Real-time AI for Accelerator Control: A Study at the Fermilab Booster

Christian Herwig, for the Accelerator AI Team

CPAD 2021

March 19, 2021

FNAL Accelerator Complex

Over the next years, FNAL is working towards a major upgrade of the accelerator complex, called the **Proton Improvement Plan-II (PIP-II)**

Goal: achieve a Megawatt proton beam, to meet the required proton per pulse density for DUNE physics

Requires: new Linac, downstream improvements to maintain luminosity

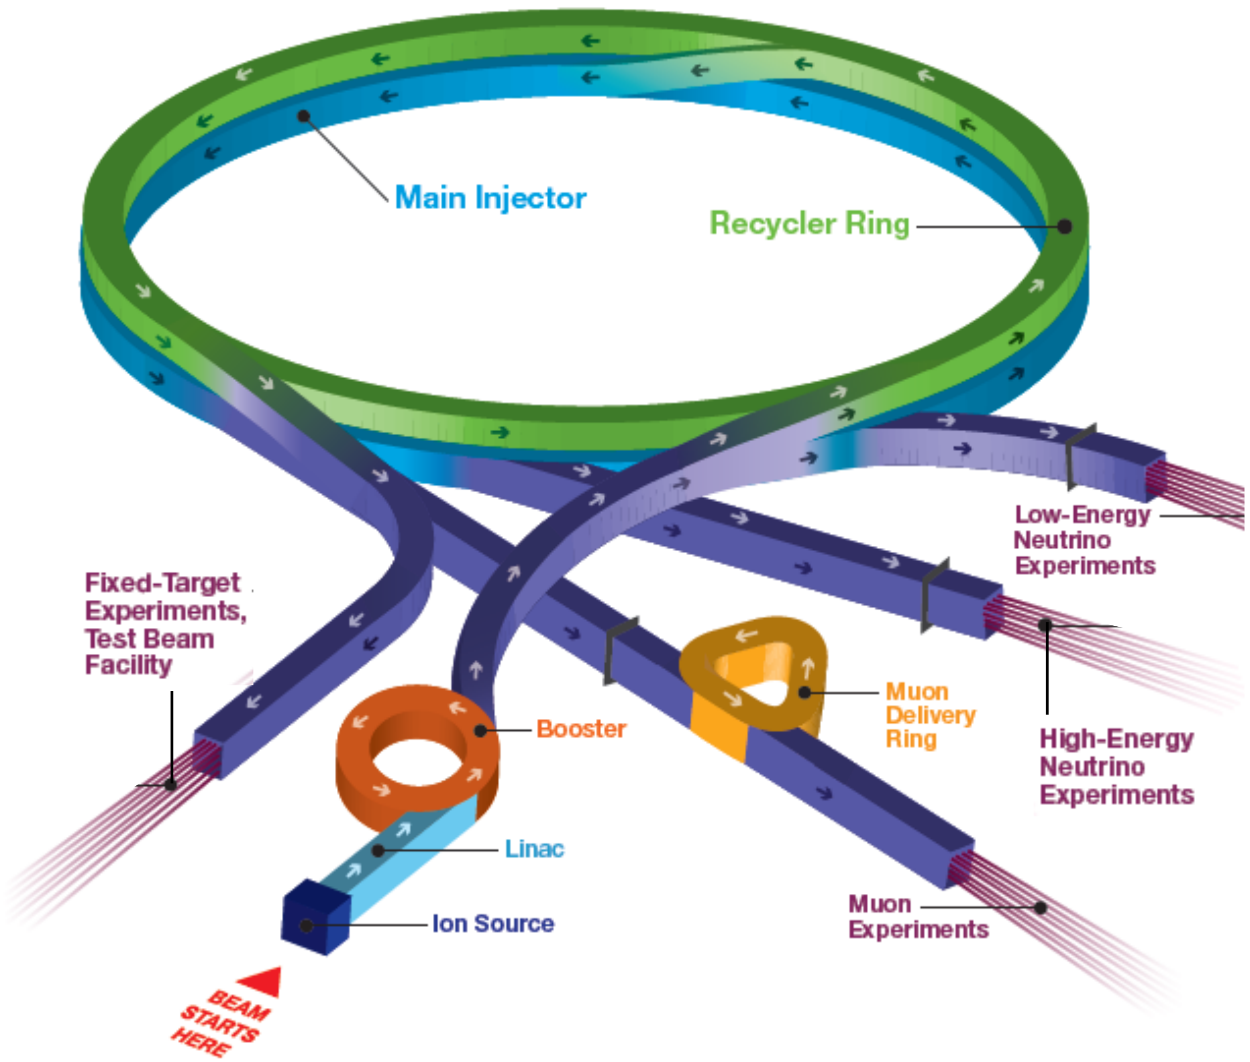

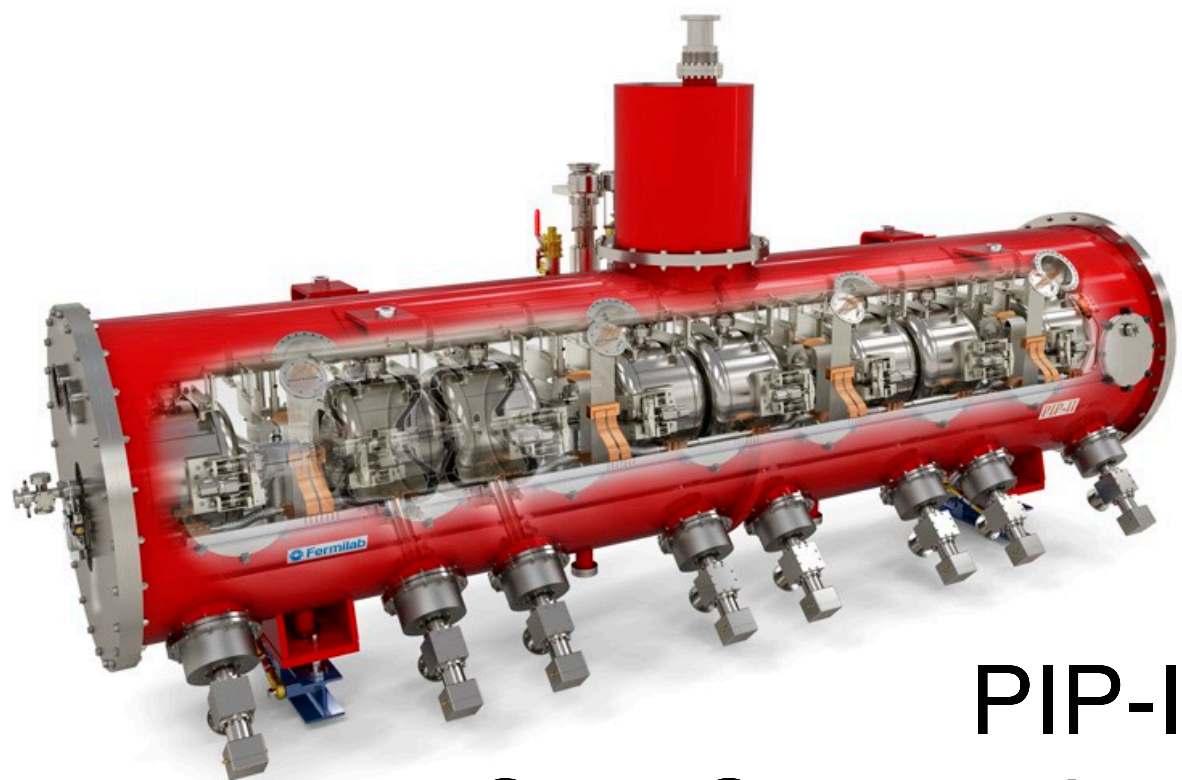

PIP-II
SRF Cryomodule

FNAL Accelerator Complex

Booster synchrotron: accelerates protons 400 MeV → 8 GeV, and delivers to Main Injector and experiments (LBNF / DUNE)

Without upgrade, Booster beam losses will limit DUNE luminosity

→ Proposal of ML regulator for enhanced beam control ([2011.07371](https://arxiv.org/abs/2011.07371))

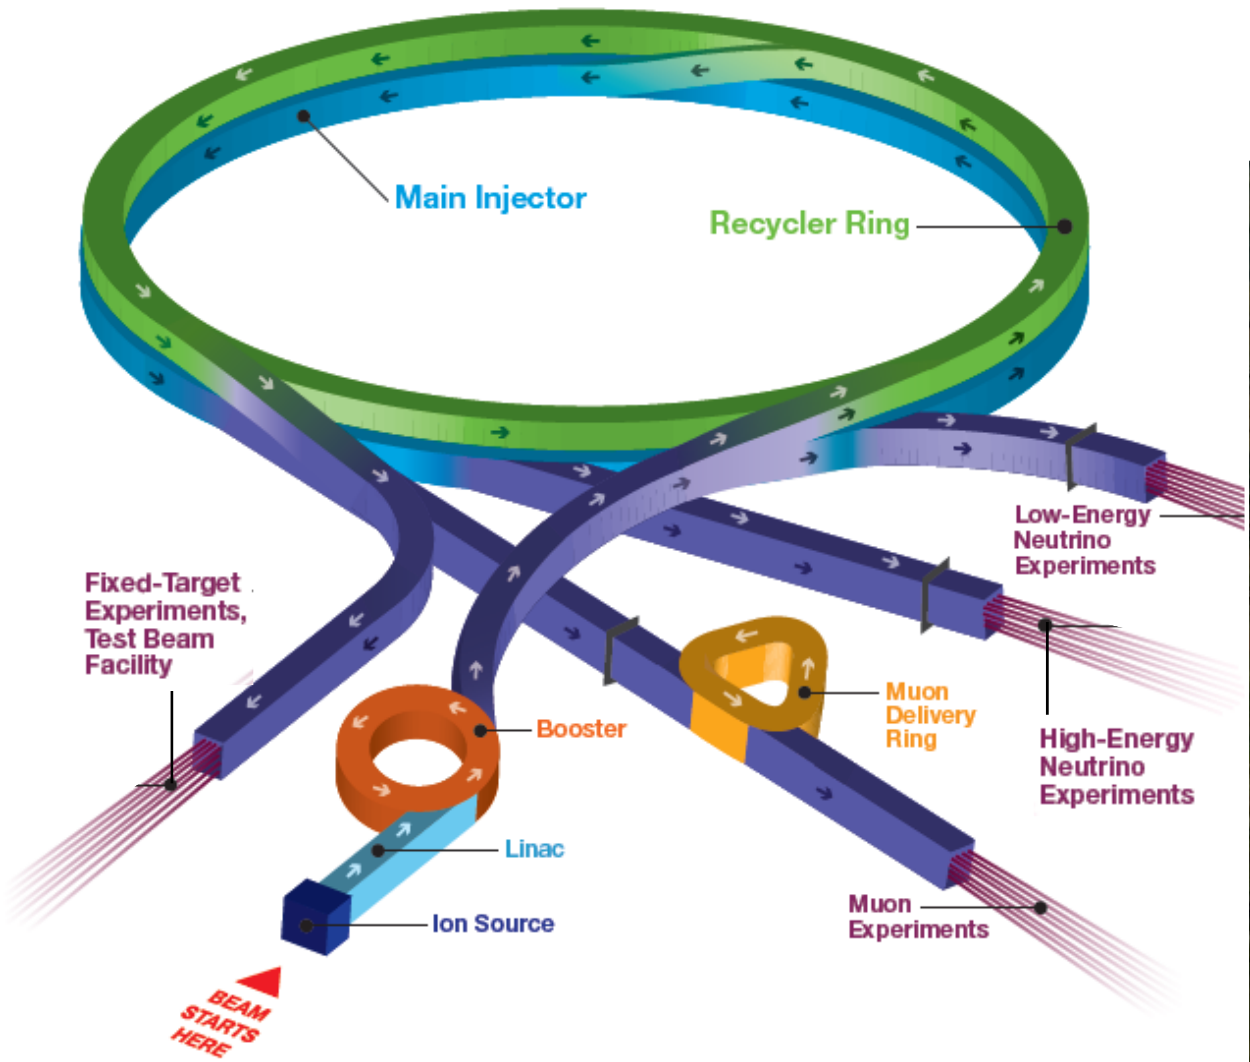

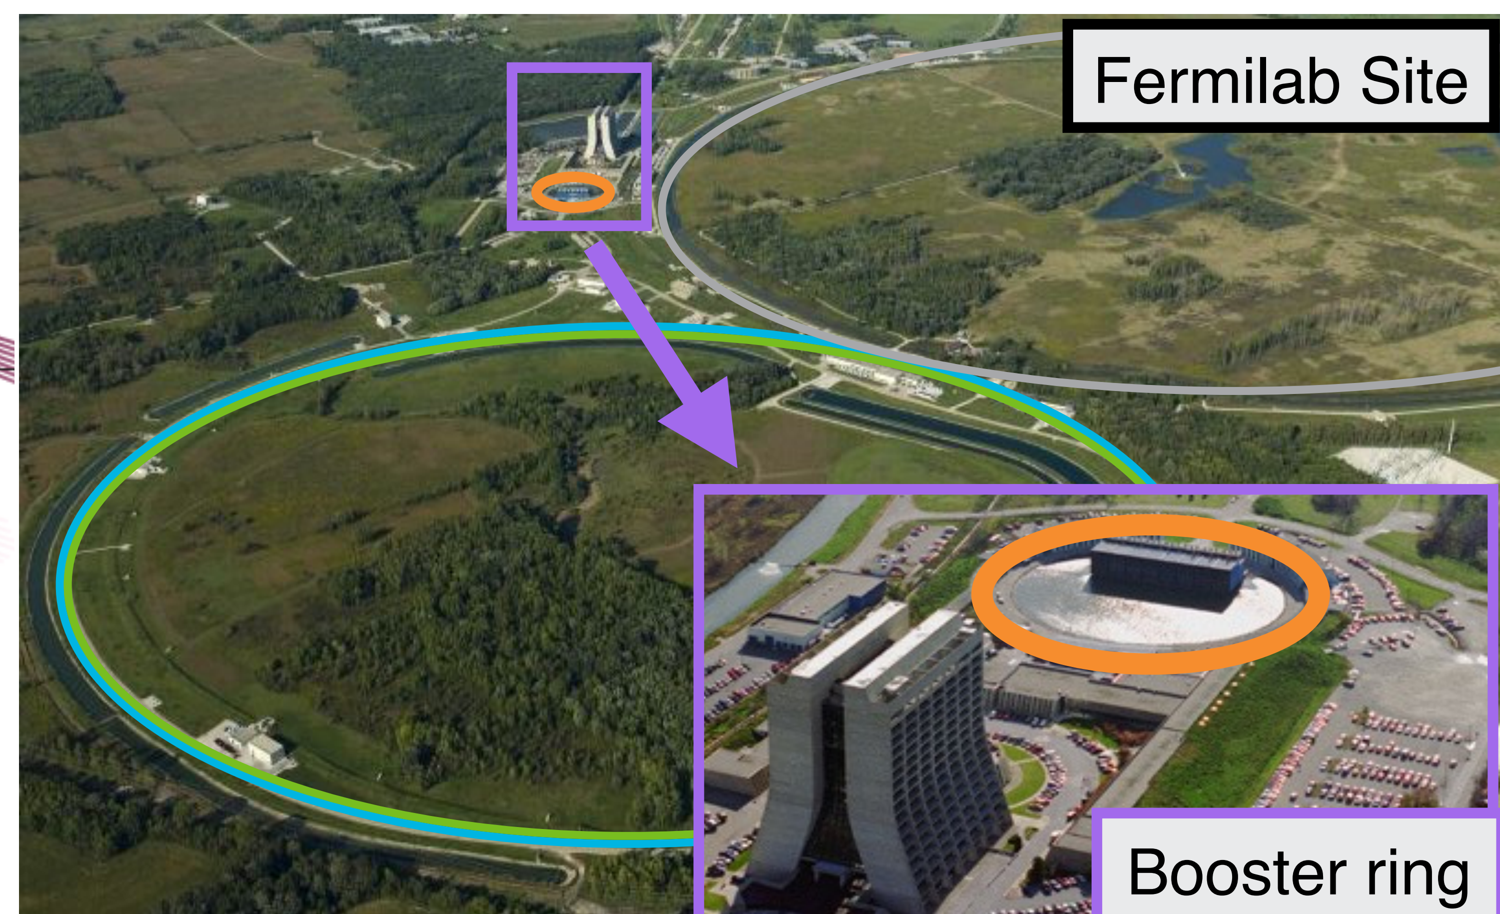

A single Booster cycle

- Combination of RF cavities 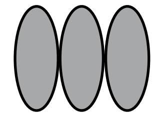 and bending magnets 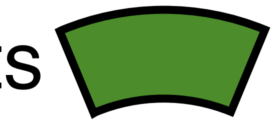
- Bending magnet current ramps in 15hz cycles to maintain the orbit of the accelerating proton beam

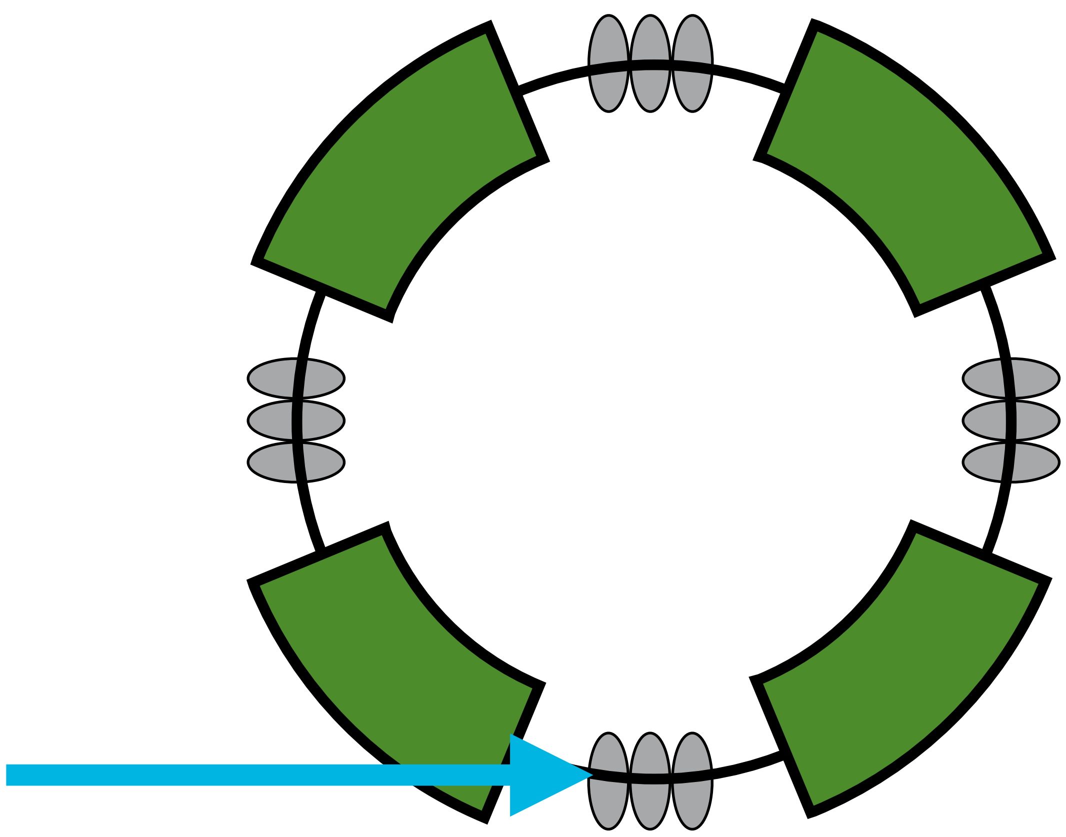

Gradient Magnet Power Supply (GMPS)

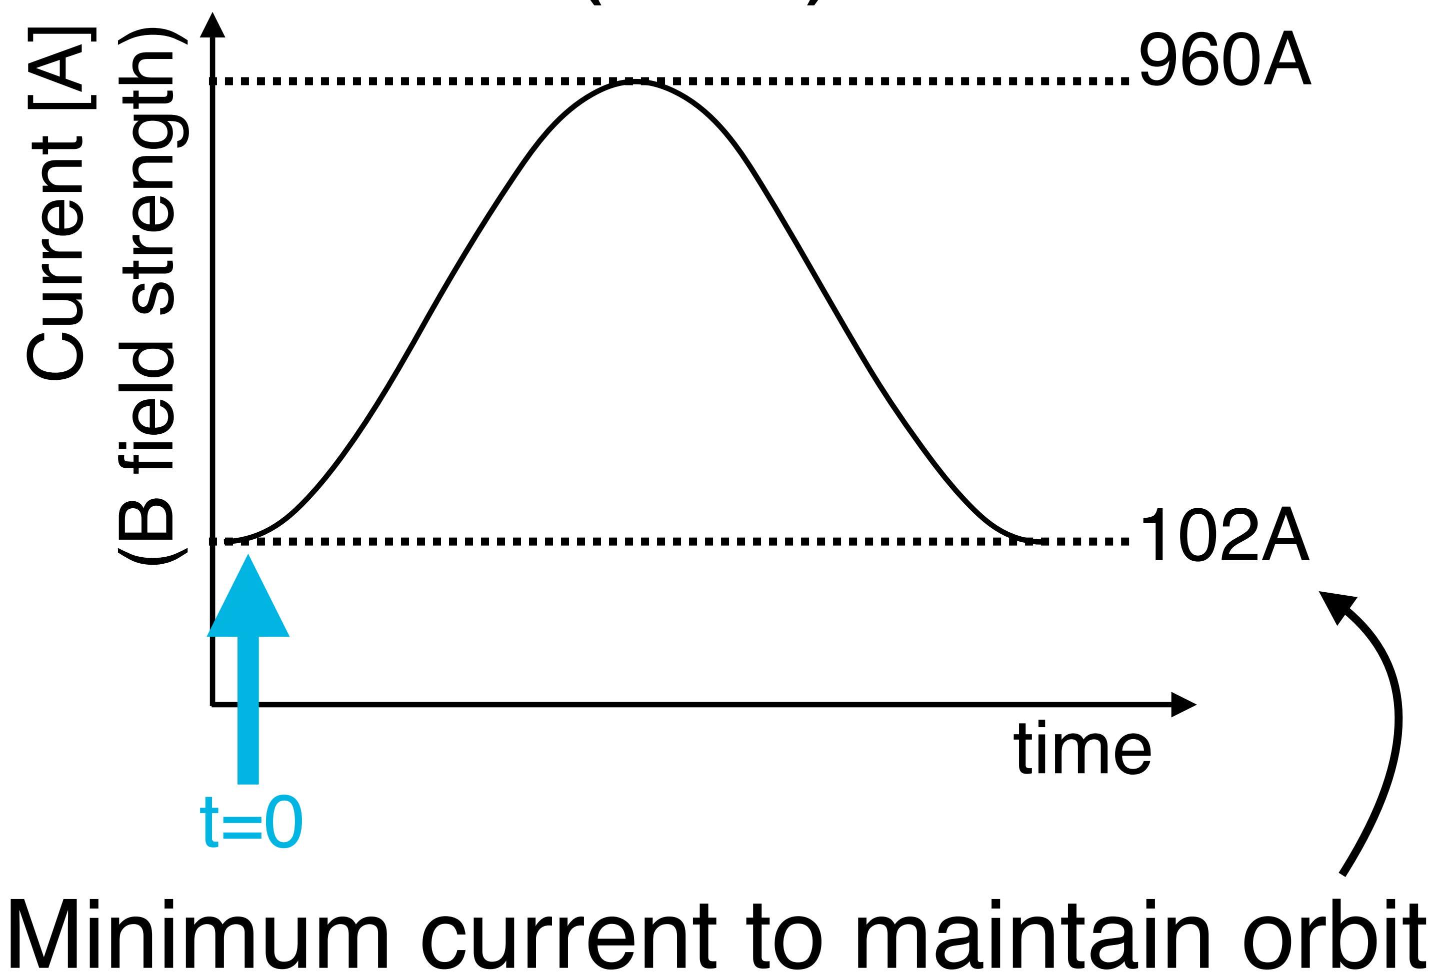

A single Booster cycle

- Combination of RF cavities 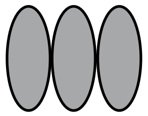 and bending magnets 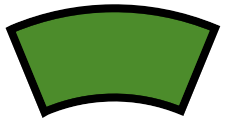
- Bending magnet current ramps in 15hz cycles to maintain the orbit of the accelerating proton beam

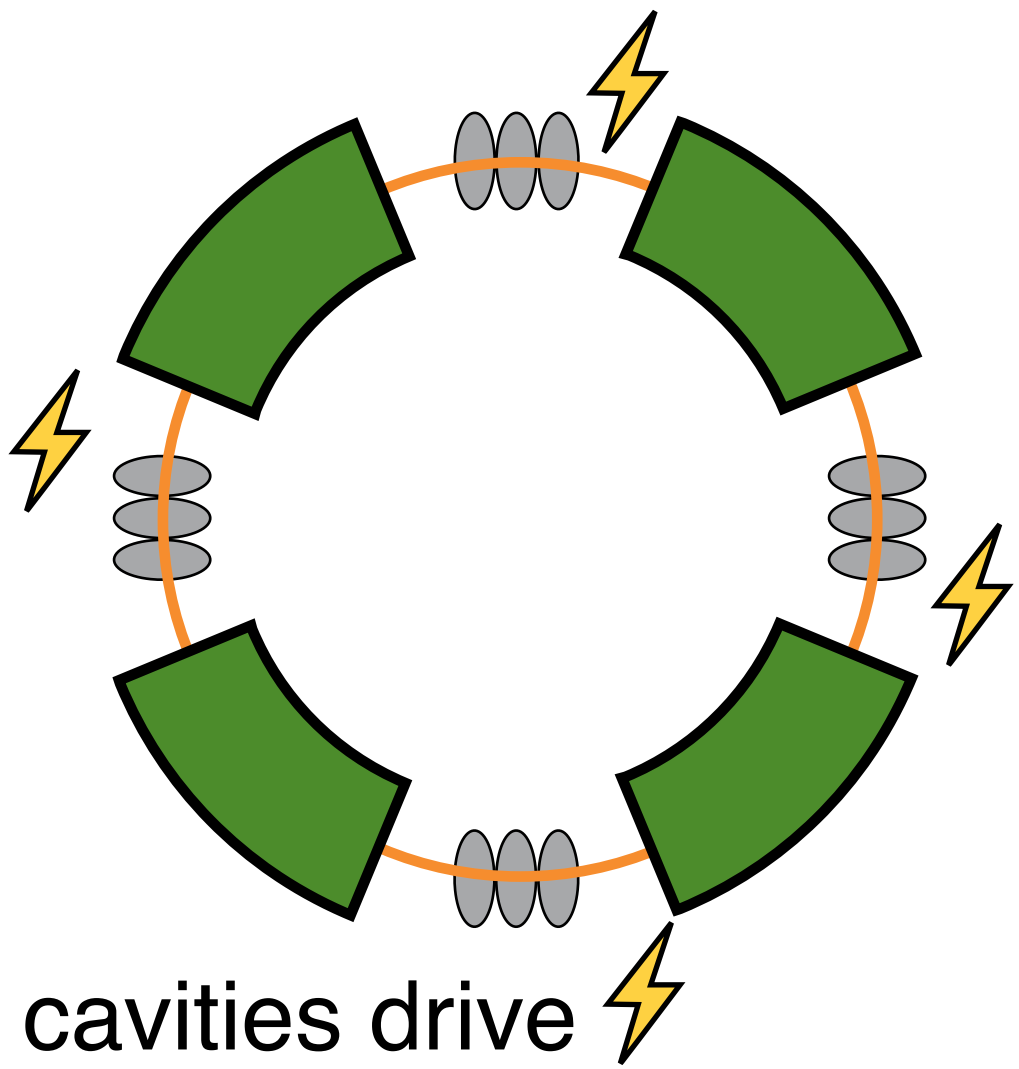

RF cavities drive further acceleration

Gradient Magnet Power Supply (GMPS)

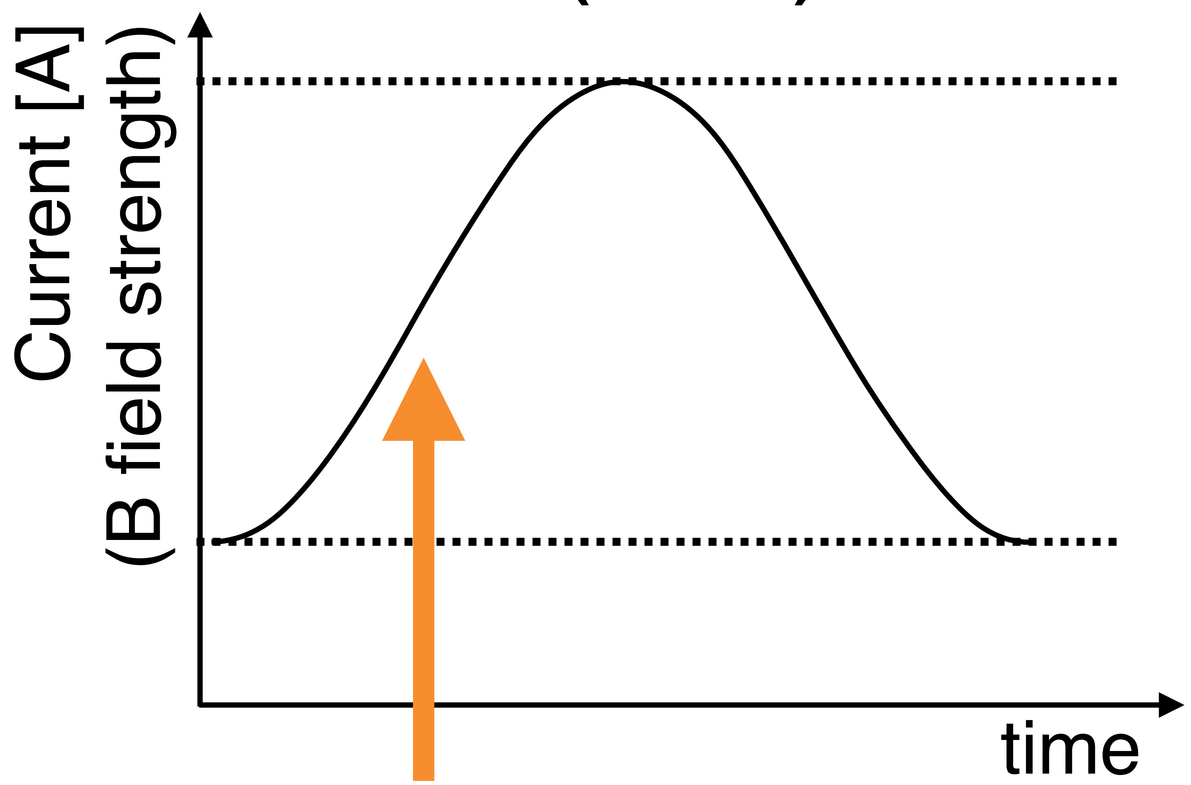

Current increases to maintain orbit

A single Booster cycle

- Combination of RF cavities and bending magnets
- Bending magnet current ramps in 15hz cycles to maintain the orbit of the accelerating proton beam

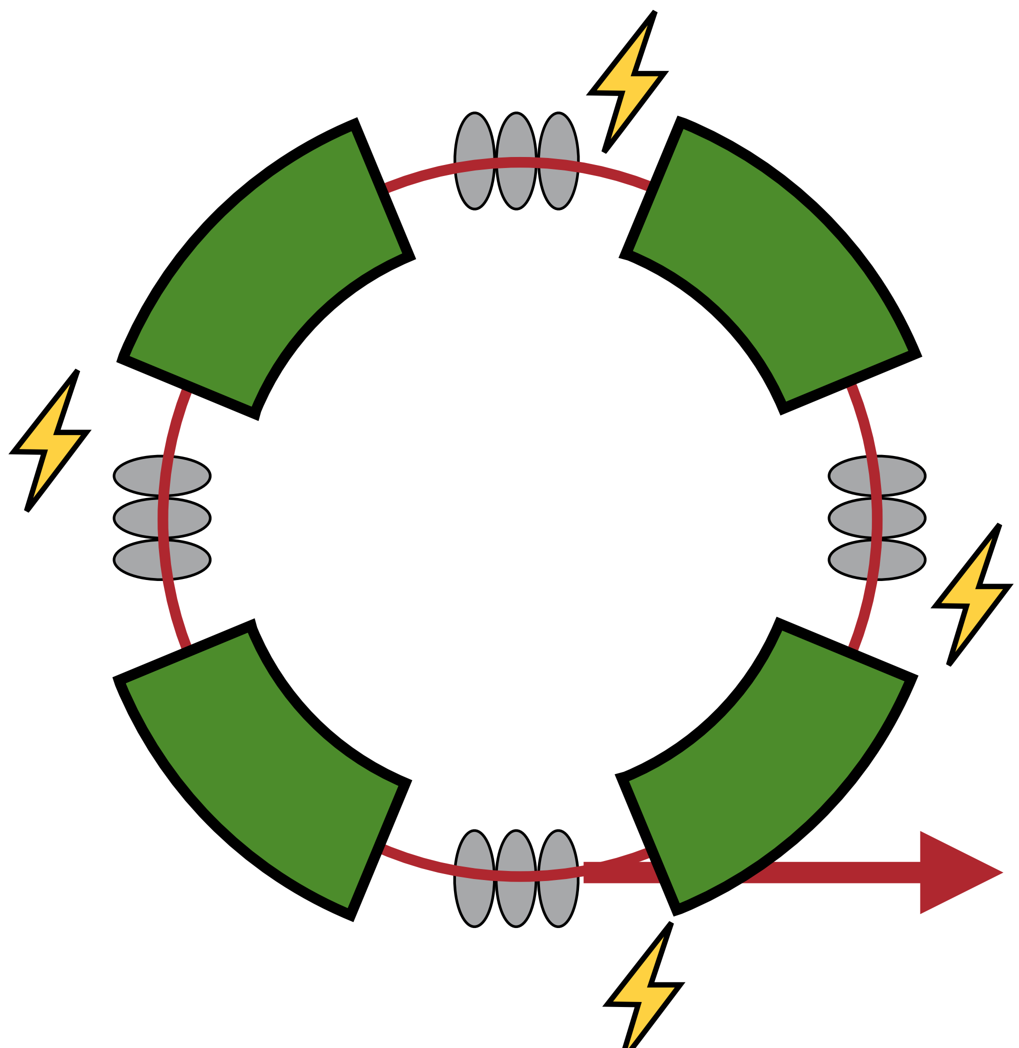

Gradient Magnet Power Supply (GMPS)

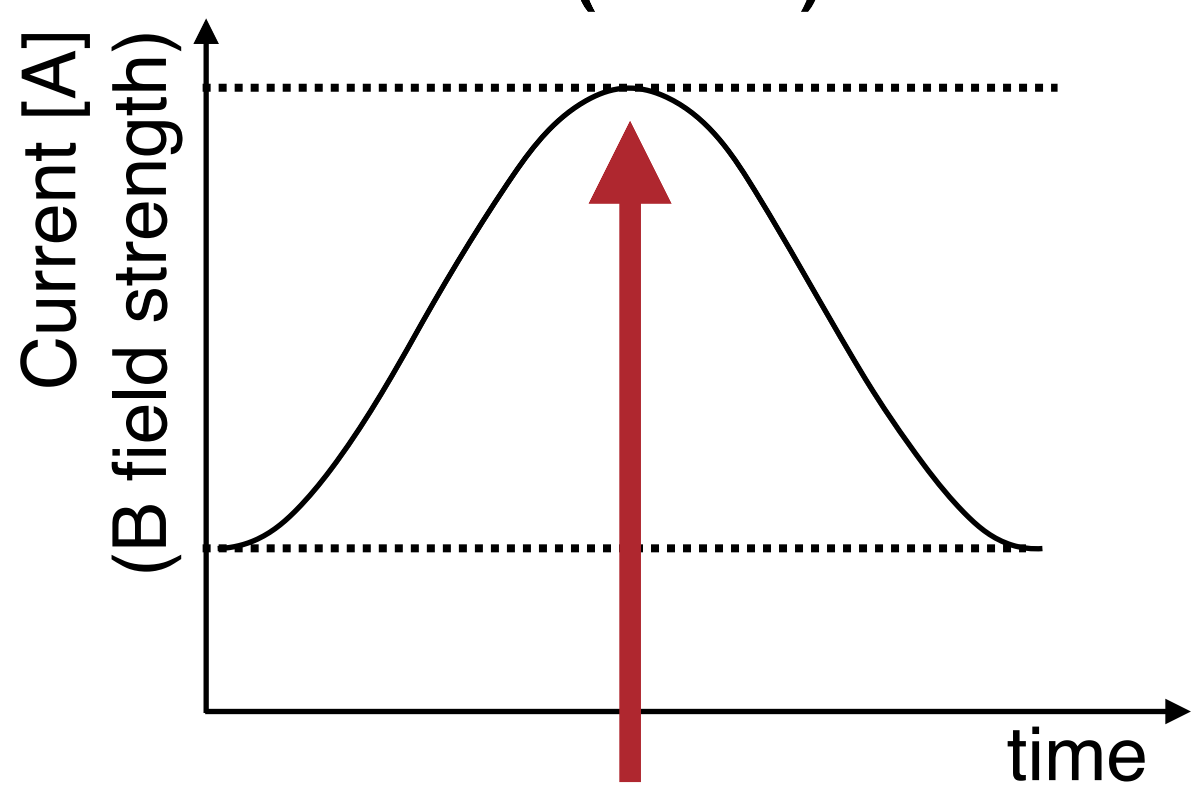

8 GeV beam extracted at maximum B-field

A single Booster cycle

- Combination of RF cavities 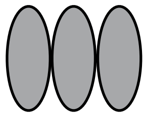 and bending magnets 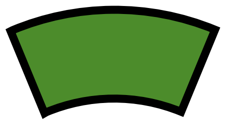
- Bending magnet current ramps in 15hz cycles to maintain the orbit of the accelerating proton beam

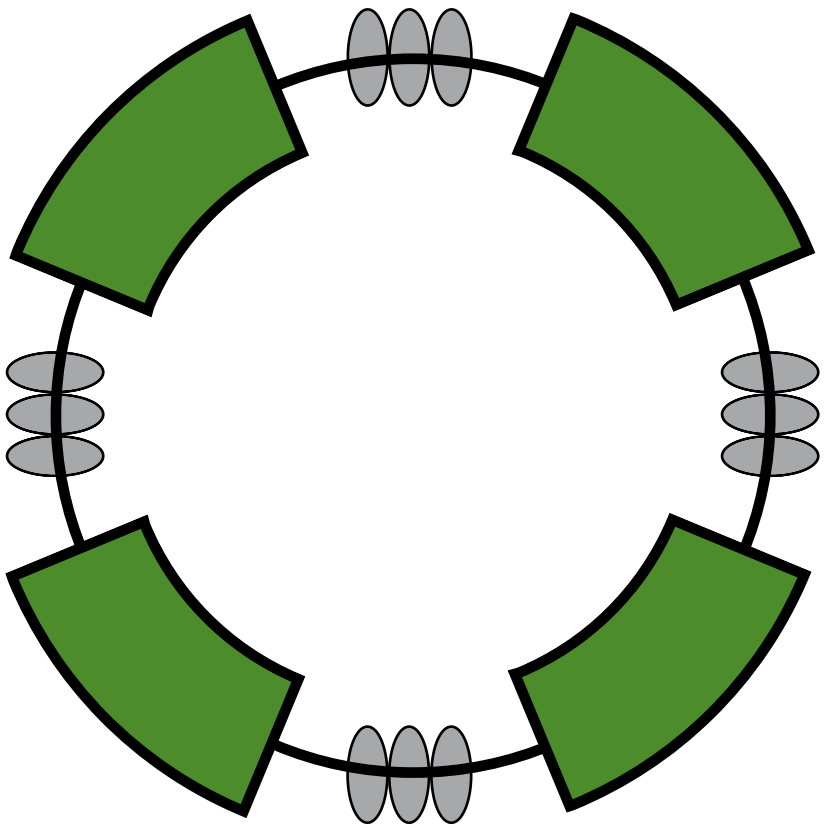

Gradient Magnet Power Supply (GMPS)

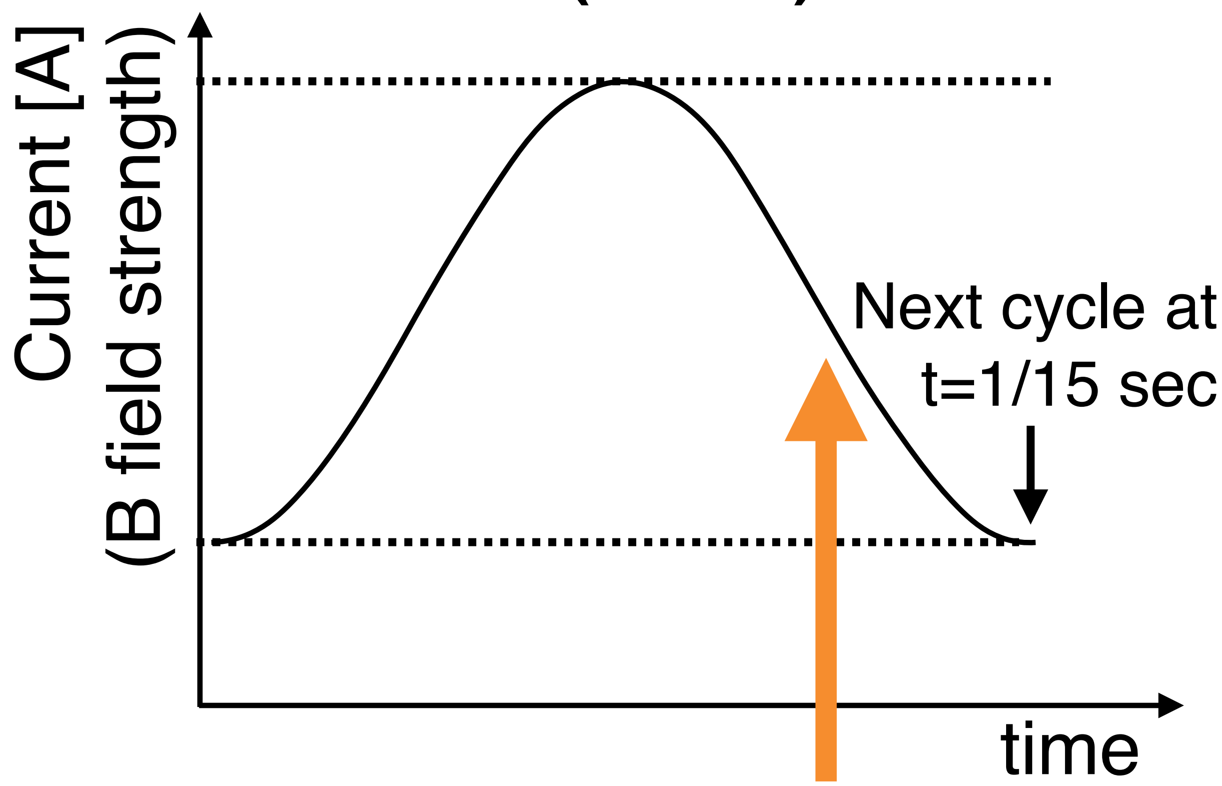

Booster unfilled for half-cycle

Ramp-down for next batch

GMPS current stability

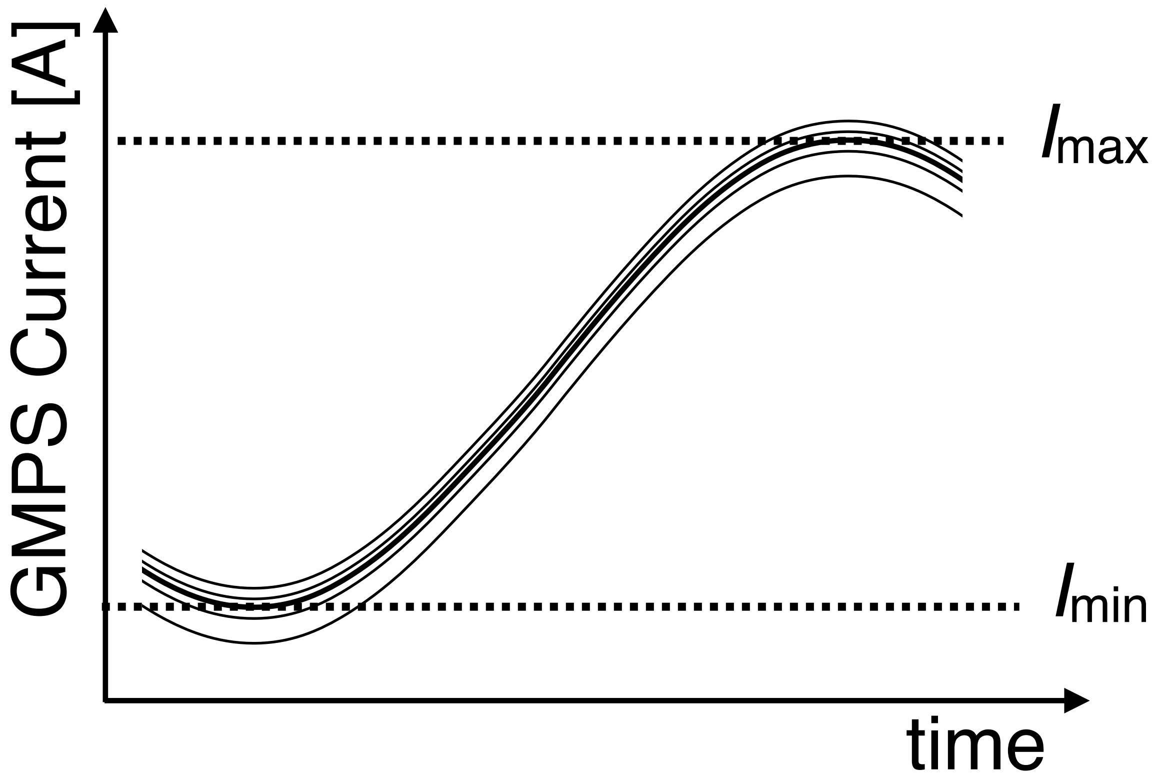

Sinusoidal waveform is prescribed for GMPS current

Measured current does not perfectly match prescription
→ Relative difference is O(%)

This spread in GMPS current (B-field) degrades the beam quality, leading to lost protons

Controls problem: How can one precisely manipulate the magnetic field to mitigate beam losses?

GMPS current stability

Current system incorporates feedback via a "PI loop".

$$I_{\mathrm{set}}(t_1) = I_{\mathrm{target}} - \alpha \cdot \mathrm{Err}(t_0) - \beta \sum_{i=-N}^{0} \mathrm{Err}(t_i)$$

Proportional, integral compensation

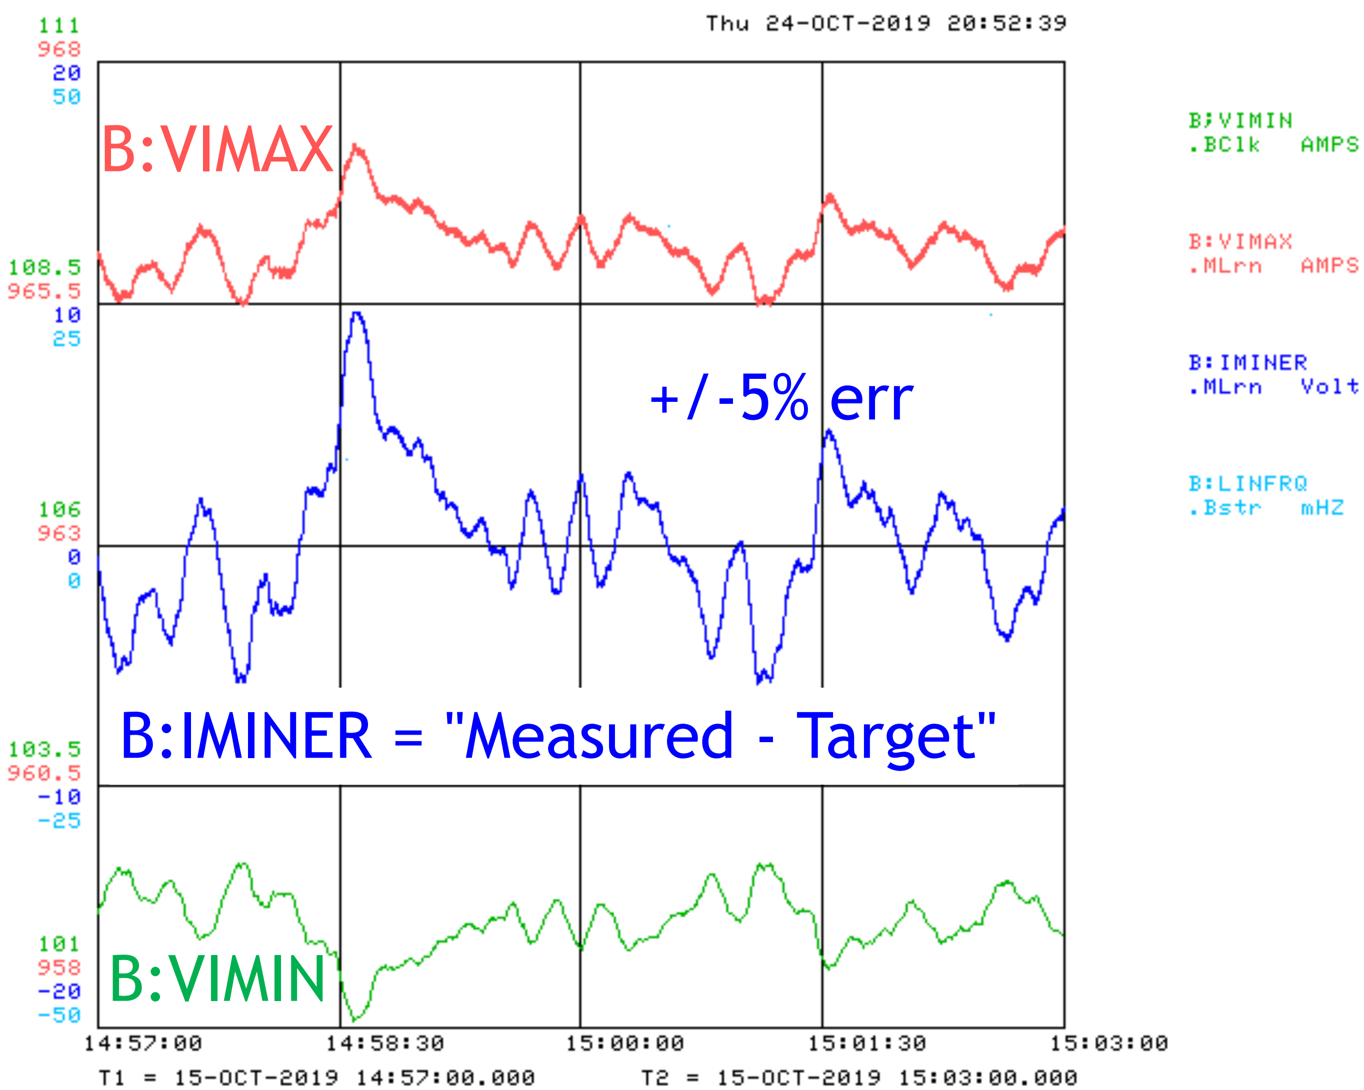

No control feedback

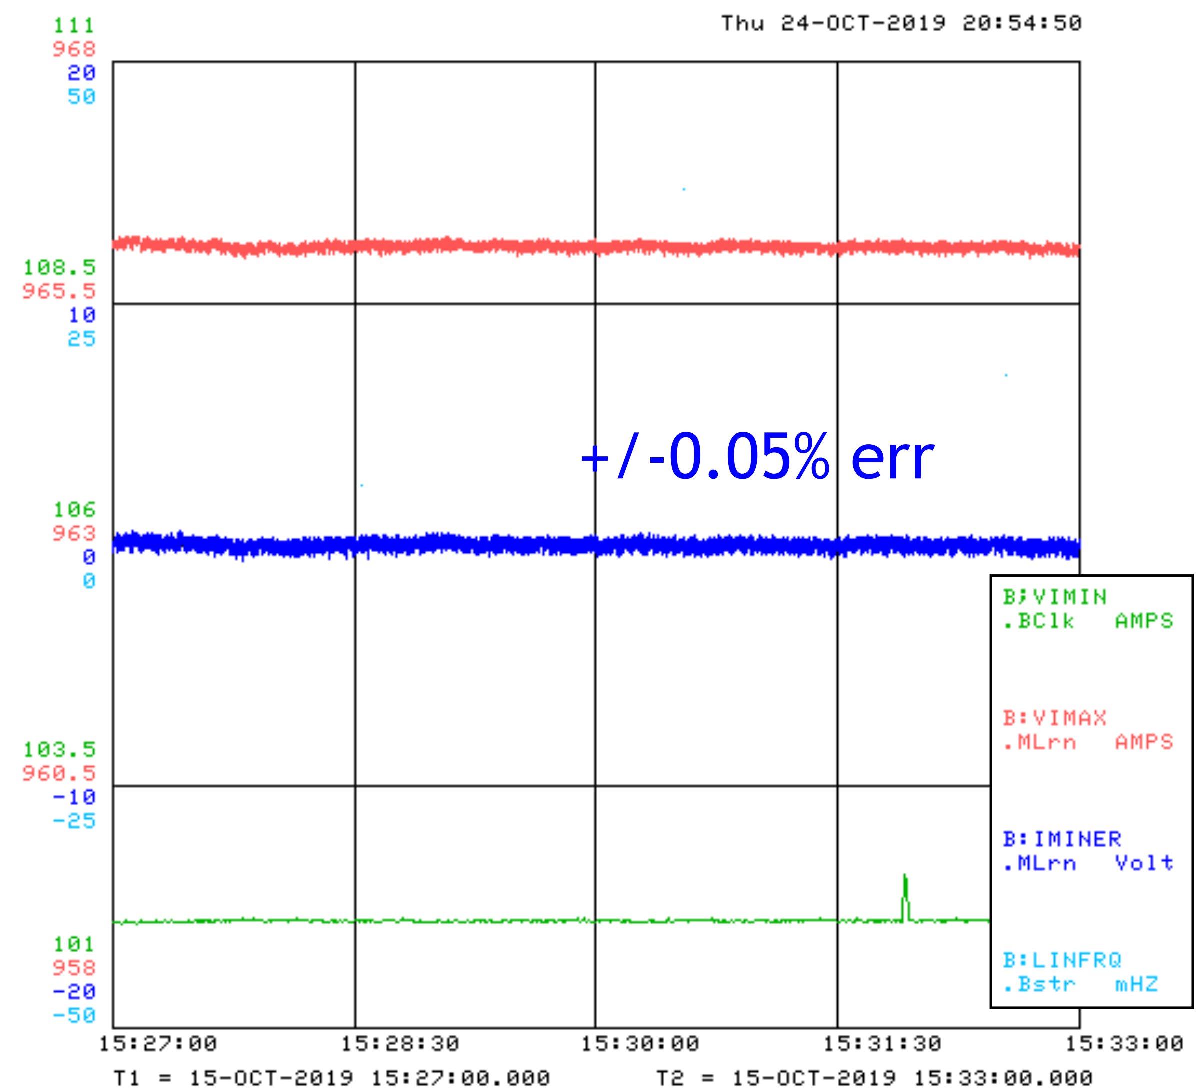

With PI controller

GMPS control schematic

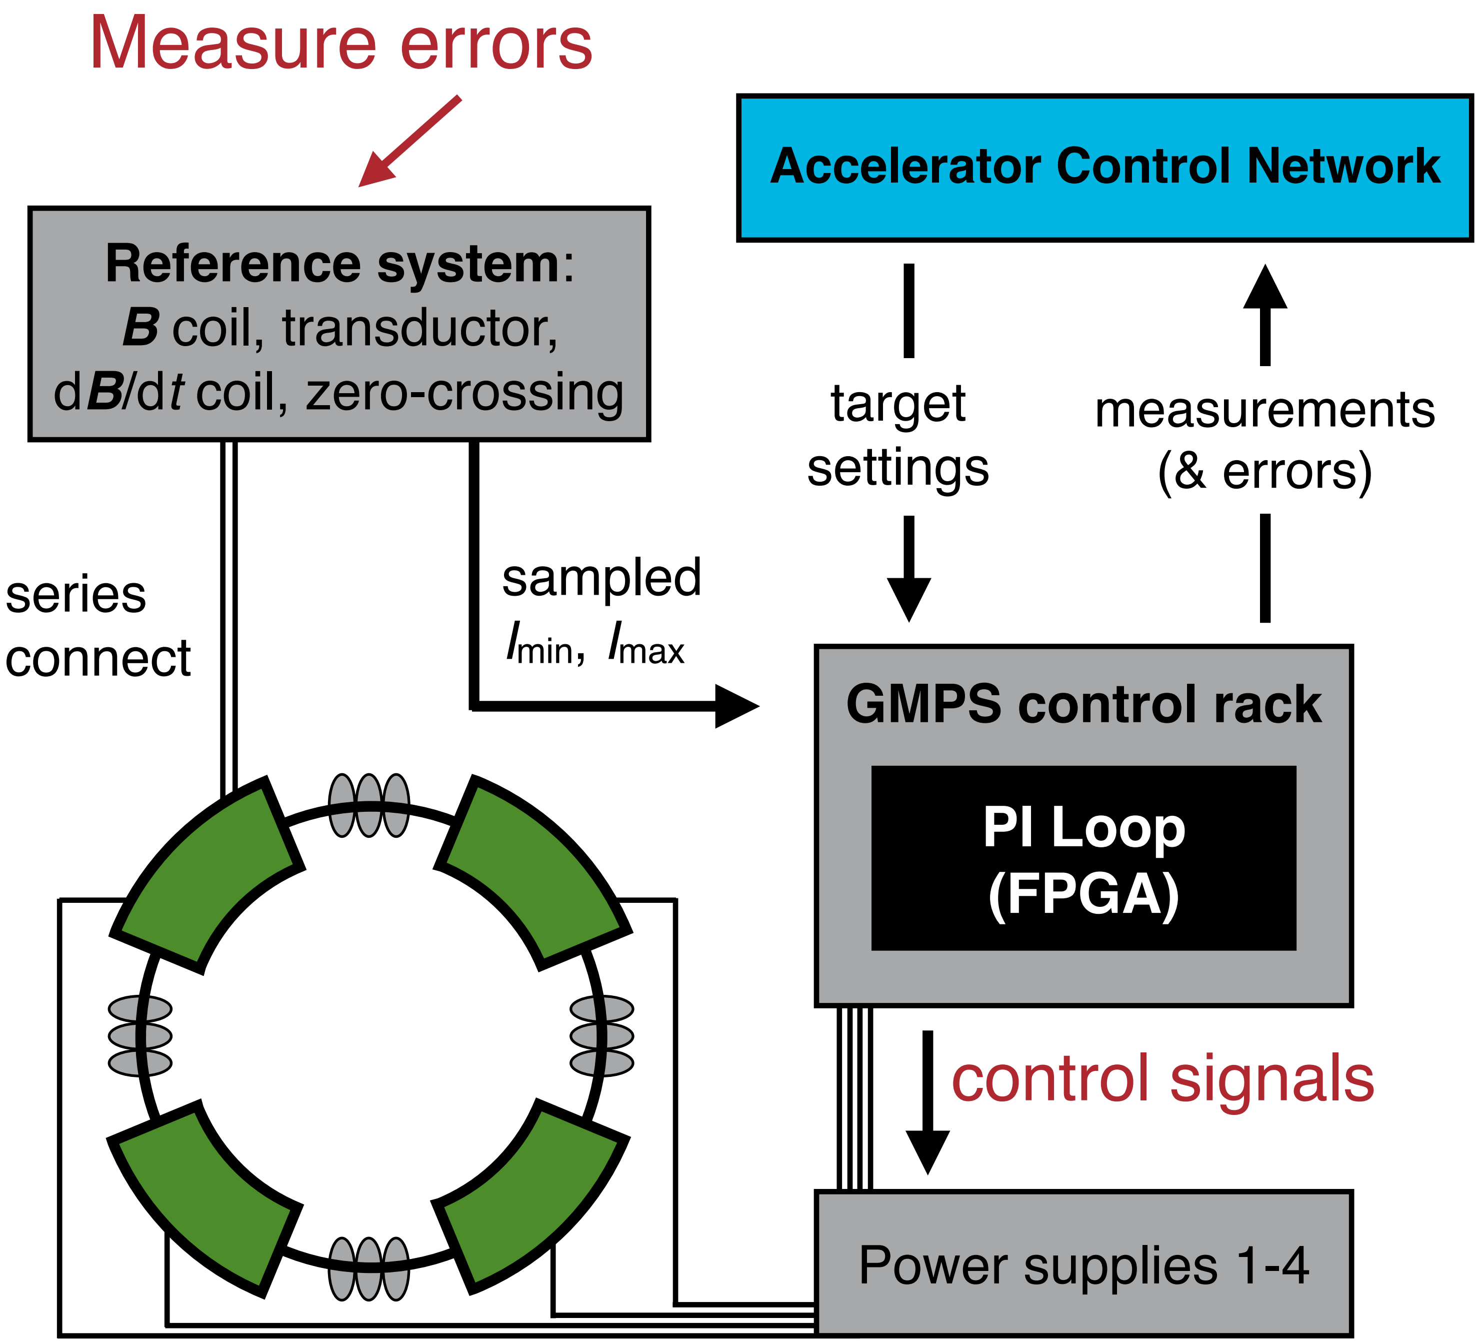

GMPS control schematic

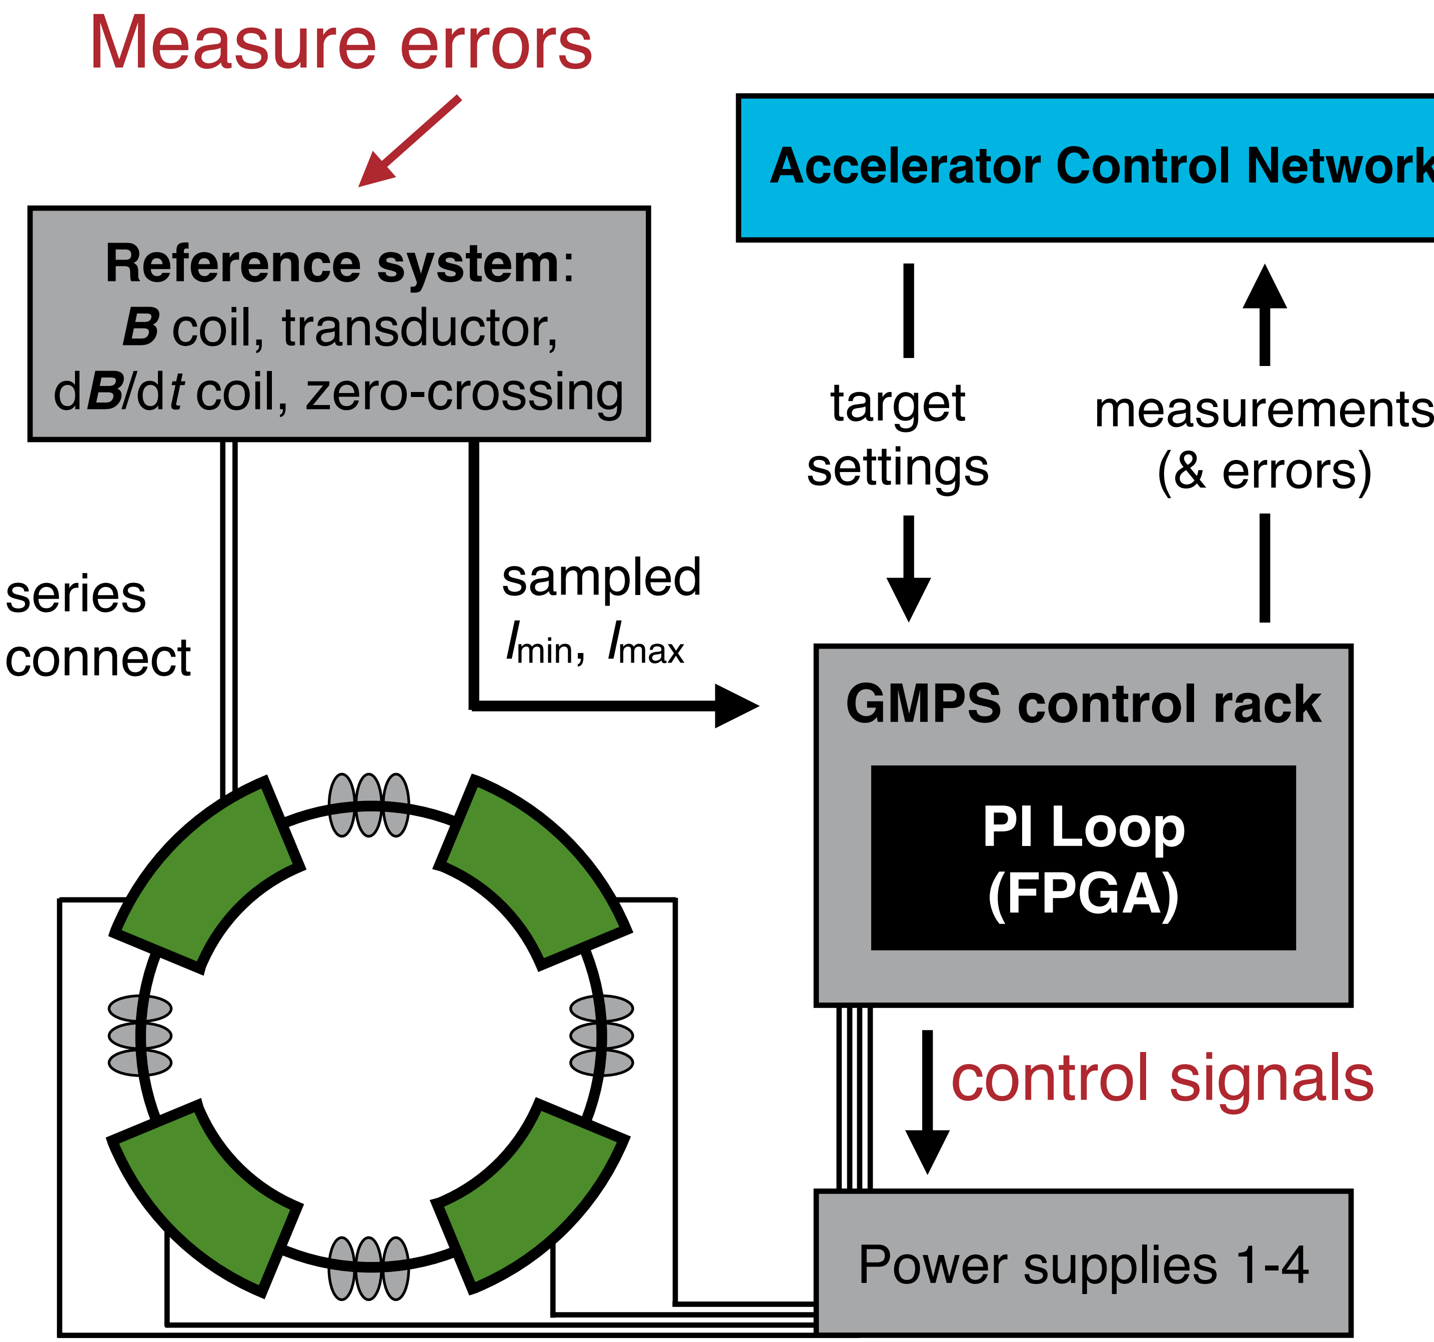

Areas for improvement:

PI loop considered I_{min} errors as the only form of feedback.

Control parameters must be selected, monitored by accelerator experts.

A Neural Network controller?

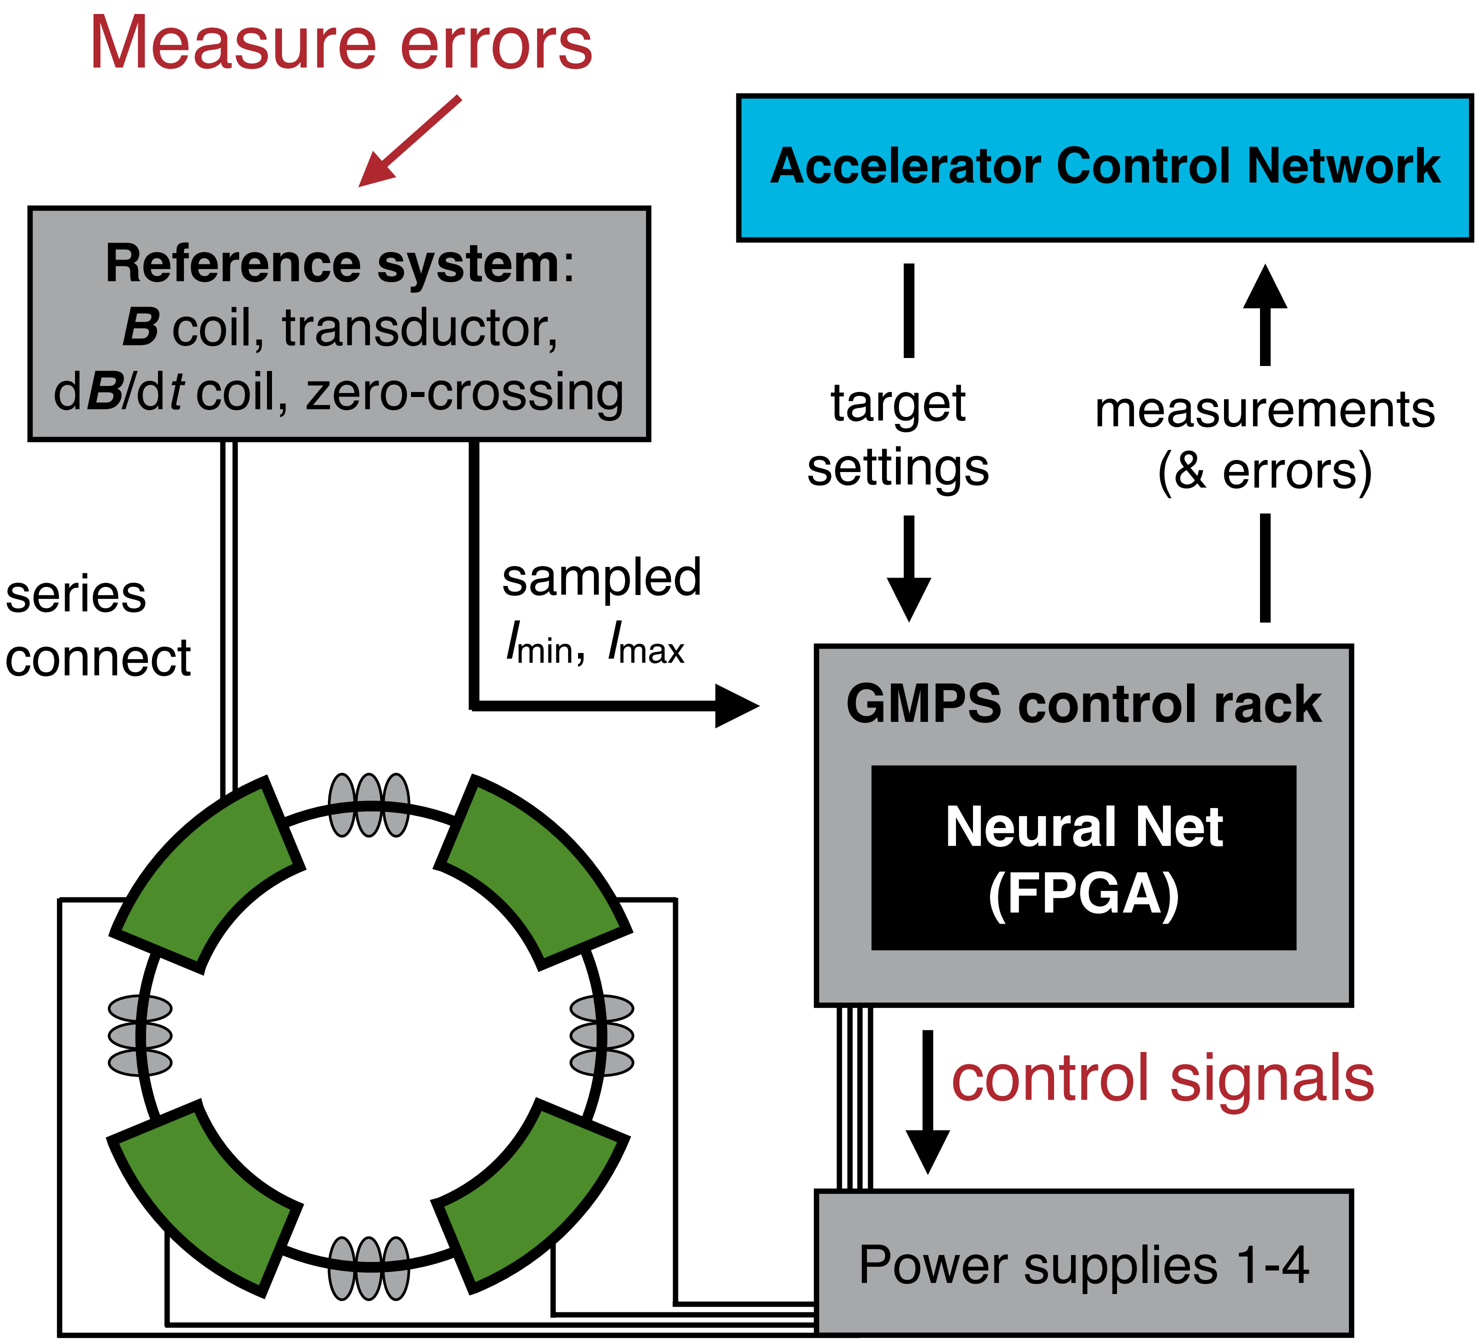

Profit from recent progress porting ML algos to FPGAs.

hls4ml See talk by J. Ngadiuba!

Can naturally incorporate many inputs.

Offers potential for "live" adjustments to the algorithm parameters while in operation.

A Neural Network controller?

Fundamental challenge of the approach: how to incorporate **realistic feedback** into the control model development process?

To begin, we cannot (should not?) test with the real Booster system

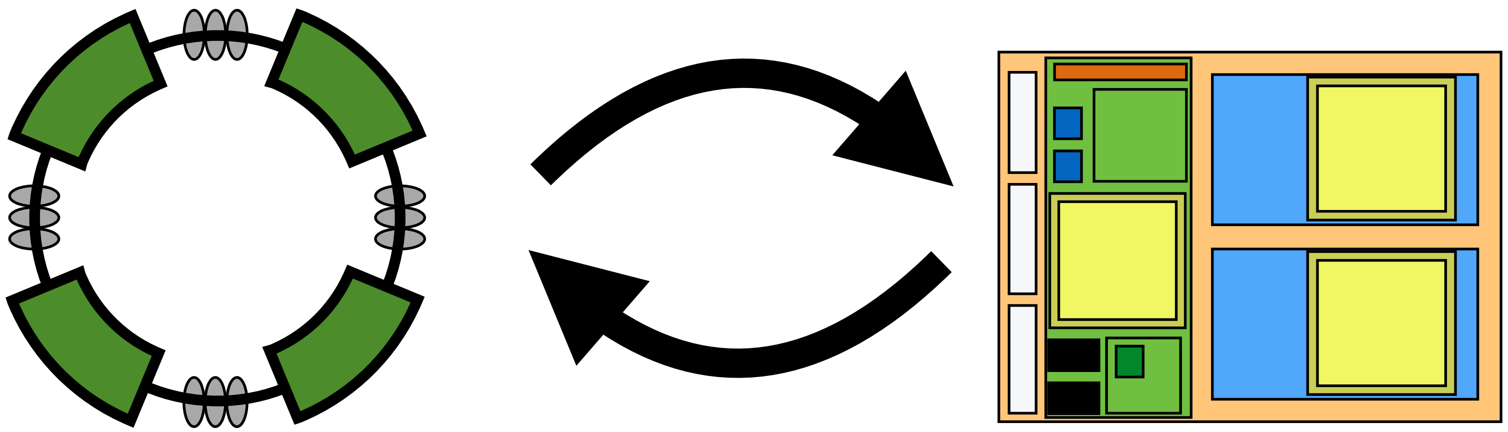

Booster system **Neural Net Control FPGA**

A Neural Network controller?

Fundamental challenge of the approach: how to incorporate **realistic feedback** into the control model development process?

To begin, we cannot (should not?) test with the real Booster system

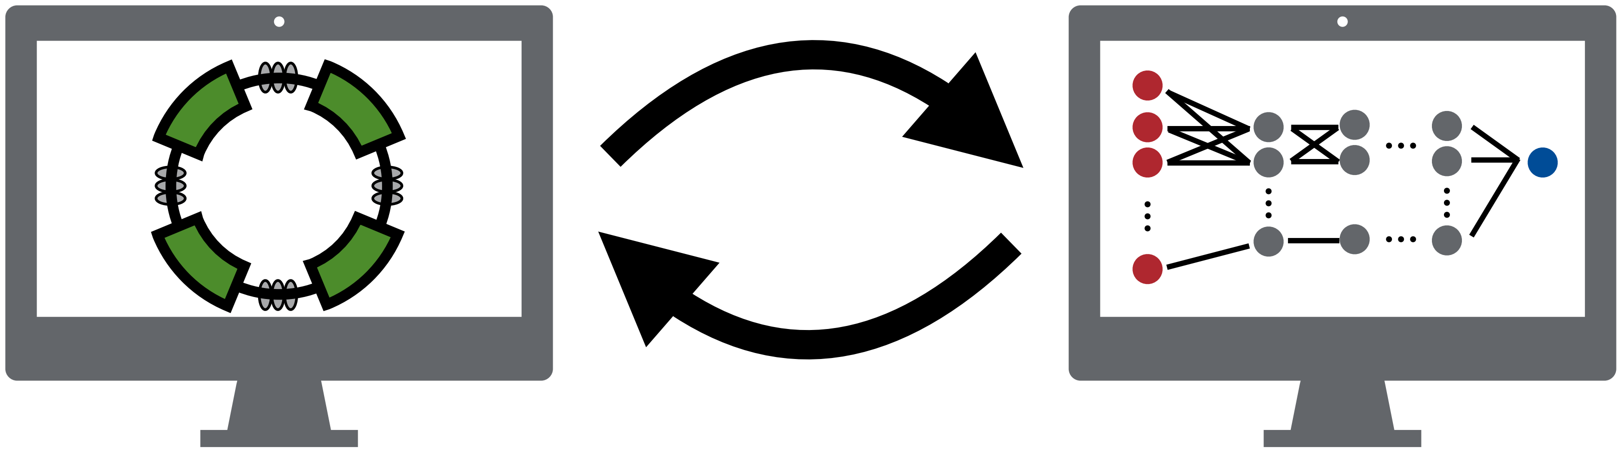

Booster Surrogate **NN Controller**

"Environment" "Agent"

Booster's digital twin

Fundamental challenge of the approach: how to incorporate **realistic feedback** into the control model development process?

To begin, we cannot (should not?) test with the real Booster system

Parameter	Details [Units]
B:IMINER	Setting-error discrepancy at injection [A]
B:LINFRQ	60 Hz line frequency deviation [mHz]
B:VIMIN	Compensated minimum GMPS current [A]
I:IB	MI lower bend current [A]
I:MDAT40	MDAT measured MI current [A]

Last 150 sampled values of predictive signals

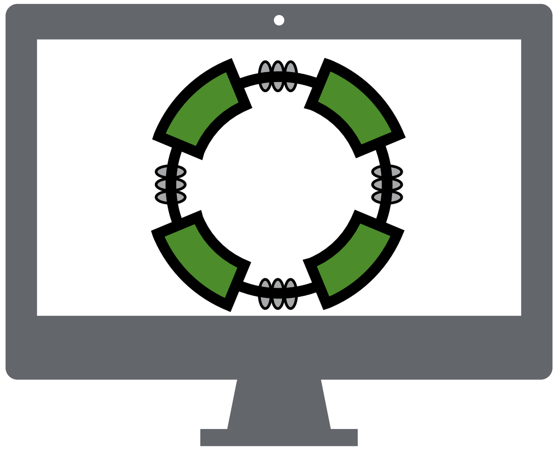

Booster Surrogate

B:IMINER
"Measured - Target"
B-field at minimum

Find that an LSTM recurrent NN can reproduce the historical Booster response quite well.

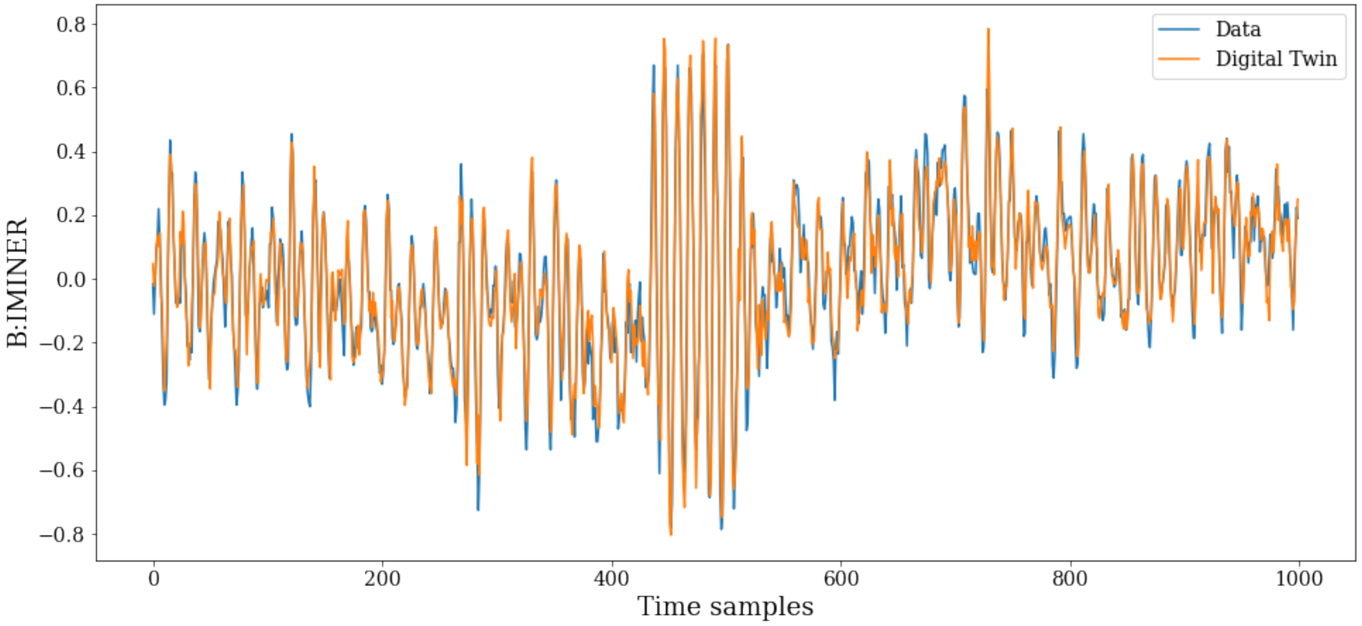

Control NN development

A simple Neural Network controller is ideal for a first demonstration

- Facilitates straightforward comparisons with the PI loop decision
- A small NN allows for maximum flexibility in our initial FPGA design

Inputs: current values for the five important signals

Parameter	Details [Units]
B:IMINER	Setting-error discrepancy at injection [A]
B:LINFRQ	60 Hz line frequency deviation [mHz]
B:VIMIN	Compensated minimum GMPS current [A]
I:IB	MI lower bend current [A]
I:MDAT40	MDAT measured MI current [A]

Control NN

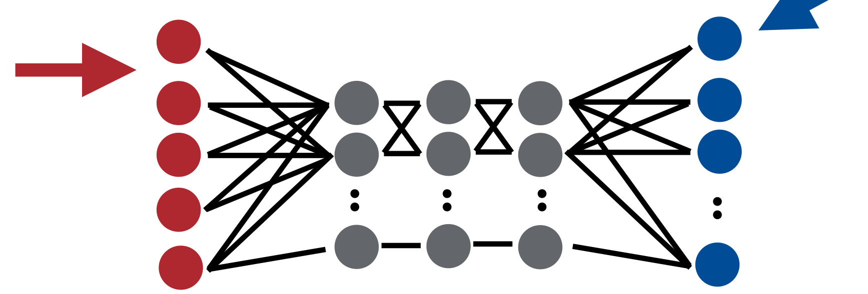

3 hidden layers of 56 nodes

Outputs: 7 actions

- No change
- ± 1 mV
- ± 0.5 mV
- ± 0.1 mV

Core of each NN "layer" is an N→M matrix multiplication

$$y_i = \sigma(w_{ij} x_i + b_i)$$

Non-linearity, e.g. $\sigma(x_i)=\max(x_i,0)$

Matrix multiplication

Prescribe that GMPS takes the action corresponding to the largest output node

Reinforcement learning

Controller interacts with Booster, accumulating rewards by minimizing errors

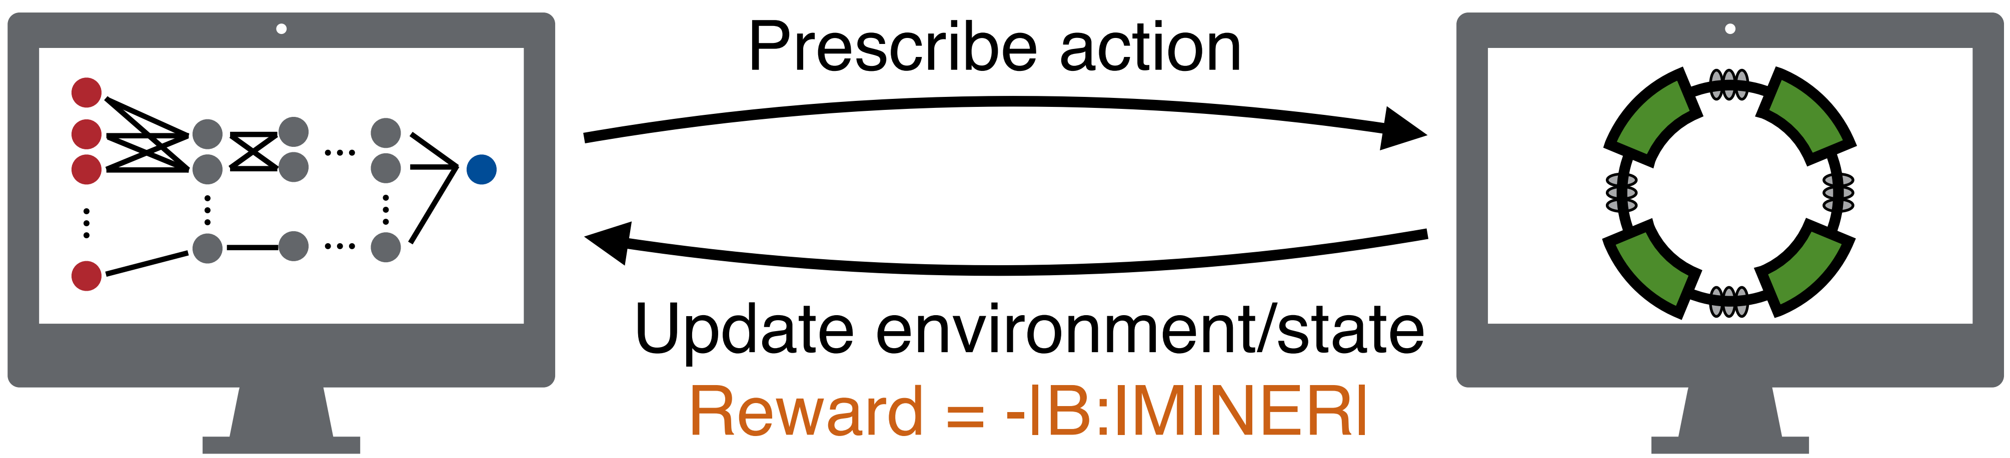

Reinforcement learning

Controller interacts with Booster, accumulating rewards by minimizing errors

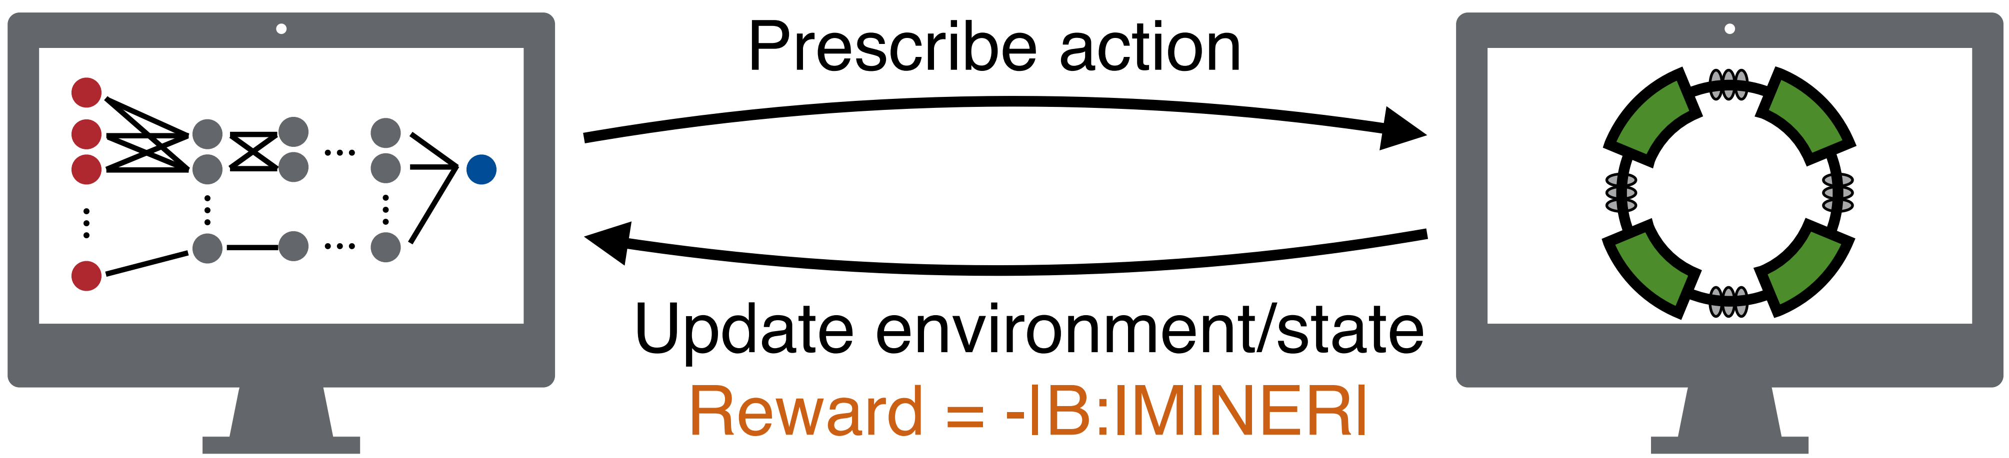

Rewards inform updates to the NN's 7k configurable weights, using the "Double Deep Q-Network" paradigm (1509.06461).

$$Q(S_t, A_t) = \sum_{t'=t}^{T} \mathbb{E}\left[\gamma^{t'-t} R(S_{t'}, A_{t'}) | S_t, A_t\right]$$

Q-value: expected sum of all rewards R, given a state S, action A, and discount factor γ

Feedback adjusts parameters so that:

$$Q(s_t,a_t) = R(s_t,a_t) + \gamma^* Q(s_{t+1},a_{t+1})$$

Estimated rewards: t'≥t

Actual reward @t

Estimated rewards: t'>t

Reinforcement learning

Average errors appear to be significantly reduced with DQN approach

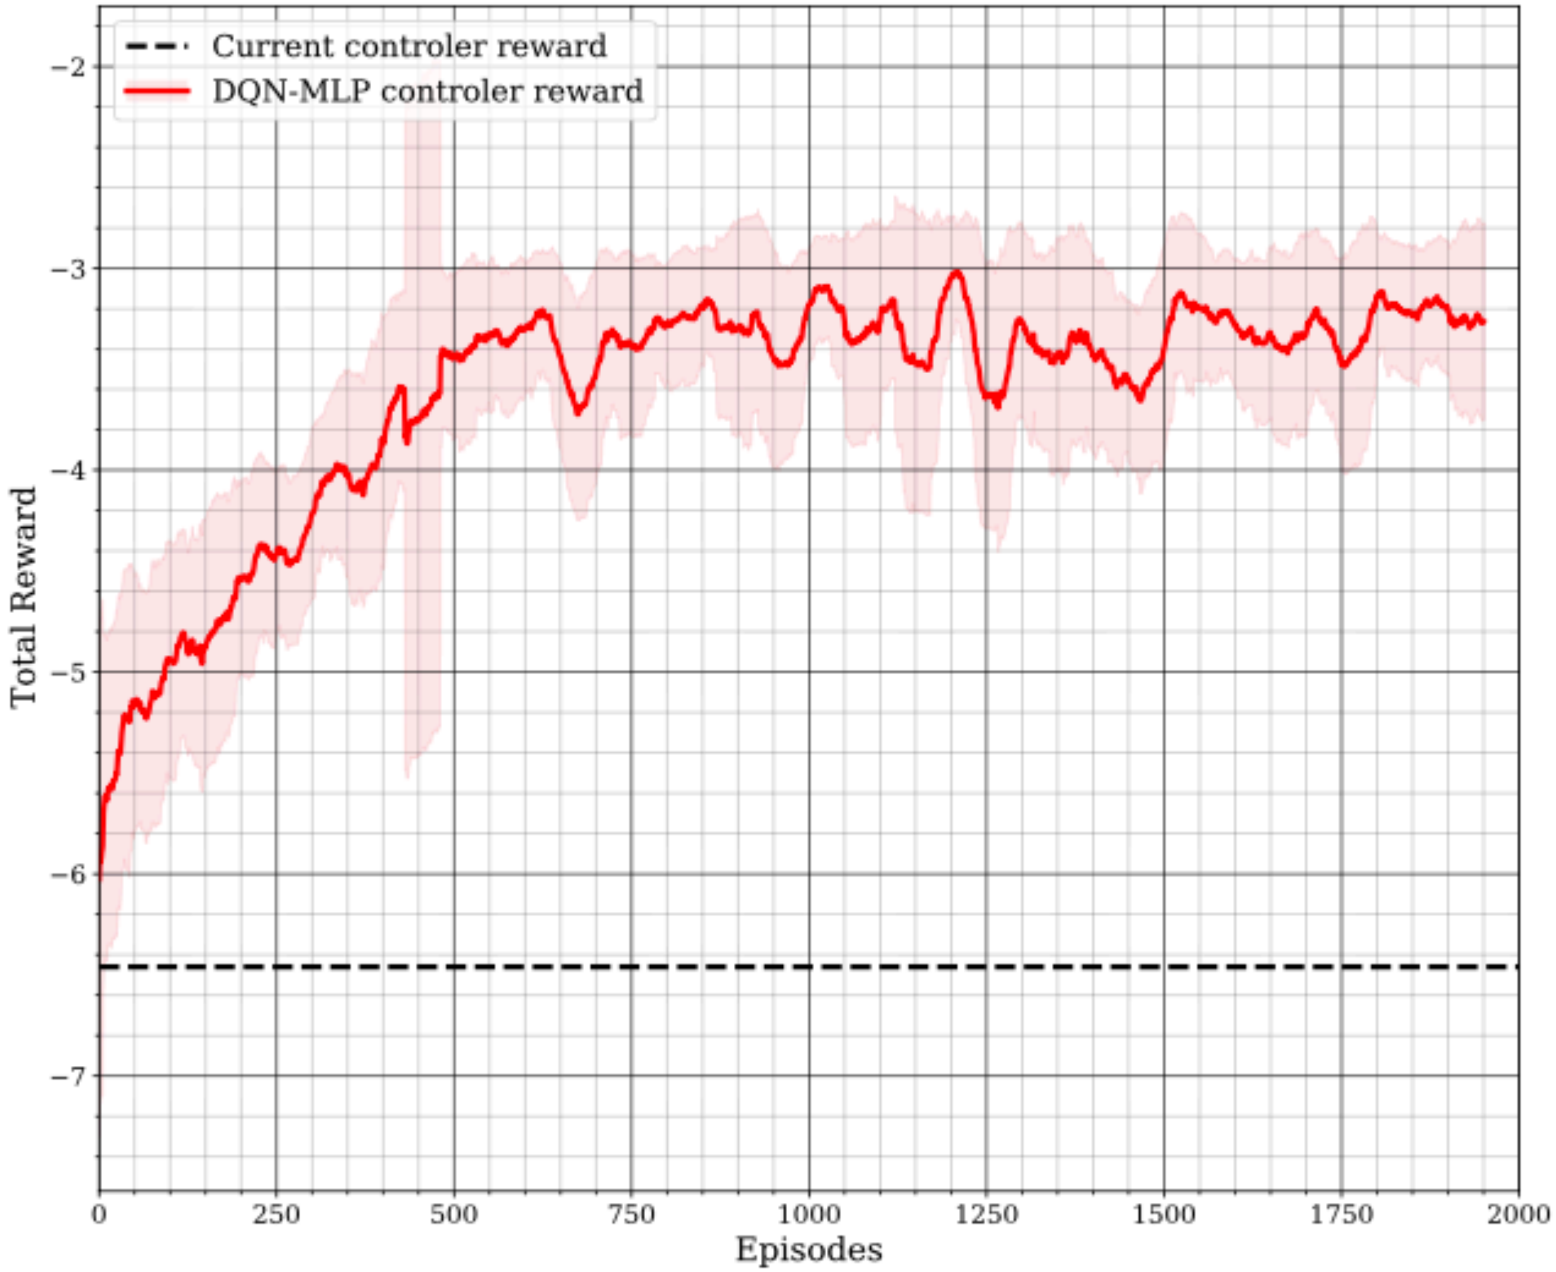

Reward accumulated by the DQN model from the Booster surrogate

Mean accumulated error observed in the historical data.

"Episodes" initialize surrogate with different historical data.

Implementing control NN on an FPGA

Benefit from significant past work in the Fast Machine Learning community

Some novel aspects for the Booster control application include:

Intel FPGA implementation:

Extended hls4ml to the Quartus HLS toolkit, establishing fine control over network implementation details for a range of resource constraints.

"Live" model updates in Booster operation:

Instead of fixing NN parameters, store in the embedded system's shared memory to push periodic improvements.

Incorporation of "guardrails":

Monitoring logic should cross-check NN controller decisions, to disable predictions outside a specified range.

Implementing control NN on an FPGA

Minimize design footprint by optimizing the precision of configurable parameters and NN calculations.

~10 bits sufficient to store all weights

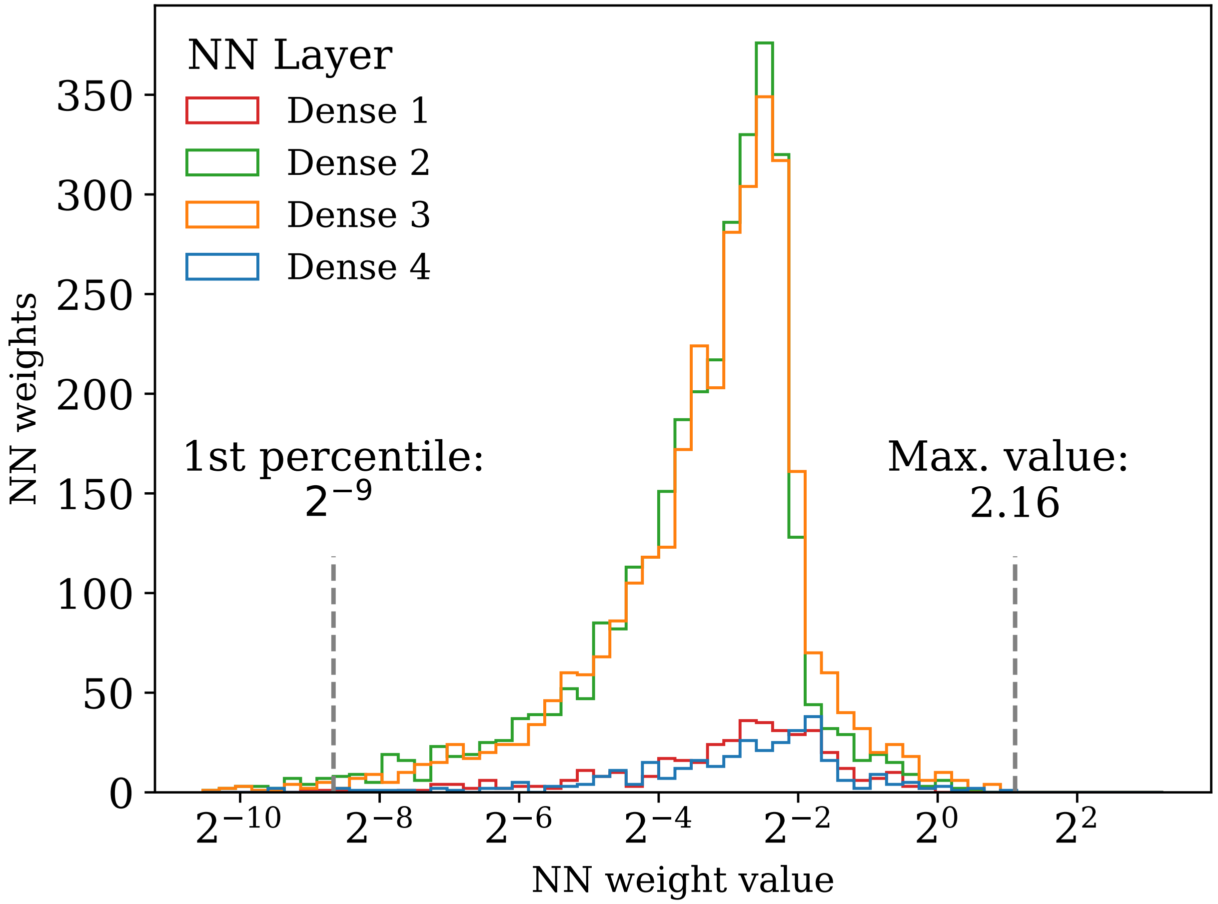

Agreement between floating-point and fixed-point model decisions

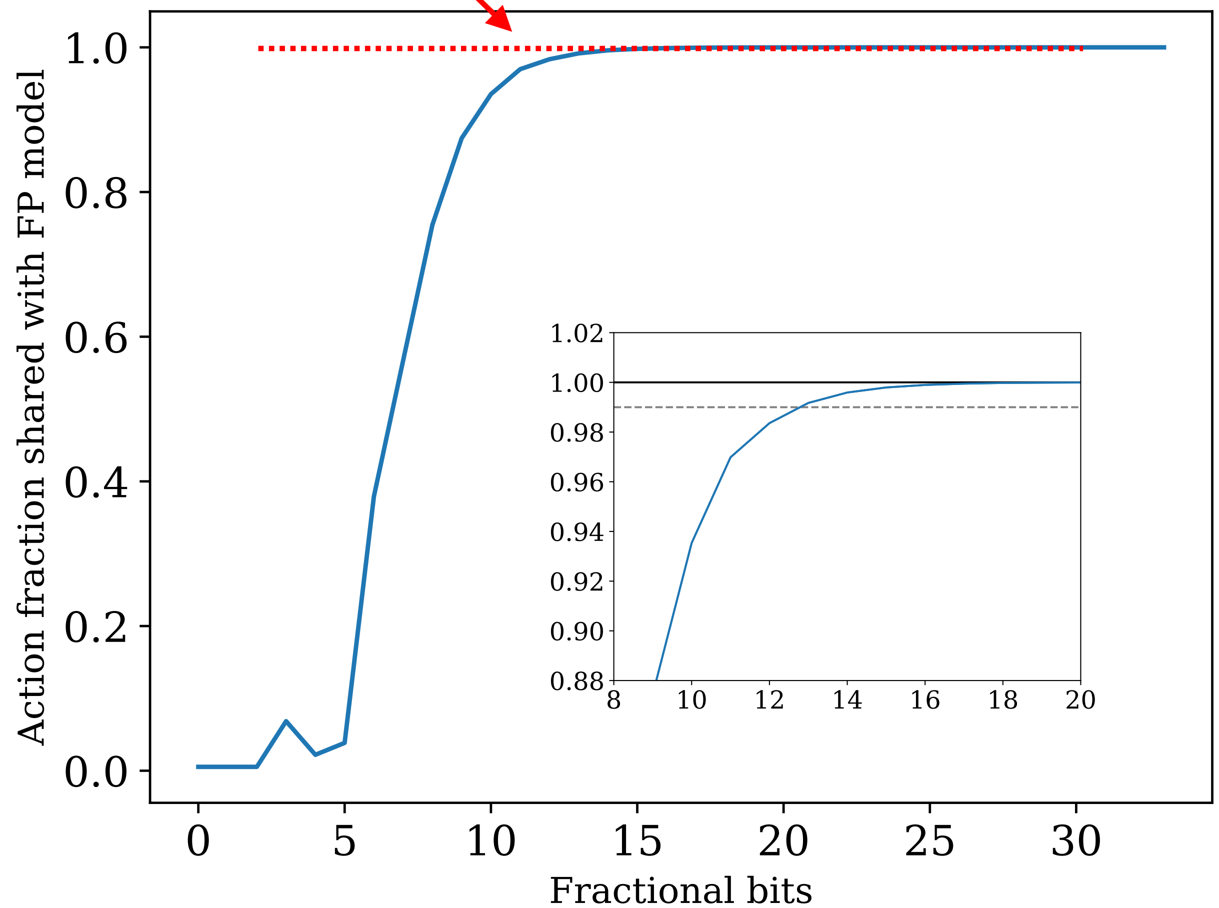

Fixed-point precision

Implementing control NN on an FPGA

Minimize design footprint by optimizing the precision of configurable parameters and NN calculations.

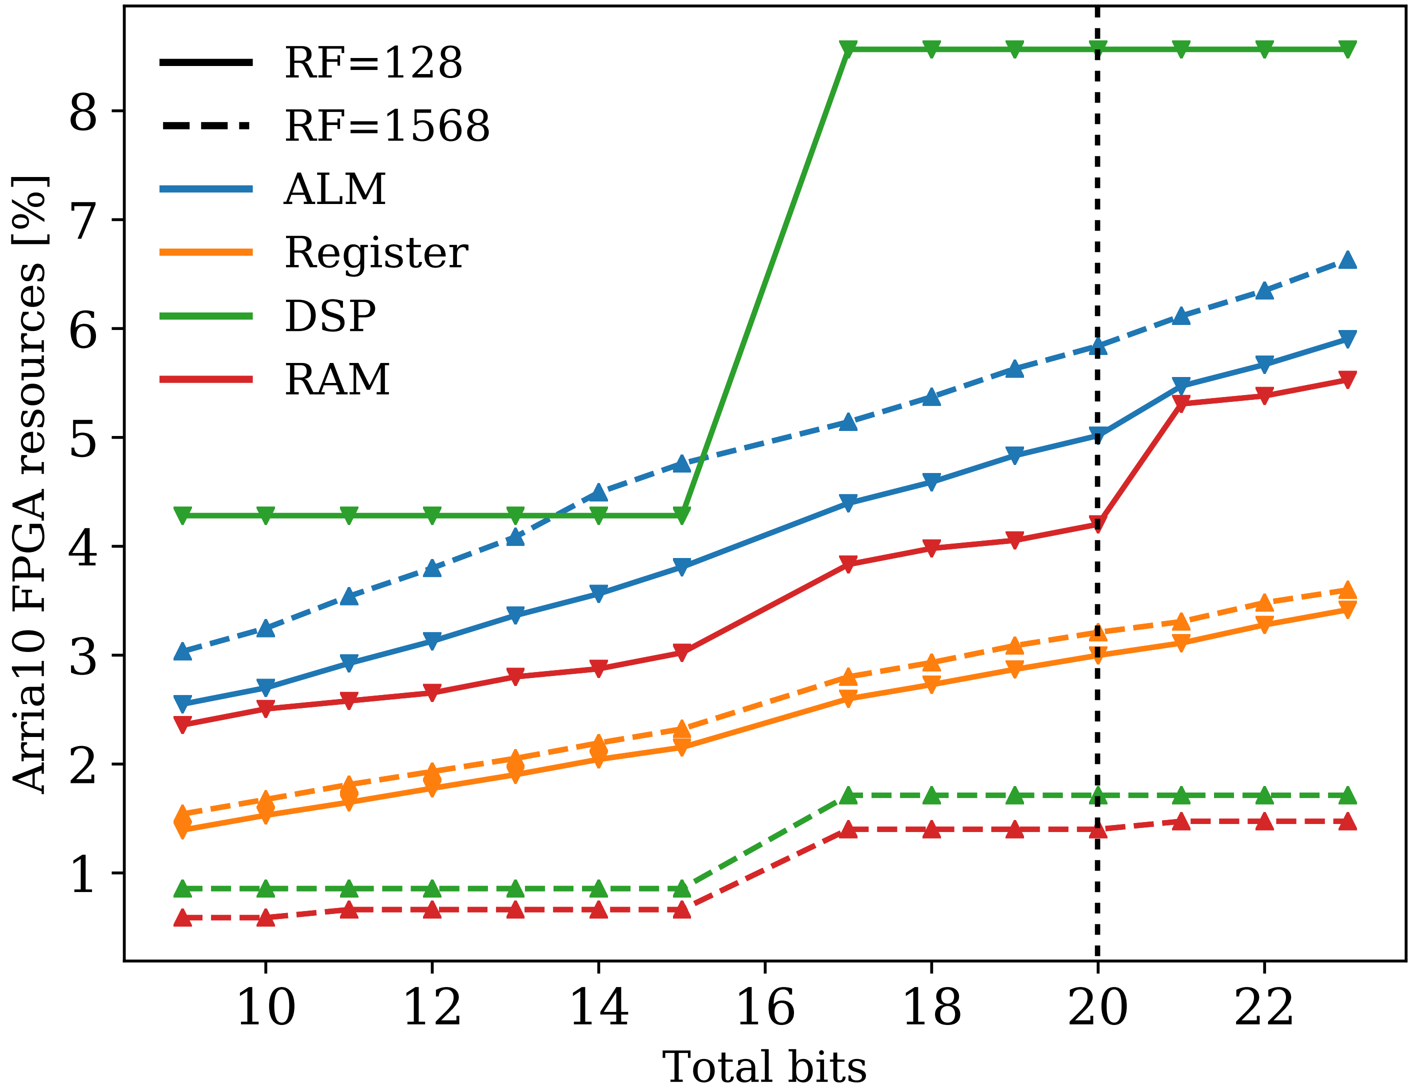

(For fixed-point operands)

Comfortably fit within 6% of the target Arria10 FPGA's resources.

Can trade serial / parallel designs to trade resources for latency.

reuse factor	DSP	BRAM	MLAB	ALM	Register	Latency
128	130	114	229	21.4 k	51.2 k	2.8 μs
224	74	100	1420	40.2 k	78.3 k	4.1 μs
1568	26	38	357	24.9 k	54.9 k	17.2 μs
Available	1518	2713	···	427 k	1.7 M	···

Looking ahead

- Simulation studies indicate that GMPS performance may be improved by a significant factor.
- Aim to deploy the new control board this spring, after Covid delay.
  - Can immediately test NN controller, running as a spy
  - Accumulate improved dataset with all signals measured *in situ*.
- In parallel, investigating new control model ideas: architectures (Larger MLPs and RNNs) and schemes (ensembles with decision by majority)

Looking ahead

- Simulation studies indicate that GMPS performance may be improved by a significant factor.
- Aim to deploy the new control board this spring, after Covid delay.
  - Can immediately test NN controller, running as a spy
  - Accumulate improved dataset with all signals measured *in situ*.
- In parallel, investigating new control model ideas: architectures (Larger MLPs and RNNs) and schemes (ensembles with decision by majority)

Looking forward to installation — thank you to all collaborators!

Pacific Northwest NATIONAL LABORATORY

Lucy Huang, Malachi Schram

UC San Diego

Javier Duarte

Columbia University

Rachel Keller

Fermilab

CH, Jason St. John, Diana Kafkes, Bill Pelico, Gabe Perdue, Andres Quintero-Parra, Brian Schupbach, Kiyomi Seiya, Nhan Tran