Fermilab

An ML control system for the Fermilab Booster

Christian Herwig, for the Accelerator AI Team

Fast ML for Science Workshop

December 1, 2020

FNAL Accelerator Complex

- Booster synchrotron: 400 MeV H^- from Linac accelerated to 8 GeV protons for delivery to Main Injector, experiments
  - Batches delivered to MI/Recycler @15 hz ('rapid cycling')
- Efficient operation critical for PIP-II goal of MW beam

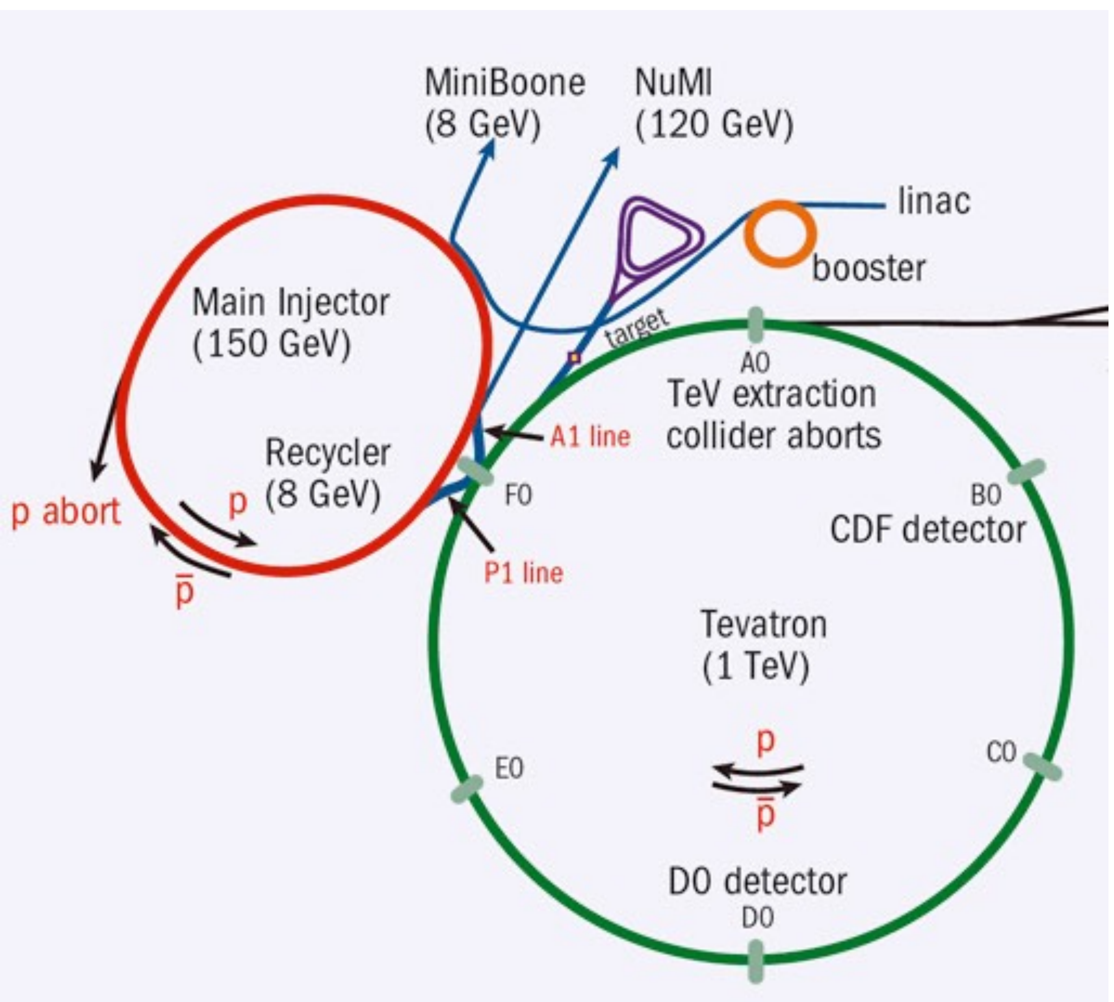

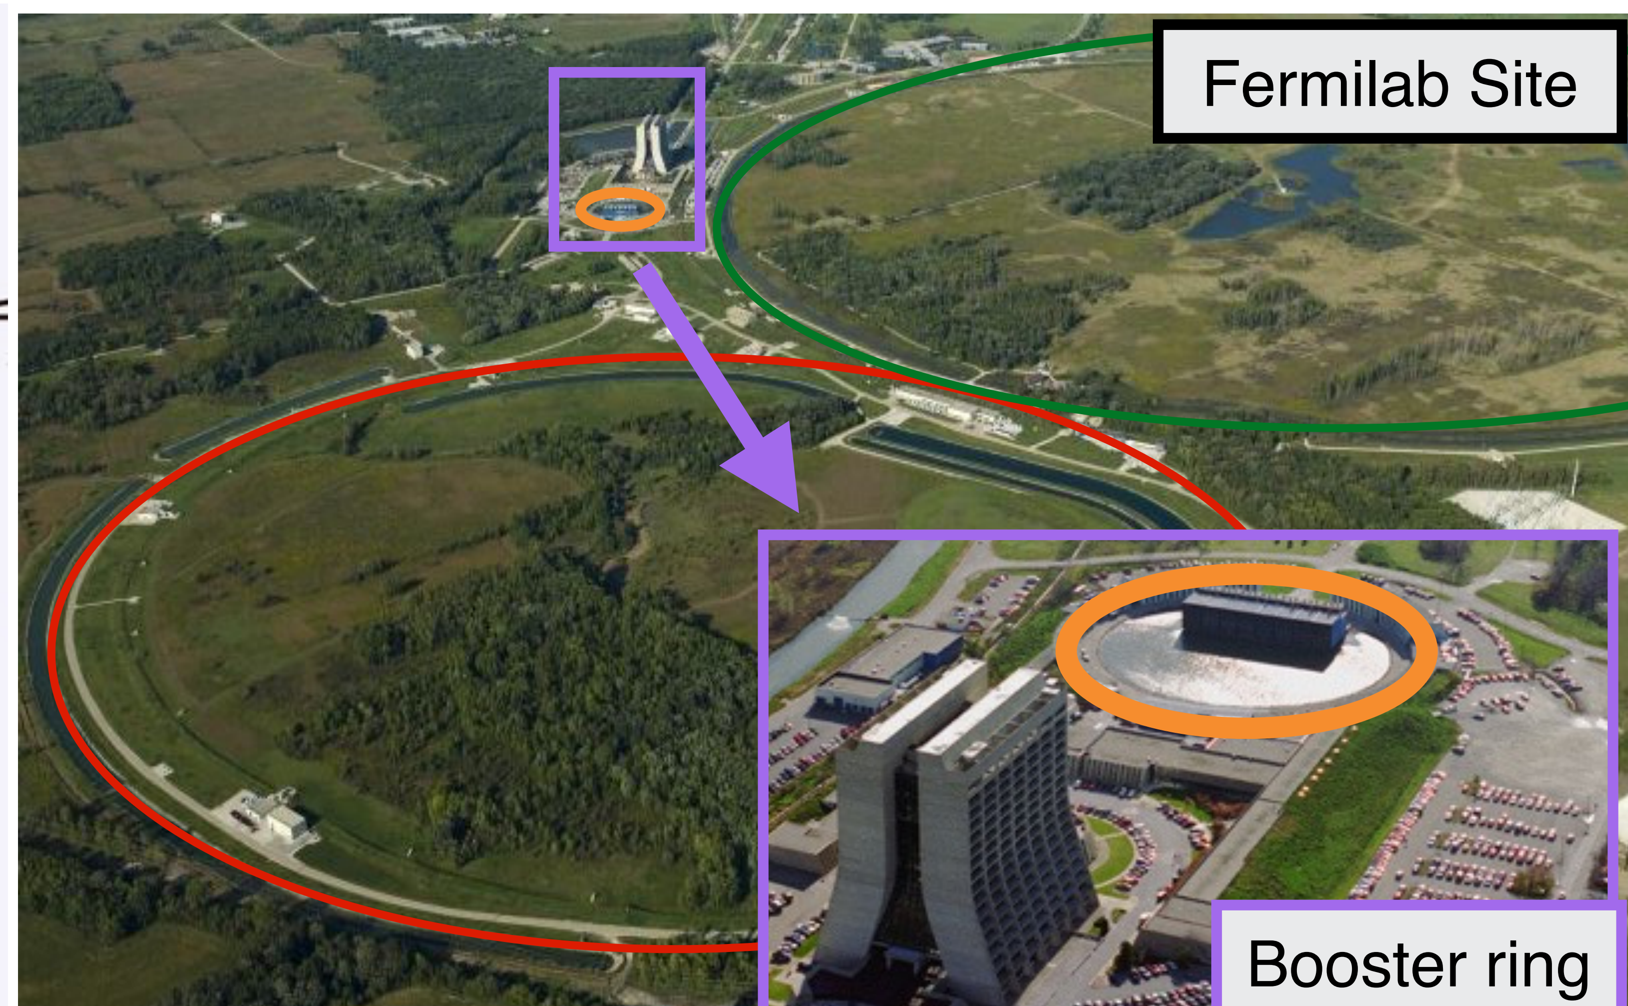

A single Booster cycle

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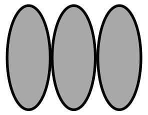 and bending magnets 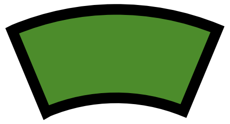
- Bending magnet current ramps in 15hz cycles to maintain the orbit of the accelerating beam

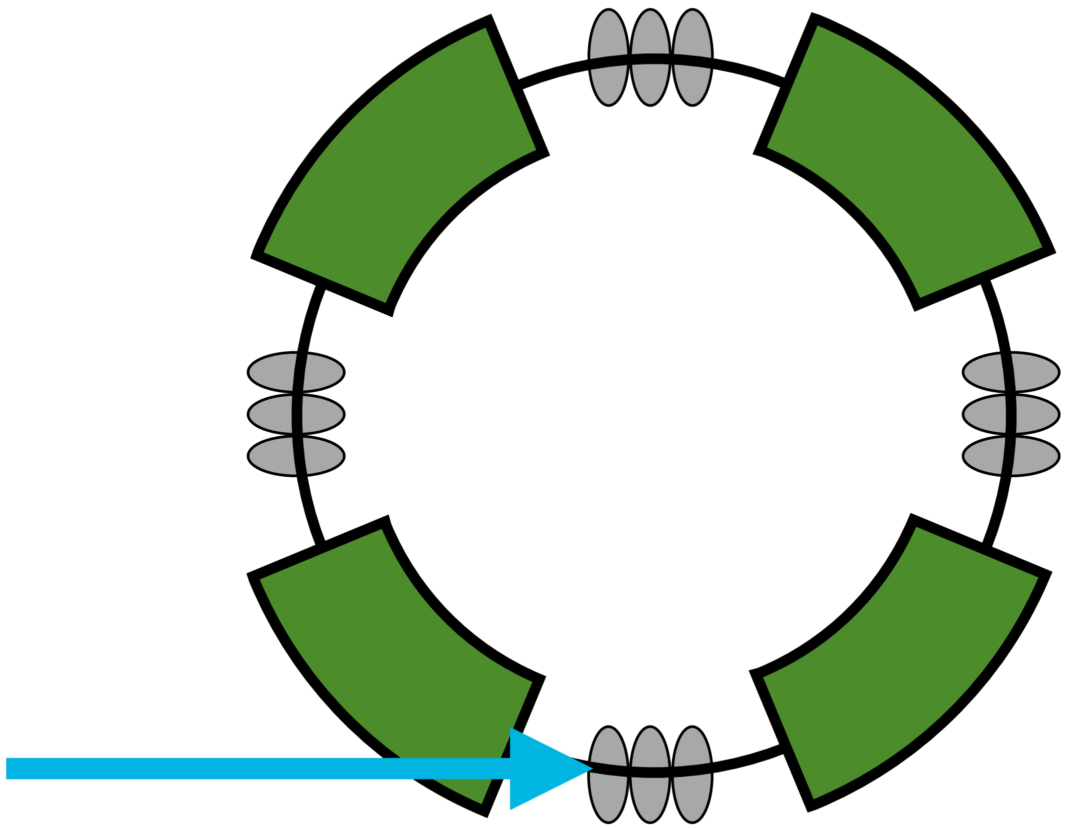

400 MeV ions
from Linac

Gradient Magnet Power Supply

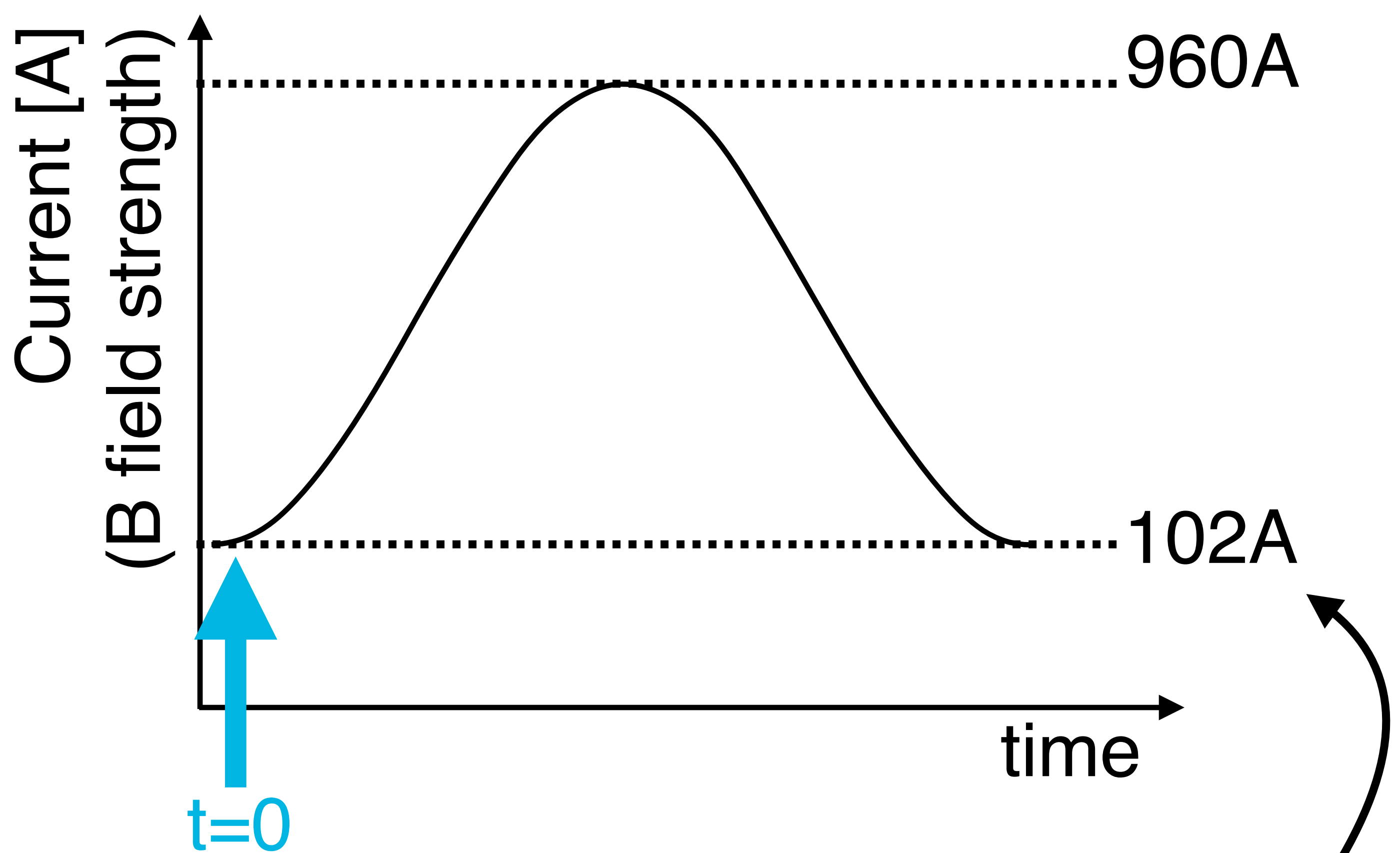

Minimum current to maintain orbit

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 beam

Credit: J. St. John

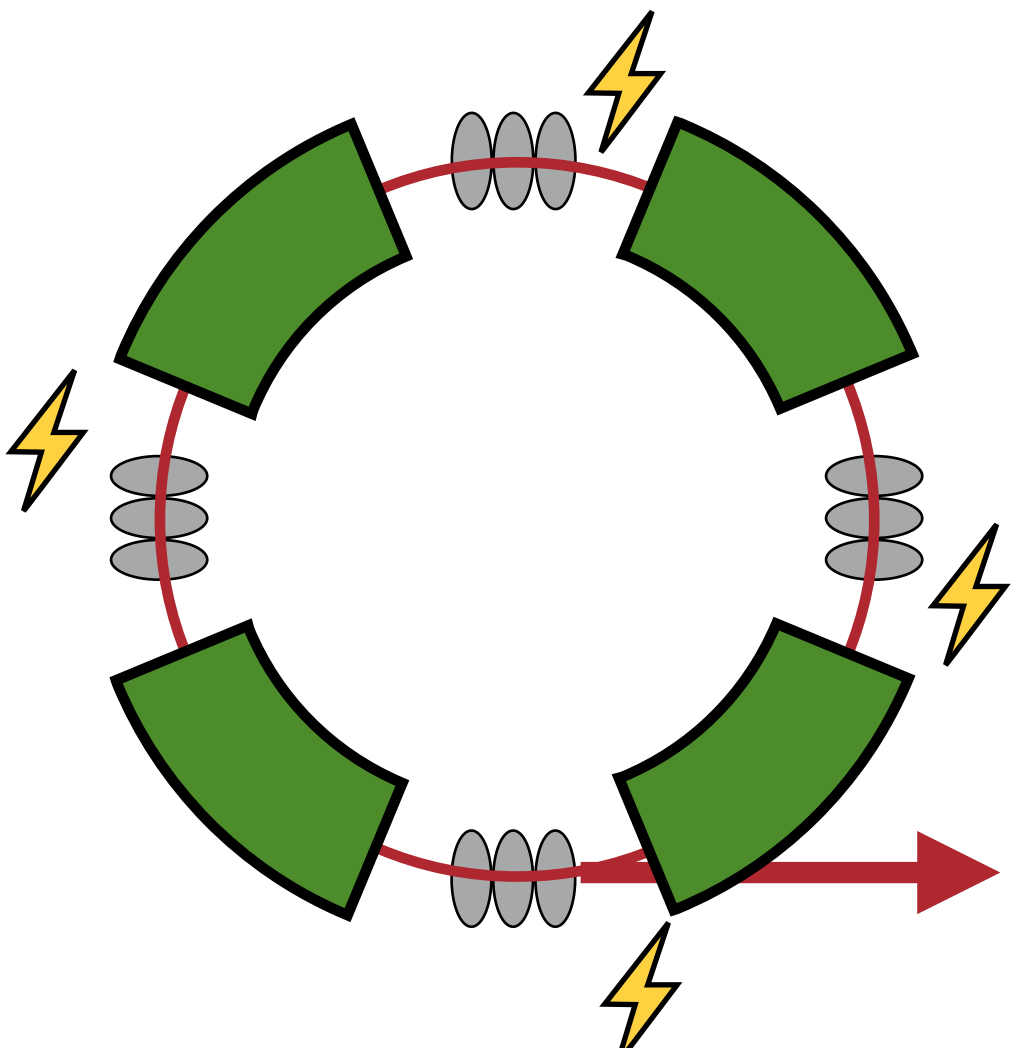

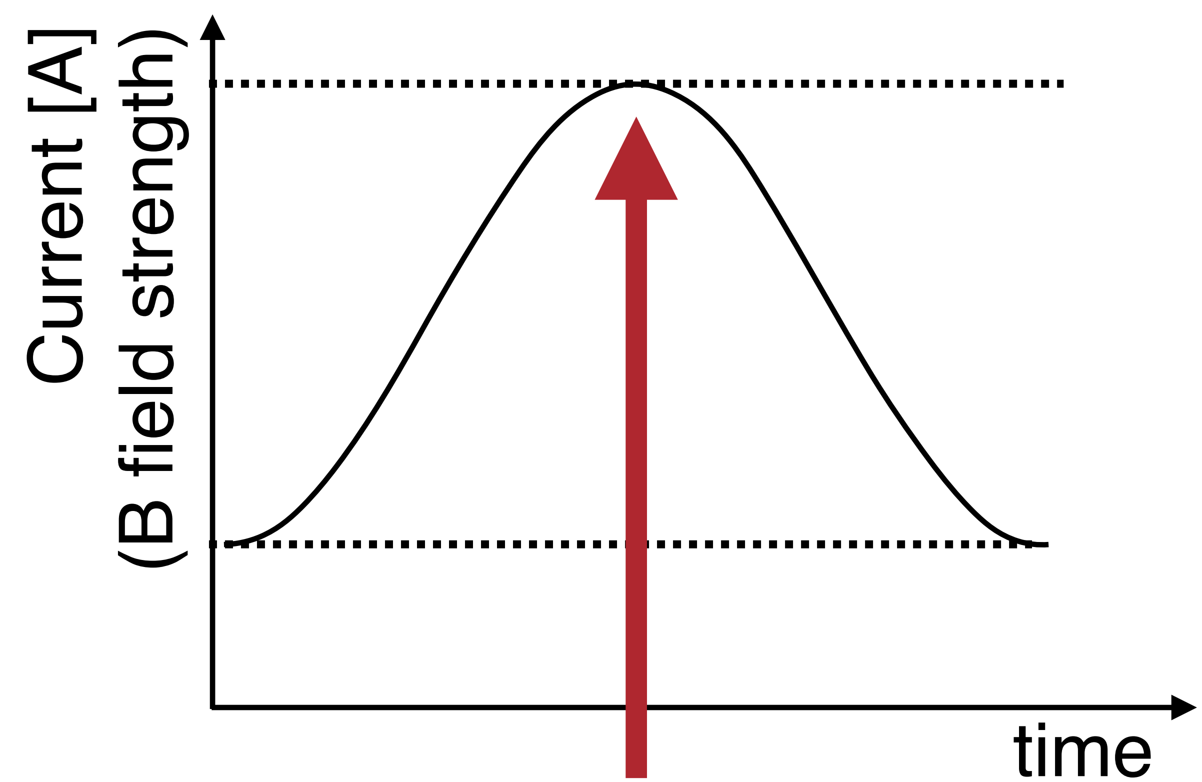

8 GeV beam extracted at maximum B-field

A single Booster cycle

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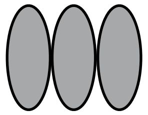 and bending magnets 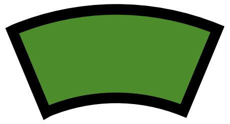
- Bending magnet current ramps in 15hz cycles to maintain the orbit of the accelerating beam

Credit: J. St. John

Spread in current*
(B field) degrades beam quality, contributing to losses!

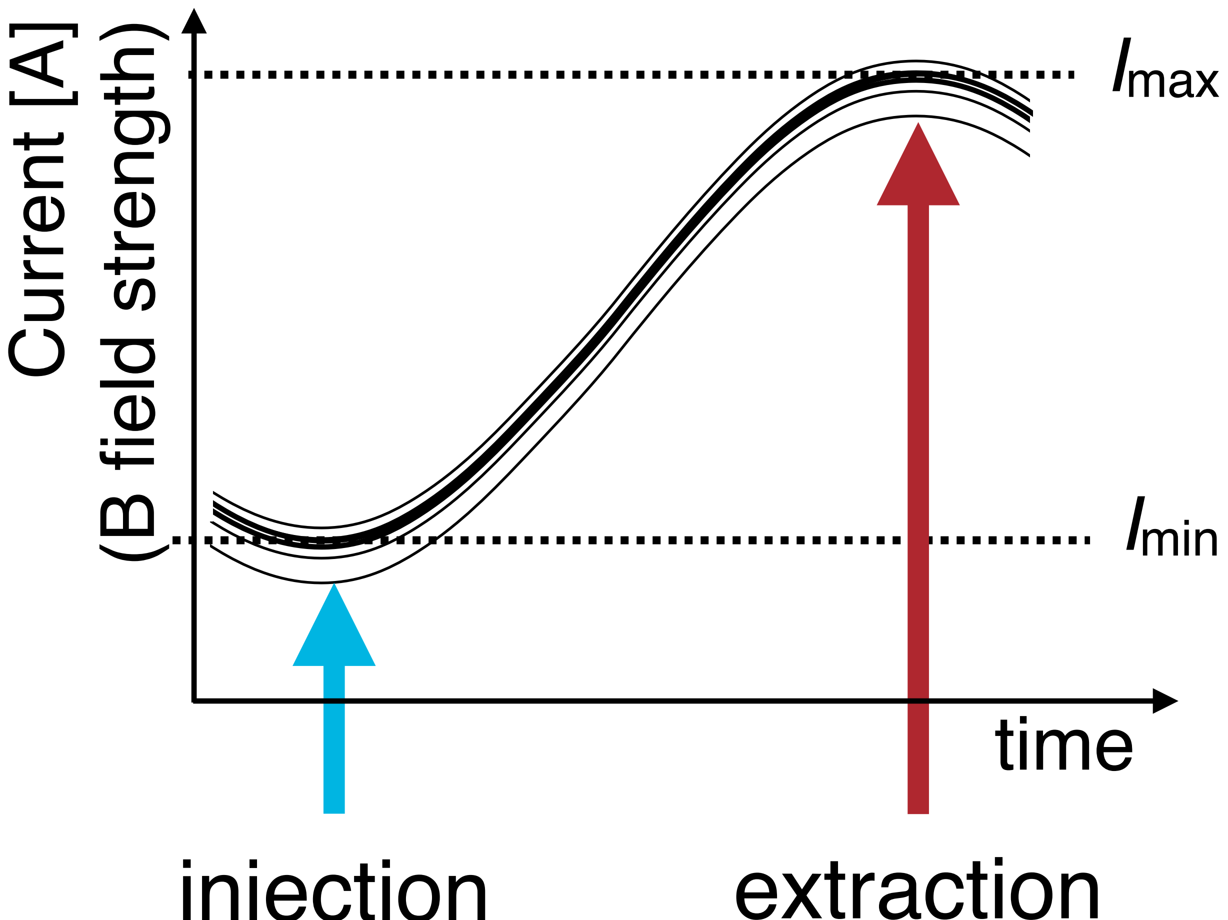

*Due to, e.g. the many nearby high-current, high-power loads

Booster control environment

- At present, currents are adjusted w/ PID loops comparing using measurements from a reference magnet

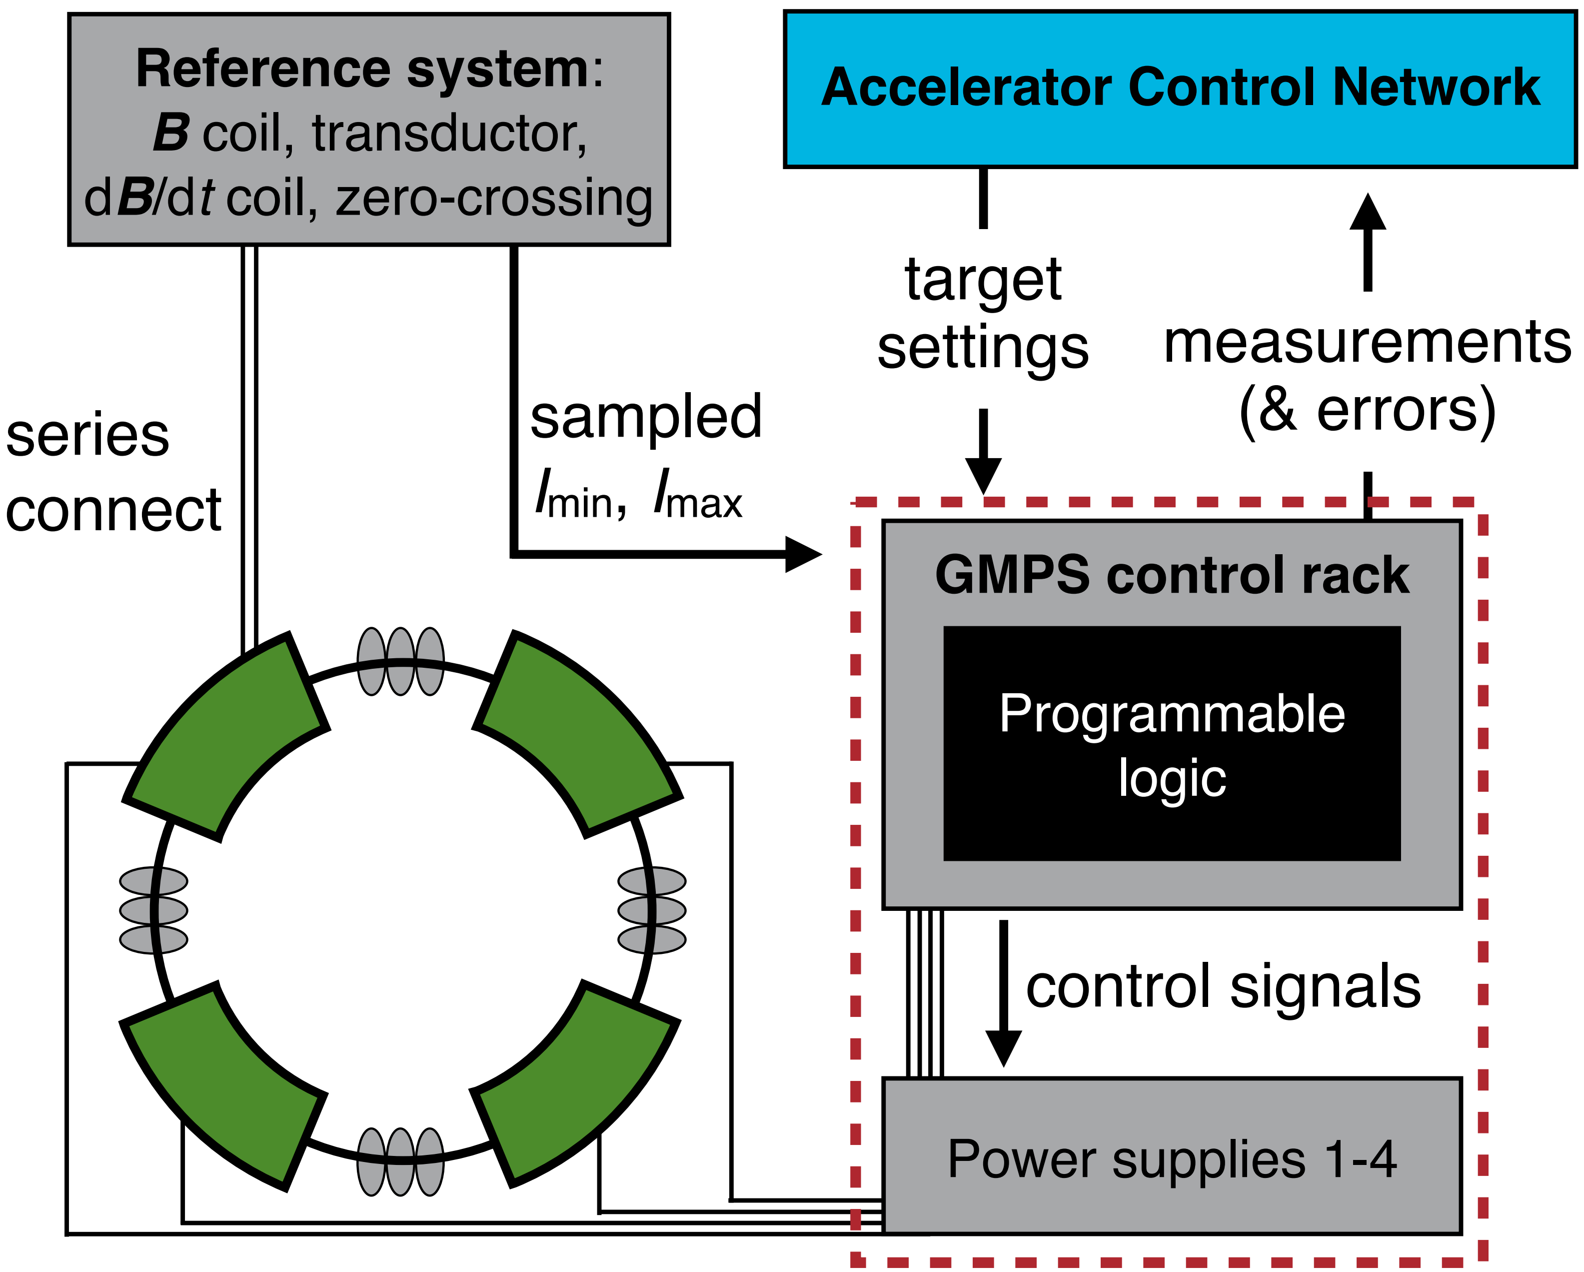

Aim is to replace PID w/ ML controller

Allow for online training and reconfiguration

Will use Reinforcement Learning paradigm

Include more features (e.g. gallery temp, line voltage, MI program)

Reinforcement Learning Setup

- Agent (the ML control algo) prescribes I_{GMPS} adjustments; interacts with environment, accumulates performance-based rewards
  - Environment: FNAL Booster or synthetic "Surrogate Model"

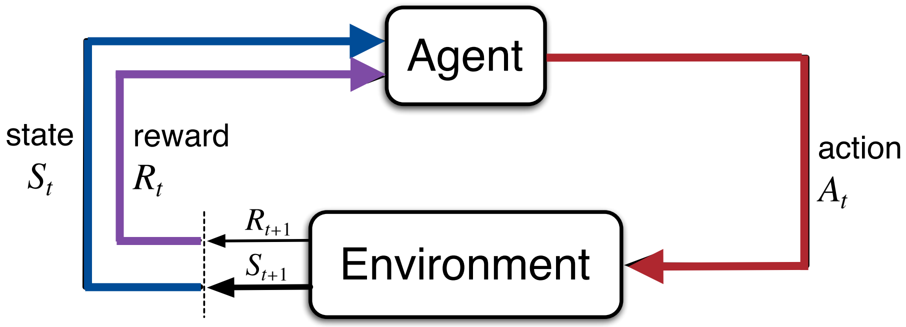

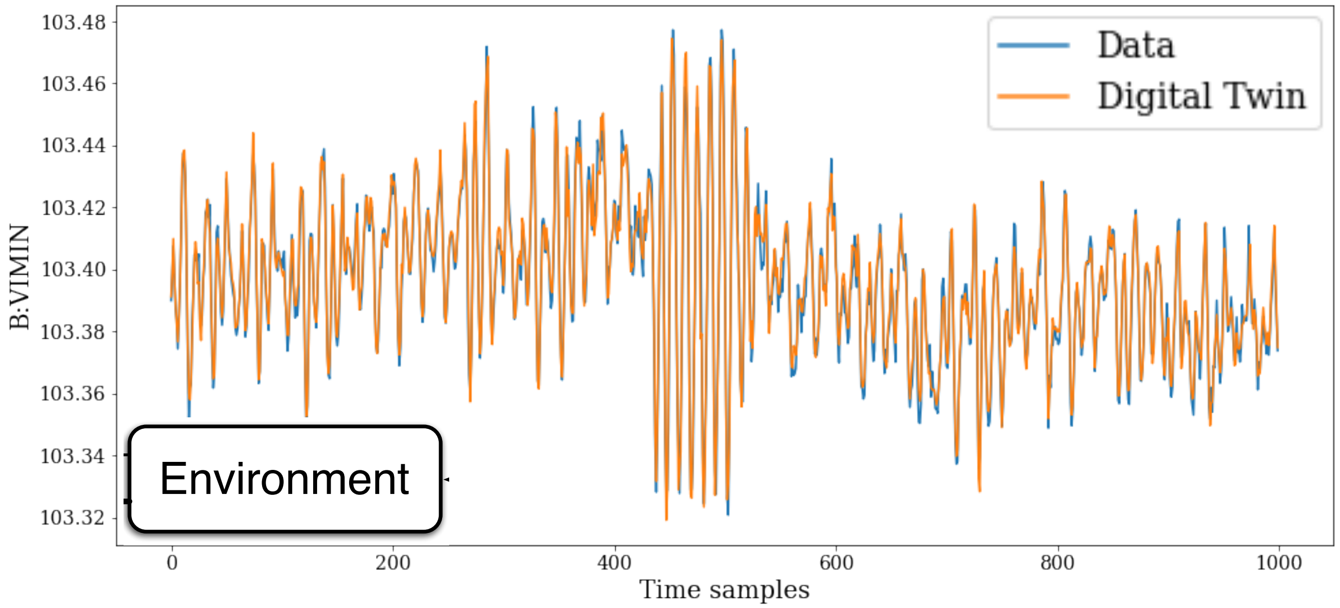

The **surrogate model** replicates I_{GMPS} well from past data

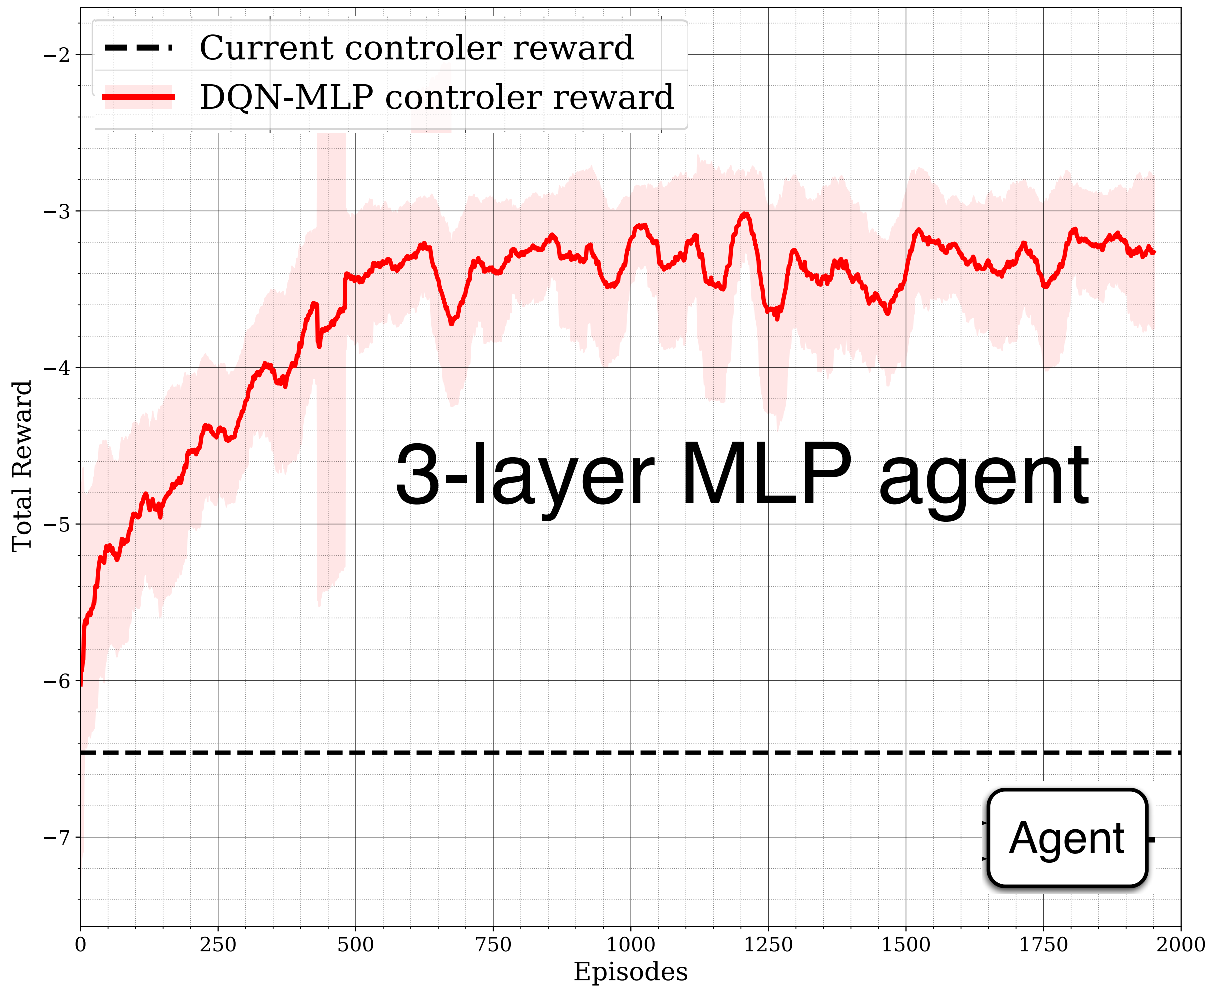

Average **agent** reward increases after many 'episodes' of training

Control Agent implementation

- Ensemble of agent models (same arch, different weights) allows a more robust controller → pref for many, lightweight models
  - Target board: Intel Arria10
  - Quartus implementation using hls4ml developed w/ much help from H. Javed + Intel folks
  - Initially: 1 model, frozen weights

Layer	Outputs	Activation	Parameters	MACs
1	56	ReLU	336	280
2	56	ReLU	3192	3136
3	56	ReLU	3192	3136
4	7	Linear	399	392
Total	...	...	7119	6944

3-layer MLP agent arch

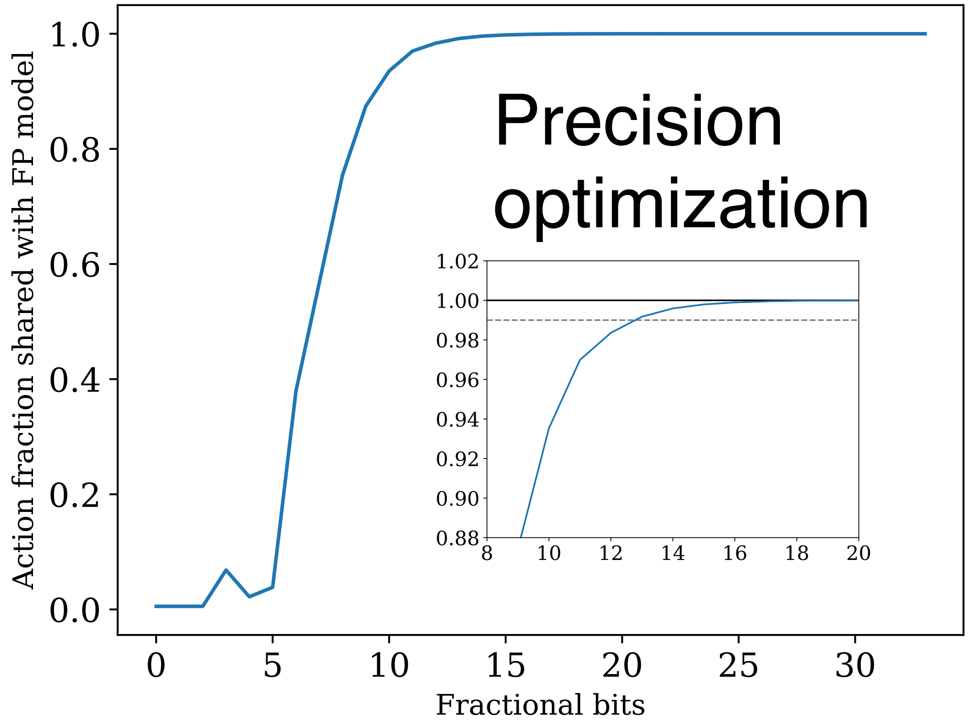

Optimized weight, operand precision comparing agent's decisions to predictions using floats

Control Agent implementation

- Ensemble of agent models (same arch, different weights) allows a more robust controller → pref for many, lightweight models
  - Target board: Intel Arria10
  - Quartus implementation using hls4ml developed w/ much help from H. Javed + Intel folks
  - Initially: 1 model, frozen weights

Layer	Outputs	Activation	Parameters	MACs
1	56	ReLU	336	280
2	56	ReLU	3192	3136
3	56	ReLU	3192	3136
4	7	Linear	399	392
Total	...	...	7119	6944

3-layer MLP agent arch

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	...

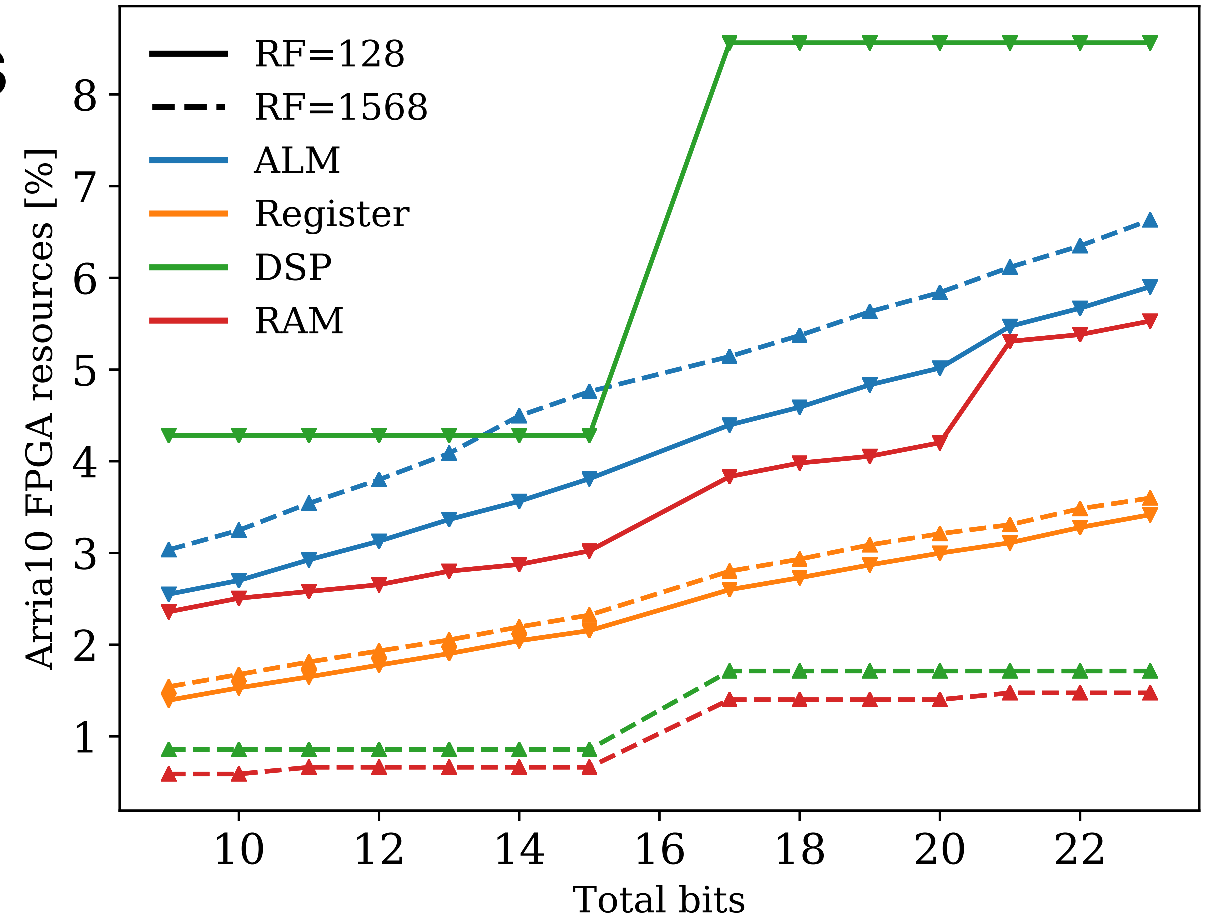

Studied range of choices, with emphasis on 'highly-serial' designs. ~5% usage →O(10) MLP ensemble

Next steps + Collaborators

- For many more results & details, see arXiv:2011.07371
- Exciting possibilities for ensembles, live-training and reconfiguration to explore.
- Due to the closure of accelerator facilities in 2020, deployment in Booster complex so far impossible
  - Looking ahead to data results in 2021!

Pacific Northwest NATIONAL LABORATORY

Lucy Huang,
Malachi Schram

UC San Diego

Javier Duarte

COLUMBIA UNIVERSITY

Rachel Keller

Fermilab

Jason St. John,
Christian Herwig,
Diane Kafkes,
Bill Pelico,
Gabe Perdue

Andres Quintero-Parra,
Brian Schupbach,
Kiyomi Seiya,
Nhan Tran,

Backup

Performance of present system

Without feedback

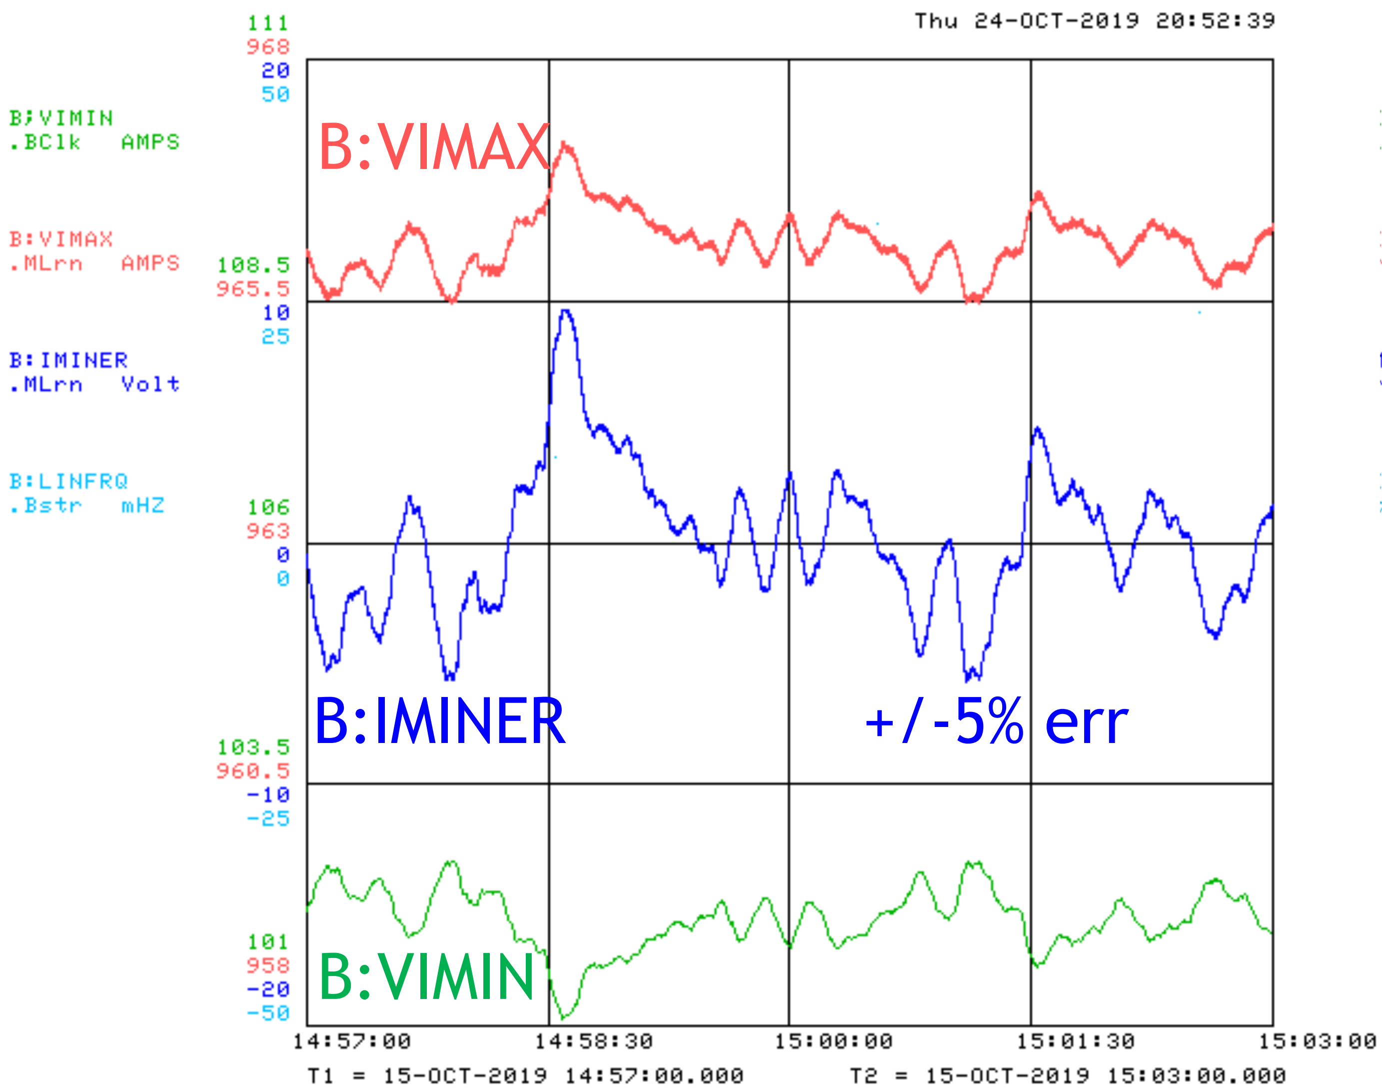

With present feedback system

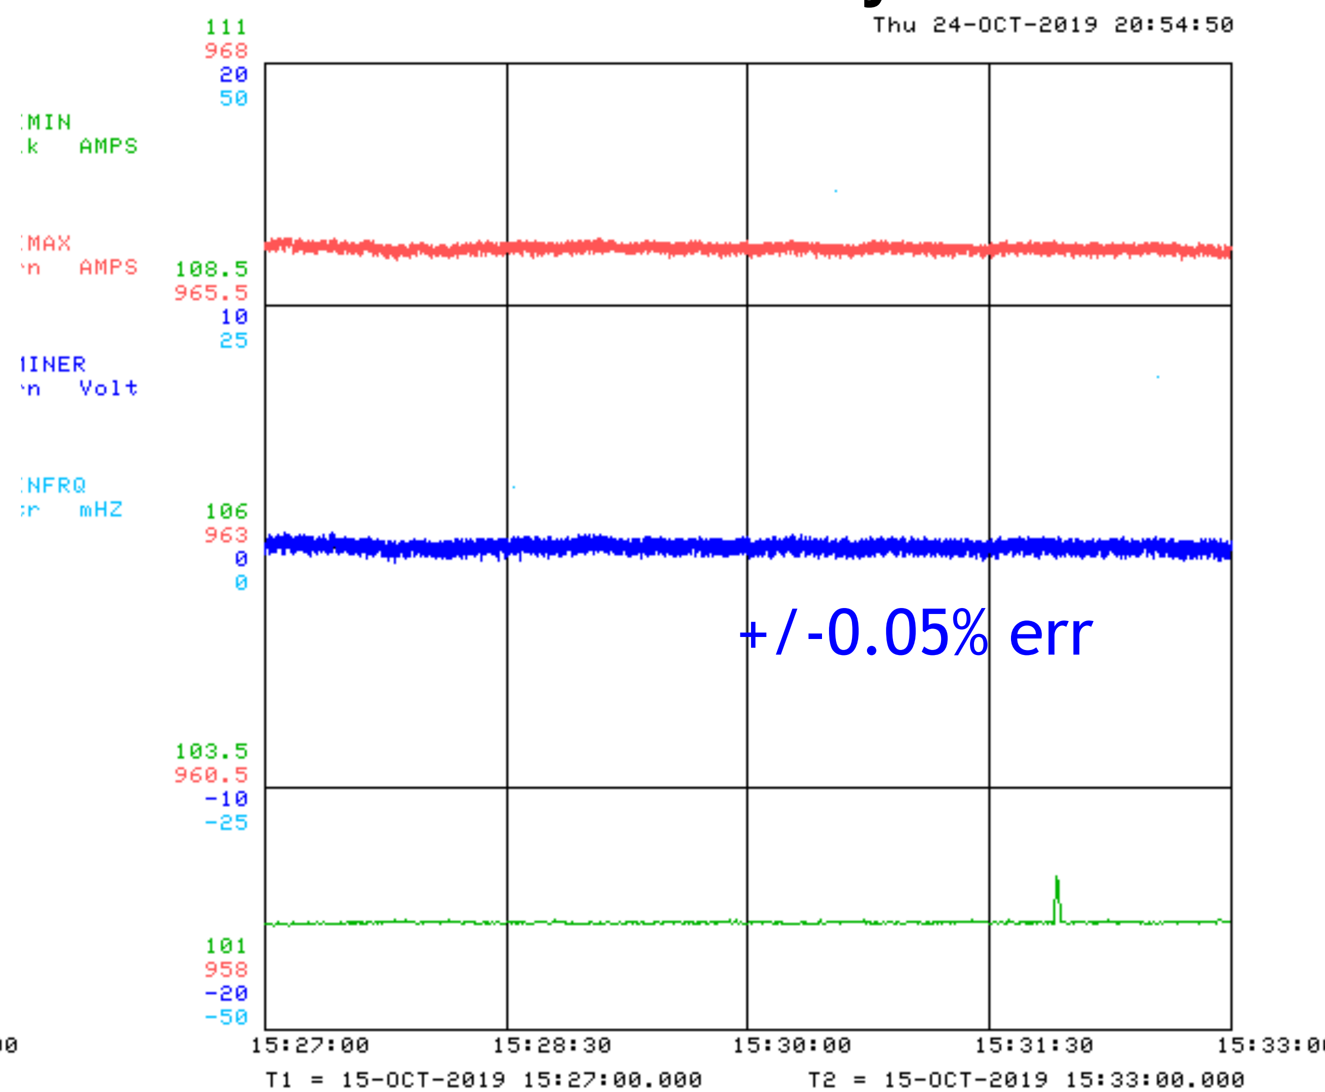

- Goal is to reduce the error to +/-0.025%
- Beam orbit change will be less than 1mm at injection

Reinforcement Learning paradigm

- Agent: Model prescribes $I_{min}(t_1)$ = Func($I_{min}(t_0)$, Temp,…)
- Environment: Booster reacts to GMPS prescription, changes the measured reference current, etc…

Could be PID controller or arbitrary model

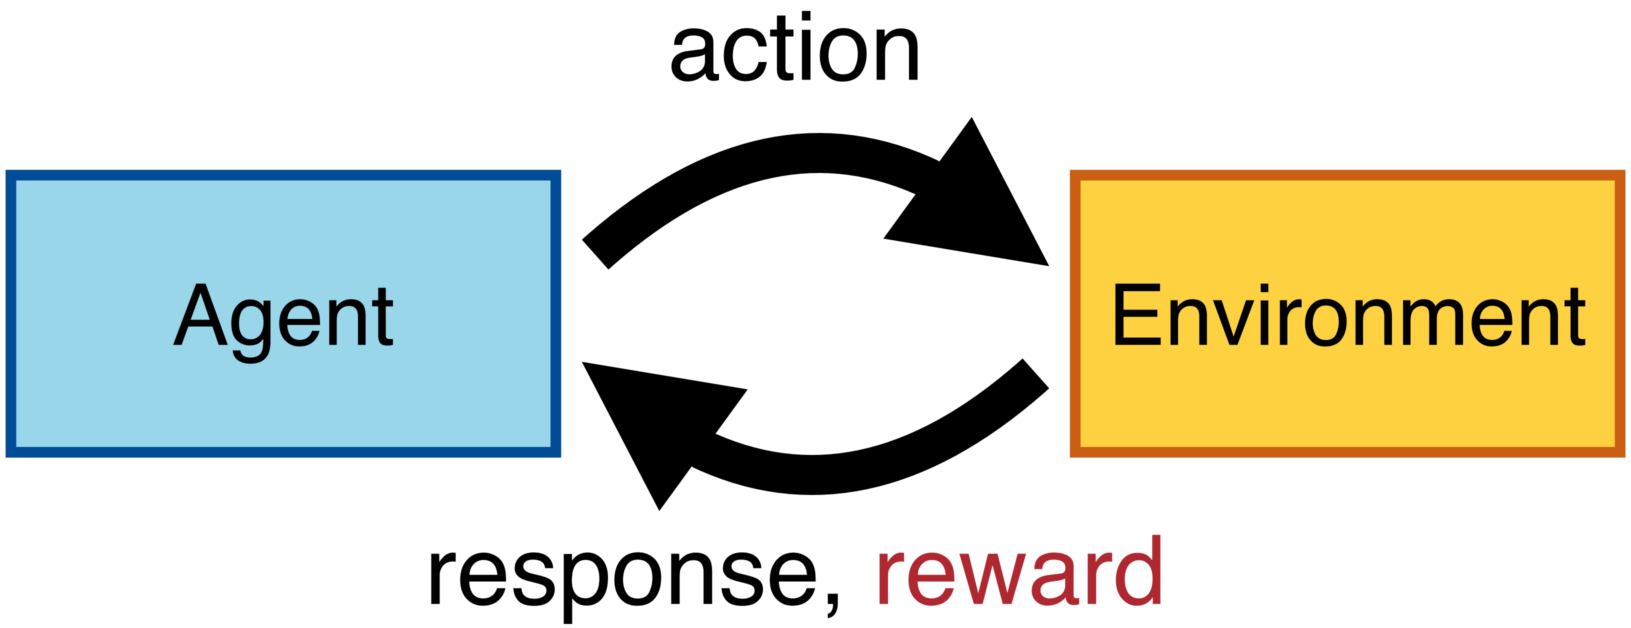

- Reward can simply be -Δ(target setting, measurement)
  - But may also include penalties, i.e. for making a prescription that wildly differs from PID controller model

Surrogate environment

- Without board, no live training in Booster environment
  - In mean time, replace with a surrogate model

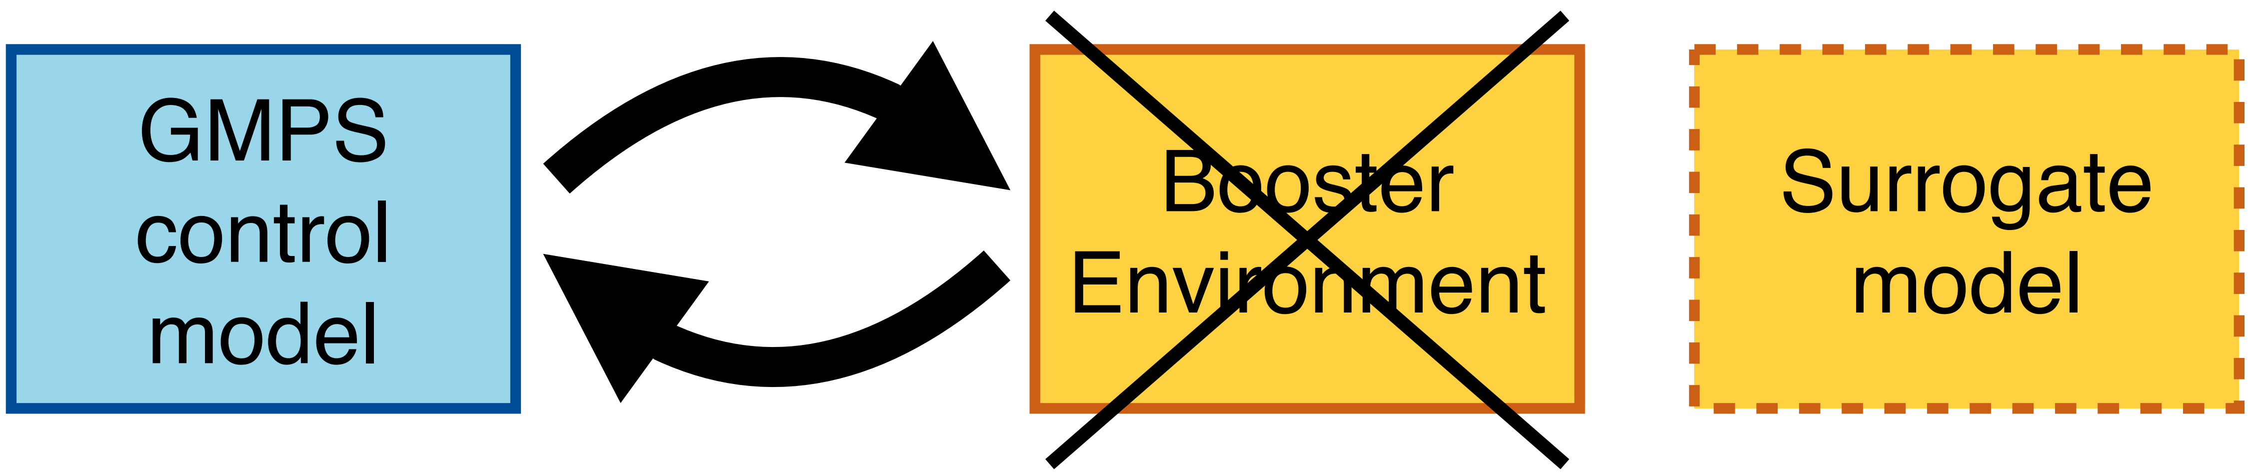

- Use recurrent NN (LSTM) as surrogate
  - feedback / state memory mechanism
- Maps: $I_{\min}(t_{\leq 0}) \rightarrow I_{\min}(t_1)$
- Train surrogate w/ Deep-Q Network

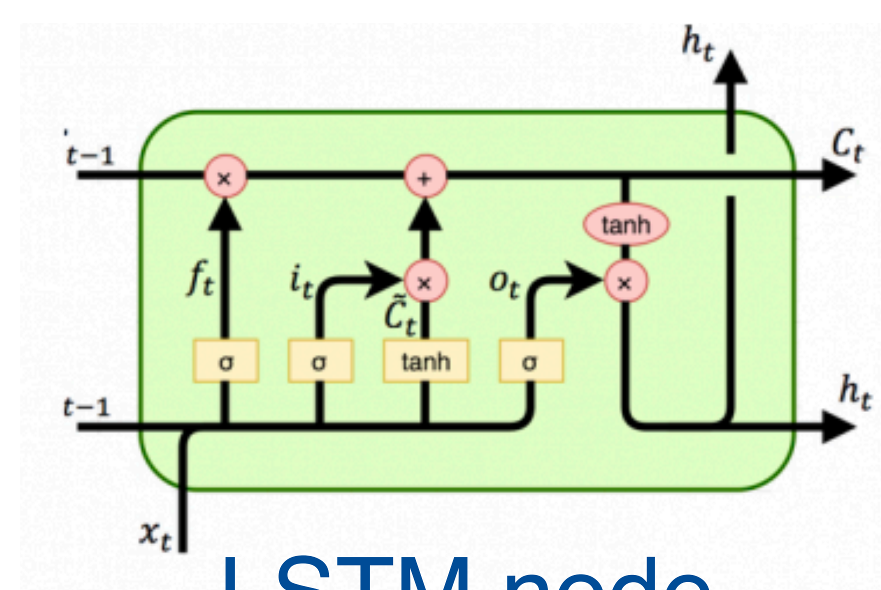

LSTM node