BNL-230131-2026-TECH

CA/AP/757

Simulation of AGS Proton Bunch Compression for Muon Collider Applications

R. Beckwith, S. Brooks

August 2026

Collider Accelerator Department

Brookhaven National Laboratory

U.S. Department of Energy

USDOE Office of Science (SC), Nuclear Physics (NP)

DISCLAIMER

Simulation of AGS Proton Bunch Compression for Muon Collider Applications

R. Beckwith, S. Brooks

August 2026

Abstract
The short lifetime of muons makes rapid production and capture essential for a future muon collider. One proposed approach uses Brookhaven National Laboratory's Alternating Gradient Synchrotron (AGS) to produce proton bunches short enough to meet the bunch length requirements needed for downstream muon production and cooling. This project investigates the feasibility of compressing AGS proton bunches to an RMS bunch length of 2 ns using a Python-based longitudinal beam dynamics simulation. The model tracks particle motion in longitudinal phase space, incorporating nonlinear RF bucket dynamics and several bunch compression techniques: non-adiabatic, unstable fixed-point, and resonant. The dominant control parameter of each method was optimized by parameter sweep. Under the initial simulation assumptions, a non-adiabatic voltage jump gave the shortest bunch, at 1.97 ns RMS, and was the only method to meet the target; unstable fixed-point and resonant bunching followed at 2.29 ns RMS and 2.32 ns RMS. Subsequent studies of longitudinal beam stability motivated a revised estimate of the initial longitudinal emittance, while comparisons with measured AGS bunch profiles provide an updated initial phase-space distribution. Because both the emittance and distribution influence the achievable compression, these refinements are expected to modify the quantitative optimization results. The results presented here therefore establish benchmarks for bunch compression under the original single-particle simulation assumptions, while identifying the initial beam conditions, longitudinal stability, and realistic RF operating constraints as important considerations for further validation.

1 Introduction

A muon collider offers a path to multi-TeV lepton collisions in a comparatively compact facility. Because the muon is approximately 207 times heavier than the electron, synchrotron radiation is strongly suppressed, allowing high-energy collisions in circular rings. As point-like particles, muons also deliver the full center-of-mass energy to the hard interaction, making muon colliders a promising platform for precision Standard Model measurements and searches for physics beyond the Standard Model [1].

A central challenge is the muon's short lifetime of approximately 2.2 μs. Muons are produced indirectly by striking a target with a high-power proton beam to generate pions, which subsequently decay into muons. The proton bunch time structure therefore affects the longitudinal distribution entering the downstream capture, phase-rotation, and cooling systems. Short, high-intensity proton bunches, on the order of a few nanoseconds RMS, are consequently an important requirement for a muon collider proton driver [2].

Brookhaven National Laboratory's Alternating Gradient Synchrotron (AGS) is being investigated as a potential proton driver for a muon collider demonstrator. The AGS delivers high-intensity proton beams at 24 GeV kinetic energy using an $h = 6$ RF system, but its nominal flattop bunches are longer than the required target. Achieving the necessary compression is constrained by available RF voltage, longitudinal emittance, and intensity-dependent beam stability. At fixed longitudinal emittance, Liouville's theorem requires that shortening the bunch increase its momentum spread. Compression methods therefore use synchrotron motion to redistribute the bunch in longitudinal phase space, with different requirements on RF voltage, timing, and beam stability.

This work develops a Python-based longitudinal beam dynamics simulation to evaluate three AGS compression techniques: non-adiabatic bunching, unstable fixed-point bunching, and resonant bunching. Adiabatic bunching is included as a verification case. The model tracks macroparticles in $(\Delta t, \Delta E)$ phase space through the RF bucket, with machine kinematics, slip factor, synchrotron frequency, and bucket geometry derived from the AGS parameters.

2 Methods

Particle motion was modeled in $(\Delta t, \Delta E)$ phase space relative to the synchronous particle, using a turn-by-turn symplectic Euler map. On turn n, each macroparticle receives an energy kick from the RF cavity and is then drifted according to its resulting energy offset:

$$\Delta E_{n+1} = \Delta E_n + qV_n \left[\sin(\phi_s + h\omega_{\mathrm{rev}}\Delta t_n) - \sin\phi_s\right], \tag{1}$$

$$\Delta t_{n+1} = \Delta t_n + \frac{\eta T_{\mathrm{rev}}}{\beta^2 E_s}\, \Delta E_{n+1}, \tag{2}$$

where $\eta = \gamma_t^{-2} - \gamma^{-2}$ is the slip factor, T_{rev} the revolution period, E_s the synchronous energy, and V_n the cavity voltage on turn n. The applied RF voltage and phase programs distinguish the four compression methods; all other machine parameters were held fixed. Because the AGS operates above transition at flattop, $\eta > 0$ and particles with positive energy offset arrive progressively later.

Machine parameters were derived from first principles rather than assumed, with kinematic quantities, the slip factor, the synchrotron frequency, and the RF bucket geometry all computed from the AGS operating point summarized in Table 1. Simulations were run at flattop with no acceleration, so $\phi_s = 0$ and the bucket is stationary.

Table 1: Representative AGS machine parameters used throughout the simulations.

Parameter	Value
Beam Energy	24 GeV
Circumference	807.1 m
Harmonic Number	6
Momentum Compaction	0.0133
Maximum RF Voltage	320 kV

Four RF programs were implemented, distinguished by how the RF voltage V_n and, where applicable, the RF phase evolve over the simulation:

- **Adiabatic bunching** raises the RF voltage slowly enough that the bunch remains matched to the evolving bucket throughout. Because the matched-ellipse condition is approximately preserved, this program produces relatively modest compression and is included primarily as a verification case.
- **Non-adiabatic bunching** applies a rapid voltage jump and allows the mismatched bunch to rotate for a quarter of a synchrotron period, reaching a minimum length before filamentation sets in.
- **Unstable fixed-point bunching** shifts the RF phase to place the bunch near an unstable fixed point, allowing the strongly nonlinear motion near the separatrix to fold and compress the longitudinal distribution.
- **Resonant bunching** modulates the RF voltage at twice the synchrotron frequency, parametrically driving the phase-space ellipse.

For each method, the RMS bunch length was recorded throughout the simulation, and the minimum RMS bunch length achieved was used as the primary performance metric for comparing the three compression candidates against one another and against the adiabatic baseline. For this comparison, all methods used the same initial bunch conditions, with a longitudinal emittance of 1.35 eV s and a uniformly filled elliptical distribution in longitudinal phase space.

2.1 Adiabatic Bunching

Adiabatic bunching compresses the proton bunch by gradually increasing the RF voltage over many turns while approximately preserving the phase-space distribution. Because the voltage changes slowly relative to the synchrotron motion, the beam remains approximately matched to the evolving RF bucket. This reduces mismatch-driven oscillations, emittance growth, and particle loss.

The voltage was ramped according to

$$V_n = V_0 + (V_{\max} - V_0)\left(\frac{n}{N}\right), \tag{3}$$

where V_0 is the initial voltage, $V_{\max}$ the final voltage, n the turn number, and N the total number of turns over which the ramp is applied. The ramp duration N was chosen to satisfy the adiabaticity condition shown in Equation 4. The fractional change in synchrotron frequency per turn remains small compared to the synchrotron motion itself, ensuring the bunch tracks the matched-ellipse contour rather than being shaken loose from it.

$$\left|\frac{1}{\omega_s}\frac{d\omega_s}{dn}\right| \ll 1, \tag{4}$$

In the simulation, the RF voltage was increased smoothly from its initial value to the maximum operating voltage. Figure 1 shows the applied RF voltage program, while Figure 2 shows the corresponding longitudinal phase-space evolution.

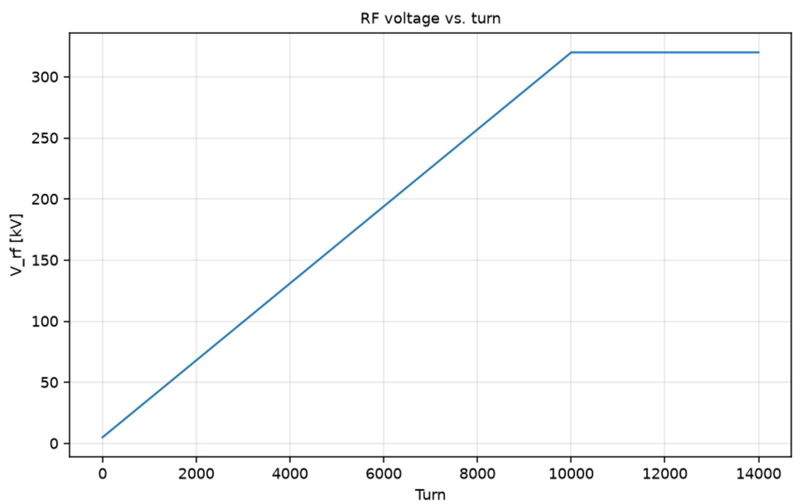

Figure 1: RF voltage program used for adiabatic bunching.

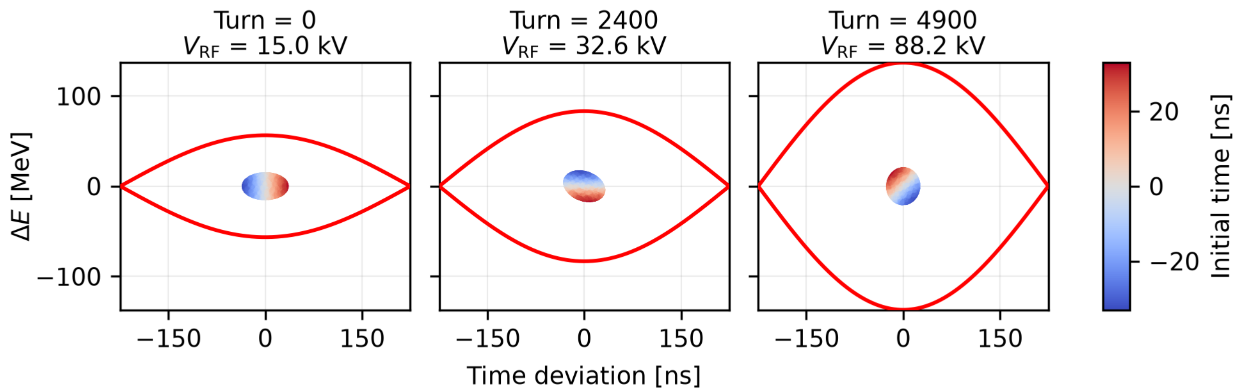

Figure 2: Representative longitudinal phase-space evolution during adiabatic bunch compression. The red curve indicates the RF separatrix, while the particle distribution shows the evolution of the proton bunch.

The resulting compression was on the order of 7 ns RMS bunch length. The final bunch distribution was found to agree with the theoretical matched distribution expected at the final RF voltage.

2.2 Non-Adiabatic Bunching

Unlike adiabatic bunching, non-adiabatic bunching deliberately violates the adiabaticity condition of Eq. 4 by applying the voltage change over a time short compared to the synchrotron period, rather than slow compared to it. The sudden voltage increase produces a mismatch between the beam distribution and the new, deeper RF bucket. This mismatch initiates phase-space rotation: the mismatched ellipse rotates about the matched contour, causing the bunch to alternate between being narrow in time and narrow in energy as it rotates.

The voltage program was a step function,

$$V_n = \begin{cases} V_0, & n < n_{\text{jump}} \\ V_{\max}, & n \geq n_{\text{jump}} \end{cases} \tag{5}$$

where n_{jump} is the turn at which the voltage is switched, chosen to be short compared to the synchrotron period $T_s = 2\pi/\omega_s$ so that the beam distribution has no time to adjust to the new bucket before the jump is complete.

In the simulation, the RF voltage was rapidly increased from its initial value to the operating voltage. The bunch was then tracked through the resulting phase-space rotation, and the RMS bunch length was calculated at every turn to identify the minimum.

Figure 3 shows the RF voltage program, and Figure 4 shows the corresponding phase-space evolution.

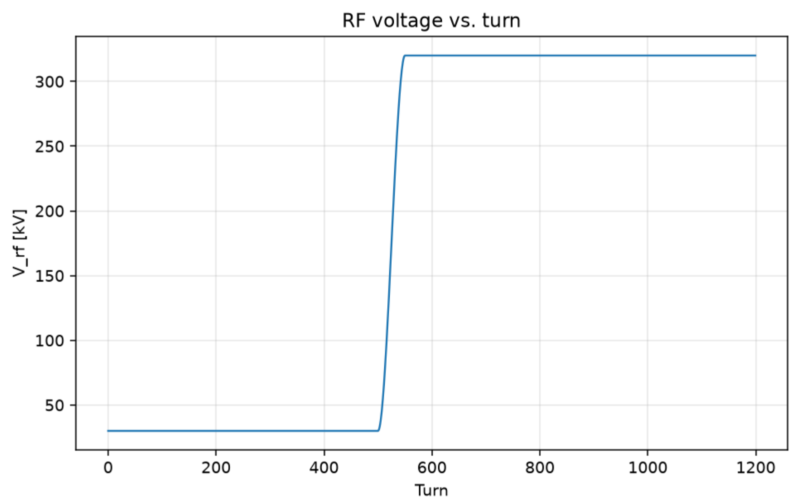

Figure 3: RF voltage program used for non-adiabatic bunching.

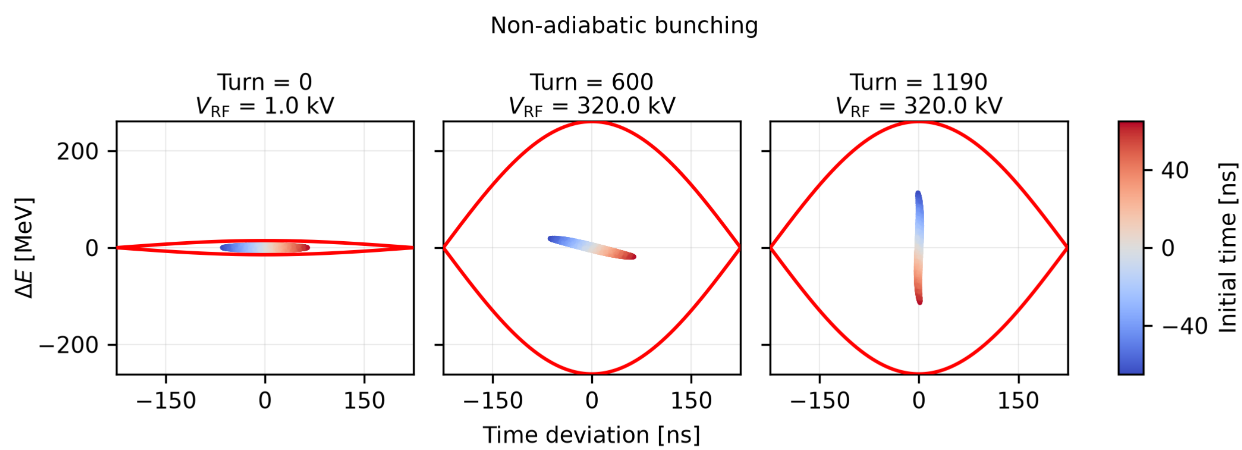

Figure 4: Representative longitudinal phase-space evolution during non-adiabatic bunching. The voltage jump induces phase-space rotation, causing the bunch to reach a temporary minimum in temporal width.

2.3 Unstable Fixed-Point Bunching

The unstable fixed-point (UFP) method exploits the nonlinear structure of the RF bucket rather than a change in voltage amplitude. Away from the stable synchronous phase ϕ_s, the bucket separatrix, computed via Brent's method, contains an unstable fixed point at $\phi_s + \pi$. The RF phase is shifted to place the initially matched bunch near the unstable fixed point, where the local phase-space flow is hyperbolic rather than elliptical: small deviations grow rather than oscillate, causing particles to move away from the fixed point and stretch along the separatrix.

After the distribution becomes sufficiently elongated along the separatrix, an RF phase shift repositions the bucket so that the now-elongated distribution finds itself displaced from the stable fixed point rather than sitting on the unstable one. The RF voltage V_n is held fixed throughout.

Figure 5 shows the RF phase program applied in the simulation. Figure 6 shows the evolution of the beam as it stretches near the unstable fixed point and subsequently rotates into a shorter bunch.

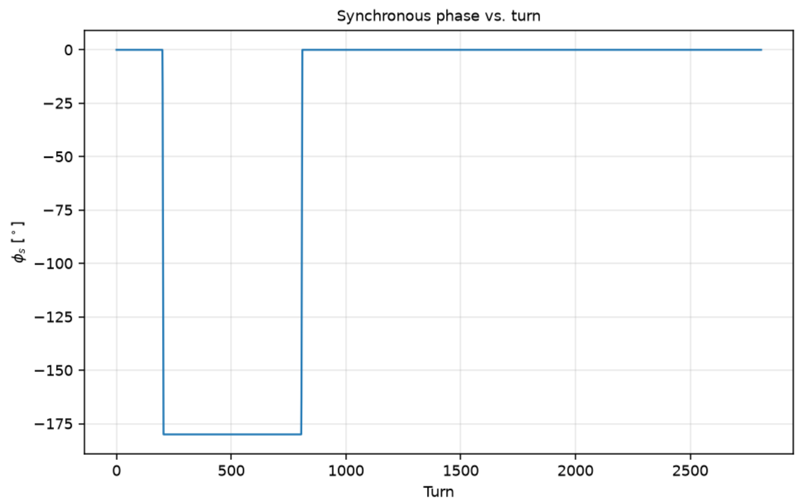

Figure 5: RF phase program used for unstable fixed-point bunching.

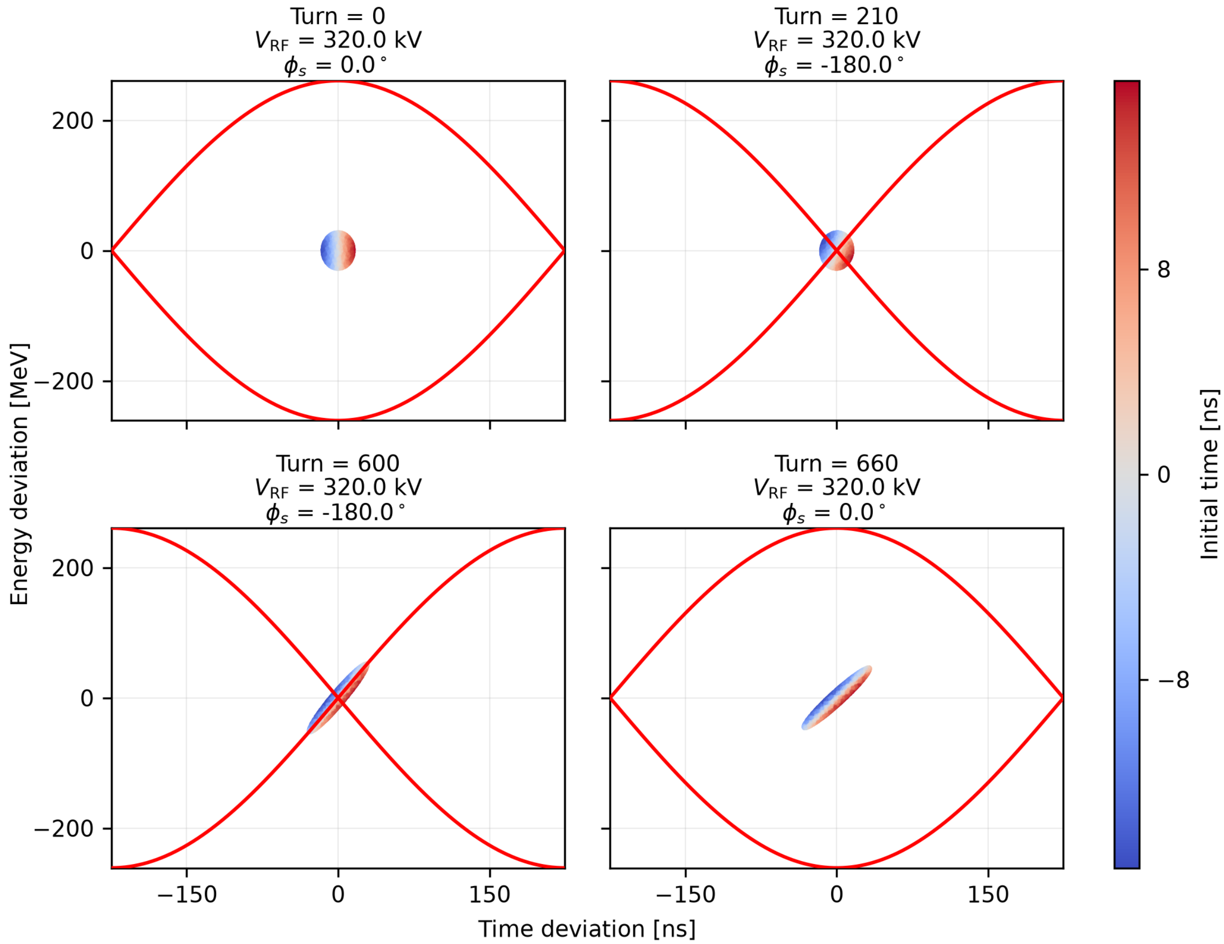

Figure 6: Representative longitudinal phase-space evolution during unstable fixed-point bunching. The beam first stretches near the unstable fixed point and then rotates within the stable RF bucket to produce a shorter bunch.

2.4 Resonant Bunching

Resonant bunching uses periodic RF voltage modulation to parametrically excite longitudinal oscillations. Unlike the adiabatic and non-adiabatic programs, which apply a single change in voltage, the resonant program modulates the voltage continuously at twice the synchrotron frequency, driving the phase-space ellipse through parametric resonance in the sense of the Mathieu/Hill equation: because the quadrupole mode of the bunch oscillates at $2\omega_s$, an amplitude-modulated drive at that same frequency resonantly pumps energy into the bunch shape rather than its centroid, stretching and compressing the phase-space ellipse a little further on each successive synchrotron oscillation. The compression therefore builds up over many turns rather than in a single rotation, in contrast to the single quarter-period rotation of the non-adiabatic method.

The RF voltage was modeled as

$$V_n = \bar{V}_n + V_0 d_n \cos\left[\omega_{\text{mod}}(n - n_{\text{start}}) + \phi_{\text{mod}}\right], \tag{6}$$

where V_0 is the nominal mean RF voltage, d_n is the turn-dependent modulation depth, and

$$\omega_{\text{mod}} = R\omega_s(1 + \delta), \tag{7}$$

with $R = 2$ corresponding to the nominal quadrupole parametric resonance condition. During the initial ramp, both the mean voltage and modulation depth are increased smoothly using a raised-cosine function,

$$r(n) = \frac{1}{2}\left[1 - \cos\left(\pi\frac{n - n_{\text{start}}}{N_{\text{ramp}}}\right)\right]. \tag{8}$$

The mean voltage and modulation depth are then

$$\bar{V}_n = V_{\text{start}} + (V_0 - V_{\text{start}})r(n), \tag{9}$$
$$d_n = d\, r(n). \tag{10}$$

After the ramp is complete, $\bar{V}_n = V_0$ and $d_n = d$.

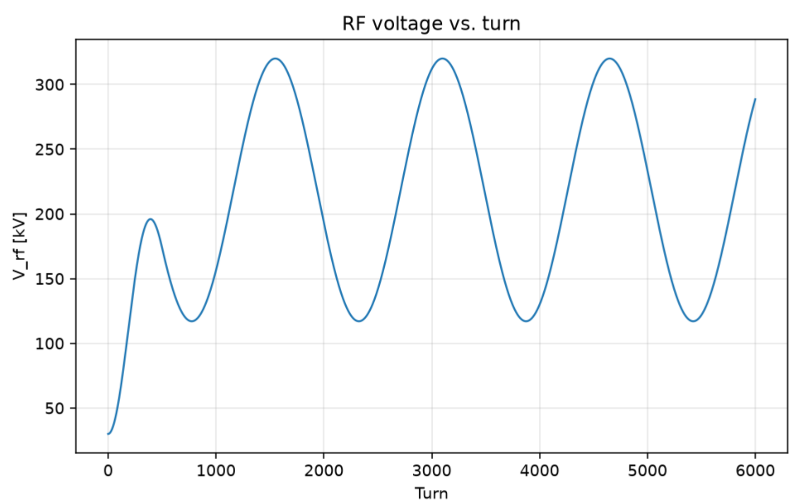

Figure 7: RF voltage program used for resonant bunching. Periodic voltage modulation at twice the synchrotron frequency drives coherent longitudinal oscillations in the proton bunch.

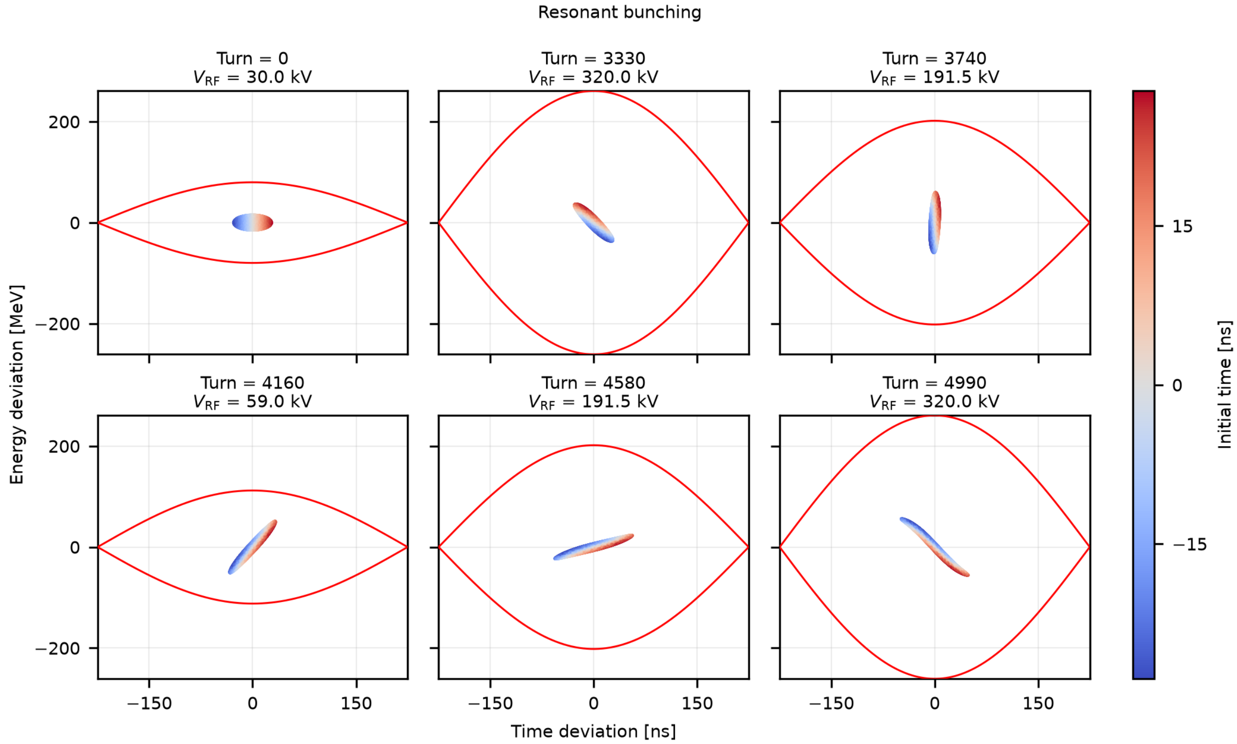

Figure 8: Representative longitudinal phase-space evolution during resonant bunching. The modulated RF voltage drives phase-space oscillations that periodically compress the bunch in time.

2.5 Initial Bunch Distribution

To refine the initial longitudinal distribution for the tracking simulations, measured AGS bunch profiles were compared with several candidate analytical distributions. Oscilloscope traces of the AGS wall current monitor signal were obtained from archived AGS scope data at RF voltages of approximately 207 and 211 kV. The traces were digitized in MATLAB by isolating the cyan oscilloscope signal with an RGB color mask, extracting the median vertical position at each horizontal pixel location, interpolating small gaps, and applying light Savitzky–Golay smoothing.

The digitized profiles were compared with a Gaussian and distributions of the form

$$I_J(u) = \left[1 - \left(\frac{u}{T_J}\right)^2\right]^J, \qquad |u| \leq T_J, \tag{11}$$

with $I_J = 0$ outside the bunch boundary. Here,

$$u = \frac{t - t_0}{t_{\text{HWHM}}}, \tag{12}$$

so that the measured half-maximum points occur at $u = \pm 1$. The scale factor

$$T_J = \frac{1}{\sqrt{1 - 2^{-1/J}}} \tag{13}$$

therefore constrains each candidate distribution to the same measured FWHM. The $J = 3/2$ and $J = 5/2$ distributions were compared with a Gaussian having the same center, amplitude, and FWHM, so that the comparison tests differences in bunch shape rather than independently fitted widths.

Agreement with the digitized data was first examined by overlaying the candidate distributions with the measured AGS profiles, as shown in Figure 9.

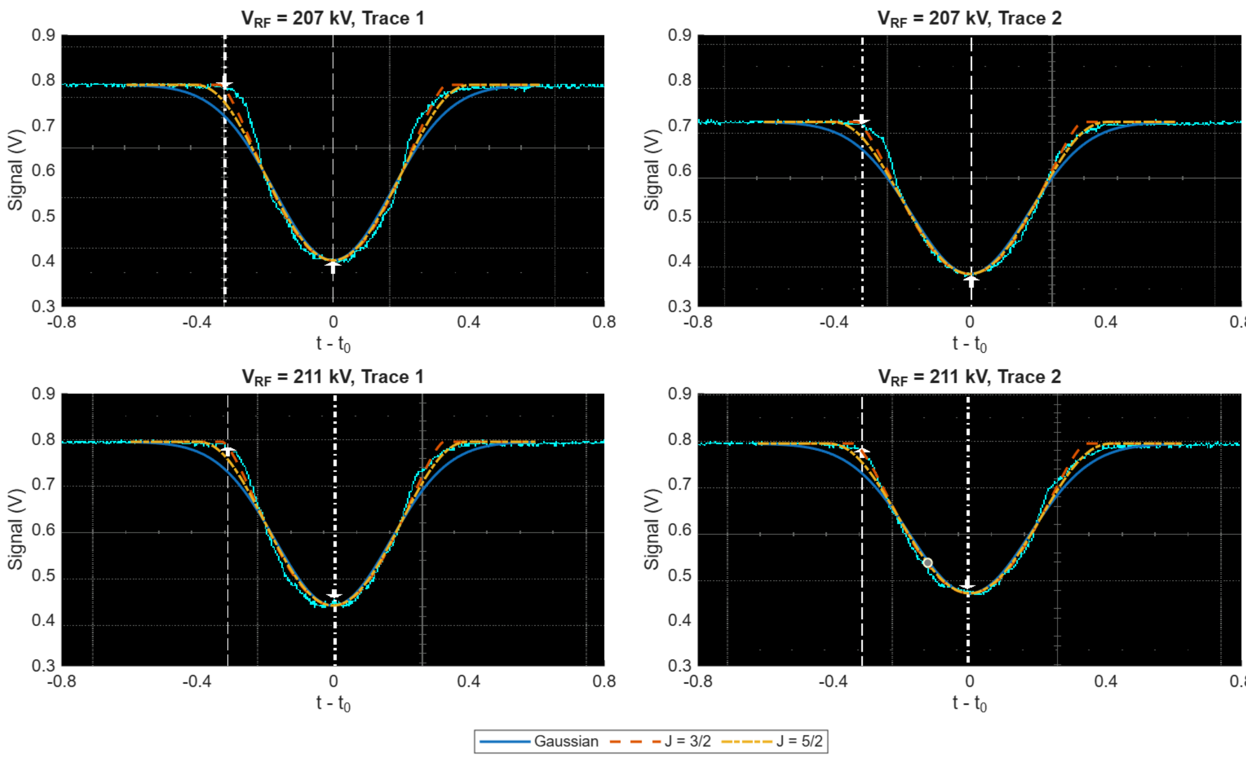

Figure 9: Comparison of digitized AGS wall current monitor profiles with Gaussian, $J = 3/2$, and $J = 5/2$ longitudinal distributions at approximately 207 and 211 kV RF voltage. Each candidate distribution was constrained to the measured bunch center, amplitude, and FWHM.

This analysis was quantified using the root mean square error (RMSE),

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} [I_{\text{data}}(u_i) - I_{\text{model}}(u_i)]^2}, \tag{14}$$

evaluated over $|u| \leq 2$. The resulting errors are summarized in Table 2.

Table 2: Normalized RMSE between the digitized AGS bunch profiles and candidate longitudinal distributions.

Measured profile	**Gaussian**	$J = 3/2$	$J = 5/2$
207 kV, Trace 1	0.0967	0.0474	0.0646
207 kV, Trace 2	0.0833	0.0423	0.0531
211 kV, Trace 1	0.0822	0.0365	0.0498
211 kV, Trace 2	0.0671	0.0352	0.0378

Both J distributions reproduce the measured profiles more closely than the Gaussian in all four cases. The $J = 3/2$ distribution gives the lowest RMSE for each measured profile. Based on this consistent agreement, $J = 3/2$ was selected as the nominal one-dimensional longitudinal bunch distribution.

A two-dimensional phase-space distribution consistent with the selected $J = 3/2$ longitudinal profile was then required to initialize the macroparticles. Defining the normalized elliptical radius

$$r^2 = \left(\frac{t}{T}\right)^2 + \left(\frac{\Delta E}{\Delta E_{\max}}\right)^2, \tag{15}$$

a parabolic phase-space density was used,

$$f(r) = C(1 - r^2), \qquad r \leq 1, \tag{16}$$

with zero density outside the ellipse. Distributions of this form are used in accelerator-physics treatments of longitudinal beam distributions and collective instabilities [4].

The connection between this two-dimensional distribution and the measured one-dimensional profile can be verified by integrating over the energy coordinate. At fixed t, the phase-space boundary is

$$|\Delta E| \leq \Delta E_{\max}\sqrt{1 - \left(\frac{t}{T}\right)^2}. \tag{17}$$

Defining $a = 1 - (t/T)^2$, the longitudinal line density becomes

$$\begin{aligned} \lambda(t) &= C \int_{-\Delta E_{\max}\sqrt{a}}^{+\Delta E_{\max}\sqrt{a}} \left[a - \left(\frac{\Delta E}{\Delta E_{\max}}\right)^2\right] d(\Delta E) \\ &= \frac{4}{3} C \Delta E_{\max} \left[1 - \left(\frac{t}{T}\right)^2\right]^{3/2}. \end{aligned} \tag{18}$$

Thus, the parabolic two-dimensional phase-space density projects directly onto the $J = 3/2$ longitudinal distribution identified from the measured AGS profiles.

As a final check of the initialization procedure, the current profile obtained from the generated macroparticles was compared with the analytical $J = 3/2$ profile. Figure 10 shows close agreement between the two distributions, confirming that the phase-space sampling reproduces the intended longitudinal bunch shape. The small bin-to-bin fluctuations arise from finite macroparticle statistics.

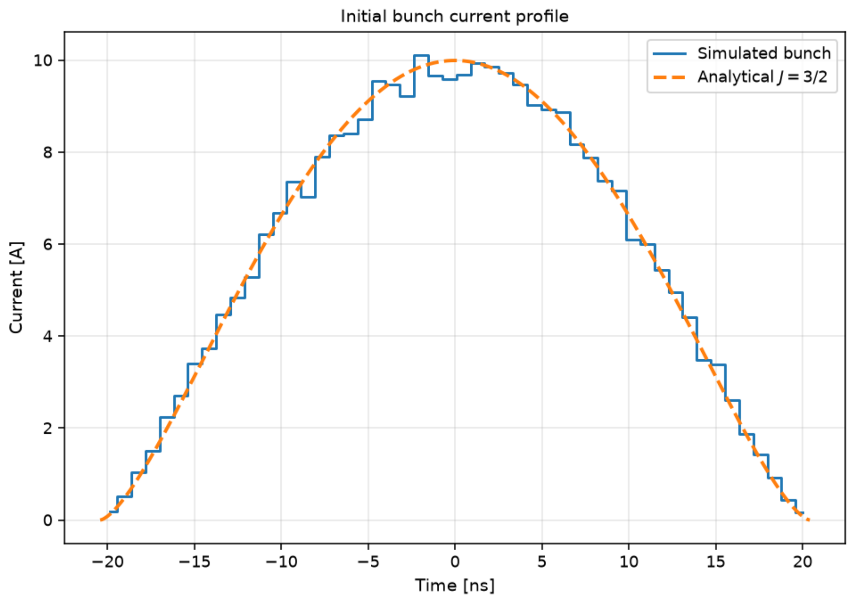

Figure 10: Initial simulated bunch current profile compared with the analytical $J = 3/2$ distribution.

The comparison with measured AGS profiles indicates that the $J = 3/2$ distribution provides a more representative description of the longitudinal bunch shape. However, the primary compression studies presented in the following sections were performed using the simpler $J = 1/2$ distribution, corresponding to a uniformly populated elliptical region in longitudinal phase space. This distribution was used throughout the initial development and optimization of the compression methods and was retained as a common baseline so that the methods could be compared under identical initial conditions.

2.5.1 Effect of the Initial Distribution

To evaluate the sensitivity of the resonant compression results to the assumed initial bunch distribution, the simulation was repeated using the more representative $J = 3/2$ distribution identified from the measured AGS bunch profiles. All other parameters were held fixed between the two cases, including an initial RF voltage of 30 kV, a modulation depth of 0.9, and a longitudinal emittance of 1.35 eV·s. The only change was therefore the initial phase-space distribution.

For the simplified $J = 1/2$ baseline, the minimum RMS bunch length was 2.314 ns at turn 5385. Under the same resonant parameters, the $J = 3/2$ distribution reached a minimum RMS bunch length of 1.843 ns at turn 5378, corresponding to an approximately 20.4% reduction in the predicted minimum bunch length.

This indicates that the assumed initial distribution can have a significant effect on the quantitative prediction of the achievable bunch length. The compression results presented in this work should therefore be interpreted as benchmarks for comparing the relative performance of the different methods, rather than as concrete predictions of achievable bunch lengths. Future studies should use the $J = 3/2$ distribution identified from the measured AGS profiles to provide a more realistic representation of the initial bunch and refine the predicted compression performance.

3 Results

Each of the three fast-compression methods depends on a single dominant control parameter whose value strongly affects the achievable compression. A one-dimensional parameter sweep was therefore performed for each method, with all remaining machine and RF settings held at the values listed in Table 1. In every sweep the RMS bunch length was tracked over the full simulation so that both the minimum value and the turn at which it occurred could be recorded, since a deep minimum that occurs at an impractical turn number is of limited use for extraction. Each subsection below presents the sweep as a heat map of RMS bunch length over turn number and the swept parameter, followed by the bunch-length evolution of the single best-performing run.

3.1 Non-Adiabatic Method Optimization

The non-adiabatic method rapidly increases the RF voltage, leaving the bunch mismatched to the new bucket and causing it to rotate in longitudinal phase space. The bunch reaches maximum compression approximately one-quarter of a synchrotron period after the jump. Lower starting voltages increase the initial bunch length and therefore the available compression, while longer voltage ramps reduce the mismatch and approach the adiabatic limit.

The starting voltage $V_{\mathrm{rf,low}}$ was swept from 5 to 200 kV in 16 steps of 13 kV, and the ramp duration N_{jump} from 5 to 3000 turns in 7 steps of 500 turns. The final voltage was fixed at 320 kV and the longitudinal emittance at 1.35 eV·s. Figure 11 shows the minimum RMS bunch length obtained for each configuration.

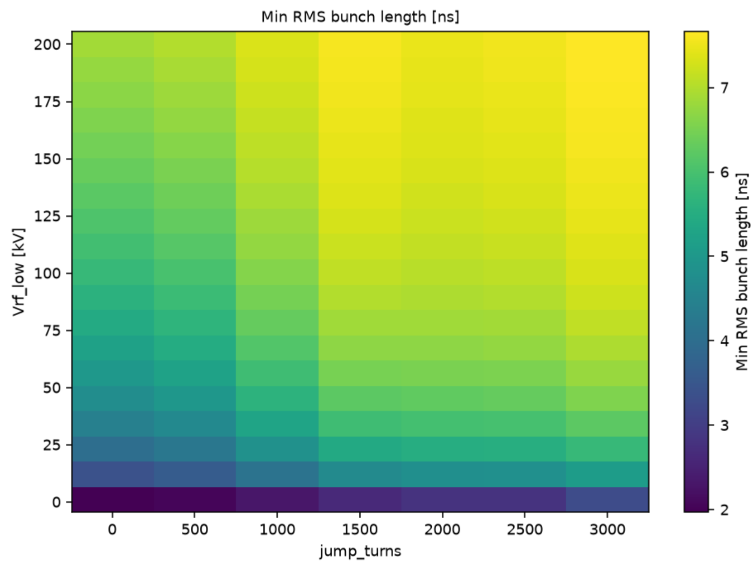

Figure 11: Minimum RMS bunch length attained during the voltage jump as a function of ramp duration N_{jump} and starting voltage $V_{\text{rf,low}}$.

Compression improves toward lower starting voltages and shorter ramp durations. For a near-instantaneous jump, the minimum bunch length decreases from approximately 7.5 ns at 200 kV to 1.97 ns at 5 kV. As the ramp duration increases, the voltage change becomes slow relative to the synchrotron motion and the bunch increasingly follows the changing bucket adiabatically, reducing the mismatch responsible for compression. The performance therefore approaches the adiabatic limit for sufficiently long ramps.

The best-performing configuration was the lowest starting voltage and fastest ramp, with $V_{\text{rf,low}} = 5$ kV and an effectively instantaneous voltage jump, reaching a minimum RMS bunch length of approximately 1.97 ns.

3.2 Unstable Fixed-Point Method Optimization

The unstable fixed-point method uses a 180° RF phase jump to place the bunch at the unstable fixed point, where the local phase-space flow is hyperbolic. The distribution shears along the separatrix, increasing in energy spread while narrowing in time. After a controlled dwell period, the RF phase is restored and the sheared distribution is recaptured in the stable bucket. The dwell time controls the extent of this evolution and was therefore selected as the primary optimization parameter.

The dwell time was swept from 100 to 1000 turns in steps of 100, with the RF voltage (320 kV), phase-jump amplitude (180°), phase-jump transition time (5 turns), and longitudinal emittance (1.35 eV s) held fixed. Figure 12 shows the resulting RMS bunch length as a function of turn number and dwell time.

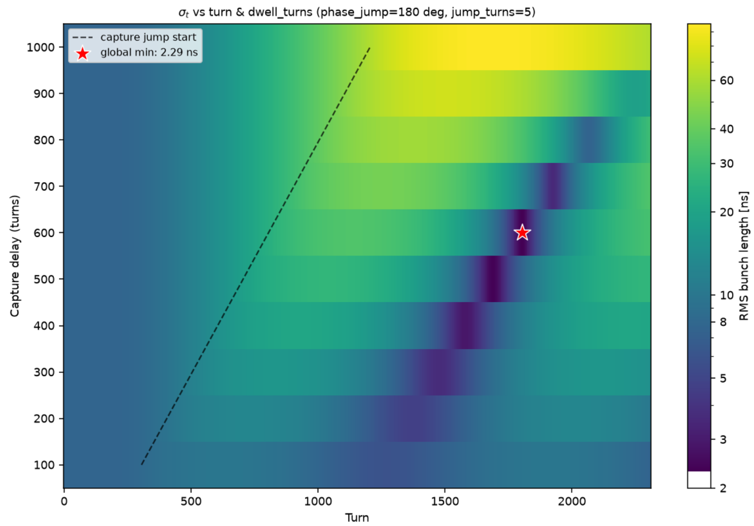

Figure 12: RMS bunch length as a function of turn number and dwell time at the unstable fixed point for a 180° phase jump executed over 5 turns. Darker regions indicate stronger compression. The dashed line marks the start of recapture for each dwell time, and the star marks the global minimum of 2.29 ns.

The heat map shows that the compression minimum shifts to later turns as the dwell time increases, following the corresponding delay in recapture. The compression also strengthens as the dwell increases toward approximately 500–600 turns, but rapidly degrades at longer dwell times as the bunch undergoes increasingly nonlinear motion near the separatrix.

Figure 13 summarizes this behavior by showing the minimum RMS bunch length attained for each dwell time and the turn at which that minimum occurs. The minimum bunch length decreases from 5.9 ns at 100 turns to approximately 2.3 ns near 500–600 turns, producing a relatively shallow optimum. Beyond this region, performance degrades rapidly, reaching approximately 3.5 ns at 700 turns and 7.5 ns at 800 turns. At still longer dwell times, the bunch does not recompress below its initial length, indicating that excessive evolution near the unstable fixed point prevents effective recapture and compression. The lower panel also shows the expected shift of the compression minimum to later turns as the dwell period is increased.

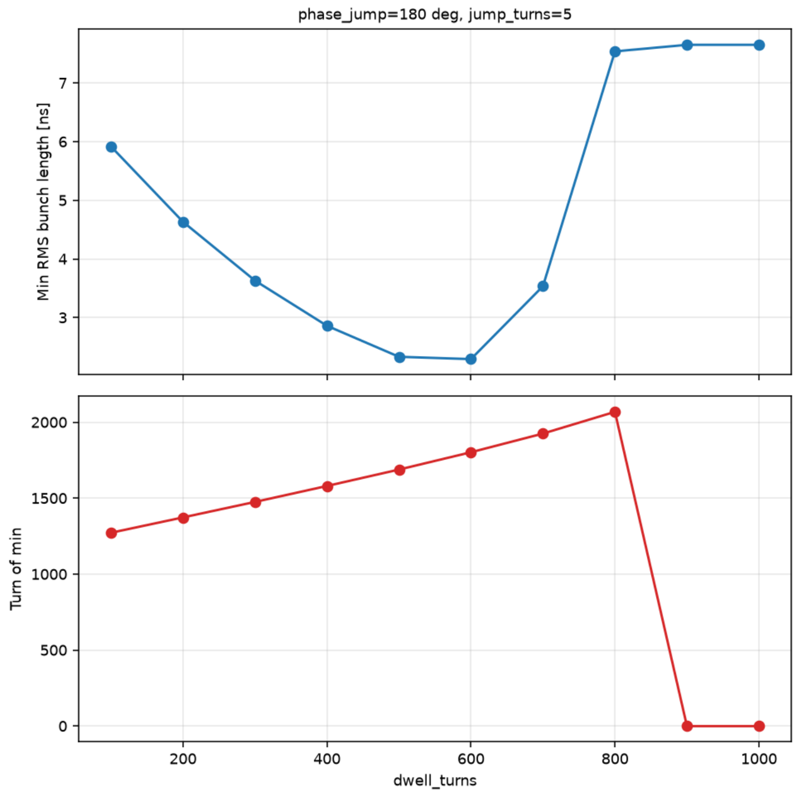

Figure 13: Summary of the dwell-time sweep. Top: minimum RMS bunch length attained for each dwell time. Bottom: turn at which the minimum occurs, increasing with dwell time until effective recompression is lost at the longest dwell times.

The best-performing case used a dwell time of 600 turns and reached a minimum RMS bunch length of 2.29 ns at turn 1803, corresponding to a compression factor of approximately 3.3 relative to the initial 7.6 ns bunch. Figure 14 shows the corresponding bunch-length evolution. The bunch initially lengthens as it evolves near the unstable fixed point and then compresses sharply following recapture as the distribution rotates in the stable bucket. The resulting minimum is transient, making the timing of extraction an important consideration for this method.

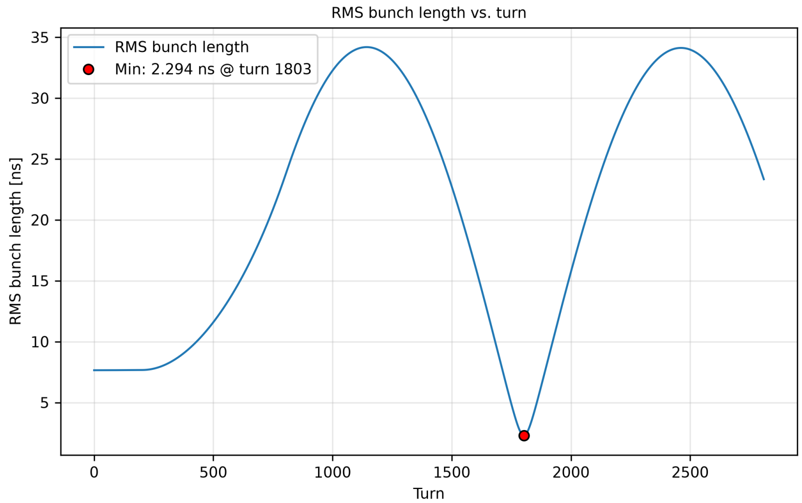

Figure 14: RMS bunch-length evolution for the best-performing unstable fixed-point case, using a dwell time of 600 turns. The minimum of 2.29 ns occurs at turn 1803.

3.3 Resonant Method Optimization

The resonant method modulates the RF voltage sinusoidally at twice the synchrotron frequency, driving the quadrupole oscillation of the bunch and progressively compressing the longitudinal distribution over successive synchrotron periods. The modulation depth controls the strength of this excitation and was therefore selected as the primary optimization parameter.

The modulation depth was swept from 0.05 to 0.95 in steps of 0.05, with the starting RF voltage (30 kV), voltage ramp duration (3000 turns), and modulation frequency ($2f_s$) held fixed. Figure 15 shows the resulting RMS bunch length as a function of turn number and modulation depth.

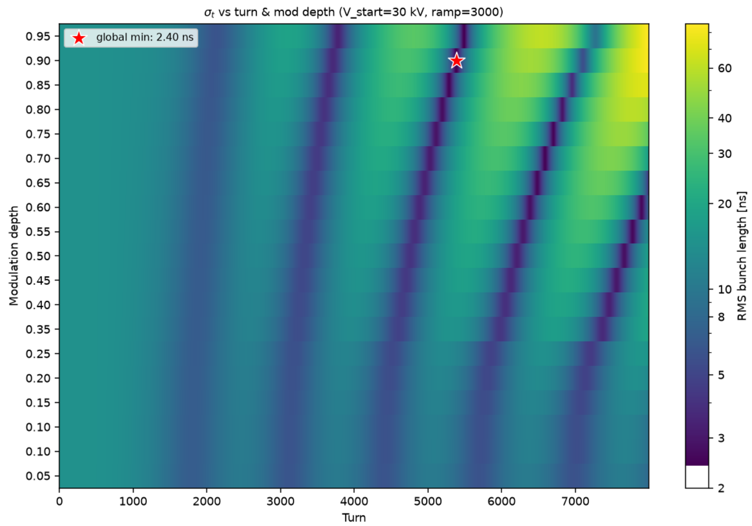

Figure 15: RMS bunch length as a function of turn number and RF modulation depth for the resonant bunching method. The bands correspond to successive minima of the driven bunch oscillation, and the red star marks the global minimum of the scan.

The heat map shows successive compression minima as the resonant drive excites the bunch over multiple oscillations. Compression strengthens with increasing modulation depth before showing diminishing improvement at higher depths, while the timing of the minima also shifts as the bunch explores the nonlinear RF bucket.

Figure 16 summarizes this behavior by showing the minimum RMS bunch length attained for each modulation depth and the turn at which that minimum occurs. The upper panel shows rapid improvement in compression up to approximately $d = 0.55$, after which the minimum bunch length remains approximately 2.4–2.7 ns. The lower panel shows that the turn of maximum compression generally shifts later with increasing modulation depth. The abrupt drops near $d = 0.65$ and $d = 0.80$ occur when an earlier oscillation minimum becomes the global minimum, producing a shift between successive compression cycles rather than a discontinuity in the bunch dynamics.

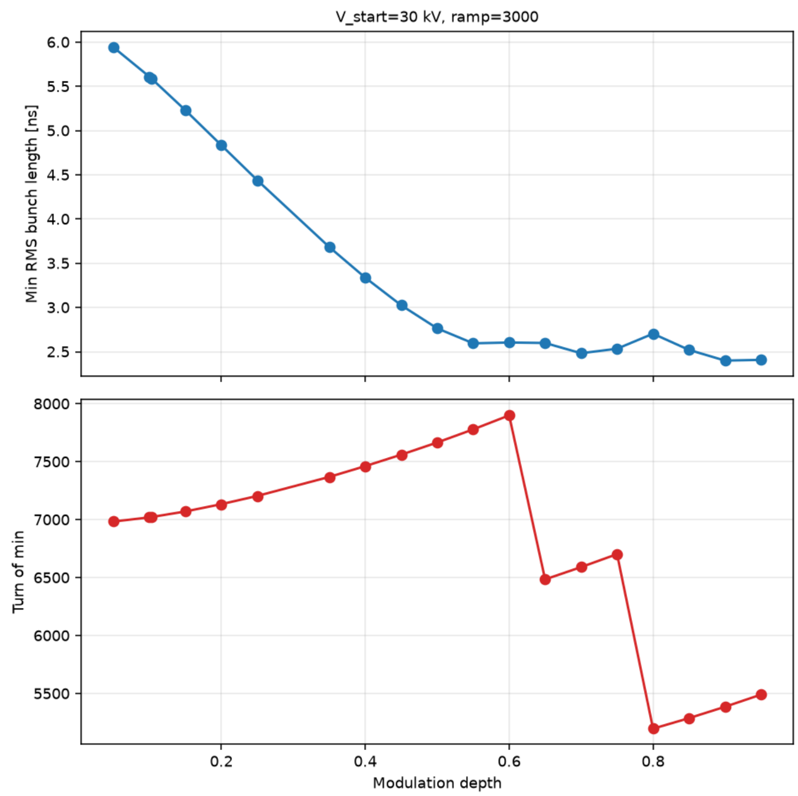

Figure 16: Summary of the modulation-depth sweep. Top: minimum RMS bunch length attained for each modulation depth. Bottom: turn at which the minimum occurs, with abrupt shifts corresponding to changes in which successive compression minimum becomes the global minimum.

The best-performing case used a modulation depth of 0.90 and reached a minimum RMS bunch length of approximately 2.32 ns at turn 5386. Figure 17 shows the corresponding bunch-length evolution, illustrating how successive oscillation minima deepen as the resonant drive progressively compresses the bunch.

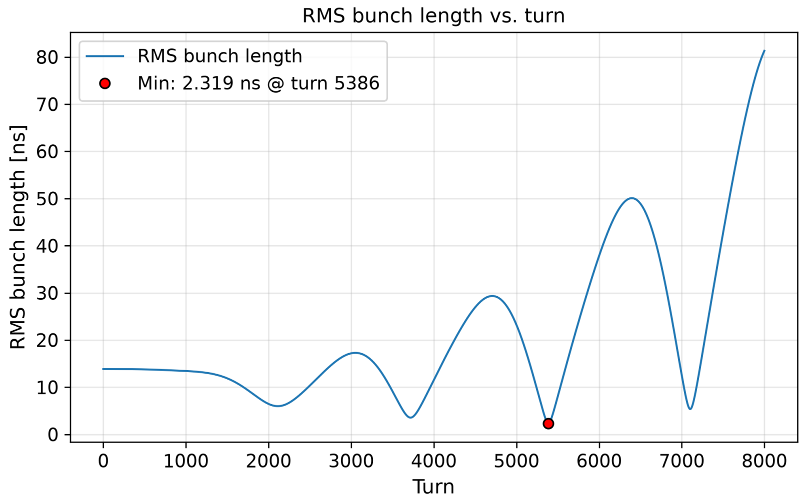

Figure 17: RMS bunch-length evolution for the best-performing resonant case, at a modulation depth of 0.90.

3.4 Comparison of Methods

The optimized performance of the three compression methods is collected in Table 3. The non-adiabatic voltage jump reached the shortest bunch, 1.97 ns, obtained at the lowest starting voltage scanned (5 kV) and the fastest jump (5 turns), and is the only method to meet the 2 ns target within the parameter ranges explored here. The unstable fixed-point and resonant methods followed, at 2.29 and 2.32 ns respectively.

The best-performing non-adiabatic case occurs at the edge of the scanned parameter range, at the lowest starting voltage and shortest voltage jump considered. The result therefore demonstrates that the 2 ns target is achievable within the simulation, but should not be interpreted as a fully optimized prediction of the achievable bunch length.

Table 3: Optimized performance of each compression method. All values are RMS bunch lengths at the minimum of the evolution.

Method	Optimized parameter	Min. σ_t (ns)	Turn of min.
Non-adiabatic	$V_{\text{start}} = 5$ kV, jump $= 5$ turns	1.97	660
Resonant	Modulation depth $= 0.90$	2.32	5386
Unstable fixed point	Dwell $= 600$ turns	2.29	1803

4 Discussion

The comparison demonstrates that non-adiabatic manipulation can provide substantially stronger compression than an adiabatic RF ramp. Of the methods studied, the non-adiabatic voltage jump produced the shortest bunch and was the only case to reach the 2 ns target. However, its optimum occurred at the lowest starting voltage and fastest transition considered, where the idealized voltage program is least representative of realistic AGS RF operation. The unstable fixed-point and resonant methods produced slightly longer bunches but provide alternative compression mechanisms with different RF and timing requirements.

The absolute bunch lengths should be interpreted as benchmarks rather than predictions of achievable AGS performance. The comparison used a common longitudinal emittance of 1.35 eV s and a uniformly filled elliptical ($J = 1/2$) initial distribution. Subsequent comparison with measured AGS bunch profiles identified $J = 3/2$ as a more representative distribution, while longitudinal stability studies indicate that the assumed emittance should also be refined. In addition, the model does not yet include collective effects or realistic RF system limitations. These effects should be incorporated before drawing quantitative conclusions about achievable machine performance.

5 Conclusion

A Python-based longitudinal beam dynamics simulation of the AGS was developed and used to evaluate adiabatic, non-adiabatic, unstable fixed-point, and resonant proton bunch compression. Under common baseline conditions, the non-adiabatic voltage jump produced the strongest compression, reaching 1.97 ns, while the unstable fixed-point and resonant methods reached 2.29 ns and 2.32 ns, respectively. These results establish comparative benchmarks for the compression methods and demonstrate that bunch lengths near the 2 ns target are possible within the simulated single-particle dynamics.

Measured AGS wall current monitor profiles were also used to refine the initial beam model, identifying a $J = 3/2$ longitudinal distribution as more representative than the uniformly filled ellipse used in the primary compression studies. Together with the longitudinal stability analysis, this indicates that the initial distribution and emittance should be updated before the simulated minima are interpreted as predictions of machine performance. Future work should therefore incorporate these revised beam conditions, realistic AGS RF constraints, and collective effects, followed by re-optimization and benchmarking against established longitudinal beam dynamics codes.

References

[1] J. P. Delahaye *et al.*, “Muon Colliders,” arXiv:1901.06150 [physics.acc-ph], 2019.

[2] M. Tanaka *et al.*, “The AGS Accelerator Complex with the New Fast Extraction System,” in *Proceedings of the 1996 Particle Accelerator Conference (PAC’96)*, Minneapolis, MN, USA, 1996.

[3] M. Bai *et al.*, "Adiabatic Excitation of Longitudinal Bunch Shape Oscillations," in *Proceedings of the 1999 Particle Accelerator Conference (PAC'99)*, New York, NY, USA, 1999.

[4] K.-Y. Ng, "Introduction to Collective Instabilities—Longitudinal and Transverse," Fermi National Accelerator Laboratory, FERMILAB-TM-2053, 1998.

Acknowledgments
This project was supported in part by the Brookhaven National Laboratory (BNL), Collider-Accelerator Department under the BNL Supplemental Undergraduate Research Program (SURP). Additional support was provided by Northeastern University's AJC Merit Research Scholars Program and the µCATALYST Summer Research Program. The author gratefully acknowledges Dr. Stephen Brooks for his mentorship and guidance throughout this project and Professor Toyoko Orimoto for her continued support and mentorship.