IMPLICIT RUNGE-KUTTA METHODS FOR UNCERTAINTY PROPAGATION

Jeffrey M. Aristoff, Joshua T. Horwood, and Aubrey B. Poore

Numerica Corporation, 4850 Hahns Peak Drive, Suite 200, Loveland CO, 80538

ABSTRACT

Accurate and efficient orbital propagators are critical for space situational awareness because they drive uncertainty propagation which is necessary for tracking, conjunction analysis, and maneuver detection. Existing sigma point- or particle-based methods for uncertainty propagation use *explicit* numerical integrators for propagating the closely spaced orbital states as part of the prediction step of the nonlinear filter (e.g. the unscented Kalman filter, Gaussian sum filter, or particle filter). As such, these methods cannot exploit the proximity of the orbital states, and each orbit is propagated independently. To remove this limitation and enable the orbital states to be propagated together, we have developed an *implicit* Runge-Kutta-based method for uncertainty propagation, and consider the propagation of the 13 sigma points needed to represent uncertainty (of a six-dimensional Gaussian state) in the unscented Kalman filter. In some cases, we can propagate uncertainty using the new propagator at about the same computational cost compared to that of propagating a single orbital state, even before the algorithm is potentially parallelized. The new propagator is applicable to all regimes of space, and additional features include its ability to estimate and control the truncation error, exploit analytic and semi-analytic methods, and provide accurate ephemeris data via built-in interpolation.

1. INTRODUCTION

Systematic and robust uncertainty quantification is a necessary part of future space surveillance and reconnaissance efforts. New computational methods are needed to mitigate the cost of uncertainty propagation, one of the major computational demands in a tracking system. In a companion paper [1], we demonstrated the use of implicit Runge-Kutta methods for accurate and efficient *orbit propagation*. Here, we demonstrate the use of implicit Runge-Kutta methods for accurate and efficient *uncertainty propagation*. The use of implicit, rather than explicit/multistep numerical methods for orbit and uncertainty propagation marks a *paradigm shift* in astrodynamical algorithms.

The general framework for the problem of estimating the state of an object and its uncertainty (i.e. its probability density function) in the presence of uncertainty in the dynamics, model parameters, and initial conditions is that of the general Bayesian nonlinear filter [2–4]. The Kalman filter optimally characterizes the state and uncertainty in the linear Gaussian case [5] and the extended Kalman filter approximately characterizes the state in the Gaussian and (mildly) nonlinear case [2, 3]. Higher order filters can also be employed to improve the characterization of the state. The uncertainty propagation component of these filters uses the Ricatti equation or a higher order Ricatti equation (e.g. the covariance matrix associated with the extended Kalman filter is obtained as a solution of the Ricatti equation).

A different class of methods starts from the representation of the initial probability density function by a set of particles or states, propagates the ensemble of particles through the differential equation defining the dynamics over some period of time, and then recovers the characterization of uncertainty at the end time*. The particles can be chosen randomly through sampling techniques as in Monte Carlo methods (e.g. particle filters) [6, 7] or deterministically as in the unscented transform of the unscented Kalman filter [8, 9] or higher order methods based on Gauss-Hermite quadrature [10], or as in the new GVM filter of Horwood and Poore [11]. Gaussian sum filters and generalizations using mixtures can also utilize the unscented Kalman filter as the basis for propagating the individual probability density functions [10,12–14]. Grid methods [15] represent the uncertainty through the propagation of particles and the subsequent recovery of the uncertainty. Ultimately, either the Fokker-Planck equation (for the continuous time dynamics) or Chapman-Kolmogorov equation (for the discrete time dynamics)

*The propagation of a given state is equivalent to solving an initial value problem (IVP).

govern the evolution of the probability density function. In the former, this parabolic partial differential equation can be solved by the method of lines in which the spatial variables are discretized with the time variables remaining continuous, resulting in an ensemble of IVPs [2]. Finally, in addition to the general nonlinear Bayesian filter, the Probability Hypothesis Density filter and its generalization, the Cardinalized Probability Hypothesis Density filters [16] can make use of Monte Carlo methods with the corresponding propagation of an ensemble of particles. Thus, a large class of methods for propagating a state and its uncertainty require the propagation of an ensemble of particles or states through the nonlinear dynamics. Once recovered and characterized, the uncertainty can be used to update a nonlinear filter, to compute a likelihood function or ratio for data association [17, 18], to compute a probability of collision in conjunction analysis [19], to assist with the detection of anomalies such as maneuvers [4], and to provide state uncertainty for sensor resource management [20]. It is to this class of methods to which the present approach applies.

Traditional algorithms for propagating an ensemble of states are based on explicit numerical methods [7, 21]. As such, they solve each IVP independently from one another, regardless of the proximity of the initial conditions, by calculating the state of a system at a later time from the state of the system at an earlier time. Hence, even if two initial conditions differed by an infinitesimal amount and the solution to the first IVP is known because it had already been computed, current approaches would not be able to exploit this information in order to solve the second IVP in a more computationally efficient manner. Since the IVPs that arise in uncertainty propagation usually contain initial conditions that are closely spaced about some estimate of the current state, new approaches, like the one presented in this paper, are needed that can reduce the computational cost of uncertainty propagation by exploiting such proximity.

The organization of this paper is as follows. In Section 2 we introduce a new technique for uncertainty propagation, one that uses Gauss-Legendre implicit Runge-Kutta methods. In Section 3 we provide a proof-of-concept and compare the performance of the new propagator to existing methods for uncertainty propagation. Finally, Section 4 provides concluding remarks.

2. METHODS

Unlike explicit/multistep numerical methods for solving initial value problems (IVPs), implicit Runge-Kutta (IRK) methods solve an IVP by calculating the state of a system at a later time from the state of the system at the current time, together with the state of the system at future times. This gives IRK methods superior stability properties and allows them to take larger (and fewer) steps than taken by their explicit/multistep counterparts [22, 23]. However, a nonlinear system of equations must be solved on each step, meaning that more work is done per step[†]. While this is a disadvantage for implicit methods in general, for the two-body problem this disadvantage is largely mitigated by the availability of approximations to the solution [21, 24] that can be computed efficiently and used to warm-start the iterations [1, 25–27], thereby reducing the number of iterations needed for convergence and lowering the computational cost of orbit propagation.

The ability to provide an initial guess for the solution of the nonlinear system of equations that arise during orbit propagation can also be exploited in the context of uncertainty propagation, wherein multiple orbital states in close proximity must be propagated. Rather than propagating each particle or state independently, information from the first propagation can be used to reduce the computational cost of each remaining propagation. Specifically, the new propagator solves one of the IVPs using an adaptive Gauss-Legendre IRK (GL-IRK) method, and the resulting trajectory (i.e. the state of the object throughout the propagation) is recorded. The system of equations that arise on each time step of the GL-IRK method are solved using iterative methods. The remaining IVPs are solved using the same GL-IRK method and time steps, together with the recorded trajectory from the first propagation, which is used to hot-start the iterations. Note that the first propagation need not be complete before the remaining propagations are initiated. Rather, after the first step of the first propagation is accepted, the remaining propagations can take that the same step. Hence, the states can be propagated together.

There are three reasons why the remaining states can be propagated more efficiently using the present technique. First, because the recorded trajectory from the first propagation and the unknown trajectories are

[†]Depending on the implementation of IRK, the work required to solve the nonlinear system of equations can be distributed across multiple processors.

close in proximity, convergence of the iterations within the remaining propagations is faster than had a trivial guess for the trajectory been provided. Second, the error estimates may be re-used, and third, steps will not be rejected. Hence, the computational cost of propagating each of the remaining states is significantly reduced from that of propagating the first state. Uncertainty in the object's state can be characterized at a fraction of the cost. The decision-maker has more time to act, or alternatively, the decision-maker can use the tracking or resource management system to handle more complex scenarios in the same amount of time.

Three additional comments regarding the performance of the new propagator should be made. Since each of the remaining propagations is faster than the first propagation, the new propagator strives to propagate the first state as efficiently as possible. Details are provided in Aristoff and Poore [1]. Another observation is that the cost of uncertainty propagation on a per particle basis decreases with increasing particle number, provided that the particles remain in close proximity. Finally, the new propagator is adaptive in that the numerical (truncation) error arising on each step of the propagations is estimated and controlled by adjusting the time step. In doing so, the propagator does just the right amount of work for a given accuracy requirement, and accuracy can be traded for speed and visa-versa.

3. RESULTS

In this section, we analyze the performance of the new propagator by propagating objects in low Earth orbit (LEO), geosynchronous orbit (GEO), and highly elliptic orbit (HEO). We first demonstrate the extent to which the computational cost of uncertainty propagation can be reduced by exploiting the proximity of the states within an initial value problem ensemble. We then compare the performance of the new propagator to that of existing methods for uncertainty propagation. In each of the tests, truth is generated using a high-accuracy[‡] 50-stage Gauss-Legendre implicit Runge-Kutta (GL-IRK) method, and the number of high-fidelity force-model evaluations, the dominant computational cost of orbit propagation, is used to quantify the cost of orbit propagation. A degree and order 32 gravity model is used for objects in LEO and HEO, and a degree and order 8 gravity model is used for objects in GEO. On each step of a given method, we require that the relative error in the 2-norm of the position of the object fall below a user-specified relative tolerance. Hence, the propagation is adaptive, variable-sized time steps are taken in order to control the error per step. For the first propagation, J_2-gravity is used to warm-start the iterations for the stage equations. For the remaining propagations, the recorded trajectory from the first propagation is used to hot-start the iterations.

3.1 Proof of Concept

As proof of concept, we propagate the uncertainty in the state of an object in LEO, the orbital elements of which are given in Table 1 (Scenario 1). The orbit is propagated for 4.5 hours, or approximately 3 orbital periods. Figure 1 demonstrates, for a given global accuracy (i.e. error in final position), the cost of propagating the first state and each remaining state (the number of which depends upon the way in which uncertainty is characterized) using an adaptive GL-IRK method. A quadratic-least-squares-fit to each of the data sets is shown. By exploiting the proximity of the orbital states, the computational cost of propagating each remaining state is less than that of propagating the first state (solid curve). The dashed and dashed-dotted curves provide bounds on the cost of propagating each remaining state. At best (dashed curve), only a single iteration is necessary to solve the system of equations for the internal stages. At worst (dash-dotted curve), multiple iterations are required, but no more than were used in the first propagation over the corresponding time step. Regardless, step rejection does not occur, and the error estimates from propagating the first state can be reused. As a result, the overall cost of propagating each remaining state is less than that of propagating the first state, even in the worst case. Had proximity of the orbital states not been exploited, the cost of propagating each remaining state would equal that of propagating the first state (solid curve). Hence, the present technique reduces the overall cost of propagating uncertainty.

[‡]twelve digits of accuracy per time step

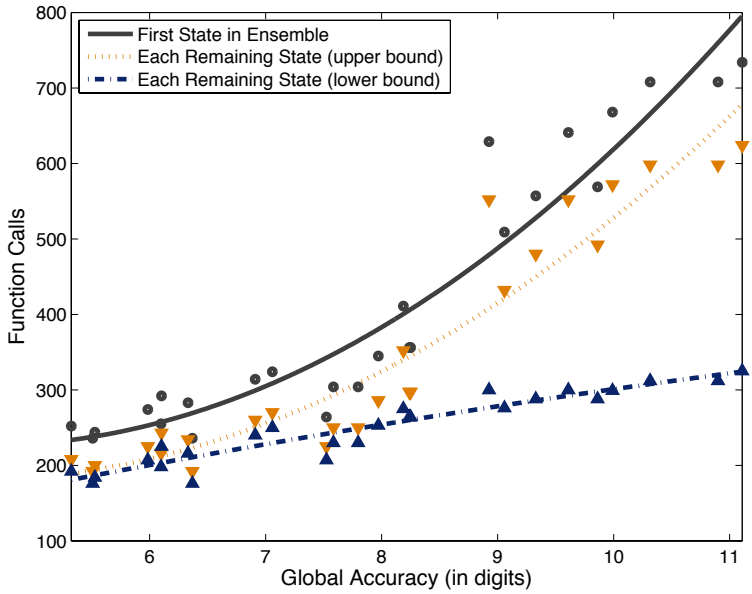

Figure 1. Computational cost of propagating closely spaced objects in low Earth orbit for a given accuracy using the new GL-IRK-based uncertainty propagator. The cost to propagate the first state in the ensemble exceeds that of each of the remaining states, bounds of which are given.

Scenario #	Orbit Type	a (km)	e	i (°)	Ω (°)	ω (°)	M (°)	orbital periods
1	LEO	6640	0.009500	72.9	116	57.7	105	3
2	LEO	6640	0.009500	72.9	116	57.7	105	30
3	GEO	42164	0	0	0	0	250	3
4	GEO	42164	0	0	0	0	250	30
5	HEO	26628	0.7416	63.4	120	0	144	3
6	Re-Entry	6518	0.0003875	53.0	145	267	94.0	3

Table 1. Initial orbital elements for the orbit propagation scenarios. The first five are taken from [24], the sixth from [28].

3.2 Performance of the New Orbital Uncertainty Propagator

Six realistic orbit propagation scenarios are used to quantify the performance of the new GL-IRK-based uncertainty propagator. The scenarios are replicas of those found in [1], and are listed in Table 1. Although the new technique for uncertainty propagation is applicable to other particle- or sigma-point-based filtering methods, we shall consider the propagation of uncertainty via the unscented Kalman filter (UKF), in which 13 sigma points (or states) are used to represent uncertainty (of a six-dimensional Gaussian state). As in [1], we use GL-IRK methods because these methods are A-stable, B-stable, parallelizable, super-convergent, symmetric, and symplectic. Gold-standard implementations[§] of Dormand-Prince 8(7) (DP8(7)) and Adams-Bashforth-Moulton (ABM), both adaptive step numerical integrators, are used for comparison. DP8(7) is an explicit Runge-Kutta method, and ABM is a multistep method. For additional information, see Aristoff and Poore [1].

In each scenario, we consider the relationship between the global accuracy and the computational cost in a serial computing environment. (An additional speedup of 5 to 15 times, corresponding to the number of internal stages, would be observed for GL-IRK in a parallel computing environment.) As in Section 3.1, rather than specifying an initial covariance and propagating the corresponding 13 sigma points, we deduce upper and lower bounds on the cost of uncertainty propagation based on the propagation of the first state in the ensemble. For reference, 7 digits of accuracy correspond to approximately one-meter accuracy for the object in LEO, and 7.5

[§]courtesy of MATLAB

digits of accuracy correspond to approximately one-meter accuracy for the object in GEO. For the object in HEO, 7 digits of accuracy correspond to approximately one-meter accuracy at perigee. A quadratic-least-squares-fit to each of the data sets is shown.

- **Scenario 1.** An object in low-altitude LEO, together with its uncertainty, is propagated for 3 orbital periods, or approximately 4.5 hours, using adaptive step-size control. The performance of the new GL-IRK-based propagator is summarized in Figure 2(a), and compared to that of DP8(7) and ABM. IRK is 70-80% more efficient than DP8(7) and ABM over the range of accuracies considered.

- **Scenario 2.** An object in low-altitude LEO, together with its uncertainty, is propagated for 30 orbital periods, or approximately 45 hours, using adaptive step-size control. The performance of the new GL-IRK-based propagator is summarized in Figure 2(b), and compared to that of DP8(7) and ABM. IRK is 75-80% more efficient than DP8(7) and ABM over the range of accuracies considered.

- **Scenario 3.** An object in GEO, together with its uncertainty, is propagated for 3 orbital periods, or approximately 3 days, using adaptive step-size control. The performance of the new GL-IRK-based propagator is summarized in Figure 2(c), and compared to that of DP8(7) and ABM. IRK is 90-95% more efficient than DP8(7) and 75-85% more efficient than ABM over the range of accuracies considered.

- **Scenario 4.** An object in GEO, together with its uncertainty, is propagated for 30 orbital periods, or approximately 30 days, using adaptive step-size control. The performance of the new GL-IRK-based propagator is summarized in Figure 2(d), and compared to that of DP8(7) and ABM. IRK is 90-95% more efficient than DP8(7) and 80-90% more efficient than ABM over the range of accuracies considered.

- **Scenario 5.** An object in HEO, together with its uncertainty, is propagated for 3 orbital periods, or approximately 36 hours, using adaptive step-size control. The performance of the new GL-IRK-based propagator is summarized in Figure 2(e), and compared to that of DP8(7) and ABM. IRK is 70-80% more efficient than DP8(7) and 55-70% more efficient than ABM over the range of accuracies considered.

- **Scenario 6.** A satellite that re-entered the atmosphere, known as ROSAT, the orbital elements of which are given in Peat [28], is propagated, together with its uncertainty, for 3 orbital periods, or approximately 4.5 hours, using adaptive step-size control. The performance of the new GL-IRK-based propagator is summarized in Figure 2(f), and compared to that of DP8(7) and ABM. Using the published weight of ROSAT (2400 kg) we estimate the area-to-mass ratio to be 4.7×10^{-4} m^2 kg^{-1}. This results in a nearly 1 km decay in the altitude of the spacecraft per orbital period when using a Harris-Priester drag model. IRK is 70-75% more efficient than DP8(7) and 75-80% more efficient than ABM over the range of accuracies considered.

4. CONCLUSIONS

This paper described a new technique for reducing the time or amount of computation required for solving an ordinary differential equation having an ensemble of initial conditions as arises in, for example, sequential state estimation, conjunction analysis, maneuver detection, and sensor resource management. Particular attention was given to the use of this technique within the prediction step of nonlinear filtering in the context of space situational awareness, i.e., for orbital uncertainty propagation.

Even before parallelization, the new adaptive-step Gauss-Legendre-implicit-Runge-Kutta-based uncertainty propagator was found to be significantly more efficient in our test scenarios than adaptive-step explicit and multistep methods often used for uncertainty propagation, specifically, Dormand-Prince 8(7) and Adams-Bashforth-Moulton (ABM). In some cases, we can propagate uncertainty (via the unscented Kalman filter) using the new propagator at about the same computational cost compared to that of propagating a single orbital state, even when restricted to a serial processor. Specifically, a 92.3% savings would reduce the cost of propagating 13 orbital states to the cost of propagating a single orbital state. Table 2 lists the computational savings (relative to ABM) for uncertainty propagation in LEO (Scenarios 1–2), GEO (Scenarios 3–4), HEO (Scenario 5), and

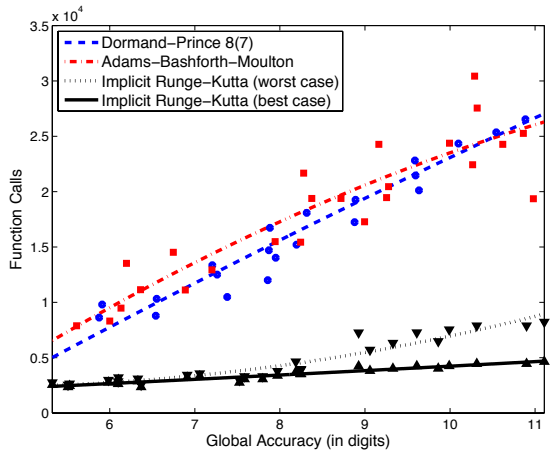

(a) Scenario 1: LEO

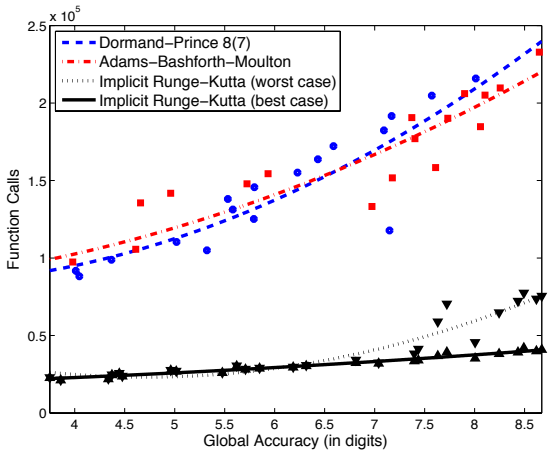

(b) Scenario 2: LEO

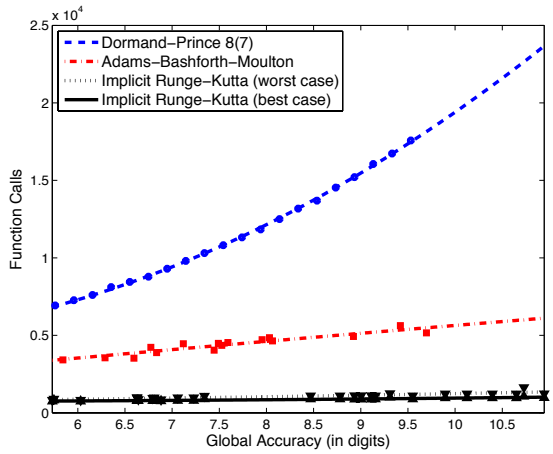

(c) Scenario 3: GEO

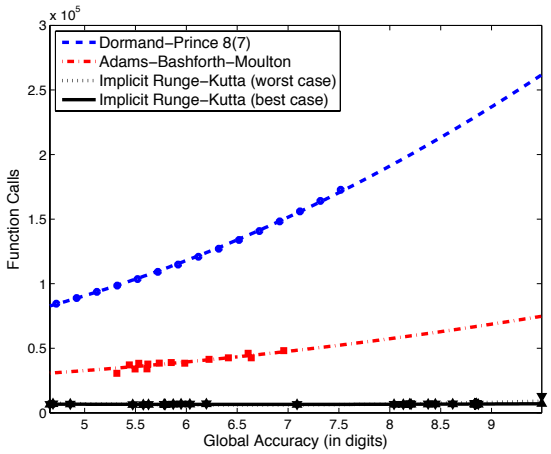

(d) Scenario 4: GEO

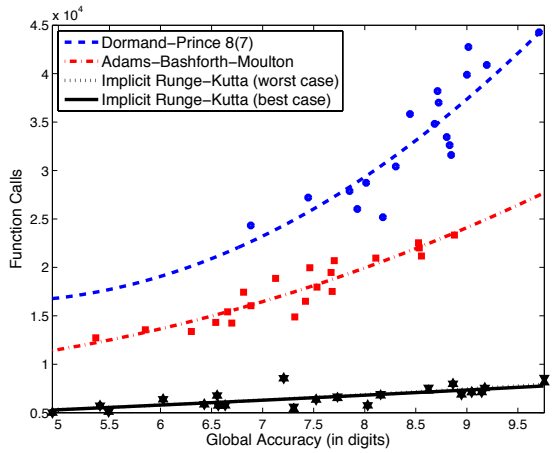

(e) Scenario 5: HEO

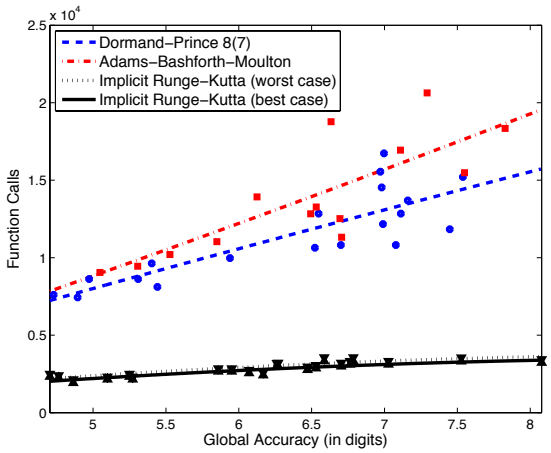

(f) Scenario 6: Re-Entry

Figure 2. Performance of the new Gauss-Legendre-implicit-Runge-Kutta-based uncertainty propagator vs. Dormand-Prince 8(7) (DP8(7)) and Adams-Bashforth-Moulton (ABM). Computational cost of uncertainty propagation via the unscented Kalman filter, measured by the number of high-fidelity force-model evaluations (function calls) needed to propagate the 13 sigma points, versus the accuracy in the final state of the mean sigma point (global accuracy) in a serial computing environment. An additional speed-up of 5-15 times would be observed for GL-IRK in a parallel computing environment. Note that the performance of ABM and DP8(7) is comparable in Scenarios 1, 2, and 6, and that ABM outperforms DP8(7) in Scenarios 3, 4, and 5.

	Scenario 1 LEO	Scenario 2 LEO	Scenario 3 GEO	Scenario 4 GEO	Scenario 5 HEO	Scenario 6 Re-Entry
SCE	70-80%	75-80%	75-95%	80-95%	55-80%	70-80%
PCE	94-99%	95-98%	95-98%	96-99%	91-99%	94-99%

Table 2. Summary of computational savings (relative to ABM) for *uncertainty propagation* in a serial computing environment (SCE) or a parallel computing environment (PCE). The computational savings tend to increase as the accuracy of the propagation increases or the initial uncertainty in the orbit decreases (i.e. the states become more closely spaced).

	Scenario 1 LEO	Scenario 2 LEO	Scenario 3 GEO	Scenario 4 GEO	Scenario 5 HEO	Scenario 6 Re-Entry
SCE	70-80%	60-70%	60-65%	75-85%	30-55%	65-75%
PCE	94-99%	92-98%	92-98%	95-99%	86-97%	93-98%

Table 3. Summary of computational savings (relatve to ABM) for *orbit propagation* in a serial computing environment (SCE) or a parallel computing environment (PCE). The computational savings tend to increase as the accuracy of the propagation increases.

Satellite Re-Entry (Scenario 6) over the accuracy ranges specified in Figures 2(a)-2(f). For comparison, Table 3 lists the computational savings for the propagation of a single orbit (see [1]).

The new uncertainty propagator builds upon the work of Aristoff and Poore [1], who developed an implicit Runge-Kutta-based method for orbit propagation. In conjunction with the Gauss von Mises framework of Horwood and Poore [11] (a statistically rigorous treatment of a space object's uncertainty), the new orbital and uncertainty propagator can provide a next-generation *uncertainty management package* for space surveillance.

ACKNOWLEDGMENTS

This work was funded, in part, by a Phase II STTR from the Air Force Office of Scientific Research (FA9550-12-C-0034) and a grant from the Air Force Office of Scientific Research (FA9550-11-1-0248).

REFERENCES

1. J. M. Aristoff and A. B. Poore, "Implicit Runge-Kutta methods for orbit propagation," in *Proceedings of the 2012 AIAA/AAS Astrodynamics Specialist Conference*, AIAA 2012-4880, (Minneapolis, MN), pp. 1–19, August 2012.
2. A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. New York: Dover, 1970.
3. A. Gelb, *Applied Optimal Estimation*. Cambridge, MA: M.I.T. Press, 1974.
4. Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley & Sons, 2001.
5. R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
6. N. G. A. Doucet, N.F. Freitas and A. Smith, "Sequential Monte Carlo methods in practice," in *Statistics for Engineering and Information Sciences*, Springer, 2001.
7. B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston: Artech House, 2004.
8. S. J. Julier and J. K. Uhlmann, "Method and apparatus for fusing signals with partially known independent error components." U.S. Patent Number 6,829,568 B2, Issued on 7 December 2004.
9. S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, pp. 401–422, 2004.
10. K. Ito and K. Xiong, "Gaussian filters for nonlinear filtering problems," *IEEE Transactions on Automatic Control*, vol. 45, no. 5, pp. 910–927, 2000.
11. J. T. Horwood and A. B. Poore, "Orbital state uncertainty realism," in *Proceedings of the 2012 Advanced Maui Optical and Space Surveillance Technologies Conference*, (Wailea, HI), September 2012.

12. J. T. Horwood and A. B. Poore, "Adaptive Gaussian sum filters for space surveillance," *IEEE Transactions on Automatic Control*, vol. 56, no. 8, pp. 1777–1790, 2011.
13. J. T. Horwood, N. D. Aragon, and A. B. Poore, "Gaussian sum filters for space surveillance: theory and simulations," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1839–1851, 2011.
14. K. J. DeMars, M. K. Jah, Y. Cheng, and R. H. Bishop, "Methods for splitting Gaussian distributions and applications within the AEGIS filter," in *Proceedings of the 22nd AAS/AIAA Space Flight Mechanics Meeting*, (Charleston, SC), February 2012. Paper AAS-12-261.
15. O. Le Maitre and O. M. Knio, *Spectral Methods for Uncertainty Quantification: with Applications to computational fluid dynamics*. Springer, 2010.
16. R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Boston: Artech House, 2007.
17. A. B. Poore, "Multidimensional assignment formulation of data association problems arising from multitarget tracking and multisensor data fusion," *Computational Optimization and Applications*, vol. 3, pp. 27–57, 1994.
18. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Boston: Artech House, 1999.
19. F. K. Chan, *Spacecraft Collision Probability*. El Segundo, CA: The Aerospace Press, 2008.
20. A. O. Hero, D. Castañón, D. Cochran, and K. Kastella, *Foundations and Applications of Sensor Management*. New York: Springer, 2008.
21. O. Montenbruck and E. Gill, *Satellite Orbits: Models, Methods, and Applications*. Berlin: Springer, 2000.
22. J. C. Butcher, *Numerical Methods for Ordinary Differential Equations*. West Sussex, England: John Wiley & Sons, 2008.
23. E. Hairer and G. Wanner, "Solving ordinary differential equations II: Stiff and differential-algebraic problems," in *Springer Series in Computational Mathematics*, Springer, second ed., 2010.
24. J. P. Vinti, "Orbital and celestial mechanics," in *Progress in Astronautics and Aeronautics* (G. J. Der and N. L. Bonavito, eds.), vol. 177, Cambridge, MA: American Institute of Aeronautics and Astronautics, 1998.
25. G. Beylkin and K. Sandberg, "ODE solvers using bandlimited approximations," *arXiv:1208.3285v1 [math.NA]*, 2012.
26. B. K. Bradley, B. A. Jones, G. Beylkin, and P. Axelrad, "A new numerical integration technique in astrodynamics," in *Proceedings of the 22nd Annual AAS/AIAA Spaceflight Mechanics Meeting*, (Charleston, SC), pp. 1–20, AAS 12-216, Jan. 30 - Feb. 2 2012.
27. B. A. Jones, "Orbit propagation using Gauss-Legendre collocation," in *Proceedings of the 2012 AIAA/AAS Astrodynamics Specialist Conference*, AIAA 2012-4967, (Minneapolis, MN), pp. 1–16, August 2012.
28. C. Peat, "Heavens above," July 2012.