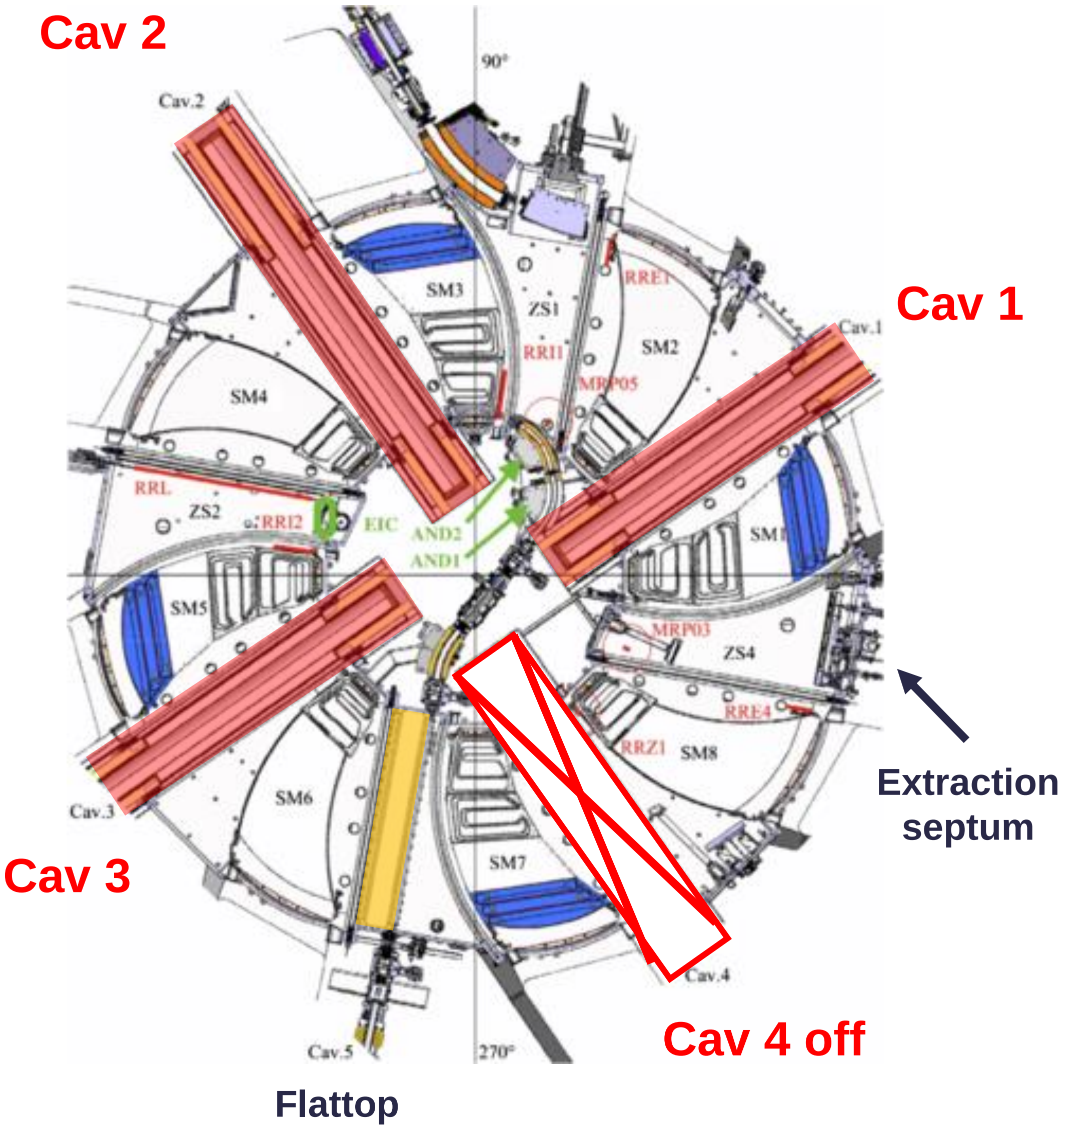

Fault Compensation

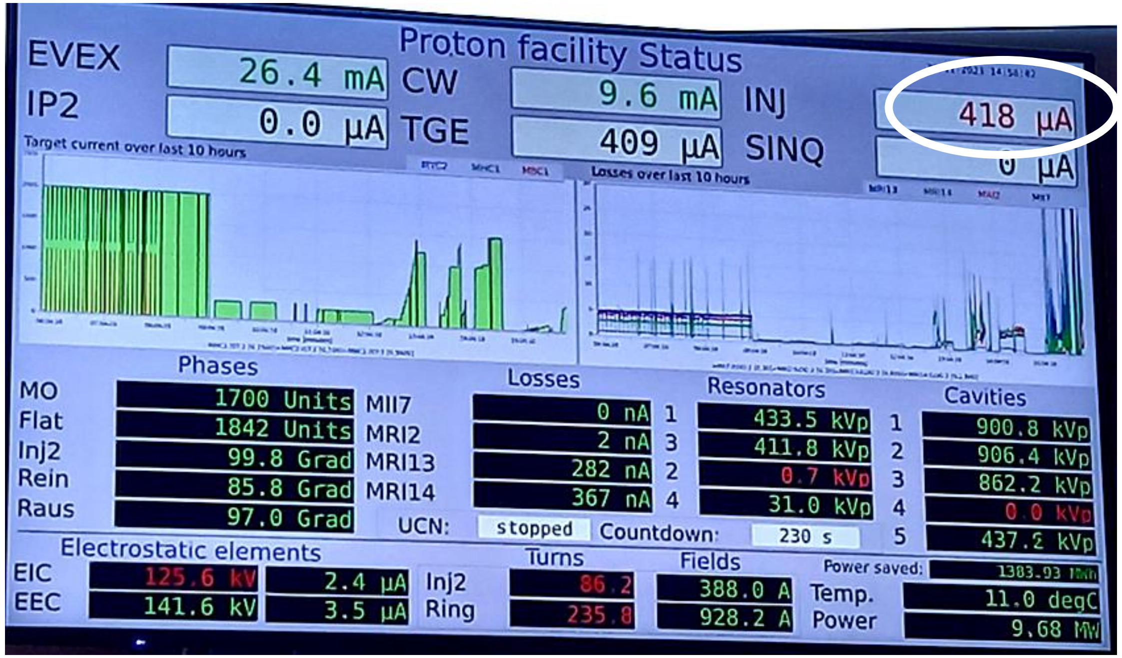

Losses at extraction around 940 nA ≡ 0.2 %

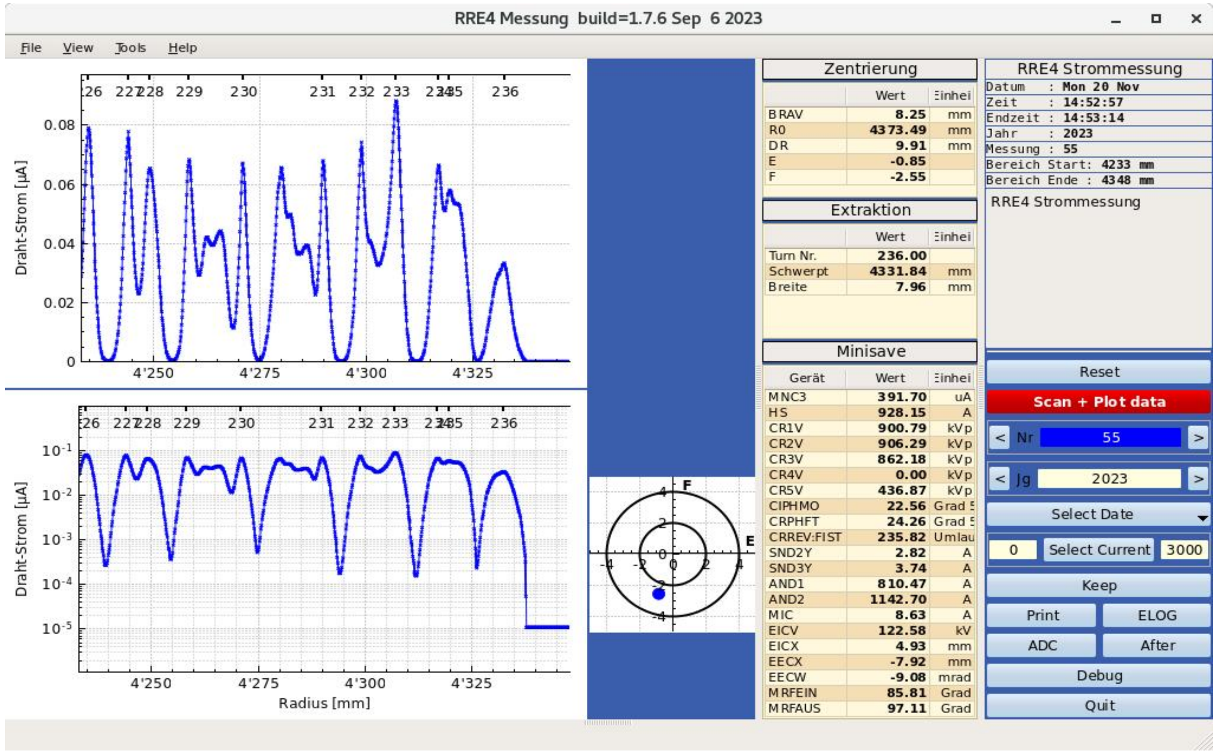

Measured radial profile last 15 turns in the cyclotron @ 418 µA
Courtesy Christian Baumgarten

- Maximum current achieved 418 μA with 236 turns
- **Yet it took experienced operators + physicists several minutes to find the optimal setting!**

Operations & Control

- **Predictive Maintenance**
  → Anticipate component degradation to minimize downtime

- **Beam diagnostics & Monitoring:**
  - Fast orbit & phase measurement
  - Loss monitors with high resolution and wide dynamic range

- **Automation & Machine Learning-Based Tuning:**
  - Real-time optimization required: relying solely on expert operators (in early operation) is insufficient to meet reliability goals
  - Machine Learning offers a promising path to achieve robust, automated beam tuning
  - PSI Injector 2 is an excellent testbed

Machine-Learning based Tuning

(Publication soon)

ML-based tuning: Motivation

- Transform cyclotron operation from manual tuning to intelligent self-optimization, using Machine Learning (ML):
  - ➢ Rapid beam tuning during start-up
  - ➢ Continuous adaptation to changing machine conditions
  - ➢ Enhanced beam quality, stability, and reliability: shall adapt to different turn numbers (5 in this experiment) and various resonators setups (2, 3 or 4 resonators)
  - ➢ Maximized energy efficiency and system robustness

  → Lower reliance on expert operators and significant time savings in setup and retuning

- **PSI Injector 2 identified as an ideal testbed for investigating ML-based tuning strategies**

Injector 2 Cyclotron

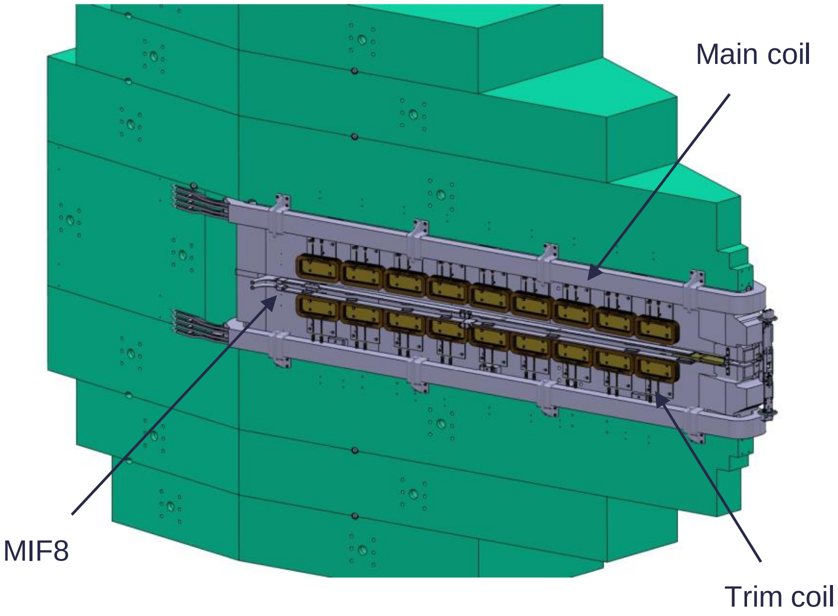

Cross section of one magnet of PSI Injector 2 cyclotron

- 14 actuators are used at once to modify the magnetic field and energy gain per turn

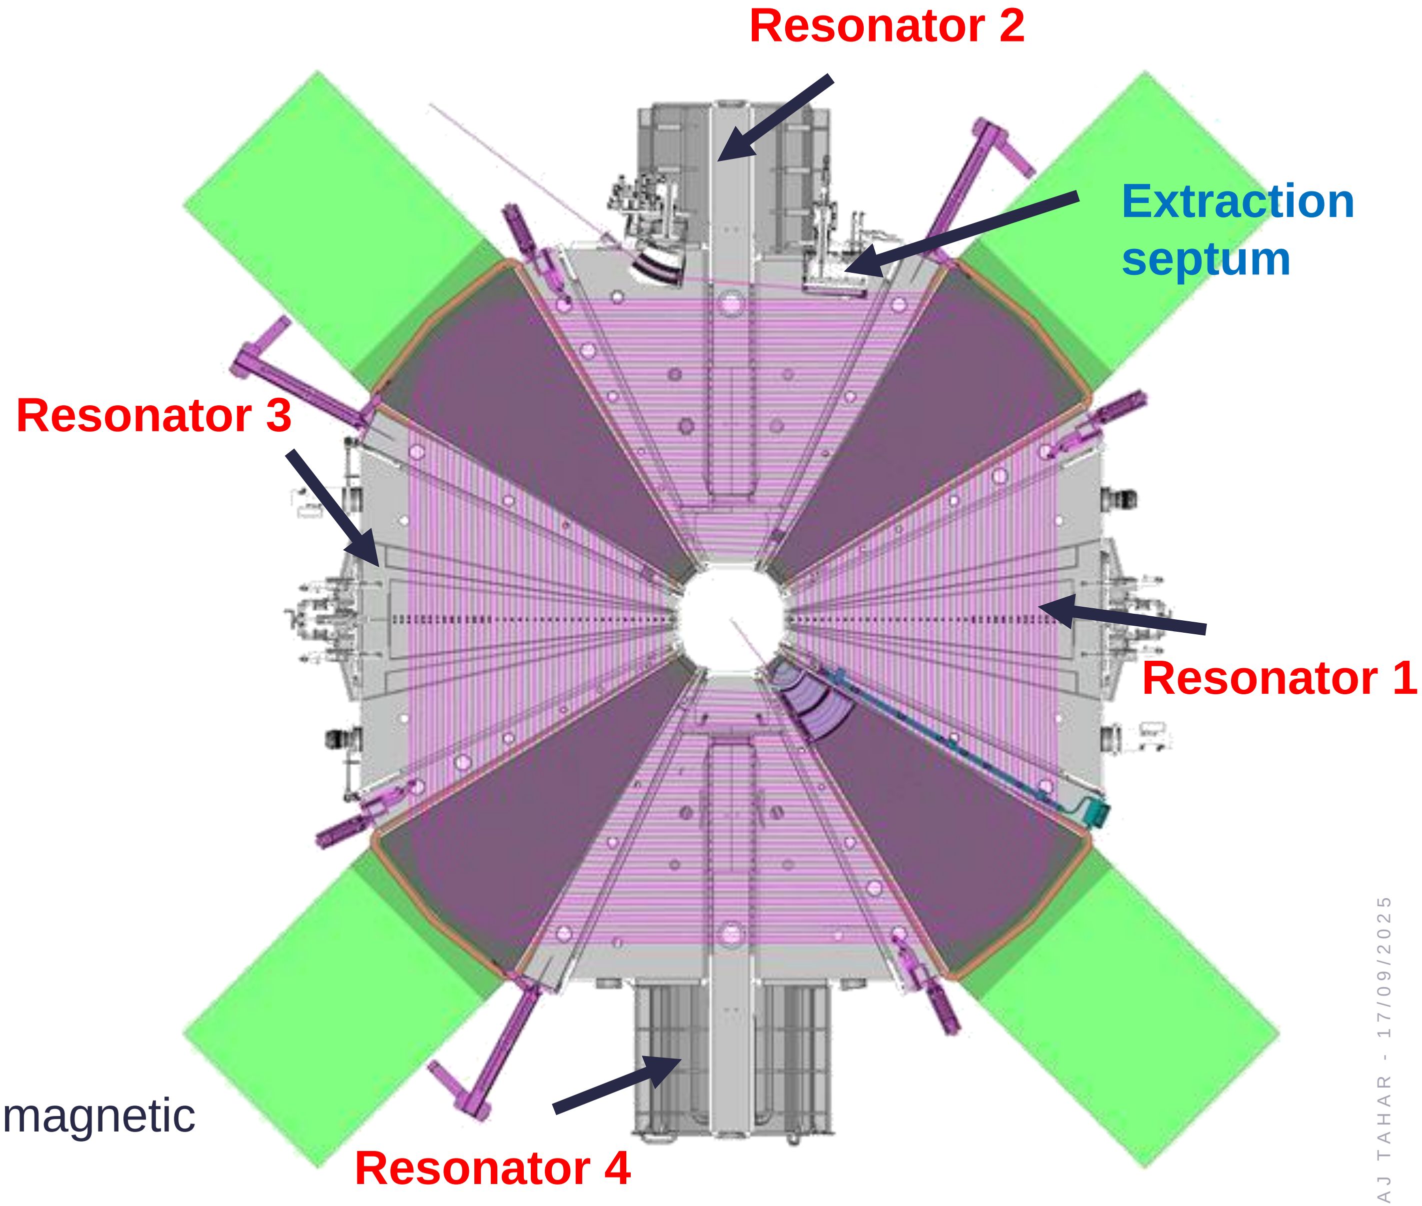

M. HAJ TAHAR - 17/09/2025

Strategic Questions

- **Convergence:** How quickly can an ML model reach optimal control in a real machine environment? Is convergence possible within a few hours?
- **Pretraining Utility:** How effective is surrogate-based or simulation-based pretraining for speeding up convergence?
- **Generalization Across Turns:** Does a model trained on one turn number generalize to others? If not, can transfer learning reduce retraining time?
- **Drift Compensation:** How frequently must a trained model be retrained to counteract thermal or magnetic drifts?
- **Scalability to High Current:** Does the ML policy learned at low beam current remain valid at high current operation?

Can a machine be taught to operate a cyclotron with limited human intervention?

Phase Breakdown of Executed Experiment

- 12 full days of beam development granted:

Stage	Turn number	Resonators setup (Res 1, Res 2, Res 3, Res 4) kVp	
0	72	430, 429, 451, 0	3 resonators
1	73	430, 401, 449, 0	
2	74	430, 371, 448, 0	
3	89	430, 0, 449, 0	2 resonators
4	60	430, 428, 448, 428	4 resonators

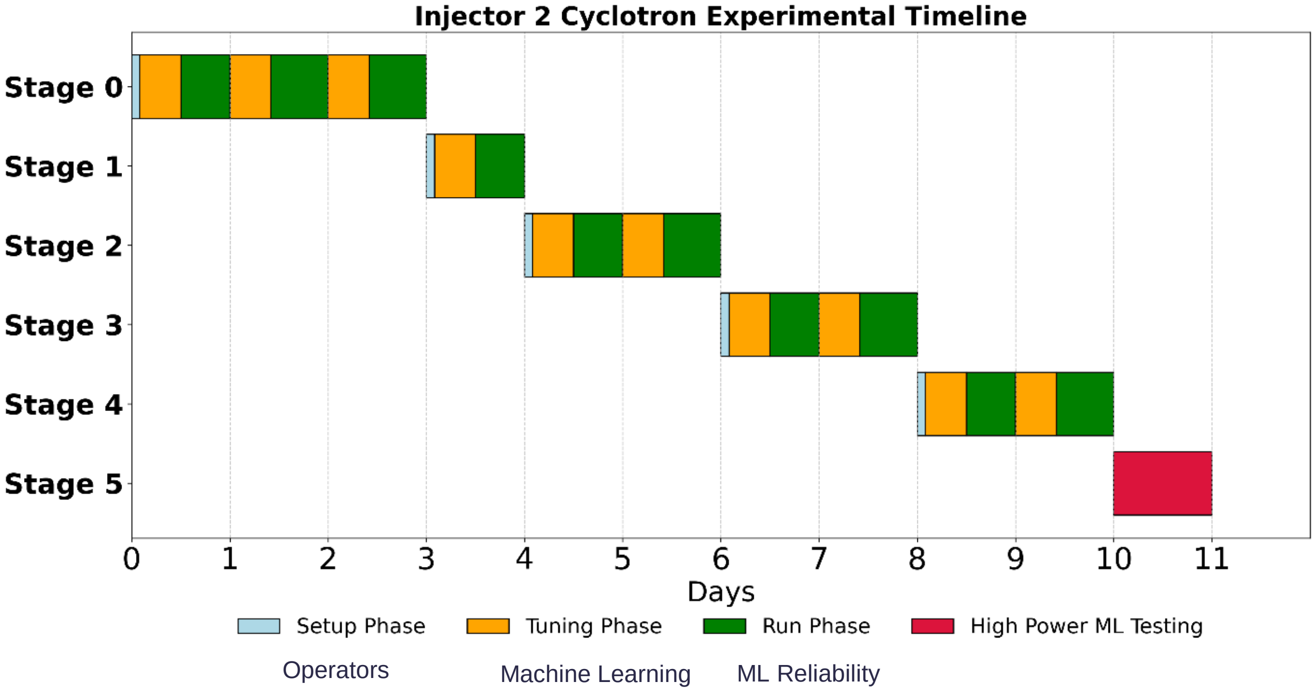

- **Setup Phase (8:00–10:00):** Operators configured the machine to the desired turn number using the resonator settings listed in Table
- **Tuning Phase (Daytime):** Machine learning agents (RL or BO) were trained at low beam currents (~20 μA)
- **Run Phase (Evening/Night):** After tuning, the ML agent was left in autonomous control during nighttime to assess long-term reliability and safety

M. HAJ TAHAR - 17/09/2025

Operator Tuning Across Turns

- Turn 72 was already setup and running
- Tuning complexity **increased significantly** with reduced number of resonators (e.g. 2 resonators at Turn 89)
- The 89-turn setup required **7 interlock recoveries** and took nearly **5x longer** to tune

Turn change	Number of interlocks activated during tuning	Equipment adjusted (except resonators and coils)	Tuning time
72 -> 73	3	-	10 min
73 -> 74	3	-	9 min
74->89	7	AXA, KIP4 collimator, KIR1L collimator	47 min
89->60	0	AXA, KIP4 collimator, KIR1L collimator	10 min

Summary of initial set-up fine tuning procedures across turns

Setup phase went very smoothly thanks to very experienced operators (we assumed 2 hours tuning time)
→ Gave us precious time to perform our experiments (next)

Measured Jacobian

- Empirical Jacobian measurement using finite differences
- What is shown: $\frac{\partial(MIF\ Phase)}{\partial(AIHS\ current)}$ across turns 60 to 89
- Key Observation:
  - Sensitivity increases with turn number, especially for outer phase probes

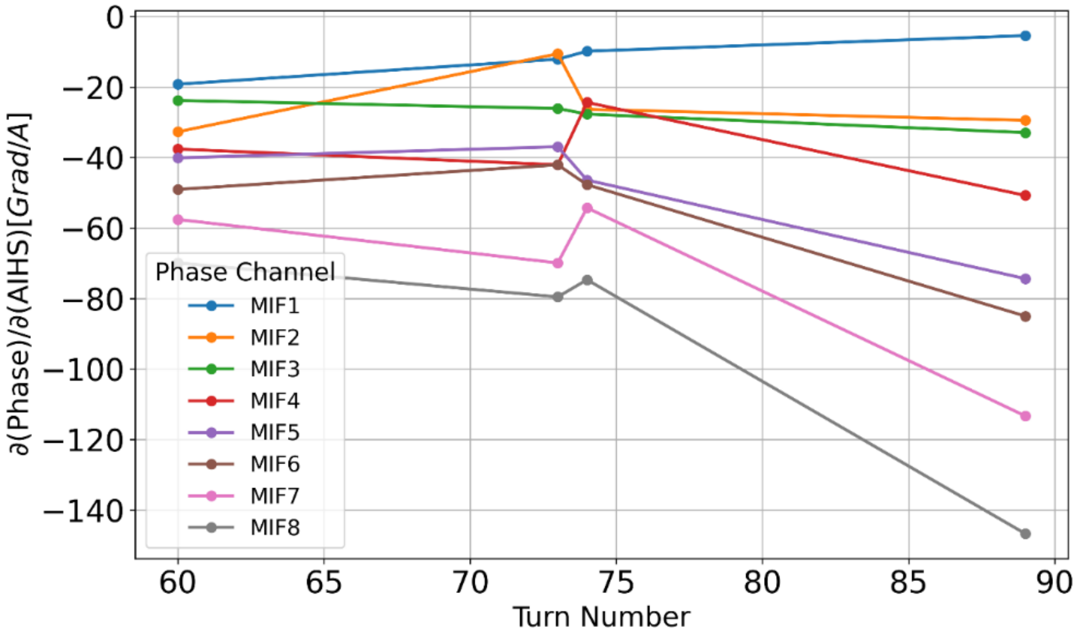

Measured partial derivatives of phase monitors with respect to AIHS current. Sensitivity increases sharply for outer probes at higher turn counts.

<u>Implication</u>:
Control must adapt to turn-specific beam dynamics
Reinforces the need for adaptive, turn-dependent tuning strategy

M. HAJ TAHAR - 23/07/2025

Beam Size Evolution with Current

- Characterize the radial size of the extracted beam using the RIE1 probe for various turn configurations and beam currents

- Consistent σ increase with current ($20\ \mu A$ to $2\ mA$) across all turn numbers (60 to 89)

→ Turn-dependent radial beam size increase with current

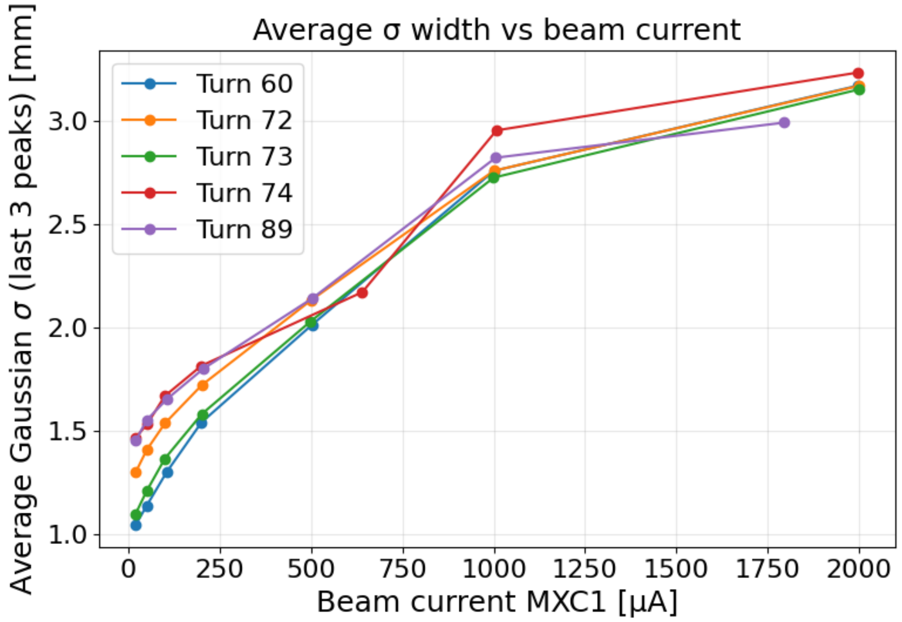

M. HAJ TAHAR - 23/07/2025

Beam Size Evolution with Current

- Characterize the radial size of the extracted beam using the RIE1 probe for various turn configurations and beam currents

- Consistent σ increase with current ($20\ \mu A$ to $2\ mA$) across all turn numbers (60 to 89)

  → Turn-dependent radial beam size increase with current

- Excellent agreement between measured and simulated values → gives us trust on our projected machine performance at 5 mA

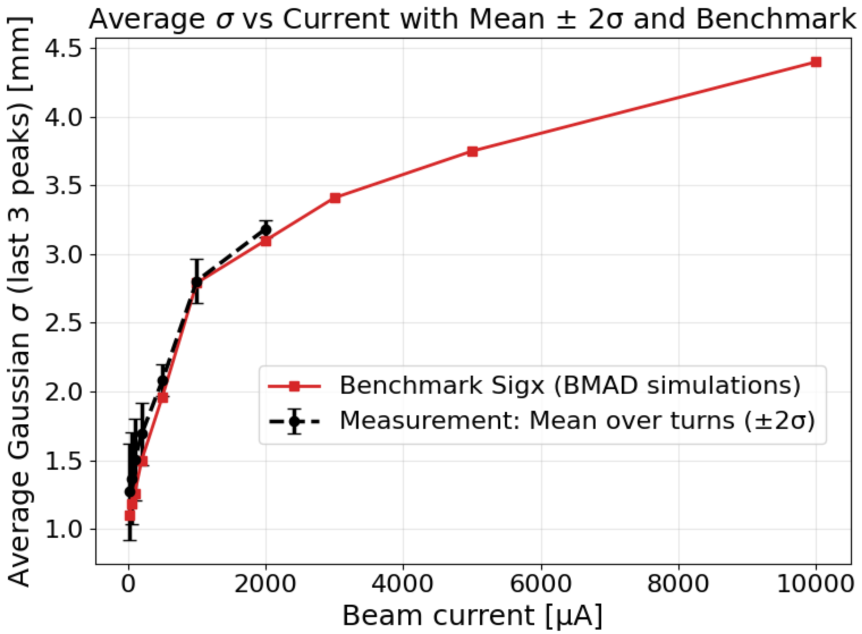

M. HAJ TAHAR - 23/07/2025

From Simulation to Reality - Approach for safe offline testing

- **After discussions with CERN colleagues, two strategic decisions made**:
  - Dual Strategy Testing:
    - Reinforcement Learning (RL)
    - Bayesian Optimization (BO)
  - Simulation-First Validation:
    - Testing conducted on in-house beam tracking simulations of our injector cyclotron
- **Why it matters:**
  - **Over 6 weeks** of RL development and tuning confirmed the feasibility of online learning
  - **Real-time training converges in < 1000 timesteps**
  - Any new feature introduced is first tested on simulations before including in the real-time software

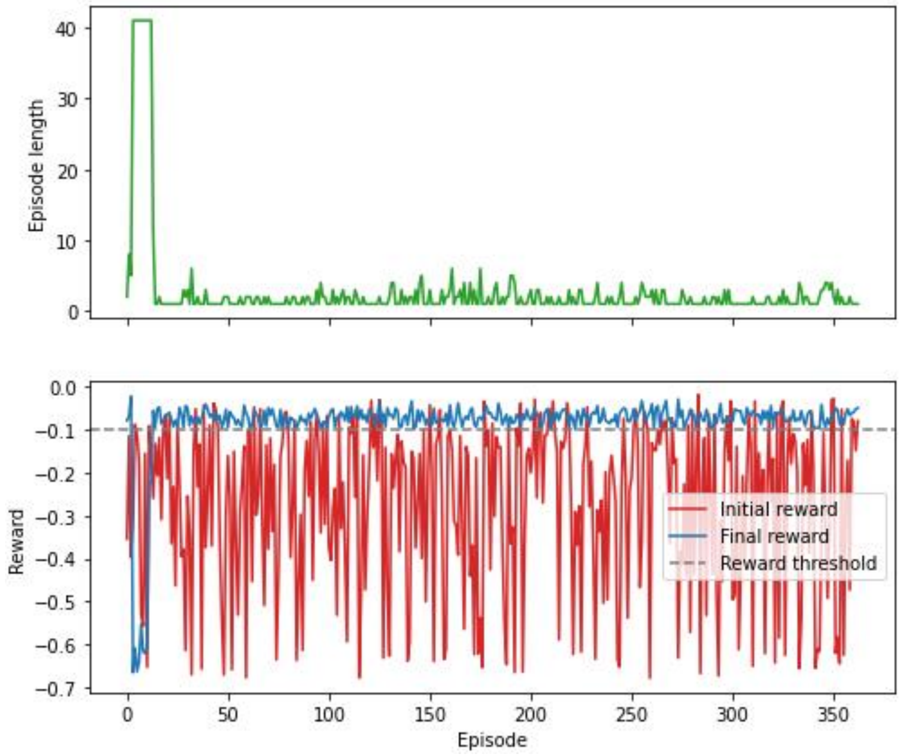

Training from scratch in tracking simulations
Convergence within less than 1000 timesteps

M. HAJ TAHAR - 17/09/2025

Reinforcement Learning Results: Turn-by-Turn Progression

Summary of RL Agent Training Performance

Turn	Resonators Active	Pretraining Used	Convergence Time	Avg. Final Reward	Interlocks	Notable Outcome
72	3 (R1, R2, R3)	No	535 timesteps ~ 4 h	-0.06	21 (3 after conv)	First full training from scratch
73	3 (R1, R2, R3)	Yes	291 timesteps ~ 2 h 20 min	-0.06	12 (0 after conv)	Fast convergence from surrogate
74	3 (R1, R2, R3)	Yes	1117 timesteps ~ 5 h 50 min	-0.06	51 (3 after conv)	Transfer Learning from Turn 73 inadequate
89	2 (R1, R3)	Yes	114 timesteps ~ 52 min	-0.055	0	Most degraded config; still successful
60	4 (R1, R2, R3, R4)	No	217 timesteps ~ 2 h 3 min	-0.06	4 (1 after conv)	Fast convergence under nominal conditions

Convergence is defined as the time required for the reinforcement learning agent to consistently solve any episode in fewer than 10 steps.

Objective: final reward > -0.08

Turn 72 – Fast Convergence, Stable Operation

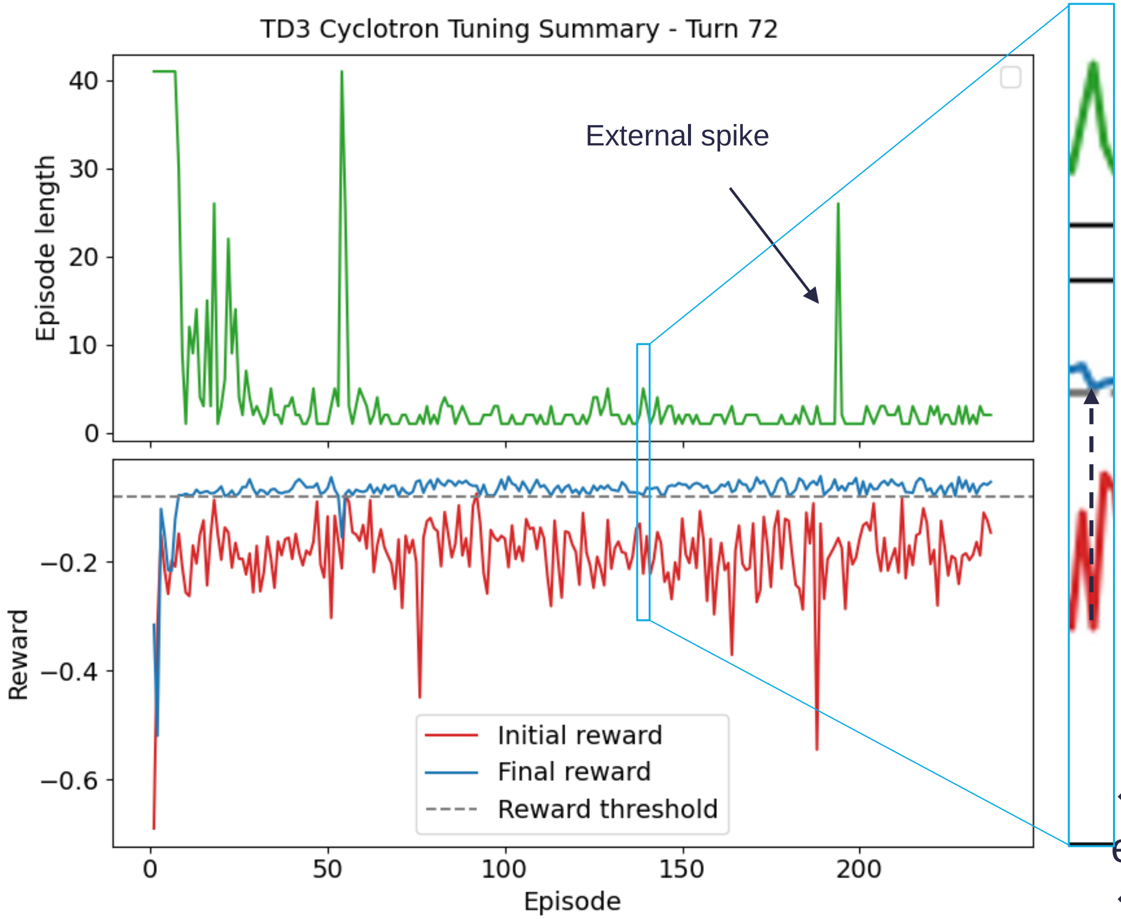

Training from scratch (7 hours in total)

Initial = machine state before action; Final = post-action result by the agent.

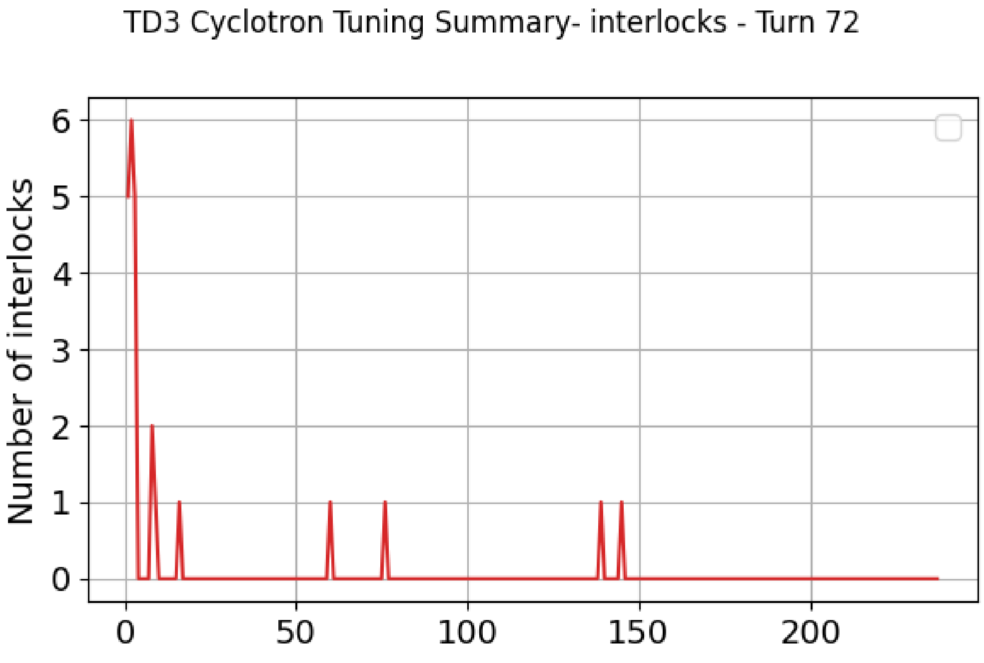

- ✔ Agent learns to solve the tuning task in < 5 steps after ~60 episodes
- ✔ Reward improves from ~–0.7 to –0.05 (×10 better)
- ✔ Episode length drops and stabilizes
- ✔ Convergence time consistent with simulations!
- ✔ Beam interlocks are avoided after early exploration

Summary of RL Agent Training Performance

Turn	Resonators Active	Pretraining Used	Convergence Time	Avg. Final Reward	Interlocks	Notable Outcome
72	3 (R1, R2, R3)	No	535 timesteps ~ 4 h	-0.06	21 (3 after conv)	First full training from scratch
73	3 (R1, R2, R3)	Yes	291 timesteps ~ 2 h 20 min	-0.06	12 (0 after conv)	Fast convergence from surrogate
74	3 (R1, R2, R3)	Yes	1117 timesteps ~ 5 h 50 min	-0.06	51 (3 after conv)	Transfer Learning from Turn 73 inadequate
89	2 (R1, R3)	Yes	114 timesteps ~ 52 min	-0.055	0	Most degraded config; still successful
60	4 (R1, R2, R3, R4)	No	217 timesteps ~ 2 h 3 min	-0.06	4 (1 after conv)	Fast convergence under nominal conditions

Convergence is defined as the time required for the reinforcement learning agent to consistently solve any episode in fewer than 10 steps.

Turn 73

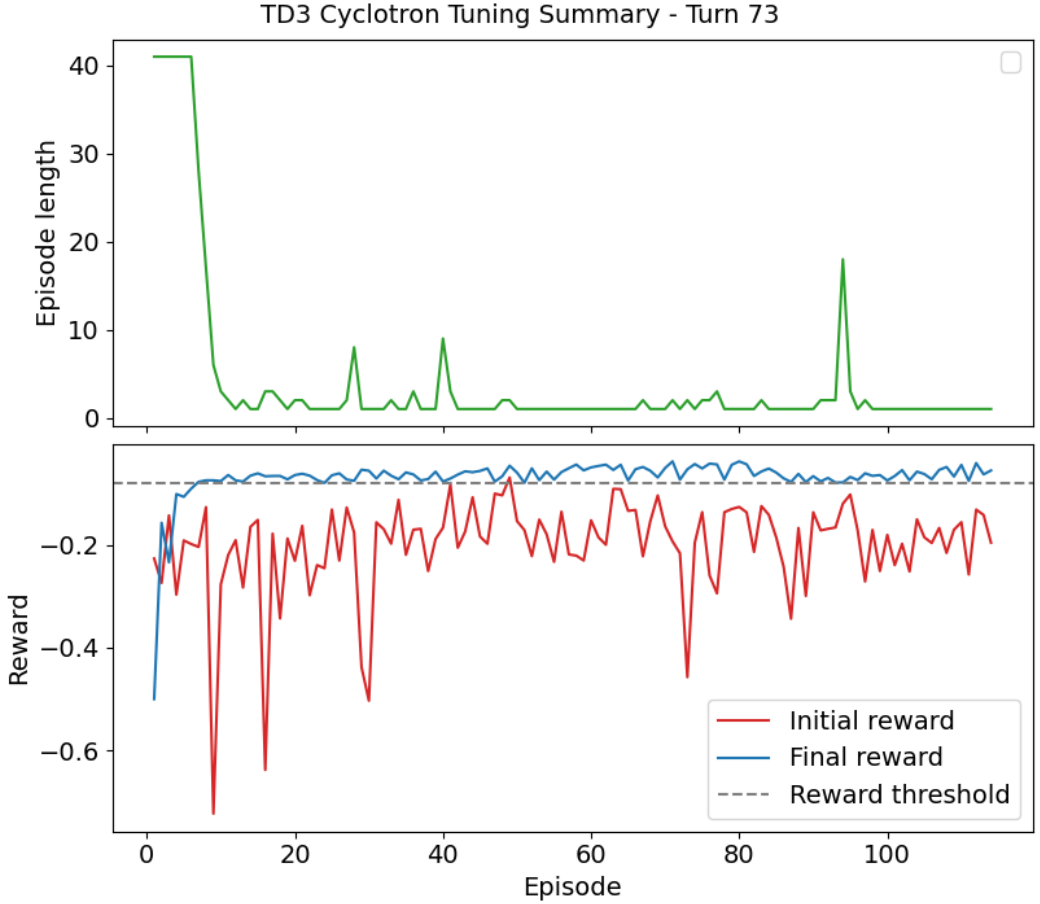

Pre-trained RL model from historical data (3 h 53 min)

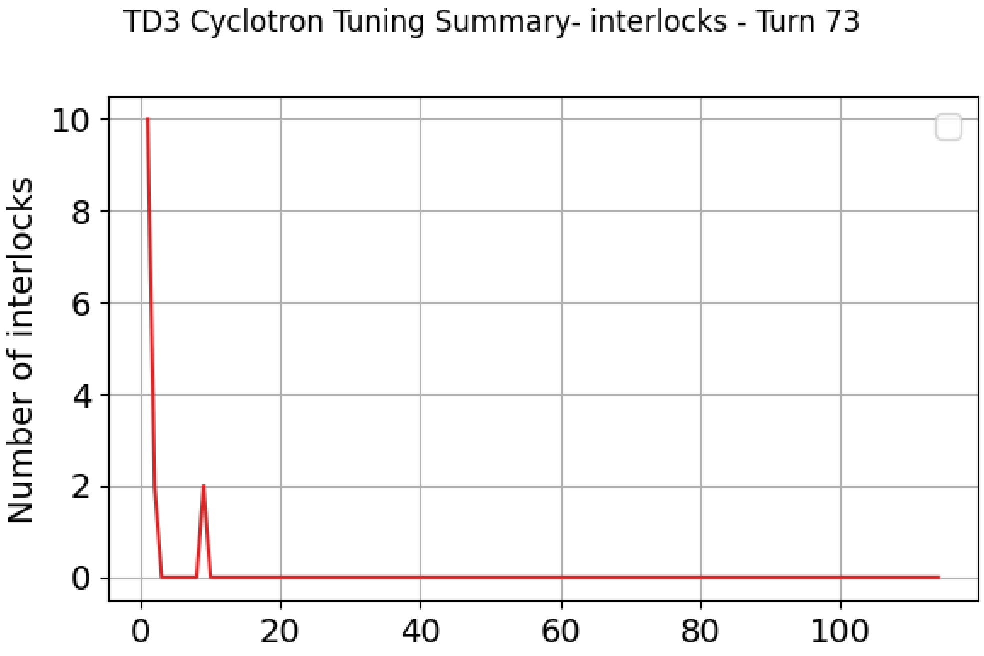

✔ Beam interlocks are avoided after early exploration

- Two models tested on evaluation mode 12 hours later: pre-trained model outperforms the one from scratch

Turn 73: training vs no training

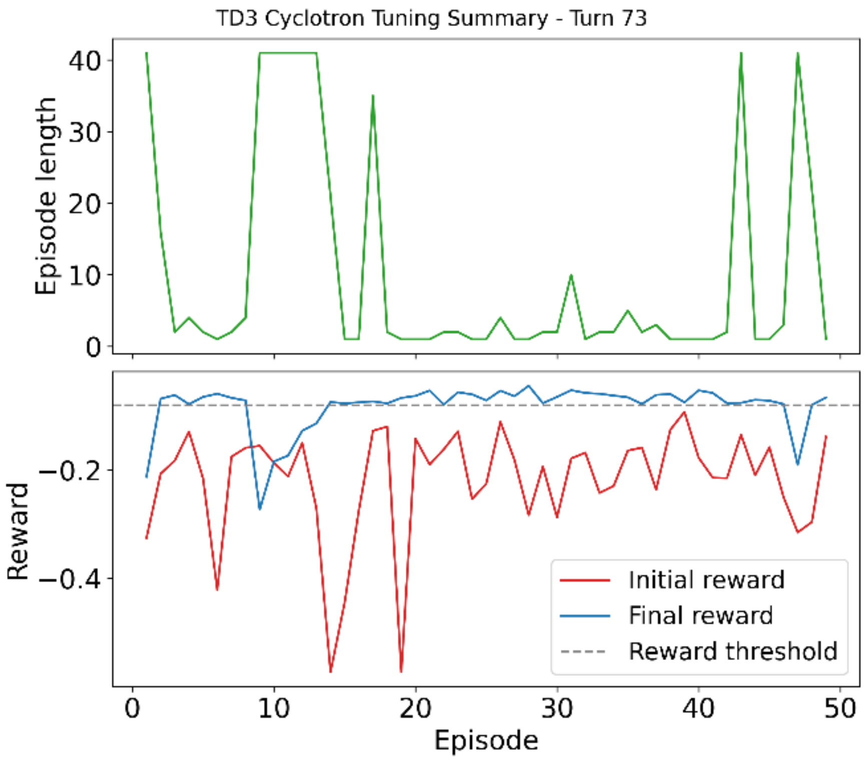

Results without pretraining

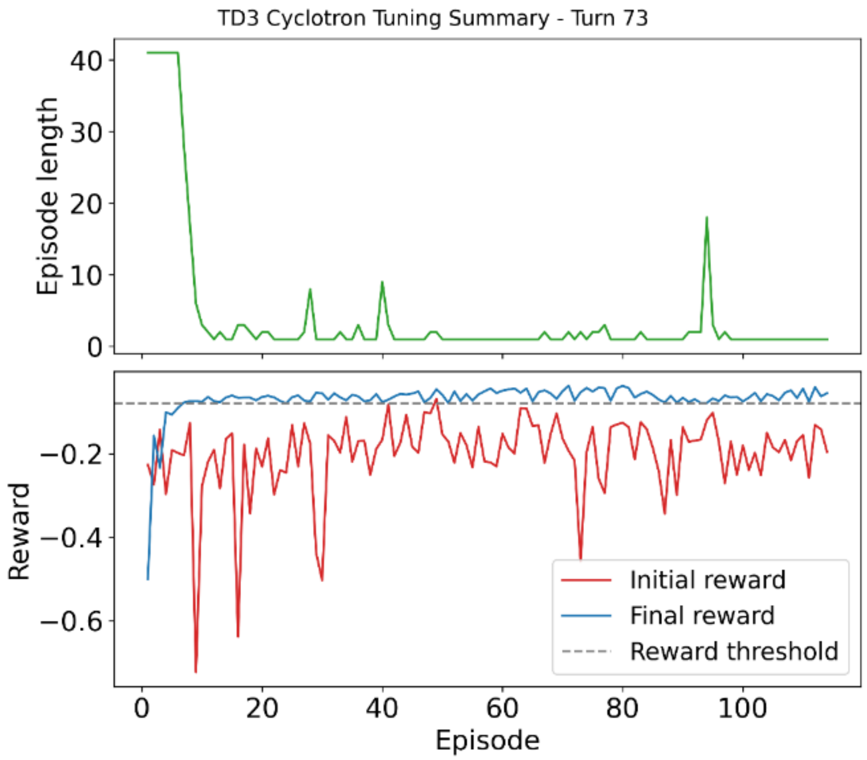

Results with pretraining

- Pretraining from a surrogate model significantly accelerated learning and provided a strong initialization

Summary of RL Agent Training Performance

Turn	Resonators Active	Pretraining Used	Convergence Time	Avg. Final Reward	Interlocks	Notable Outcome
72	3 (R1, R2, R3)	No	535 timesteps ~ 4 h	-0.06	21 (3 after conv)	First full training from scratch
73	3 (R1, R2, R3)	Yes	291 timesteps ~ 2 h 20 min	-0.06	12 (0 after conv)	Fast convergence from surrogate
74	3 (R1, R2, R3)	Yes	1117 timesteps ~ 5 h 50 min	-0.06	51 (3 after conv)	Transfer Learning from Turn 73 inadequate
89	2 (R1, R3)	Yes	114 timesteps ~ 52 min	-0.055	0	Most degraded config; still successful
60	4 (R1, R2, R3, R4)	No	217 timesteps ~ 2 h 3 min	-0.06	4 (1 after conv)	Fast convergence under nominal conditions

Convergence is defined as the time required for the reinforcement learning agent to consistently solve any episode in fewer than 10 steps.

Turn 74: Transferability of Pretrained RL Actor

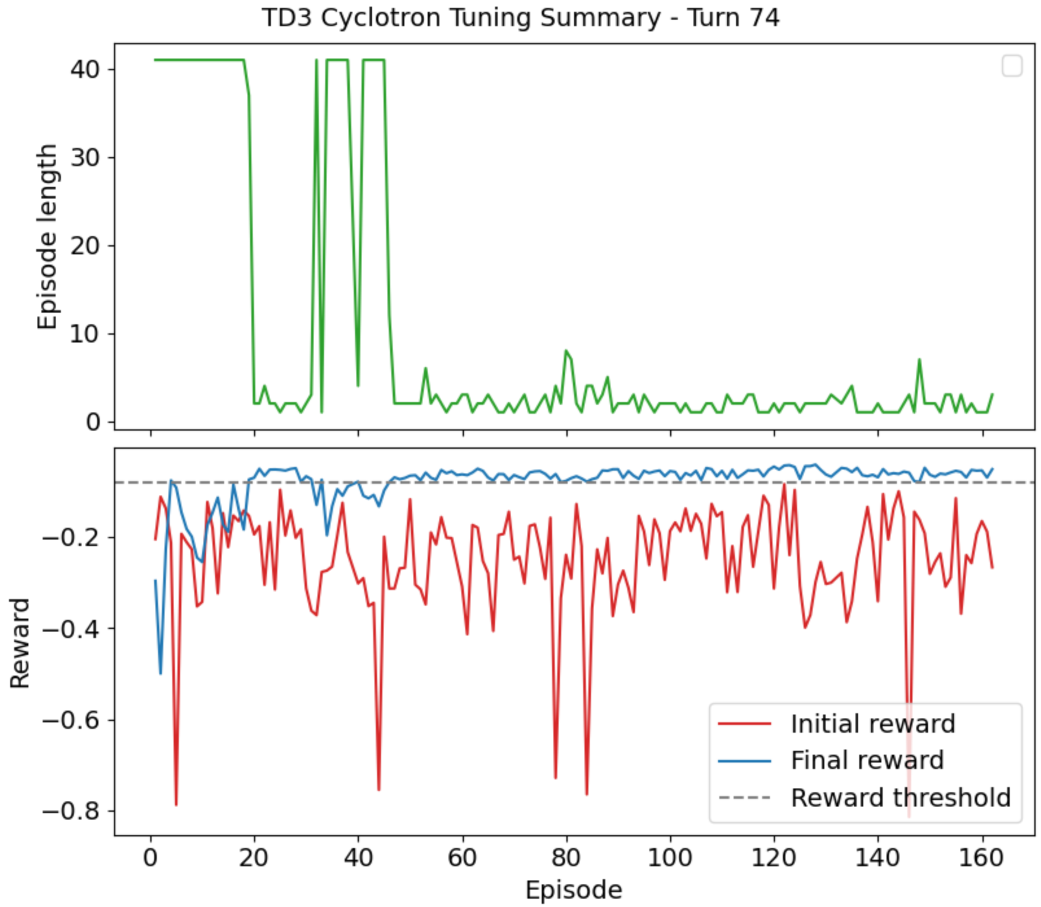

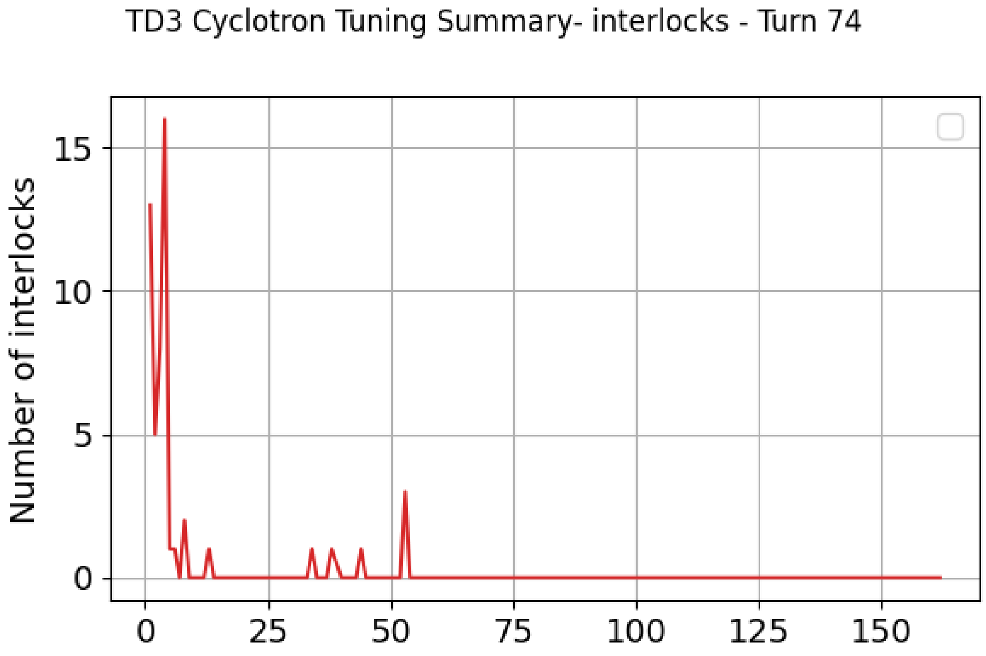

Pre-trained RL actor from turn 73 (7 h 50 min)

- Test conducted: RL actor trained at turn 73 was applied to Turn 74
- Performance degraded compared to training from scratch → turn-specific training is essential

Lessons Learned from Real-Time RL Control – Turn-by-Turn Deployment

- **Simulation results provided an accurate starting point**, helping shape a realistic RL training environment for the live cyclotron
- **Training a model for a single turn typically takes 3 to 6 hours** (less than 1000 timesteps), depending on exploration noise chosen, can be accelerated with another approach
- **We systematically validated across different turn numbers** to confirm the robustness of the RL strategy in all machine configurations
- **A model trained on one turn is not valid for another,** hence the multi-turn training effort
- **The agent quickly learns to avoid beam losses (interlocks)** by internalizing strong penalties during training, a key success factor for safe & reliable deployment

Evaluating RL Model Reliability Under Real Conditions

- RL agent operates in evaluation mode, continuously controlling the machine
- Tested overnight for robustness to drifts and perturbations
- Agent monitors reward and reacts when performance drops

Evaluation mode – Turn 74 (3 resonators)

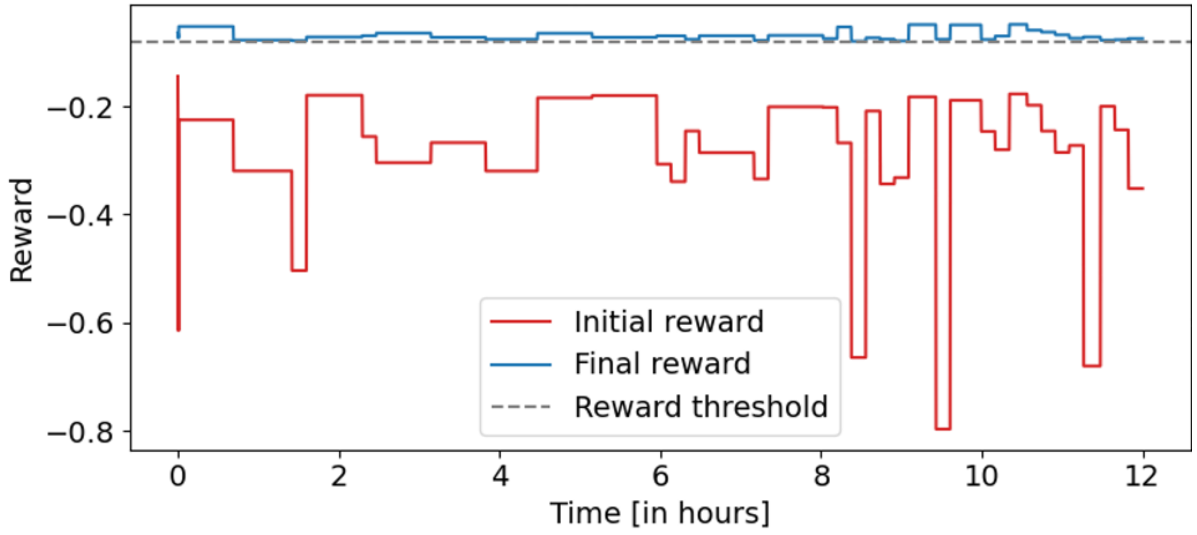

- Trained model in the morning, launched in evaluation mode from 7pm to 7am
- Throughout the night, the machine was **intermittently perturbed** with random coil kicks to test model robustness
- The RL agent consistently **restored beam quality**:
  - ➢ No interlocks triggered
  - ➢ Final reward remained high, losses stayed low

- A 30-minute retraining corrected this drift → performance fully restored → **importance of intermittent retraining**

Evaluation mode – Turn 60 (4 resonators)

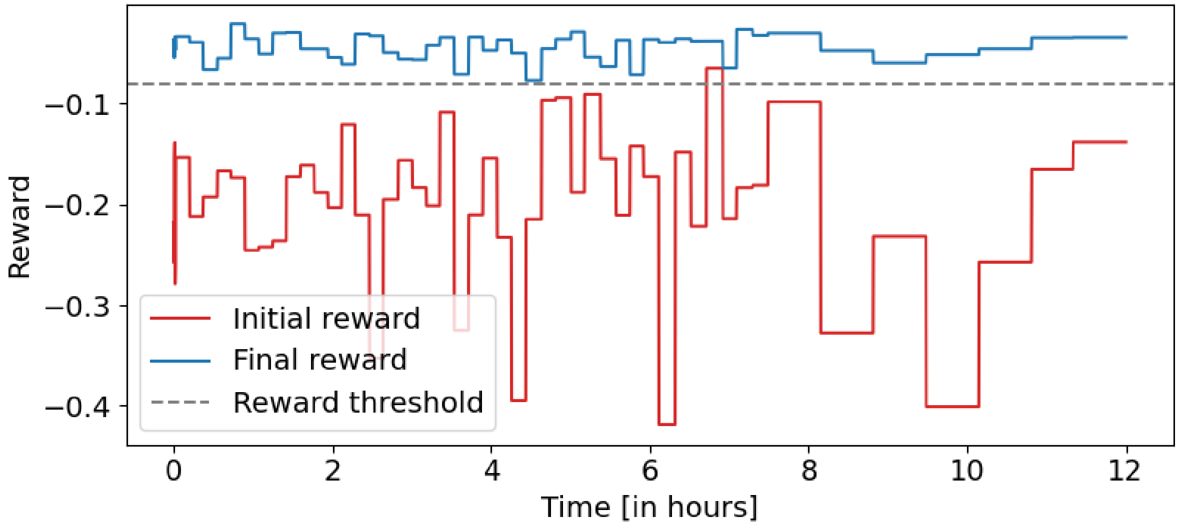

- Trained model during the day, launched in evaluation mode from 9pm to 9am:
  - ➢ No interlocks triggered
  - ➢ Succeeds to exceed the reward threshold level and bring losses to nearly vanishing levels

Key Takeaways – Model Reliability in Real Operation

- **First live deployment of an RL-based control loop** on a cyclotron, running autonomously for **12 hours**
  → Continuously tuned in real time, adapting to drift and fluctuations

- Successfully tested under **three machine configurations**:
  - All 4 resonators (Turn 60)
  - Only 3 resonators (Turn 74)
  - Only 2 resonators (Turn 89)

  → The RL agent consistently recovered and adapted

- Demonstrated that RL-based tuning can become the **default control strategy** for future systems like the **TMX cyclotron**

- **Intermittent retraining** boosts robustness: overnight tests showed that occasional retraining helps recover optimal performance

- **Current limitation**: validated at **20 µA**, i.e., relatively low current
  → How will this scale to high-current, high-power regimes? **Results at the cyclotron conference** ☺

Conclusion

- **Reliable high-power cyclotrons** are essential for Accelerator-Driven Systems (ADS).

- **Progress across hardware, design optimization, and operations** forms the foundation for meeting reliability targets.

- **Machine Learning-assisted tuning** offers a promising path to fast fault compensation and robust operation.

- PSI Injector 2 experiment demonstrated feasibility; **HIPA will be the next step** toward ADS-scale deployment.

Acknowledgements

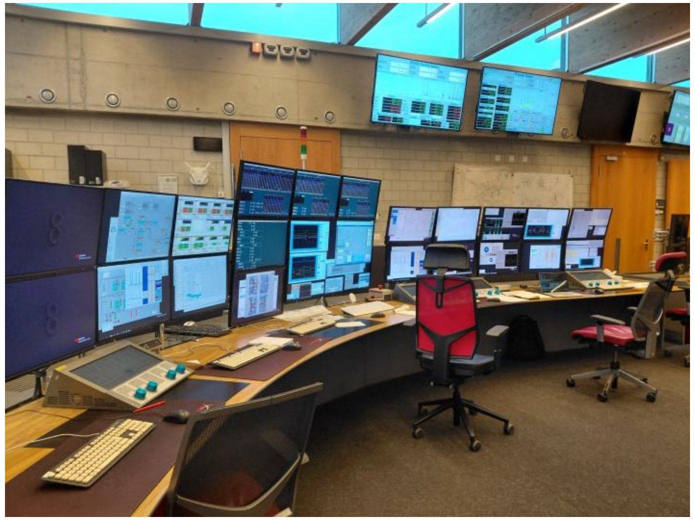

Acknowledgements:

The success of this RL-based tuning project at PSI Injector 2 was made possible by the contributions, support, and collaboration of many talented individuals and teams, from **PSI, CERN and TRANSMUTEX**, in particular:

- **Evgeny Solodko** for his continuous involvement from the beginning, for his work on the historical data, surrogate model and on-site deployment
- **Marco Bocchio** for elaborating an alternative approach (BO) and contributing to help refine our method
- **Serge Marquie** for his invaluable support and stimulating discussions regarding the ML models
- **Jochem Snuverink** for setting up the TMX server infrastructure and continuous help and support with technical aspects
- The Cyclotron team at PSI: **Antonio Barchetti** for his availability, advice and support during critical phases of deployment,

Christian Baumgarten, Joachim Grillenberger, Markus Schneider, Mariusz Sapinski and Rudolf Dolling for their impactful discussions, suggestions and support

- **Franklin Servan-Schreiber (CEO), Donovan Maire (CTO) and Marco Busch (head Acc)** for their continuous support

Opportunities at TRANSMUTEX

- We are hiring in the field of **accelerator physics and technology**
- Interested? Visit our website: Transmutex | Innovators in Sustainable Energy
- Or contact me directly (e.g. via LinkedIn) - I will forward your request to the HR

TRANSMUTEX

www.transmutex.com