Pre-Flare-Up Fusion Control

Concept paper draft

Working purpose: capture what we think is already true, what we suspect may be true, and a rough control architecture for preventing plasma flare-ups before they become damaging bursts, crashes, or disruptions.

Status note. This is a concept paper, not a proof paper. It does not claim a new plasma-physics discovery. It frames an engineering-control idea: treat pre-burst plasma behavior as an evolving spectral state and intervene before the machine crosses into a bad regime.

1. Problem statement

Fusion plasma is not only a problem of ignition or bulk confinement. In practice, a major bottleneck is keeping a high-performance plasma from drifting into a bad state and dumping energy through bursty edge events, island coupling, tearing activity, sawtooth crashes, or larger disruptions. Existing work already shows that real-time prediction, resonant magnetic perturbation control, and machine-learning-assisted control can change outcomes. The unresolved engineering opportunity is to move from mainly reactive suppression toward a more explicit pre-flare-up prevention layer.

2. What we think is already true

1) Many harmful plasma events show precursors. Depending on the event class, these can appear as changes in frequency, amplitude, phase relationship, rotation speed, or mode coupling before the visible flare-up.

2) Different event classes live in different bands. There is no single universal outburst frequency. ELM repetition often sits in the Hz-to-tens-of-Hz range, tearing/rotating island tracking often lives in the roughly 100 Hz to 10 kHz band, and some ELM precursors have been reported in the hundreds-of-kHz range.

3) What matters is often not one absolute frequency, but a trajectory: drift, narrowing, matching, locking, or fast growth in a multi-signal state.

4) Magnetic control already works in pieces. RMPs, rotating perturbations, localized current drive, and real-time feedback can suppress or mitigate specific instabilities. AI/ML control has also already been demonstrated in real-time tokamak settings.

5) Because the machine already has control levers, the practical next question is not whether intervention is possible in principle, but whether a controller can detect the right precursor state early enough and choose the least-damaging intervention quickly enough.

3. What we suspect but cannot yet prove

1) Plasma flare-ups may be better controlled if treated as a spectral-state transition problem rather than a simple threshold alarm problem.

2) The earliest useful warning may be a bundle of signals rather than a single precursor: frequency drift plus amplitude growth plus rotation change plus phase tightening plus mode coupling risk.

3) A predictive controller may outperform static recipes if it can forecast the near-term evolution of the instability spectrum and de-tune the plasma away from bad relationships before locking or runaway growth occurs.

4) The most useful control may not be brute suppression. The best intervention may be the smallest actuation that changes the trajectory away from the bad state while preserving as much confinement and performance as possible.

5) There may be invention space in a control architecture that explicitly manages instability as an evolving spectral field rather than as a generic disruption-probability score.

4. Event-class starter map

ELMs: likely two important clocks — the repetition rate of the edge event itself and any precursor magnetic activity before it.

Sawteeth: likely better described by precursor relationships and frequency ratios than by one absolute canonical number.

Tearing / neoclassical tearing / rotating islands: especially control-relevant because frequency, phase, and rotation matching can matter directly to coupling and locking.

Larger disruptions: often not one frequency problem but a late-stage outcome of earlier instability evolution; the usable design question is precursor window length and reliability, not a single disruption tone.

5. Rough architecture

A. Sensing layer
Collect real-time measurements from magnetic probes, Mirnov coils, soft X-ray channels, edge diagnostics, density/temperature fluctuation monitors, rotation estimates, and any already-available control-oriented plasma state estimates.

B. Spectral state engine
Build a live picture of what the plasma is doing: active bands, mode amplitudes, growth rates, phase relationships, rotation speeds, coupling signatures, radial location, and time derivatives.

C. Risk engine
Convert the live spectral picture into event-class risks, for example: ELM risk, tearing-locking risk, sawtooth crash risk, coupling risk, and global disruption risk.

D. Predictive engine
Forecast the short-horizon evolution of the current state. Ask: if no action is taken, is the plasma drifting toward locking, mode merger, sudden amplitude growth, or loss of confinement?

E. Intervention planner
Choose the smallest plausible intervention that can change the trajectory. Candidate levers include perturbation amplitude, perturbation phase, perturbation rotation, localized current-drive timing/location, edge-control actuators, or a controlled bleed-off path.

F. Safety override
If the machine is already too close to a cliff edge, protection outranks optimization. The system must fall back to conservative machine-protection logic.

6. Core design philosophy

The proposed controller should be predictive rather than purely reactive.

Reactive philosophy: something bad has started, suppress it.

Predictive philosophy: the plasma is drifting toward a dangerous mode relationship; de-tune it now.

This matters because once a flare-up or locked configuration is fully formed, the available interventions become fewer, harsher, and more damaging to performance.

7. What should be tested first

1) Earliest reliable precursor window by event class.
2) Which signals lead and which lag.
3) Whether frequency ratios / matching / phase locking outperform simple amplitude thresholds.
4) Whether one integrated spectral-state score predicts bad events better than isolated single-signal alarms.
5) Whether archived shots suggest plausible counterfactual interventions.
6) Whether a reduced simulator can learn a safer detuning policy than fixed recipes.

8. Prototype scope we could actually support

Without a fusion lab, the realistic deliverable is not a reactor-grade controller. The realistic deliverable is a prototype stack:

- a formal architecture,
- a reduced simulation environment,
- a multi-signal precursor detector,
- a short-horizon risk estimator,
- and a policy engine that recommends the least-disruptive intervention.

That would be enough to make the concept concrete, pressure-test the framing, and prepare better technical questions for a fusion company, lab, or university group.

9. Main uncertainties

1) Whether available diagnostics give enough signal quality and latency margin.
2) Whether different event classes require very different models.
3) Whether interventions can act fast enough compared with the precursor window.
4) Whether the controller would create new instability problems while solving old ones.
5) Whether the strongest improvement comes from AI itself or from better state representation and control framing.

10. Working conclusion

The strongest current version of the idea is not "AI solves fusion." It is narrower and more defensible: a pre-flare-up control layer that watches the evolving spectral state of the plasma, predicts when it is approaching a dangerous instability regime, and applies the least-damaging intervention before burst behavior begins.

The concept is grounded enough to justify a technical question set, a reduced-model prototype, and a future collaboration conversation. It is not a proof of novelty or success. It is a coherent control hypothesis worth testing.

Appendix A. Short list of technical questions for outside experts

• Which precursor quantities are operationally most useful on your machine: absolute frequency, frequency drift, phase, amplitude growth, rotation matching, or coupling onset?

• What precursor window length is typical by event class, and how much of that window is actually usable for intervention?

• Do your present controllers act on a multi-signal spectral state, or mainly on class-specific thresholds and model-based recipes?

• Which actuators are fast enough to matter before a flare-up: RMP phase/amplitude/rotation, ECCD, edge fueling, or only mitigation after the fact?

• Where does real-time AI help today: prediction, tuning, optimization, or control authority selection?

Appendix B. Source categories already known to be relevant

• Real-time disruption prediction and control-oriented prediction work

• RMP and rotating-RMP control studies

• Real-time tearing / NTM tracking and control

• AI / machine-learning plasma-control demonstrations

• ITER and tokamak magnetic control architecture papers

Appendix C. Immediate next actions

1) Build a machine-agnostic signal dictionary: frequency, drift, amplitude, phase, rotation, coupling, location, and event-class labels.

2) Gather archived example shots and classify precursor windows by event class rather than mixing all bursts together.

3) Create a reduced simulator that lets a controller compare reactive suppression against predictive detuning.

4) Use the technical question list in expert conversations before making any novelty claims.