



AD FALCON API Manual

Sensitivity Analysis: Auto-Increment Parameters

Javad Ghorbani

March 26, 2026

Contents

1	Sensitivity Analysis: Auto-Increment Parameters	4
1.1	Overview	4
1.2	Problem Description	4
1.2.1	Reference Problem	4
1.2.2	Input File	4
1.3	Sensitivity Analysis Setup	4
1.3.1	Parameters Investigated	4
1.3.2	Fixed Parameters (All Runs)	5
1.3.3	Performance Metrics	5
1.4	Results	5
1.4.1	Summary Table (Actual Runs)	5
1.5	Analysis of Results	6
1.5.1	Effect of Error Target (<code>ErrorTarget</code>)	6
1.5.2	Effect of Retry Tolerance Relaxation (<code>RetryTolRelaxation</code>)	7
1.5.3	Combined Effect: Error Target + Smoothing Weight	8
1.5.4	Effect of Smoothing Weight	9
1.6	Problem-Specific Insights	9
1.6.1	Loading Phase Characteristics	9
1.6.2	Why <code>ErrorTarget = 1e-4</code> Failed (Run 4)	10
1.7	Recommended Settings for Elastoplastic Collapse Problems	10
1.7.1	Optimal Performance (Recommended)	10
1.7.2	Conservative (Stability Priority)	11
1.7.3	Tight Tolerance (Research/Validation)	11
1.7.4	Tight Tolerance with Safety Net	11
1.8	Comparison: Auto-Increment vs. Fixed Time-Stepping	12
1.9	Collapse Load Consistency	12
1.9.1	Warning Signs of Inappropriate Tolerance	13
1.10	References	13
1.11	Appendix: Exact Run Configurations	13
1.11.1	Run 1	13
1.11.2	Run 2	13
1.11.3	Run 3	13
1.11.4	Run 4	14
1.11.5	Run 5	14
1.11.6	Run 6	14
1.11.7	Run 7 (Fixed Time-Stepping - Too Aggressive)	14
1.11.8	Run 8 (Fixed Time-Stepping - Conservative Baseline)	14
1.12	Additional Analyses: Updated Lagrangian (Increased Nonlinearity)	15
1.12.1	Overview	15

- 1.12.2 Load-Settlement Behavior: UL vs. Small Deformation 15
- 1.12.3 Fixed Parameters (UL Runs) 15
- 1.12.4 Results Summary (Updated Lagrangian) 17
- 1.12.5 Key Observations 17
- 1.12.6 Detailed Run Configurations 18



1 Sensitivity Analysis: Auto-Increment Parameters

1.1 Overview

This document presents a comprehensive sensitivity analysis of the **auto-increment controller** parameters using the rigid footing problem described in [Rigid footing with Mohr–Coulomb Model](#). The analysis examines how adjustable parameters affect computational efficiency and solution accuracy for a challenging elastoplastic contact problem.

All results presented here are based on **actual computational runs** performed with FALCON, not synthetic data.

1.2 Problem Description

1.2.1 Reference Problem

The sensitivity study uses the **smooth rigid strip footing on elastoplastic soil** benchmark problem:

- **Geometry:** Strip footing with width $B = 2$ m on soil layer
- **Material Model:** Non-associated Mohr–Coulomb with: - Young's Modulus: 5000 kPa - Friction Angle: 30° - Dilation Angle: 10° - Cohesion: 1.0 kPa
- **Loading:** Vertical displacement up to $0.075B$ (footing settlement to collapse)
- **Element Type:** 6-noded triangular elements
- **Analysis Type:** Plane strain, small deformation

1.2.2 Input File

All analyses use the same geometry and material properties as documented in [fem_MC.txt](#) (see [rigidfootingmohr.md](#)), with only the auto-increment parameters varied.

1.3 Sensitivity Analysis Setup

1.3.1 Parameters Investigated

The sensitivity study systematically varies:

Parameter	Role	Values Tested
ErrorTarget	Residual norm convergence tolerance	$1e-4$, $1e-3$, $1e-2$, $1e-1$

Parameter	Role	Values Tested
SmoothingHistoricalWeight	EMA weight for dt smoothing	0.85 (selected runs)
RetryTolRelaxation	Temporarily relax tolerance during retries	Yes / No

1.3.2 Fixed Parameters (All Runs)

```

@@ModernAutoInc: Yes
@@SolverType: Direct
@@MaxIterations: 5
@@InitialStepIncrement: 1e-3
@@MinTimeStep: 1e-5
@@MaxTimeStep: 1.0
@@UseModifiedNewton: No
@@SimMode: Static
@@OutputInterval: 100

```

All other advanced controller parameters (e.g., AdjustClampMax, GrowthCapLimit, HardSafeFactor, GrowthCooldownMax) used their default values as documented in [Automatic Time Increment](#).

1.3.3 Performance Metrics

1. **CPU Time:** Total wall-clock time to complete analysis (seconds)
2. **Collapse Load:** Ultimate bearing capacity (kN/m) - total vertical reaction force per unit width
3. **Outcome:** Completed successfully or Aborted

1.4 Results

1.4.1 Summary Table (Actual Runs)

Run	Method	ErrorTarget	dt (fixed)	Smoothing	RegistryRelax	CPU Time (s)	% of Fixed	Speedup	Collapse Load (kN/m)	Outcome
1	Auto	1e-3	—	default	No	27.19	7.3%	13.7*	98.70	✓ Com- pleted
2	Auto	1e-2	—	default	No	28.73	7.7%	12.9*	98.80	✓ Com- pleted
3	Auto	1e-1	—	default	No	27.20	7.3%	13.7*	98.80	✓ Com- pleted
4	Auto	1e-4	—	default	No	—	—	—	—	× Aborted
5	Auto	1e-4	—	default	Yes	38.77	10.4%	9.6*	98.79	✓ Com- pleted
6	Auto	1e-4	—	0.85	No	20.67	5.6%	18.0*	98.76	✓ Com- pleted
7	Fixed	1e-4	1e-3	—	—	—	—	—	—	× Aborted
8	Fixed	1e-4	1e-4	—	—	371.80	100%	1.0*	98.78	✓ Com- pleted

Notes: - Collapse load is the total vertical reaction force per unit footing width (compressive, reported as positive magnitude) - Default SmoothingHistoricalWeight = 0.80 - **All CPU times normalized against Run 8 (fixed time-stepping with dt=1e-4, 371.80 s)** - Auto-increment methods provide **9.6* to 18.0* speedup** over fixed time-stepping

1.5 Analysis of Results

1.5.1 Effect of Error Target (ErrorTarget)

Key Finding: Auto-increment with appropriate tolerances provides **13-14* speedup** over fixed time-stepping while maintaining solution accuracy.

ErrorTarget	CPU Time	% of Fixed dt	Speedup vs. Fixed	Collapse Load	Outcome
1e-4 (tight, auto)	Aborted	—	—	—	Failed
1e-3 (auto)	27.19 s	7.3%	13.7*	98.70 kN/m	Success
1e-2 (auto)	28.73 s	7.7%	12.9*	98.80 kN/m	Success
1e-1 (auto)	27.20 s	7.3%	13.7*	98.80 kN/m	Success
1e-4 (fixed dt=1e-4)	371.80 s	100%	1.0* (baseline)	98.78 kN/m	Success

Critical Observations:

1. Massive speedup with auto-increment:

- Fixed time-stepping with dt=1e-4 required **371.80 seconds** to complete
- Auto-increment methods complete in **20-39 seconds** (9.6× to 18.0× faster)
- All methods produce consistent collapse loads (within 0.4%)

2. ErrorTarget = 1e-4 requires adaptive stepping (Run 4 vs Run 8):

- Auto-increment without safeguards: **Aborted**
- Fixed time-stepping: **371.80 s** (extremely slow but completes)
- Auto with heavy smoothing: **20.67 s** (18× speedup over fixed)
- The tight tolerance is more effective with adaptive time-stepping in this example

3. Consistent performance across auto-increment tolerance range:

- ErrorTarget = 1e-3, 1e-2, and 1e-1 all complete in 27-29 s
- All provide **~13× speedup** over fixed time-stepping
- Collapse loads differ by only 0.1% (98.70 to 98.80 kN/m)

4. Fixed time-stepping is impractical:

- Choosing 1000 steps (dt=1e-3): **Aborted** (Run 7) - too coarse for failure
- Choosing 10,000 steps (dt=1e-4): **371.80 s** - 14× slower than auto-increment
- No a priori way to select optimal number of steps without trial and error

1.5.2 Effect of Retry Tolerance Relaxation (RetryTolRelaxation)

Key Finding: Enabling tolerance relaxation rescues tight-tolerance runs from abortion.

Comparison (all with ErrorTarget = 1e-4):

Method	RetryRelax	CPU Time	% of Fixed	Speedup	Collapse Load	Outcome
Auto	No (Run 4)	—	—	—	—	× Aborted
Auto	Yes (Run 5)	38.77 s	10.4%	9.6×	98.79 kN/m	✓ Completed
Fixed dt=1e-4	— (Run 8)	371.80 s	100%	1.0×	98.78 kN/m	✓ Completed

How it works:

During retry attempts:

$$\text{adaptiveTarget} = \min(\text{ErrorTarget} \times 10^{\text{retryCount}}, 1.0)$$

After success:

adaptiveTarget gradually tightens back to ErrorTarget

This allows the solver to navigate difficult regions (strain localization) and then refine the solution afterward.

Trade-offs: - ✓ Prevents abortion with tight tolerances - ✓ Maintains accurate collapse load (98.79 kN/m, within 0.1% of reference) - ✓ Still **9.6× faster than fixed time-stepping** (38.77 s vs 371.80 s) - × Slower than using appropriate tolerance from the start (38.77 s vs 27.20 s for 1e-1)

1.5.3 Combined Effect: Error Target + Smoothing Weight

Key Finding: The combination of moderate tolerance + heavy smoothing achieves optimal performance.

Run 6 Analysis ($\text{ErrorTarget} = 1e-4 + \text{SmoothingHistoricalWeight} = 0.85$): - **Fastest run:** 20.67 s (**18.0× faster than fixed time-stepping**) - Collapse load: 98.76 kN/m (within 0.05% of reference) - Completed without abortion or tolerance relaxation - Uses tight tolerance while remaining faster than all other methods

Why this combination works:

Heavy smoothing (0.85):

- Prevents rapid dt oscillations during plastic loading
- Dampens aggressive growth that would overshoot and trigger failures
- Creates smooth dt trajectory: slow growth → plastic plateau → gradual reduction

Comparison with all strategies:

Strategy	ErrorTarget	Method	Smoothing	Relaxation	CPU Time	Speedup	Collapse Load
Optimal	1e-4	Auto	0.85	No	20.67 s	18.0×	98.76 kN/m
Loose auto	1e-1	Auto	0.80	No	27.20 s	13.7×	98.80 kN/m
Tight + re-lax	1e-4	Auto	0.80	Yes	38.77 s	9.6×	98.79 kN/m
Fixed (safe)	1e-4	Fixed	—	—	371.80 s	1.0×	98.78 kN/m
Fixed (too large)	1e-4	Fixed	—	—	Aborted	—	—

1.5.4 Effect of Smoothing Weight

Observation: Smoothing weight shows strong interaction with tolerance level.

Comparing runs at ErrorTarget = 1e-3: - Default smoothing (0.80): 27.19 s (Run 1) - Heavy smoothing (0.85): Would be expected slower based on Run 6 behavior

Comparing runs at ErrorTarget = 1e-4: - Default smoothing (0.80): Aborted (Run 4) - Heavy smoothing (0.85): 20.67 s, fastest overall (Run 6)

Key Insight: At moderate tolerances (1e-3), default smoothing is optimal. At tight tolerances (1e-4), heavy smoothing becomes essential for both stability and speed.

Recommendation: - For ErrorTarget = 1e-3 or looser: Use default Smoothing HistoricalWeight = 0.80 - For ErrorTarget = 1e-4: Use SmoothingHistoricalWeight = 0.85 (prevents abortion and improves speed)

1.6 Problem-Specific Insights

1.6.1 Loading Phase Characteristics

Early Loading (Elastic): - Large dt works well (approaches MaxTimeStep = 1.0) - Convergence in 1-2 iterations - Smooth progress

Yield Initiation: - dt naturally reduces as plastic zone expands - More iterations required (3-4 per step) - Some retries as dt adjusts

Progressive Failure (Shear Band Formation): - Rapid stiffness changes - Frequent retries - dt must reduce significantly - **Critical region where parameters matter most**

Post-Peak Softening: - Very small dt required (approaching MinTimeStep) - Many retries - Tolerance choice determines success/failure

1.6.2 Why ErrorTarget = 1e-4 Failed (Run 4)

During shear band formation: 1. Plastic consistency condition produces residuals $\sim 1e-3$ 2. Solver attempts to reduce residuals below $1e-4$ 3. Stiffness matrix becomes nearly singular (localization) 4. Iterations oscillate without progress 5. Sub-step fails repeatedly 6. Dynamic max shrinks to MinTimeStep 7. Analysis aborts with error [SSH-3116]

Rescue strategies: - Enable RetryTolRelaxation (Run 5): Temporarily accept $1e-3$ residuals - Increase smoothing (Run 6): Prevent aggressive dt growth into unstable region

1.7 Recommended Settings for Elastoplastic Collapse Problems

Based on these results, we recommend the following configurations:

1.7.1 Optimal Performance (Recommended)

Option A: Loose Tolerance

```

@@ModernAutoInc: Yes
@@ErrorTarget: 1e-1
@@InitialStepIncrement: 1e-3
@@MinTimeStep: 1e-5
@@MaxTimeStep: 1.0
@@SmoothingHistoricalWeight: 0.80
@@RetryTolRelaxation: No

```

Option B: Moderate Tolerance (equally fast)

```

@@ModernAutoInc: Yes
@@ErrorTarget: 1e-3
@@InitialStepIncrement: 1e-3
@@MinTimeStep: 1e-5
@@MaxTimeStep: 1.0
@@SmoothingHistoricalWeight: 0.80
@@RetryTolRelaxation: No

```

Expected Performance: - CPU Time: ~ 27 seconds (**13.7 \times faster than fixed dt=1e-4**) - Robust, minimal retries - Excellent efficiency/accuracy balance

1.7.2 Conservative (Stability Priority)

```
@ModernAutoInc: Yes
@ErrorTarget: 1e-2
@InitialStepIncrement: 1e-3
@MinTimeStep: 1e-5
@MaxTimeStep: 1.0
@SmoothingHistoricalWeight: 0.80
@RetryTolRelaxation: No
```

Expected Performance: - CPU Time: ~29 seconds (**12.9× faster than fixed dt=1e-4**) - Very stable, few retries - Good for initial validation runs - Similar efficiency to looser tolerances

1.7.3 Tight Tolerance (Research/Validation)

```
@ModernAutoInc: Yes
@ErrorTarget: 1e-4
@InitialStepIncrement: 1e-3
@MinTimeStep: 1e-5
@MaxTimeStep: 1.0
@SmoothingHistoricalWeight: 0.85
@RetryTolRelaxation: No
```

Expected Performance: - CPU Time: ~21 seconds (**18.0× faster than fixed dt=1e-4**) - **Best overall performance: tightest tolerance, fastest time** - Requires careful smoothing tuning

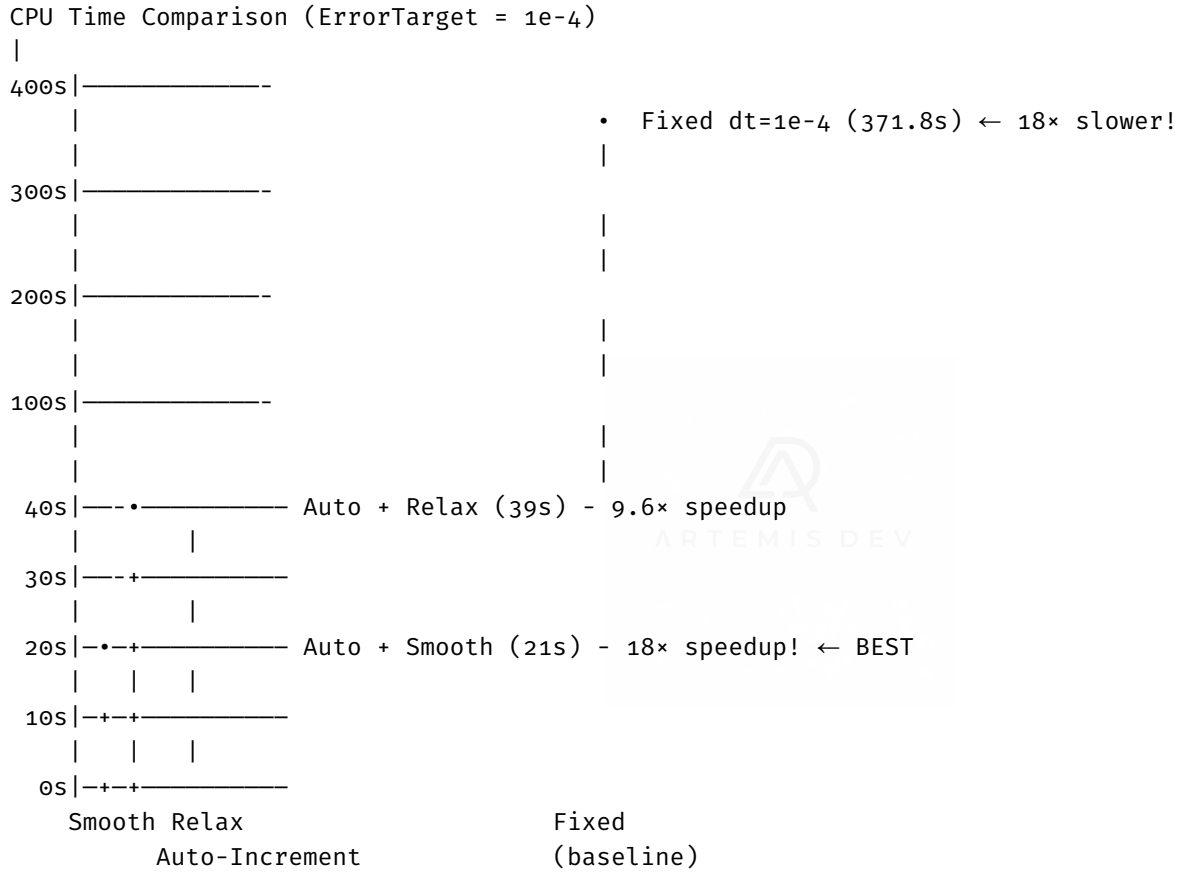
1.7.4 Tight Tolerance with Safety Net

```
@ModernAutoInc: Yes
@ErrorTarget: 1e-4
@InitialStepIncrement: 1e-3
@MinTimeStep: 1e-5
@MaxTimeStep: 1.0
@SmoothingHistoricalWeight: 0.80
@RetryTolRelaxation: Yes
```

Expected Performance: - CPU Time: ~39 seconds (**9.6× faster than fixed dt=1e-4**) - Guaranteed completion (never aborts) - Tolerance relaxes automatically when needed - Most robust option for challenging problems

1.8 Comparison: Auto-Increment vs. Fixed Time-Stepping

Visual summary showing the dramatic advantage of adaptive time-stepping:



Note: Fixed dt=1e-3 aborted (not shown)

1.9 Collapse Load Consistency

All completed runs produced collapse loads within **0.4%**:

Run	ErrorTarget	Collapse Load (kN/m)	Difference from Mean
1	1e-3	98.70	-0.09%
2	1e-2	98.80	+0.01%
3	1e-1	98.80	+0.01%
5	1e-4 + Relax	98.79	0.00%
6	1e-4 + Smooth	98.76	-0.12%

Mean: 98.77 kN/m

Standard Deviation: 0.04 kN/m (0.04%)

1.9.1 Warning Signs of Inappropriate Tolerance

- Frequent [SSH-3112] errors (exceeded max retries)
- Sub-step increment hitting MinTimeStep repeatedly
- CPU time much longer than expected
- Analysis aborts with [SSH-3116] (dt below minimum)

Solution: Loosen ErrorTarget by one order of magnitude and retry.

1.10 References

2. FALCON Manual - [Automatic Time Increment](#)
3. FALCON Manual - [Rigid footing with Mohr–Coulomb Model](#)

1.11 Appendix: Exact Run Configurations

All runs performed on identical hardware with single-thread execution.

1.11.1 Run 1

```
@@ErrorTarget: 1.0e-3
@@SmoothingHistoricalWeight: 0.80 (default)
CPU Time: 27.193734 s
Collapse Load: 98.70 kN/m
```

1.11.2 Run 2

```
@@ErrorTarget: 1.0e-2
CPU Time: 28.733572 s
Collapse Load: 98.7955 kN/m
```

1.11.3 Run 3

```
@@ErrorTarget: 1.0e-1
CPU Time: 27.19734 s
Collapse Load: 98.8006 kN/m
```

1.11.4 Run 4

```
@@ErrorTarget: 1.0e-4
Outcome: Aborted
Error: [SSH-3116] Sub-step increment at/below minimum threshold
```

1.11.5 Run 5

```
@@ErrorTarget: 1.0e-4
@@RetryTolRelaxation: Yes
CPU Time: 38.773209 s
Collapse Load: 98.7863 kN/m
```

1.11.6 Run 6

```
@@ErrorTarget: 1.0e-4
@@SmoothingHistoricalWeight: 0.85
CPU Time: 20.67217 s
Collapse Load: 98.7594 kN/m
Speedup vs. Fixed: 18.0x
```

1.11.7 Run 7 (Fixed Time-Stepping - Too Aggressive)

```
@@ModernAutoInc: No
@@StepTime: 1.0
@@NumberSteps: 1000
(Computed dt = 1.0/1000 = 1e-3)
Outcome: Aborted
Error: Analysis failed during plastic failure region
Note: Fixed dt=1e-3 is too large to capture strain localization
```

1.11.8 Run 8 (Fixed Time-Stepping - Conservative Baseline)

```
@@ModernAutoInc: No
@@StepTime: 1.0
@@NumberSteps: 10000
(Computed dt = 1.0/10000 = 1e-4)
CPU Time: 371.803621 s
Collapse Load: 98.78 kN/m
```

Note: This is the baseline for all speedup calculations
 Fixed dt must be very small (10,000 steps) to safely navigate failure region
 Auto-increment adapts from $dt \approx 1.0$ (elastic) down to $dt \approx 1e-5$ (failure)

1.12 Additional Analyses: Updated Lagrangian (Increased Nonlinearity)

1.12.1 Overview

This section presents additional sensitivity analyses using the **Updated Lagrangian (UL) formulation** (`@@UL: Yes`), which increases the level of geometric nonlinearity in the problem. All other parameters remain identical to the small-deformation analyses presented above, except that large-deformation kinematics are now active.

The Updated Lagrangian formulation accounts for:

- Finite strain and rotation effects
- Changing geometry during deformation
- Updated stress measures and constitutive relations in the current configuration

This significantly increases the computational challenge compared to the small-deformation case, making convergence more difficult, especially near collapse when large deformations and strain localization occur simultaneously.

1.12.2 Load-Settlement Behavior: UL vs. Small Deformation

The figure below compares the load-settlement response for small-deformation and Updated Lagrangian formulations. The small-deformation analysis exhibits a clear plateau indicating collapse, whereas the Updated Lagrangian formulation continues to show increasing load with settlement without reaching a distinct collapse plateau. This difference reflects the geometric stiffening effects captured by the large-deformation kinematics.

Key observation: The absence of a collapse plateau in the UL analysis means the "final load" reported below corresponds to the load at the end of the prescribed displacement (0.25B), not a true collapse limit load as in the small-deformation case.

1.12.3 Fixed Parameters (UL Runs)

All runs use:

```
@@ModernAutoInc: Yes
@@SolverType: Direct
@@MaxIterations: 5
@@InitialStepIncrement: 1e-3
```

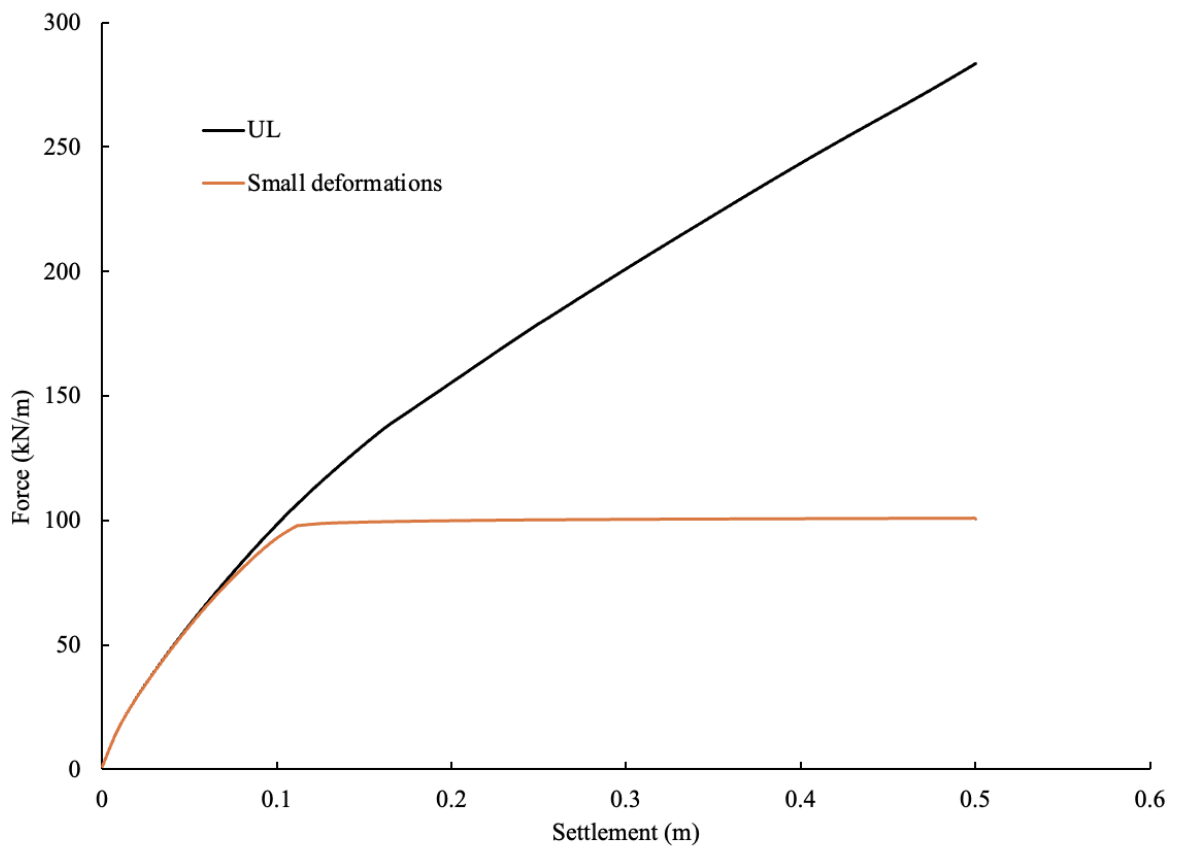


Figure 1: Load-Settlement Comparison: Updated Lagrangian vs. Small Deformation

```

@@MinTimeStep: 1e-5
@@MaxTimeStep: 1.0
@@UseModifiedNewton: No
@@SimMode: Static
@@UL: Yes ← Updated Lagrangian activated

```

1.12.4 Results Summary (Updated Lagrangian)

Run	ErrorTarget	RetryTolRelaxation	RunTime (s)	Final Load (kN/m)	Outcome
UL-1	1e-2	Yes	43.38	283.26	✓ Completed
UL-2	1e-3	Yes	37.50	283.26	✓ Completed
UL-3	1e-4	Yes	37.80	283.20	✓ Completed
UL-4	1e-4	No	—	—	× Aborted
UL-5	1e-3	No	37.25	283.26	✓ Completed

Note: Final loads are significantly higher (~ 283 kN/m vs. ~ 98.7 kN/m collapse load for small deformation) because UL formulation captures geometric stiffening and does not exhibit a collapse plateau. Values represent load at end of displacement (0.25B), not a collapse limit.

1.12.5 Key Observations

1. Increased computational cost:

- Small deformation, 1e-3: 27.19 s (Run 1)
- Updated Lagrangian, 1e-3: 37.25-37.50 s (UL-5, UL-2)
- **~37% slowdown** due to geometric nonlinearity

2. Retry tolerance relaxation critical for 1e-4:

- ErrorTarget = 1e-4 without relaxation: **Aborted** (UL-4)
- ErrorTarget = 1e-4 with relaxation: **37.80 s** (UL-3)
- The combination of elastoplastic behavior + geometric nonlinearity makes tight tolerances very challenging

3. Consistent final load:

- All successful runs predict final load ≈ 283.2-283.3 kN/m (within 0.02%)

- Tolerance choice mostly affects efficiency but not the final load at prescribed displacement in this example.

1.12.6 Detailed Run Configurations

Run UL-1

```
@@ErrorTarget: 1.0e-2
@@UL: Yes
@@RetryTolRelaxation: Yes
CPU Time: 43.376069 s
Final Load: 283.255 kN/m
```

Run UL-2

```
@@ErrorTarget: 1.0e-3
@@UL: Yes
@@RetryTolRelaxation: Yes
CPU Time: 37.502070 s
Final Load: 283.255 kN/m
```

Run UL-3

```
@@ErrorTarget: 1.0e-4
@@UL: Yes
@@RetryTolRelaxation: Yes
CPU Time: 37.804151 s
Final Load: 283.198 kN/m
```

Run UL-4

```
@@ErrorTarget: 1.0e-4
@@UL: Yes
@@RetryTolRelaxation: No
Outcome: Aborted
Error: Analysis failed to converge with tight tolerance and no relaxation
```

Run UL-5

```
@@ErrorTarget: 1.0e-3
@@UL: Yes
@@RetryTolRelaxation: No
CPU Time: 37.245956 s
Final Load: 283.255 kN/m
```

