

Speed and Accuracy Tests of the Variable-Step Störmer-Cowell Integrator

Matthew M. Berry* Liam M. Healy†

Abstract

The variable-step Störmer-Cowell integrator is a non-summed, double-integration multi-step integrator derived in variable-step form. The method has been implemented with a Shampine-Gordon style error control algorithm that uses an approximation of the local error at each step to choose the step size for the subsequent step. In this paper, the variable-step Störmer-Cowell method is compared to several other multi-step integrators, including the fixed-step Gauss-Jackson method, the Gauss-Jackson method with s -integration, and the variable-step single-integration Shampine-Gordon method, in both orbit propagation and orbit determination. The results show the variable-step Störmer-Cowell method is comparable with Gauss-Jackson using s -integration, except in high drag cases where the variable-step Störmer-Cowell method has an advantage in speed and accuracy.

Introduction

Use of numerical integration in space surveillance has grown in recent years as accuracy requirements have increased. Numerical integration requires a great deal of computation time compared to the analytic propagators previously used. An upgrade planned for the Navy's Space Surveillance System (known as the Fence) will greatly increase the number of objects being tracked, and hence significantly increase the amount of computation time required. Numerical integration methods requiring less computation time than those currently employed while maintaining or improving accuracy requirements are needed to reduce this burden.

A variable-step version of the Störmer-Cowell integrator has recently been developed (Ref. 1; Ref. 2). A study of currently available integration methods indicates that this new method has features desirable for elliptical orbits, and for orbits experiencing significant atmospheric drag (Ref. 2). The Störmer-Cowell integrator is a non-summed double-integration multi-step integrator. Multi-step integrators integrate forward using information from several previous backpoints, and are also known as predictor-corrector methods. Multi-step integrators are generally faster than single-step integrators, which use information from only a single point, because multi-step methods have fewer evaluations per step, though if a single-step method uses a larger step size this advantage is lost. Double-integration methods integrate second-order differential equations directly, so position is found directly from acceleration without first computing velocity. On the other hand, single-integration methods, known as the Adams methods, integrate first-order differential equations so they must be applied twice, first integrating to find velocity and then integrating velocity to find position. Double-integration methods are generally more accurate than single-integration methods, because removing the velocity calculation reduces the round-off error. In addition, with multi-step integration, double-integration methods are more stable than single-integration methods, and only require one evaluation per step, so double-integration is faster than single-integration. Multi-step integrators have both a non-summed and a summed form, depending on whether a summation term is used in the derivation. The non-summed form of double-integration is called Störmer-Cowell

*Astrodynamics Engineer, Analytical Graphics, Inc, 220 Valley Creek Blvd, Exton, PA 19341. E-mail: mberry@agi.com.

†Research Physicist, Naval Research Laboratory, Code 8233, Washington, DC 20375-5355. E-mail: Liam.Healy@nrl.navy.mil.

integration, while the summed form is called Gauss-Jackson integration. The summed form reduces the round-off error, so is generally more accurate than the non-summed form.

Multi-step integrators can be expressed in both fixed-step and variable-step forms. Variable-step integrators are advantageous over fixed-step methods for elliptical orbits, because variable-step methods can take fewer steps near apogee, where the satellite is moving slowly. Variable-step methods are combined with an error control algorithm to achieve this efficiency. However, variable-step methods have increased overhead costs, so may be slower than fixed-step methods for circular orbits. The fixed-step forms are often derived in terms of backward differences of the backpoints, which offer a simple form for the derivation. However, backward differences require that the backpoints be equally spaced, so the step size must remain constant. Derivations using backward differences of the Adams, Störmer-Cowell, and Gauss-Jackson methods in fixed-step form can be found in several references including Ref. 3, and codes using these methods are in wide use.

Variable-step forms of multi-step integrators can be derived by using divided differences instead of backward differences. Divided differences do not require that the backpoints be equally spaced. A variable-step version of the non-summed Adams method, known as the Shampine-Gordon method, is derived in Ref. 4, which also supplies code to implement the method. The method estimates the local error at each step by taking the difference of correctors of different orders. The size of the next step is then adjusted based on the local error estimate to meet a given tolerance. However, the method is biased towards keeping the step size constant, because there is an overhead cost associated with changing the step size. The step size is only increased when the error control algorithm indicates that it can at least be doubled. The Shampine-Gordon integrator is also variable order, in addition to being variable step. The order is adjusted at each step to maintain the error within the tolerance. Because the method is variable order it is self-starting, unlike other multi-step methods which need a separate startup procedure to get the initial set of backpoints.


The variable-step Störmer-Cowell method was derived in Ref. 1 to combine the advantages associated with double-integration and variable-step integration. The method was derived by following the derivation of the Shampine-Gordon method and applying that derivation to double-integration. The method uses the same concept of adjusting the step size based on local error estimates. However, the implementation of the variable-step Störmer-Cowell differs from that of the Shampine-Gordon method in several ways. The variable-step Störmer-Cowell method does not employ the restriction of only doubling the step size because the expense of a function evaluation for a sophisticated orbit model outweighs the overhead cost of changing the step size. Therefore, the method is biased toward changing the step size. Furthermore, the variable-step Störmer-Cowell is not variable-order, because the variable-order algorithm requires two full evaluations per step. However, the method does employ a variable-order algorithm as its startup procedure, using two evaluations per step until the initial set of backpoints is found.

Another method commonly used to handle elliptical orbits is to change the independent variable using a generalized Sundman transformation (Ref. 1). The transformation changes the independent from time, t , to a new variable s , using the relationship


$$dt = cr^n ds. \tag{1}$$

With this transformation the steps are spread more evenly about the orbit, giving an analytic step regulation. The manner in which the points are spread about the orbit is dependent on the value of the exponent n . Figure 1 shows the distribution of fixed integration points for various values of n in an orbit with an eccentricity of 0.75. In the tests in this paper n is equal to 1.5. Integrating with the transformation is known as s -integration, while integrating with time as the independent variable is known as t -integration. With s -integration, two evaluations per step are necessary for stability. However, to save computation time the second evaluation can be a partial evaluation, where only the two-body force is re-evaluated (Ref. 1). In the tests in this paper the partial evaluation is used with s -integration so it has effectively the same run-time as a method with one evaluation per step.

In this paper, the variable-step Störmer-Cowell method is compared in accuracy and speed tests in both orbit propagation and orbit determination to several other multi-step integrators to determine when it has an advantage over the other methods. The method is compared to the Shampine-Gordon method, as well as the 8th-order Gauss-Jackson with both t -integration and s -integration. Though s -integration spreads the points about the orbit, it is used with the fixed-step method Gauss-Jackson



(a) Equal Mean Anomaly ($n=0$); Step Size at Perigee Is $\Delta\nu = 1.97$ Radians
 (b) Equal Eccentric Anomaly ($n=1$); Step Size at Perigee Is $\Delta\nu = 0.97$ Radians



(c) Equal Intermediate Anomaly ($n=3/2$); Step Size at Perigee Is $\Delta\nu = 0.60$ Radians
 (d) Equal True Anomaly ($n=2$); Step Size at Perigee Is $\Delta\nu = 0.39$ Radians

Figure 1: Points Separated by Equal Values of Various Orbit Angles, with 16 Steps in each Orbit

method, taking fixed steps in s , so there is no local error control. For orbit propagation, the orbits are tested for various eccentricities and perigee heights. The accuracy of the integrators can be controlled through a tolerance for the variable-step methods and through the step size for the fixed-step methods. At each eccentricity and perigee height, the tolerance or step size is found for each integrator to meet a given accuracy requirement. Then, with the integrators set with the tolerance or step size from the accuracy test, the computation time of the integrator is measured for a 30-day propagation. This test procedure compares the speed of the integrators when they are tuned to give equivalent accuracy. To test orbit determination, the computation time to perform orbit determination on 1000 objects with Gauss-Jackson using t -integration is found. The computation time is then found using both Gauss-Jackson with t -integration and the other integrators to perform orbit determination on objects with eccentricity where the orbit propagation tests indicate the other methods are faster than Gauss-Jackson. This test shows the advantage of using t -integration for circular orbits and variable-step methods for elliptical orbits compared to using only t -integration.

Prior Studies

Many studies have compared the use of different methods to integrate satellite orbits. Merson (Ref. 5) tested different integration methods on a variety of types of orbits, and found that the eighth-order Gauss-Jackson method, with only the predictor, to be most efficient for circular orbits. Merson recommended the Gauss-Jackson method with s -integration for highly elliptical orbits. However, Merson's report was prepared in 1974, before variable-step multi-step methods such as Shampine-Gordon were available.

Fox (Ref. 6) compared the fixed-step Gauss-Jackson method, Gauss-Jackson with s -integration, a variable-step variable-order Adams method, and several variable-step single-step methods for orbits with various eccentricities. Fox used the two-body problem in his tests so an exact measure of error was available. To simulate the effects of complex perturbations on computation time, Fox added unnecessary calculations to the evaluation to waste time. Like Merson, Fox found the Gauss-Jackson method to be best for low eccentricities. Fox found that for elliptical orbits with complicated force models s -integration is the best method, because it requires fewer evaluations than the single-step methods. The main drawback of Fox's tests is that the effect of perturbations on integration error was not considered.

More recently, Montenbruck (Ref. 7) performed a study comparing several single-step and multi-step methods, with both fixed and variable steps. Again, his study did not include perturbations. His tests indicated that some of the more recently developed single-step methods are competitive with multi-step methods in some cases.

A set of tests performed by Lundberg (Ref. 8) compared integrators similar to the methods considered here, as well as single-step integrators. Lundberg's study examined both the number of function evaluations and the total computation time, so he was able to draw conclusions about the overhead cost of the integrators. For instance, integrators that have fewer evaluations but longer run-times than another method, when the evaluation is not expensive, have a greater overhead cost. Lundberg concluded that the variable-step multi-step methods have a higher overhead cost than the variable-step single-step methods. However, both the fixed and variable-step multi-step methods were more efficient than the single-step methods in terms of function evaluations. In comparing the double integration methods to the single integration methods, Lundberg found the double integration methods to be more accurate for a given step size and order.

Though Lundberg's tests demonstrate many of the differences between these methods, the tests do not reveal the best integrator to use for satellite orbits with a full force model in all cases. Lundberg found the fixed-step methods to be more efficient for satellite orbits, but he only examined circular orbits. His study did not reveal the benefit of using variable-step methods for elliptical satellite orbits. Also, Lundberg's tests did not include drag, which has a significant effect on integration error.

Though Merson and Fox both recommended variable-step methods for elliptical orbits and fixed-step methods for circular orbits, they do not indicate when to switch between the two methods. Here, tests are performed at various eccentricities to show where the variable-step methods have an advantage over the fixed step method. The tests are also performed at different perigee heights to show the effect that drag has on the results.

Orbit Propagation Tests

To compare the integrators in orbit propagation, speed tests are performed with the integrators tuned to give equivalent accuracy for a set of test case orbits. The test cases have various eccentricities and perigee heights, and all have an inclination of 40° , a ballistic coefficient of $0.01 \text{ m}^2/\text{kg}$, and an epoch of 1999-10-01 00:00:00 UT. To measure accuracy, an error ratio is defined in terms of the RMS error of the integration (Ref. 5). First define position errors as

$$\Delta r = |r_{\text{computed}} - r_{\text{reference}}|. \quad (2)$$

The RMS position error can be calculated,

$$\Delta r_{\text{RMS}} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\Delta r_i)^2}. \quad (3)$$

The RMS position error is normalized by the apogee distance and the number of orbits to find the position error ratio,

$$\rho_r = \frac{\Delta r_{\text{RMS}}}{r_A N_{\text{orbits}}}. \quad (4)$$

The integrators are tuned to give an error ratio of 1×10^{-9} over three days. The Gauss-Jackson integrator, with both t - and s -integration, is tuned by adjusting the step size. The Shampine-Gordon and variable-step Störmer-Cowell integrators are tuned by changing the relative error tolerance. The reference used to compute the error in (2) is a 14th-order Gauss-Jackson method. Testing against a higher-order integrator has been shown to give a reasonable estimate of integration error (Ref. 9). To ensure that the 14th-order Gauss-Jackson is a good reference, its step-size is chosen to give an error ratio of 1×10^{-10} in a step-size halving test. In the step size halving test, the reference used in (2) is found with the same integrator, but with half the step size (Ref. 9).

Once the integrators have been tuned to meet the accuracy requirement, they are run in speed tests. In the speed test the runtime is measured to propagate the orbit 30 days. The speed ratio is then found for the variable-step methods, where speed ratio is the run-time of the fixed-step Gauss-Jackson divided by the run-time of the variable-step methods. The variable-step methods have an advantage when the speed ratio is above one.

In both the accuracy and speed tests perturbations are considered. The perturbation forces are 36×36 WGS-84 geopotential, the Jacchia 70 drag model (Ref. 10), and lunar and solar forces. Note that all of these perturbations are continuous forces. Solar radiation pressure, a perturbation that has discontinuities at eclipse boundaries, can cause numerical integration errors, especially when using multi-step integrators, because the discontinuity violates the assumption made in the formulation of the integrators that the forces are smooth and continuous (Ref. 11). To simplify our study this perturbation is not considered.

The tests are performed using the Special-K software suite (Ref. 12), developed by the Naval Research Laboratory, which is used operationally by Naval Network and Space Operations Command (NNSOC, formerly Naval Space Command). Operationally, the software uses the Gauss-Jackson integrator with t -integration for eccentricities below 0.25, and Gauss-Jackson with s -integration for eccentricities above 0.25. These tests are performed with a research version of the software, which has been modified to allow different integrators to be used.

Figures 2 – 5 show plots of speed ratio against eccentricity for perigee heights of 300 km, 400 km, 500 km, and 1000 km, respectively. Plots for s -integration, Shampine-Gordon, and the variable-step Störmer-Cowell method are all shown on each figure. The horizontal line on the plots represent a speed ratio of one, above which the variable-step methods are more efficient than the fixed-step methods. Tables 1 – 4 show the times from the 30-day propagation used to compute the speed ratios. The 300 km circular orbit decays in less than 30 days, so the times shown for that case are the times to integrate 20 days. The tables also show the step sizes and tolerances needed for the methods to give a position error ratio, given by (4), of 1×10^{-9} against the 14th-order Gauss-Jackson method.

Figure 2 shows that the variable-step Störmer-Cowell method has an advantage over s -integration for 300 km perigee orbits with eccentricities below 0.55. This result indicates that the variable-step




Figure 2: Speed Ratios to t -integration at 300 km Perigee




Figure 3: Speed Ratios to t -integration at 400 km Perigee




Figure 4: Speed Ratios to t -integration at 500 km Perigee




Figure 5: Speed Ratios to t -integration at 1000 km Perigee

Table 1: Comparisons for Perigee Height of 300 km

e	Step Size / Tolerance				Time for 30 Day Run (sec)			
	t	s	SG	vSC	t	s	SG	vSC
0	12	7	5×10^{-12}	3×10^{-12}	33.3	58.7	82.6	45.5
0.10	26	15	4×10^{-11}	5×10^{-11}	22.9	35.6	55.3	20.6
0.15	26	23	4×10^{-11}	5×10^{-11}	22.7	21.4	46.0	17.3
0.20	14	11	1×10^{-11}	2×10^{-11}	41.3	39.5	49.2	16.9
0.30	32	30	1×10^{-11}	4×10^{-11}	18.0	12.2	40.6	13.0
0.40	35	20	1×10^{-11}	2.5×10^{-11}	16.3	14.5	31.9	11.6
0.50	31	21	1×10^{-11}	3×10^{-11}	18.4	10.8	25.7	9.32
0.60	28	25	7×10^{-11}	6×10^{-11}	20.3	6.71	15.2	6.91
0.65	26	24	4×10^{-11}	4.5×10^{-11}	22.0	5.86	13.8	6.10
0.70	28	24	1×10^{-11}	6×10^{-11}	20.2	4.74	13.5	4.99
0.80	26	23	2×10^{-11}	8×10^{-11}	21.8	2.92	7.76	3.09
0.90	22	18	1×10^{-10}	1×10^{-10}	25.6	1.53	2.93	1.37
0.95	23	18	2×10^{-10}	1.5×10^{-10}	24.4	0.65	1.21	0.62

Table 2: Comparisons for Perigee Height of 400 km

e	Step Size / Tolerance				Time for 30 Day Run (sec)			
	t	s	SG	vSC	t	s	SG	vSC
0	29	20	2×10^{-11}	2×10^{-12}	20.6	32.1	89.9	31.4
0.05	36	34	7×10^{-11}	8×10^{-11}	16.5	17.5	58.8	20.4
0.10	41	40	2×10^{-10}	2×10^{-10}	14.5	13.9	42.4	16.7
0.15	39	36	1×10^{-10}	6×10^{-11}	15.1	13.8	37.6	15.8
0.20	38	34	2×10^{-11}	7×10^{-11}	15.3	13.2	40.1	13.7
0.25	38	34	7×10^{-11}	5×10^{-11}	15.2	12.1	30.6	12.7
0.30	39	34	2×10^{-11}	6×10^{-11}	14.7	11.0	32.8	11.2
0.40	35	32	1×10^{-11}	4×10^{-11}	16.3	9.29	29.6	10.3
0.50	33	30	1×10^{-11}	3×10^{-11}	17.2	7.69	23.8	8.83
0.55	35	25	7×10^{-12}	2×10^{-11}	16.2	7.85	22.2	8.20
0.60	31	28	1×10^{-11}	6.5×10^{-11}	18.3	6.03	18.2	6.47
0.65	31	27	2×10^{-11}	8×10^{-11}	18.2	5.23	14.1	5.42
0.70	29	26	1×10^{-11}	6×10^{-11}	19.4	4.41	13.0	4.75
0.80	26	23	1×10^{-11}	8×10^{-11}	21.8	2.92	8.16	3.13
0.90	25	22	7×10^{-11}	1.5×10^{-10}	22.5	1.24	2.89	1.21
0.95	25	22	2×10^{-10}	1×10^{-10}	22.5	0.56	1.15	0.54

Table 3: Comparisons for Perigee Height of 500 km

e	Step Size / Tolerance				Time for 30 Day Run (sec)			
	t	s	SG	vSC	t	s	SG	vSC
0	45	40	5×10^{-11}	1×10^{-10}	13.3	16.1	75.0	25.8
0.10	44	39	8×10^{-11}	1.5×10^{-10}	13.7	14.2	43.3	15.9
0.15	42	37	1×10^{-10}	5×10^{-11}	14.0	13.4	35.2	15.1
0.20	42	37	2×10^{-11}	1×10^{-10}	13.8	12.3	37.6	12.7
0.30	40	35	6×10^{-11}	7.5×10^{-11}	14.4	10.6	26.8	10.9
0.50	35	32	2×10^{-11}	2.5×10^{-11}	16.2	7.23	20.6	8.61
0.55	34	31	1×10^{-11}	5×10^{-11}	16.7	6.49	19.9	7.24
0.60	34	30	4×10^{-11}	6.5×10^{-11}	16.7	5.71	14.5	6.25
0.70	31	28	2×10^{-11}	1×10^{-10}	18.2	4.12	11.4	4.42
0.80	29	26	2×10^{-11}	9×10^{-11}	19.4	2.61	7.16	2.89
0.90	26	24	2×10^{-10}	1×10^{-10}	21.7	1.15	2.51	1.27
0.95	27	24	5×10^{-10}	8×10^{-11}	20.9	0.53	1.05	0.59

Table 4: Comparisons for Perigee Height of 1000 km

e	Step Size / Tolerance				Time for 30 Day Run (sec)			
	t	s	SG	vSC	t	s	SG	vSC
0	68	60	1×10^{-10}	2×10^{-10}	8.83	10.8	49.1	16.9
0.10	65	56	1×10^{-10}	2×10^{-10}	9.19	9.78	29.9	11.0
0.15	62	55	1×10^{-10}	1×10^{-10}	9.40	9.04	25.7	10.4
0.20	60	53	1×10^{-10}	1.5×10^{-10}	9.65	8.57	23.1	9.33
0.30	57	50	1×10^{-10}	1×10^{-10}	10.1	7.48	19.8	8.41
0.50	51	46	1×10^{-11}	6×10^{-11}	11.2	5.08	18.3	6.59
0.60	48	44	5×10^{-11}	7×10^{-11}	11.8	3.93	11.7	5.21
0.70	46	41	3×10^{-11}	1.5×10^{-10}	12.3	2.88	8.98	3.63
0.90	38	34	1×10^{-10}	1×10^{-10}	14.9	0.87	2.30	1.14

Störmer-Cowell method handles the high-drag cases better than s -integration. Controlling the step size by local-error control allows the method to use appropriately small steps near perigee where drag is a factor, and still use large enough steps at apogee to decrease run-time. With analytic step regulation, s -integration cannot account for the fact that there is drag near perigee. For s -integration to give accurate results near perigee at the low altitude, it must use a step that is smaller than necessary at apogee. So this test case demonstrates the advantage of local error control over analytic step regulation. The results at the higher perigee heights indicate that s -integration does vary the step size appropriately when drag is a less significant factor. Figure 1(c) shows that s -integration does take smaller steps, in arc-length, at perigee than at apogee. The comparison between the variable-step Störmer-Cowell method and s -integration indicates that these smaller steps are appropriate to account for the increased geopotential perturbation forces near perigee, but are not small enough to handle high drag cases.

Figures 3 – 5 show that at higher perigees the variable-step Störmer-Cowell method and s -integration have roughly the same speed, though the variable-step Störmer-Cowell method is somewhat slower. The additional overhead associated with the variable-step algorithm may account for the difference between the run-times of the two methods. The variable-step Störmer-Cowell method and s -integration are always faster than the Shampine-Gordon integrator. The Shampine-Gordon method requires two evaluations per step, while the other two methods use only one full evaluation per step, so the Shampine-Gordon method should have twice the run-time as the other methods. However, the results shown in Tables 1 – 4 indicate that the Shampine-Gordon integrator has more than double the run-time as the other variable-step methods. The Shampine-Gordon method is biased toward keeping the step size constant, only increasing it when the step size can be doubled. This limitation accounts for the additional run-time of the Shampine-Gordon method beyond what is expected from the additional evaluation.

Figures 2 – 5 show that s -integration has an advantage over t -integration for orbits with eccentricities over approximately 0.15, and the Shampine-Gordon integrator is more efficient than the Gauss-Jackson method with t -integration for eccentricities over approximately 0.60. The figures indicate that these results are independent of perigee height. The variable-step Störmer-Cowell method also has an advantage over the Gauss-Jackson method for eccentricities over approximately 0.15 for perigee heights of 400 km and higher. For 300 km perigee height orbits, the variable-step Störmer-Cowell method has an advantage for eccentricities over approximately 0.10.

Orbit Determination Tests

Orbit determination testing is performed on a test set of catalogued satellites for 1999-09-29. There are 8003 objects in the catalog, of which 1000 are randomly selected for testing. The goal of the test is to find the improvement in computation time by using the variable-step methods where the orbit propagation tests show they are more efficient, as well as to validate that the variable step methods give comparable results to the fixed-step Gauss-Jackson.

The initial vectors from the catalog are fit using differential correction with a fitspan between 1.5 and 10 days. The fitspan, which is the time span from which observations are used in the fit, is determined by algorithms used by NNSOC in daily operations and depends on the mean motion and rate of change of mean motion. The fit includes observations up to 1999-10-01 00:00:00, going back through the length of the fitspan. Before the fit, the initial vector is propagated forward to the time of the last observation. The fit solves for position and velocity at the time of the last observation, and also solves for the ballistic coefficient when the perigee height is below 1200 km.

The fit is performed with a batch least-squares differential correction process. The differential correction process propagates the epoch state backwards to the time of each observation. At each observation time a residual is calculated, which is the difference between the actual observation and the observation computed from the propagation. These residuals are then used to correct the value of the state at epoch,

$$\delta\mathbf{x} = (A^TWA)^{-1} A^T\tilde{\mathbf{w}}, \quad (5)$$

where $\delta\mathbf{x}$ is the correction to the epoch state vector, $\tilde{\mathbf{w}}$ are the residuals, A is a matrix of partial derivatives, and W is a weighting matrix. The A matrix contains the partial derivatives of the observations with respect to the components of the epoch state vector. The W matrix weights observations based

on how well the sensors are known to perform. After the state is updated the process is repeated with residuals given by the updated state. This process repeats through several iterations until the epoch state converges. At each step the weighted RMS is calculated,

$$RMS = \sqrt{\frac{\tilde{\mathbf{b}}^T W \tilde{\mathbf{b}}}{N - 1}}, \quad (6)$$

where N is the number of observations. The process is converged when the percent change in weighted RMS from one iteration to the next is less than 1%, or if the change in weighted RMS is less than 1×10^{-5} . The weighted RMS is a unitless value, because the W matrix normalizes the residuals.

To get a baseline for computation time in the tests, the time is found to fit the 1000 test objects using Gauss-Jackson with t -integration. The test is performed on a 450 MHz Pentium II machine running Linux. The total user time is 11.2 hours. Of the 1000 objects, 913 update, while 87 do not update because they fail some criteria, such as having a final RMS that is too large, or not having enough observations for the object.

To test s -integration, the objects with eccentricities above 0.15 are fit with Gauss-Jackson using both t -integration and s -integration. The time using t -integration is 2.28 hrs, and the time using s -integration is 0.65 hrs. So s -integration is 3.5 times faster than t -integration for processing only the objects with eccentricities above 0.15. Operational algorithms are used to choose the step sizes for the methods, with both t -integration and s -integration having the same step at perigee. Objects with eccentricities over 0.25 use a step size of 30 seconds, and low-earth objects with lower eccentricities use a step size of 1 minute. Objects at higher altitudes use larger steps. There are 136 objects with eccentricities above 0.15, of which 102 update and 34 do not. Comparing the final states given by t - and s -integration, 71 of the objects have a final position difference of less than 1 m. The remaining 31 objects are shown in Table 5, which gives the final position difference in meters, and the difference in final weighted RMS. A negative value for RMS difference in the table indicates that the s -integration has a lower final weighted RMS.

The objects with the largest position difference in Table 5, 19622 and 21589, have a lower weighted RMS with s -integration, indicating that s -integration gives a better fit. The object with the next highest position difference, 3827, converged on a different iteration in the differential correction, and accepted a different number of observations. Though this position difference is relatively large, it is still within the accuracy of the observations and the force model. The remaining position differences are relatively minor, and well within the accuracy of the observations.

Using s -integration to perform the fits saves 1.63 hrs over t -integration. In the entire set of 1000 objects, if t -integration is used to fit the objects with eccentricities below 0.15 and s -integration is used to fit objects with eccentricities above 0.15, the total computation time would be 9.57 hrs, which is a 14.6% savings over using only t -integration.

To test the variable-step Störmer-Cowell integrator, the method is also used to fit objects with eccentricities over 0.15. A relative tolerance of 2×10^{-11} is used in the fits. One of the objects that did not update under t -integration or s -integration, object 21538, is removed from the set for the variable-step Störmer-Cowell method, because the variable-step algorithm uses a step size that is prohibitively small for the object. Under t -integration and s -integration the integrations become unstable for this object, generating an error that causes the differential correction to stop, so the object is not updated. Instead of becoming unstable, the variable-step method finds a step size small enough to perform the integration, but the step size is on the order of one millisecond, so the object can not update in a reasonable amount of time. This problem is not necessarily a flaw in the method, since the other methods cannot handle this object either. However, additional error handling to prevent situations like this one from occurring are necessary before the variable-step Störmer-Cowell method can be used operationally.

Of the remaining 135 objects with eccentricities over 0.15, 105 of the objects update with the variable-step Störmer-Cowell method, while 30 do not. Three objects that failed to update with t -integration and s -integration, 14136, 18719 and 25539, are able to update with the variable-step method. Characteristics of these objects, and their reason for failing in t -integration, are show in Table 6. Two of the objects have perigee heights below 300 km, where drag is significant. The other object has a higher perigee height of 382 km, but with ballistic coefficient $0.43\text{m}^2/\text{kg}$ can still be considered a high-drag case. These results indicate that the variable-step method is able to handle drag better than the other methods.

Table 5: Orbit Determination Differences for t - vs. s -integration

Satellite Number	Position Difference (m)	RMS Difference (s -int - t -int)
3827	112.778	0.0007
10960	2.70625	0.0000
13970	4.65988	-0.0002
19622	378.819	-0.0912
19884	1.06457	0.0000
19994	7.05259	-0.0007
19998	1.13009	-0.0004
21589	264.48	-0.0555
21591	4.22768	0.0008
21709	9.7989	0.0000
22020	13.9651	0.0000
22098	18.3051	0.0000
22238	7.41498	0.0000
22633	10.5869	0.0000
22997	1.71619	0.0007
23229	5.68604	0.0000
23332	15.019	0.0000
23403	7.89305	0.0000
23430	9.05374	0.0000
23460	15.5898	0.0000
23523	7.00024	0.0000
23616	4.74221	-0.0001
23950	3.42605	0.0000
24211	1.82044	-0.0001
24293	10.1841	0.0000
24655	14.9491	0.0000
24764	17.3145	-0.0001
25503	4.20061	0.0000
25542	18.3253	0.0000
25552	2.91824	0.0000
25805	2.00252	0.0000

With local error control the variable-step method can take smaller steps when drag is a factor. Of the 102 objects that update under both t -integration and the variable-step Störmer-Cowell method, 85 have final position differences under 1 m. The remaining 17 objects are shown in Table 7. A negative value for RMS difference indicates the fit is better with the variable-step method.

Table 6: Objects Updated by var. Störmer-Cowell But Not t -integration

Satellite	Reason for Failing in t -integration	e	Perigee Height (km)	Ballistic Coefficient (m^2/kg)
14136	Final B-term too large	0.72	272	0.0018
18719	Final RMS too large	0.51	161	0.045
25539	Final B-term too large	0.72	382	0.43

Table 7: Orbit Determination Differences for t -integration vs. var. Störmer-Cowell

Satellite Number	Position Difference (m)	RMS Difference (vSC – t -int)
5977	5.07898	0.0001
8195	1.4376	0.0000
9911	1.11593	0.0000
10946	1.01612	0.0011
10960	3.03786	0.0002
11007	231.851	-0.0407
14131	52.4079	-0.0317
19622	380.928	-0.0909
19807	1.56643	0.0000
19994	10.3783	-0.0011
21589	259.809	-0.0545
21590	30.3005	-0.2210
22020	1.48802	0.0000
22997	2.87322	0.0010
23177	1.20855	-0.0054
23824	115.035	-0.0896
23950	466.252	0.6853

The object with the largest position difference, 23950, has a larger final weighted RMS with the variable-step method. This object has only 25 observations within its fitspan. The differential correction converges on the 16th iteration for this object under the variable-step method. With t -integration the differential correction rejects one observation after the 15th iteration, and continues for six more iterations considering only 24 observations. So the difference between the fits with t -integration and with the variable-step Störmer-Cowell method for this object can be accounted for by the low number of observations for this object. All of the remaining objects with position differences greater than 10 m have lower final RMS values with the variable-step Störmer-Cowell method.

The variable-step Störmer-Cowell method takes 0.63 hours to process the objects, so it is 1.65 hours faster than t -integration. The method is 3.6 times faster than the Gauss-Jackson method with t -integration for processing only the objects with eccentricities above 0.15. Using the the variable-step Störmer-Cowell method for objects over 0.15 eccentricity and Gauss-Jackson with t -integration for the remaining objects would take 9.55 hours to process all 1000 objects, a 14.7% savings over only using t -integration for all of the objects. This savings is comparable to the savings with s -integration.

To test Shampine-Gordon, the objects with eccentricities above 0.6 are fit with both Gauss-Jackson using t -integration and Shampine-Gordon. A relative tolerance of 1×10^{-11} is used for the Shampine-Gordon integrator in the fits. Gauss-Jackson takes 1.64 hrs to process the 87 objects, while Shampine-Gordon takes 0.85 hrs. So Shampine-Gordon is 1.9 times faster than Gauss-Jackson for these objects.

The Gauss-Jackson method updates 67 of the objects, while the Shampine-Gordon integrator updates 66 of the objects. The Shampine-Gordon method fails to update object 10946 because the final RMS is too large, though the object updates with Gauss-Jackson. The final position differences after the fit between the two integrators is less than 1 m for 55 of the objects, the remaining 15 are shown in Table 8. A negative RMS difference in the table indicates that the weighted RMS is lower with Shampine-Gordon.

Table 8: Orbit Determination Differences for t -integration vs. Shampine-Gordon

Satellite Number	Position Difference (m)	RMS Difference (SG – t -int)
8195	2.15318	-0.0001
9911	2.47439	0.0000
12992	1.57563	0.0007
13999	1.28456	-0.0002
17078	1.56477	0.0000
19622	1.87689	0.0480
19884	1.3129	-0.0003
20649	1.2846	-0.0001
21589	264.477	-0.0577
22020	1.01124	0.0000
22068	1.65849	0.0002
22633	1.79357	-0.0001
22997	2.53152	0.0009
23824	115.065	-0.0877
24655	36.3626	0.5819

Again the objects with the largest position differences, 21589 and 23824, have lower weighted RMS values with Shampine-Gordon. The remaining objects have position differences that are relatively small, and well within the accuracy of the observations.

Using Shampine-Gordon on eccentricities over 0.60 saves 0.79 hrs over t -integration. If the entire set of 1000 objects is fit with Gauss-Jackson using t -integration for eccentricities below 0.60 and with Shampine-Gordon for eccentricities above 0.60, the total computation time would be 10.41 hrs, which is a 7.0% savings over using only Gauss-Jackson with t -integration.

Conclusions

The variable-step Störmer-Cowell method is compared to the variable-step Shampine-Gordon method and the Gauss-Jackson method using both the fixed-step t -integration and the analytic step regulation of s -integration. These tests show the advantage that s -integration and the variable-step methods have over fixed-step t -integration methods for elliptical orbits. Orbit propagation tests in which the integrators are set to give equivalent accuracy show that s -integration and the variable-step Störmer-Cowell method are faster than the Shampine-Gordon integrator, because Shampine-Gordon has an additional evaluation per step, and because Shampine-Gordon is biased toward keeping a constant step size. The variable-step Störmer-Cowell method is somewhat slower than s -integration in most cases, which indicates that the analytical step regulation that s -integration uses is giving appropriate step sizes. However, the variable-step Störmer-Cowell method has an advantage over s -integration at lower perigees, indicating that s -integration needs step sizes at apogee that are too small in order to have small enough steps at perigee to maintain accuracy when drag is a significant factor. The tests show that s -integration has an advantage over t -integration above an eccentricity of approximately 0.15, and the Shampine-Gordon integrator is more efficient than t -integration for eccentricities over approximately 0.60. These results are independent of perigee height. The variable-step Störmer-Cowell method also has an advantage over t -integration for eccentricities over approximately 0.15 for perigee heights of 400 km and higher. For 300 km perigee height orbits, the variable-step Störmer-Cowell method has an advantage for eccentricities over approximately 0.10.

Orbit determination tests show that s -integration and Shampine-Gordon give comparable results to t -integration when used at the eccentricities specified by the orbit propagation tests. A time improvement of 14.6% and 14.7% are achieved by using s -integration and the variable-step Störmer-Cowell method, respectively, at eccentricities above 0.15, and a time improvement of 7.0% is achieved by using Shampine-Gordon at eccentricities above 0.60. The variable-step Störmer-Cowell method is also able to update more objects than t -integration and s -integration, indicating that local error control is necessary to properly handle high drag cases.

Recommendations for Future Study

The results suggest that s -integration works well except when drag is a significant factor. To give larger steps at apogee while maintaining accuracy at perigee with drag present, a different value of n is needed in the general Sundman transformation. While Merson (Ref. 5), Nacozy (Ref. 13; Ref. 14), and others found that $n = 1.5$ works best in the transformation, they only investigated values of $n = 1, 1.5,$ and 2 . A study of other values of n is warranted to find the value that gives the best speed advantage, which may depend on perigee height and other factors. The variable-step Störmer-Cowell method, or a similar method, could be used in such a study. If a variable-step method is used with s as the independent variable, then the variable-step algorithm should never want to change the step size. The best value of n can be found by varying n and examining the change of the step size over the integration. The value of n that keeps the step size closest to constant is best. The results here indicate that a value close to $n = 1.5$ is most likely ideal for low drag, but a different value is needed with higher drag.

ACKNOWLEDGEMENTS

We thank Shannon Coffey at the Naval Research Laboratory for providing funding for this work. We thank Chris Hall for his comments and advice on the research, and we thank Vince Coppola and Jim Woodburn for comments on the paper.

REFERENCES

- [1] Berry, M. M., *A Variable-Step Double-Integration Multi-Step Integrator*, Ph.D. thesis, Virginia Tech, Blacksburg, VA, 2004, Available at: <http://scholar.lib.vt.edu/theses/available/etd-04282004-071227/>.
- [2] Berry, M. and Healy, L., “A Variable-Step Double-Integration Multi-Step Integrator,” *Advances in Astronautics*, American Astronautical Society, San Diego, CA, February 2004, AAS 04–238.
- [3] Maury, J. L. and Segal, G. P., “Cowell type numerical integration as applied to satellite orbit computation,” Tech. Rep. X-553-69-46, NASA, 1969, NTIS #N6926703.
- [4] Shampine, L. F. and Gordon, M. K., *Computer Solution of Ordinary Differential Equations*, W. H. Freeman and Company, San Francisco, 1975.
- [5] Merson, R. H., “Numerical Integration of the Differential Equations of Celestial Mechanics,” Tech. Rep. TR 74184, Royal Aircraft Establishment, Farnborough, Hants, UK, January 1975, Defense Technical Information Center number AD B004645.
- [6] Fox, K., “Numerical integration of the equations of motion of celestial mechanics,” *Celestial Mechanics*, Vol. 33, No. 2, June 1984, pp. 127–142.
- [7] Montenbruck, O., “Numerical Integration Methods for Orbital Motion,” *Celestial Mechanics and Dynamical Astronomy*, Vol. 53, 1992, pp. 59–69.
- [8] Lundberg, J., “Multistep Integration Formulas for the Numerical Integration of the Satellite Problem,” Tech. Rep. IASOM TR 81-1, Center for Space Research, The University of Texas at Austin, Austin, TX, April 1981.

- [9] Berry, M. and Healy, L., “Comparison of Accuracy Assessment Techniques for Numerical Integration,” *Advances in Astronautics*, American Astronautical Society, San Diego, CA, February 2003, AAS 03–171.
- [10] Jacchia, L. G., “New Static Models of the Thermosphere and Exosphere with Empirical Temperature Models,” Tech. Rep. 313, Smithsonian Astrophysical Observatory, 1970.
- [11] Woodburn, J., “Mitigation of the Effects of Eclipse Boundary Crossings on the Numerical Integration of Orbit Trajectories Using an Encke Type Correction Algorithm,” *AAS/AIAA Space Flight Mechanics Meeting, Santa Barbara, CA, 11–14 February 2001*, AAS/AIAA, AAS Publications Office, P. O. Box 28130, San Diego, CA 92198, 2001, Paper AAS 01-223.
- [12] Neal, H. L., Coffey, S. L., and Knowles, S., “Maintaining the Space Object Catalog with Special Perturbations,” *Astrodynamics 1997 Part II*, edited by F. Hoots, B. Kaufman, P. Cefola, and D. Spencer, Vol. 97 of *Advances in the Astronautical Sciences*, American Astronautical Society, San Diego, CA, August 1997, pp. 1349–1360, AAS 97–687.
- [13] Nacozy, P., “The intermediate anomaly,” *Celestial Mechanics*, Vol. 16, 1977, pp. 309–313.
- [14] Nacozy, P., “Time elements in Keplerian orbital elements,” *Celestial Mechanics*, Vol. 23, 1981, pp. 173–198.