

EFFICIENT, ROBUST AND UNIVERSAL COMPUTATION OF CLASSICAL ORBITAL ELEMENTS

R. Flores^a and E. Fantino^{a*}

^a Department of Aerospace Engineering, Khalifa University of Science and Technology, P.O. Box 127788, Abu Dhabi (United Arab Emirates)

* Corresponding author: elena.fantino@ku.ac.ae

Abstract

This paper reviews the standard algorithm for converting spacecraft state vectors to Keplerian orbital elements with a focus on its computer implementation. It analyzes the shortcomings of the scheme as described in the literature, and proposes changes to address orbits of arbitrary eccentricity and inclination in a robust way. Next, it presents two alternative schemes that simplify the program structure while improving the accuracy and speed of the transformation on modern computer architectures. Comprehensive numerical benchmarks demonstrate accuracy improvements by two orders of magnitude, together with a 40% reduction of computational cost relative to the standard implementation.

Keywords: state vector to orbital elements; algorithm accuracy; computational cost

1. Introduction

Transforming between spacecraft State Vector (SV) $\mathbf{x} = \{\mathbf{r}, \mathbf{v}\}$ (position and velocity) and Classical Orbital Elements (COE) $\mathbf{o} = \{a, e, i, \Omega, \omega, \theta\}$ (semimajor axis, eccentricity, inclination, longitude of the ascending node, argument of pericenter and true anomaly) [1] is a recurring task in many orbital mechanics calculations. There are many high-precision propagators that operate with the state vector, but it is often the case that the orbital elements provide a more convenient description of the trajectory. This is due to their small changes over one revolution¹, which makes them attractive for applications like station-keeping [2], long-term orbital stability [3], and orbit transfer optimization [4]. Moreover, the orbital elements provide an intuitive understanding of the orbit orientation and shape.

The algorithm to transform from SV to COE is considered an essential tool of celestial mechanics and, as such, it is ubiquitous in the reference literature (see for example [1] [5] [6] [7] [8]). Unlike the switch from COE to SV, which is straightforward and involves a fixed sequence of steps, the inverse transform must take into account the singularities in the definition of the orbital elements. For equatorial orbits, Ω is ill-defined. In circular orbits, ω also becomes degenerate. These singularities are especially relevant for

¹ Of course, the true anomaly experiences large variations, but it can be replaced easily with the time of periapsis or the true/mean anomaly at epoch, which vary slowly.

applications where the equations of motion are formulated directly in terms of the orbital elements (e.g., [9]). Different sets of elements have been proposed to work around this limitation. For a review of alternatives see [10]. The family of equinoctial elements was a first attempt to address the issues. It traces its origins to Lagrange's secular theory of planetary motion [11]. The term "equinoctial elements" was coined by Arsenault et al. [12] in 1970, and multiple variations of them can be found in the literature [13] [14] [15] [16] [17] [18] [19]. Equinoctial elements still suffer from limitations, they are singular for parabolic orbits and cannot treat simultaneously prograde and retrograde equatorial orbits. A further refinement are the modified equinoctial elements [20] [21]. This work only considers the transformation from SV to COE, and focuses on improving its accuracy and efficiency in modern computers.

Software implementations commonly handle the singularities by means of separate branches of code for each case. This approach increases program size and the likelihood of coding errors, and, crucially, can be detrimental for computational performance in modern CPUs. Moreover, identifying which code path to take involves setting tolerances for eccentricity and inclination. This may compromise accuracy if the thresholds are not chosen carefully.

Current processor architectures are pipelined [22]. Each instruction goes through several execution stages (one per clock cycle) before completion, while subsequent instructions are moving through the pipeline concurrently. It is often the case that the result of a previous instruction is required to decide which path of a branch must be taken. Waiting for the result to be ready (i.e., for the value of the condition to come out of the pipeline) would waste processor cycles (this is known as a "bubble"). Bubbles are mitigated with speculative execution assisted by a branch predictor [23]. The hardware makes a guess of the likely path to follow and starts execution before the condition can be evaluated. In case the prediction turns out erroneous, the pipeline must be flushed, wasting the processing already done for the instructions in flight [24]. An effective strategy against performance degradation due to branch mispredictions is, unsurprisingly, branchless programming [25].

This document presents two coding strategies that simplify the calculation of classical orbital elements and improve the accuracy of the results. Section 2 reviews a standard implementation, as found in the astrodynamics literature. It outlines the main shortcomings and strategies to mitigate them. Section 3 presents a branchless version that eliminates the need to contemplate coordinate system singularities. It enables robust, strictly linear (i.e., free from branching) execution, for increased accuracy and consistent performance. Section 4 compares the accuracies of the original and branchless algorithms in the context of low inclination and near-circular orbits. Section 5 introduces an alternative algorithm for computer systems where the branchless approach is not suitable. It improves accuracy and reduces the number of branches (without completely eliminating them). Section 6 presents additional benchmarks. Finally, the main conclusions are drawn in section 7.

2. Standard programming approach

A popular algorithm (which we shall denote as AL1) follows the steps below to compute the orbital elements from the state vector \mathbf{x} . All vectors are expressed in a Cartesian frame $\{x, y, z\}$ with axes along the $\{\hat{\mathbf{i}}, \hat{\mathbf{j}}, \hat{\mathbf{k}}\}$ directions. Hats denote unit vectors, for example: $\hat{\mathbf{r}} = \frac{\mathbf{r}}{r}$, where $r = \sqrt{\mathbf{r} \cdot \mathbf{r}}$. The reference plane for measuring inclination is XY and μ denotes the gravitational parameter.

1. Determine the orbital specific angular momentum $\mathbf{h} = \mathbf{r} \times \mathbf{v}$
2. The orbital inclination is $i = \arccos(\hat{\mathbf{h}} \cdot \hat{\mathbf{k}})$
3. If $i > 0$ continue with step 4, otherwise skip to step 7
4. Find the direction of the ascending node $\hat{\mathbf{n}} = \frac{\hat{\mathbf{k}} \times \mathbf{h}}{\|\hat{\mathbf{k}} \times \mathbf{h}\|}$
5. $\Omega_{aux} = \arccos(\hat{\mathbf{n}} \cdot \hat{\mathbf{x}})$
6. If $n_y \geq 0 \rightarrow \Omega = \Omega_{aux}$, otherwise $\Omega = -\Omega_{aux}$. Skip to step 8
7. Set $\Omega = 0$ and $\hat{\mathbf{n}} = \hat{\mathbf{i}}$
8. The semimajor axis is $a = \left(\frac{2}{r} - \frac{\mathbf{v} \cdot \mathbf{v}}{\mu} \right)^{-1}$
9. Compute the eccentricity vector $\mathbf{e} = \frac{\mathbf{v} \times \mathbf{h}}{\mu} - \hat{\mathbf{r}}$
10. The eccentricity is $e = \sqrt{\mathbf{e} \cdot \mathbf{e}}$
11. If $e > 0$ proceed with step 12, otherwise skip to 16
12. $\theta_{aux} = \arccos(\hat{\mathbf{e}} \cdot \hat{\mathbf{r}})$
13. If $\mathbf{r} \cdot \mathbf{v} \geq 0 \rightarrow \theta = \theta_{aux}$, otherwise $\theta = -\theta_{aux}$
14. $\omega_{aux} = \arccos(\hat{\mathbf{e}} \cdot \hat{\mathbf{n}})$
15. If $e_z \geq 0 \rightarrow \omega = \omega_{aux}$, otherwise $\omega = -\omega_{aux}$. Finished
16. Set $\omega = 0$ and $\theta_{aux} = \arccos(\hat{\mathbf{r}} \cdot \hat{\mathbf{n}})$
17. If $r_z \geq 0 \rightarrow \theta = \theta_{aux}$, otherwise $\theta = -\theta_{aux}$

Table 1 - Traditional scheme (AL1).

First and foremost, AL1 does not even work for eccentric inclined orbits in practice. Steps 12, 14 and 16 hide an insidious pitfall. In exact arithmetic, one expects $|\hat{\mathbf{a}} \cdot \hat{\mathbf{b}}| \leq 1$ for any pair of unit vectors $\{\hat{\mathbf{a}}, \hat{\mathbf{b}}\}$. However, finite precision arithmetic is subject to rounding errors. Thus, the magnitude of the dot product could exceed 1 if both vectors are quasi-parallel. This can happen, for example, in step 12 when the spacecraft is very close to the apocenter or pericenter. The computer will trigger an exception if it tries to evaluate the arccosine of a value larger than 1, disrupting the calculation. While this occurrence may seem unlikely, if the elements are evaluated a large number of times, the problem

will eventually arise². To make matters worse, the circumstances where this issue is triggered depends, on the exact sequence of operations used to compute the argument of the arccosine (this is typical with rounding-related problems). For example, these two apparently equivalent ways of computing θ_{aux} can yield different results:

$$\theta_{aux} = \arccos(\hat{\mathbf{e}} \cdot \hat{\mathbf{r}}) ; \theta_{aux} = \frac{\arccos(\mathbf{e} \cdot \mathbf{r})}{e r}. \quad (2.1)$$

A simple workaround is to replace step 12 with

$$\xi = \max(\min(\hat{\mathbf{e}} \cdot \hat{\mathbf{r}}, 1), -1) ; \theta_{aux} = \arccos(\xi), \quad (2.2)$$

which, unfortunately, introduces additional latencies because θ_{aux} cannot be evaluated until ξ is available. The same applies to steps 14 and 16. One could think that there is a potential for a similar issue in step 5. However, in this case the argument of the arccosine is just one of the components of a unit vector. This operation is very robust numerically and no special precaution is required.

The sequence of calculations requires conditional branches at steps 3 and 11 because \mathbf{n} and \mathbf{e} vanish for non-inclined and circular orbits, respectively. In those cases, it is not possible to compute the respective unit vectors. Aside from the aforementioned performance penalty, the branches are problematic because conditions such as $i > 0$ and $e > 0$ do not fare well in the domain of finite precision arithmetic³ (which always applies for numerical orbit propagation). Once the magnitude of the vectors becomes sufficiently small, the rounding errors inherent to floating point arithmetic can start dominating the calculations. While it may still be possible to normalize the vectors without triggering an overflow, the resulting direction can be erroneous due to this kind of contamination. A simple workaround is to set finite thresholds below which the alternative code path is taken. However, these tolerances must be adjusted carefully. Too low a value does not prevent the rounding issues, while an overly large threshold will incorrectly characterize normal orbits as degenerate cases. Therefore, for the sake of performance and accuracy it is highly desirable to write an algorithm that can run in a strictly linear way (i.e., branchless) and without tolerances that require tuning.

There are subtler errors in AL1 that need to be fixed before it can be used for arbitrary orbits. Note that if the inclination vanishes, the sign of ω in step 15 cannot be determined from the vertical component of the eccentricity vector. Instead, it is given by the sign of the y component (because in this case the reference direction for ω is the x axis). Similarly, in step 17, if the orbit is equatorial, the sign of the true anomaly is given by the y component of \mathbf{r} . Furthermore, one must consider if the orbit is prograde or retrograde when choosing the sign. Also, to handle retrograde orbits, the actual condition that must be tested in step 3 is $i > 0$ AND $i < \pi$. The correct classical

² In fact, it was not possible to complete the tests in section 4 without addressing this problem first.

³ Due to a lucky coincidence, explained in section 3, the condition $i > 0$ is acceptable in this case.

algorithm, including an eccentricity threshold (e_{thr}) in step 10 to detect quasi-circular orbits, is summarized in Table 2.

1. $\mathbf{h} = \mathbf{r} \times \mathbf{v}$
2. $\xi = \max(\min(\hat{\mathbf{h}} \cdot \hat{\mathbf{k}}, 1), -1) \rightarrow i = \arccos(\xi)$
3. Orbit is inclined if $i > 0$ AND $i < \pi$. Continue with step 4 for inclined orbits, otherwise skip to step 7
4. $\hat{\mathbf{n}} = \frac{\hat{\mathbf{k}} \times \mathbf{h}}{\|\hat{\mathbf{k}} \times \mathbf{h}\|}$
5. $\Omega_{aux} = \arccos(\hat{n}_x)$
6. If $n_y \geq 0 \rightarrow \Omega = \Omega_{aux}$, otherwise $\Omega = -\Omega_{aux}$. Skip to step 8
7. $\Omega = 0, \hat{\mathbf{n}} = \hat{\mathbf{i}}$
8. $a = \left(\frac{2}{r} - \frac{\mathbf{v} \cdot \mathbf{v}}{\mu} \right)^{-1}$
9. $\mathbf{e} = \frac{\mathbf{v} \times \mathbf{h}}{\mu} - \hat{\mathbf{r}}$
10. $e = \sqrt{\mathbf{e} \cdot \mathbf{e}}$
11. If $e > e_{thr}$ continue with step 12, otherwise skip to step 18
12. $\xi = \max(\min(\hat{\mathbf{e}} \cdot \hat{\mathbf{r}}, 1), -1) \rightarrow \theta_{aux} = \arccos(\xi)$
13. If $\mathbf{r} \cdot \mathbf{v} \geq 0 \rightarrow \theta = \theta_{aux}$, otherwise $\theta = -\theta_{aux}$
14. $\xi = \max(\min(\hat{\mathbf{e}} \cdot \hat{\mathbf{n}}, 1), -1) \rightarrow \omega_{aux} = \arccos(\xi)$
15. If the orbit is inclined continue to step 16, otherwise skip to step 17
16. If $e_z \geq 0 \rightarrow \omega = \omega_{aux}$, otherwise $\omega = -\omega_{aux}$. Finished
17. If $(e_y \cdot h_z) \geq 0 \rightarrow \omega = \omega_{aux}$, otherwise $\omega = -\omega_{aux}$. Finished
18. $\omega = 0$
19. $\xi = \max(\min(\hat{\mathbf{r}} \cdot \hat{\mathbf{n}}, 1), -1) \rightarrow \theta_{aux} = \arccos(\xi)$
20. Continue to step 21 for inclined orbits, otherwise proceed to step 22
21. If $r_z \geq 0 \rightarrow \theta = \theta_{aux}$, otherwise $\theta = -\theta_{aux}$. Finished
22. If $(r_y \cdot h_z) \geq 0 \rightarrow \theta = \theta_{aux}$, otherwise $\theta = -\theta_{aux}$

Table 2 – Correct traditional scheme (AL2).

While the algorithm in Table 2 (AL2 henceforth) addresses many issues from the original formulation, it is clearly more convoluted. This increases the likelihood of coding errors, and can negatively impact performance. It also requires setting an appropriate value for e_{thr} , which shall be discussed in section 4.1.

To close this section, we call to the attention of the reader that step 8 of AL1 and AL2 can trigger an exception —overflow, infinity or not a number (NaN), depending on the specifics of the computer and compiler— for parabolic orbits. A parabolic orbit, by

definition, has infinite semimajor axis, so this is more a limitation of the COE themselves than a software issue. For most applications this edge case does not require special attention. For example, in double precision floating-point arithmetic it happens if $a \geq 1.8 \cdot 10^{308}$ [26]. Note that, because the semimajor axis is not used in subsequent steps, the exception will not stop the calculation, and can be detected and treated a posteriori. For applications focusing on quasi-parabolic orbits, an alternative is to compute a^{-1} instead, which is finite. This edge case falls outside the scope of the work, and will not receive further consideration. We also note that the classical algorithm is applicable to hyperbolic orbits ($e > 1$), where it returns a negative semimajor axis, as per standard conventions.

3. A branchless general code path

It is possible to overcome the aforementioned shortcomings by using the protections built inside the ATAN2 function available in compilers. It is also a native instruction in many processor architectures⁴ [27]. A call to ATAN2(y,x) returns the argument of the complex number $x+iy$ in the $[-\pi, \pi]$ range. Crucially, even if both arguments are null, it returns well-defined values (0 or π , depending on the signs of the operands⁵). This instruction can compute Ω and the direction of the eccentricity vector to the best accuracy offered by the hardware, free from exceptions and without having to set arbitrary tolerances. With a reliable value of Ω , the $\hat{\mathbf{n}}$ vector can be uniquely determined. Then, a reference triad for the orbital plane $\{\hat{\mathbf{n}}, \hat{\mathbf{b}}, \hat{\mathbf{h}}\}$ is defined. In this frame, ω and θ can be determined reliably with ATAN2 calls, irrespective of the eccentricity and inclination.

Interestingly, use of the arctangent (and even ATAN2 calls) for transforming SV to COE is mentioned as early as 1992 [28]. However, the driving factor is completely different. The author used an old programming language, where the only inverse trigonometric function available was the arctangent. In fact, he explicitly states that additional code is required to treat circular and equatorial orbits.

As shown below, the improved algorithm (AL3 hereafter) results in a single code path for all types of orbits:

1. $\mathbf{h} = \mathbf{r} \times \mathbf{v}$
2. $i = \text{ATAN2}\left(\sqrt{h_x^2 + h_y^2}, h_z\right)$
3. $\Omega = \text{ATAN2}\left(h_x, -h_y\right)$
4. $\hat{\mathbf{n}} = \cos \Omega \hat{\mathbf{i}} + \sin \Omega \hat{\mathbf{j}}$
5. Compute $\hat{\mathbf{b}} = \hat{\mathbf{h}} \times \hat{\mathbf{n}}$ such that $\{\hat{\mathbf{n}}, \hat{\mathbf{b}}, \hat{\mathbf{h}}\}$ is right-handed triad

⁴ For example, in x86 and IA64 its known as Partial Arctangent and its opcode is FPATAN.

⁵ In standard floating-point arithmetic, zeros are signed.

6. $a = \left(\frac{2}{r} - \frac{\mathbf{v} \cdot \mathbf{v}}{\mu} \right)^{-1}$
7. $\mathbf{e} = \frac{\mathbf{v} \times \mathbf{h}}{\mu} - \hat{\mathbf{r}}$
8. $e = \sqrt{\mathbf{e} \cdot \mathbf{e}}$
9. $\omega = \text{ATAN2}(\mathbf{e} \cdot \hat{\mathbf{b}}, \mathbf{e} \cdot \hat{\mathbf{n}})$
10. $\theta = \text{ATAN2}(\mathbf{r} \cdot \hat{\mathbf{b}}, \mathbf{r} \cdot \hat{\mathbf{n}}) - \omega$

Table 3 - Branchless scheme (AL3).

Clearly, the branchless algorithm is easier to follow and requires fewer steps. This simplicity is a considerable advantage by itself, as it reduces the likelihood of coding errors and improves maintainability. In addition to eliminating branches, the use of ATAN2 removes the need to check the magnitude of the arccosine argument, and determining the sign of the angle afterwards. Of course, it can be argued that some of the tests have not disappeared, they are simply hidden inside the ATAN2 call. However, if the function is supported by the CPU, these steps are accomplished at the hardware level, which is usually more efficient than a software implementation.

Step 3 deals cleanly with vanishing inclination orbits. It computes Ω as the argument of the vector $\hat{\mathbf{k}} \times \mathbf{h}$ (which is orthogonal to the projection of \mathbf{h} over the XY plane). This way, the longitude of the ascending node is determined to the best accuracy the computer is capable of. In case the projection of the angular momentum vanishes completely (i.e., the orbit is equatorial), Ω is set to either 0 or π . The two values are equally acceptable, because step 3 ensures that Ω and the direction $\hat{\mathbf{n}}$ are always consistent⁶. Because $\hat{\mathbf{n}}$ is well-defined, the calculation of the remaining vector of the orbital plane reference triad ($\hat{\mathbf{b}}$) is straightforward.

The case of near-circular orbits is treated seamlessly by step 9. Because the eccentricity vector is never normalized, there is no risk of exceptions. Again, the argument of the pericenter is calculated to the best accuracy attainable by the computer, and any spurious component of \mathbf{e} along $\hat{\mathbf{h}}$ (which can appear due to rounding errors) is automatically cancelled out during the calculations. If the eccentricity vanishes completely, the argument of the pericenter will be set to 0 or π , but the choice is irrelevant because $\omega + \theta$ yields the correct value of the argument of latitude. Step 10 ensures that the true anomaly remains consistent with ω (i.e., the original state vector can always be recovered from the computed orbital elements, to the accuracy allowed by the approximate arithmetic, irrespective of the type of orbit).

The attentive reader has surely noticed another subtle difference in AL3. In step 2, the inclination is computed using the ATAN2 function, instead of the arccosine. Replacing

⁶ Note that, for a zero-inclination orbit, the sum $\Omega + \omega$ gives the correct longitude of the pericenter.

the arccosine with the arctangent improves the accuracy when operating in finite precision arithmetic. For small values of the inclination, we can approximate

$$\cos i \approx 1 - \frac{i^2}{2}. \quad (3.1)$$

In floating-point arithmetic, there exists $\varepsilon > 0$ such that⁷

$$\text{If } |x| < \varepsilon \rightarrow 1 + x \hat{=} 1, \quad (3.2)$$

where we used a hat over the equal sign to indicate the result of approximate arithmetic [29]. For double precision $\varepsilon \sim 10^{-16}$ [26]. Combining equations (3.1) and (3.2):

$$\text{If } |i| \leq \sqrt{\varepsilon} \rightarrow \cos i \approx 1 - \frac{i^2}{2} \hat{=} 1 \rightarrow \arccos(\hat{\mathbf{h}} \cdot \hat{\mathbf{k}}) \hat{=} 0. \quad (3.3)$$

As the inclination of the orbit approaches $\sqrt{\varepsilon}$, AL2 yields vanishing inclination, introducing additional errors. Interestingly, while this phenomenon is undesirable from the accuracy standpoint, it has a convenient side effect. Because the computed value of the inclination vanishes well before the angular momentum becomes truly parallel to the z direction, it is possible to apply directly the condition $i > 0$ (i.e., without setting a finite tolerance) to select the code path for equatorial/inclined orbits.

The ATAN2 instruction solves the accuracy issue because, for small inclinations,

$$\tan i \approx i. \quad (3.4)$$

There is no constant term summed to the inclination in Eq. (3.4). Therefore, the arctangent function computes a good numerical approximation of the inclination, no matter how small it is. This same advantage applies whenever the arccosine function is used to compute an angle close to 0 or π , so it also applies to the determination of Ω , ω and θ . To keep the presentation concise, this paper focuses on the inclination, but the discussion also applies to the other angles.

It must be mentioned that AL3 is not truly general, as it will not work for zero orbital angular momentum. However, this case is of minimal relevance in practical applications (the trajectory is a straight line passing through the primary) and can usually be ignored.

4. Comparative accuracy of AL2 and AL3

The accuracy of the two implementations was tested building families of orbits with decreasing values of eccentricity, inclination, or both. Each orbit is sampled at 100 equally-spaced values of the true anomaly (starting at the pericenter). The error due to the transformation between state vector and orbital elements is estimated as follows:

⁷ Rigorously speaking, there are two alternative definitions of ε . The first one is given by Eq. (3.2). The second one is the relative rounding error of floating-point arithmetic operations, which is half as large. For the sake of simplicity, we shall use both of them interchangeably, because it does not affect the substance of the discussion.

1. From the reference orbital elements \mathbf{o}_i^{ref} at the i -th sampling point, compute the corresponding reference state vector $\mathbf{x}_i^{ref} = \mathbf{x}(\mathbf{o}_i^{ref})$.
2. Apply the transformation to determine the approximate orbital elements $\tilde{\mathbf{o}}_i = \mathbf{o}(\mathbf{x}_i^{ref})$. Then, invert the transformation to recover an approximation of the reference state vector $\tilde{\mathbf{x}}_i^{ref} = \mathbf{x}(\tilde{\mathbf{o}}_i)$.
3. The absolute error at the i -th point is $\delta_i = \|\mathbf{x}_i^{ref} - \tilde{\mathbf{x}}_i^{ref}\|$.

For each orbit, the maximum error across all the sampling points and the RMS error are recorded:

$$\delta^{\max} = \max\{\delta_1, \dots, \delta_{100}\} ; \delta^{RMS} = \sqrt{\frac{\sum_{i=1}^{100} \delta_i^2}{100}} . \quad (4.1)$$

This error measure includes a contribution from the transformation from orbital elements to state vector. However, given that this part is the same for both AL2 and AL3, it yields valid comparisons. Note that it is not appropriate to evaluate the error comparing by \mathbf{o}_i^{ref} directly against $\tilde{\mathbf{o}}_i$, because, near the singularities of the transformation, vastly different combinations of $\{\Omega, \omega, \theta\}$ correspond to almost identical state vectors. Instead, (4.1) measures how well the original state vector can be recovered from the approximate orbital elements. This is more relevant from a physical standpoint.

For the sake of simplicity, dimensionless variables have been used. In this system the gravitational parameter is $\mu = 1$, and a circular orbit of unitary radius has an orbital velocity of one. The fixed values $a = 1, \Omega = \omega = 1 \text{ rad}$ have been used to focus on the effects of e and i . Because the eccentricity is kept small in all tests (maximum 0.01) the distance to the primary and the velocity are always close to unity. Thus, the absolute and relative errors have the same order of magnitude and it is not necessary to examine them separately. The absolute error is used for the comparisons. All calculations use double precision floating-point arithmetic (i.e., $\varepsilon \sim 10^{-16}$).

4.1 Quasi-circular orbits with nonzero inclination

The inclination is fixed at $\pi/4$, and the eccentricity is progressively reduced from 10^{-2} to 10^{-16} . The test compares the traditional scheme (AL2) against the branchless algorithm (AL3). For this initial run $e_{thr} = 0$, as in AL1. The results are shown in Fig. 1.

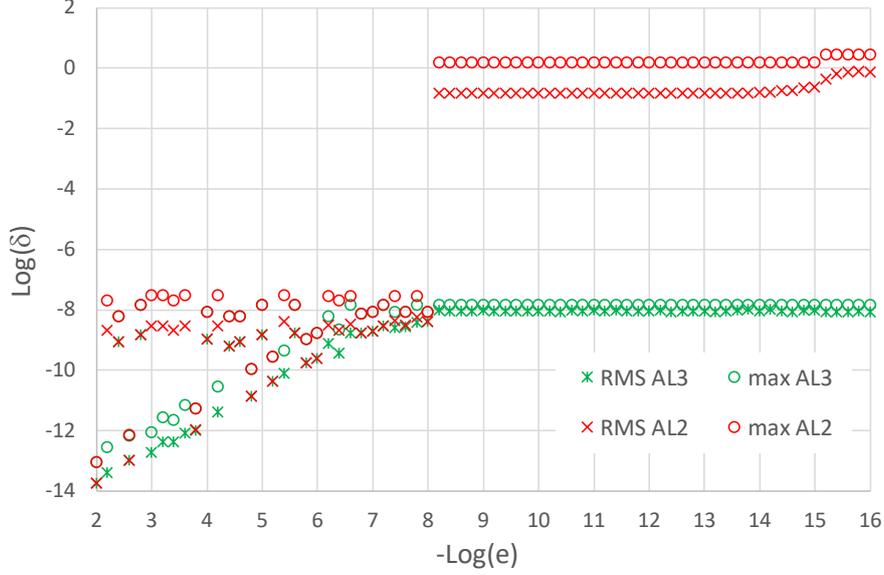


Fig. 1 – AL2 vs. AL3 error. Inclined quasi-circular orbits. $e_{thr}=0$.

As anticipated in section 2, the standard algorithm fails catastrophically for small eccentricities (below 10^{-8}). The issue arises because, for quasi-circular orbits, the calculation of the eccentricity involves the difference of two very similar vectors. It is a well-known fact in finite precision arithmetic that subtracting two close numbers is an ill-conditioned operation. The relative error of the result can be much larger than the uncertainties of the operands, a phenomenon known as catastrophic cancellation [30]. It is easier to understand the source of the problem if, instead of the eccentricity vector, we focus on its magnitude. The eccentricity can be recast as

$$e = \sqrt{1 - \frac{h^2}{\mu} \left(\frac{2\mu}{r} - v^2 \right)}. \quad (4.2)$$

With floating-point arithmetic, the evaluation of the radicand will be subject to a rounding error of order $O(\varepsilon)$ at least⁸. Therefore, the computed (approximate) eccentricity is

$$\tilde{e} = \sqrt{e^2 + O(\varepsilon)}. \quad (4.3)$$

It is clear that whenever $e < \sqrt{\varepsilon}$ the rounding error term dominates the calculation and \tilde{e} becomes unreliable. This agrees perfectly with the behavior observed in Fig. 1, where AL2 becomes unusable for $e < 10^{-8} \approx \sqrt{\varepsilon}$. The reason for the unacceptable increase in δ is not the uncertainty of e itself (the orbit is so close to circular that it does not change the result much), but the loss of accuracy of \hat{e} (the direction of the pericenter). This

⁸ Note that the radicand is the difference of two operands, the first one having the fixed value 1. Therefore, for small eccentricities, the second approaches unity. The intrinsic rounding error of the floating-point approximation of a value close to 1 is ε . Therefore, the expected uncertainty of the radicand calculation is, at a minimum, ε .

contaminates the calculation of ω and θ , to the point that they correspond to a completely erroneous position along the orbit. AL3 is impervious to this issue because, by design, it ensures that the sum $\omega + \theta$ is the angle between the orbiter and the ascending node, even if $e \ll 1$.

AL2 can be fixed by treating the orbit as circular whenever the eccentricity is sufficiently small. From the discussion above, a reasonable threshold is $e_{thr} = 10\sqrt{\varepsilon} \approx 10^{-7}$. A safety factor 10 has been included in e_{thr} to ensure that catastrophic failure never occurs.

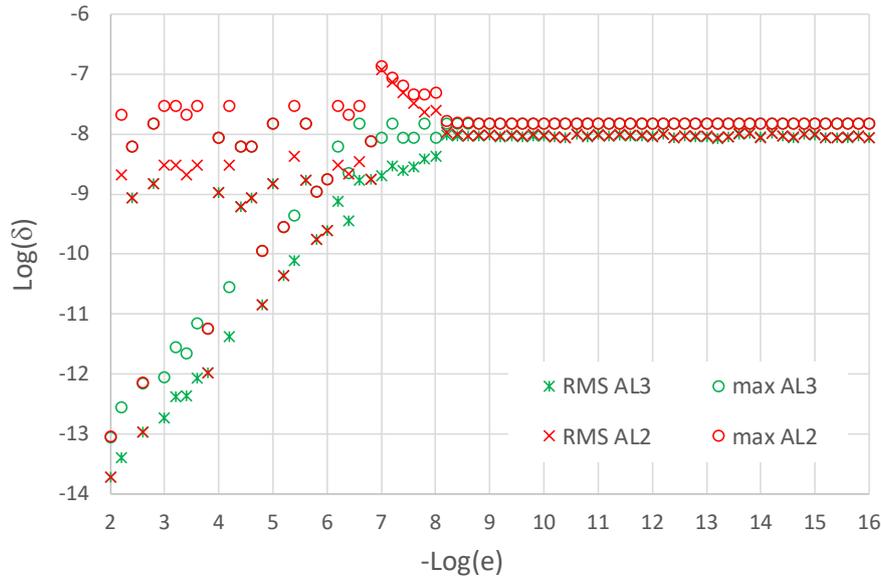


Fig. 2 – AL2 vs. AL3 error. Inclined quasi-circular orbits. $e_{thr} = 10^{-7}$.

Adding the eccentricity tolerance makes the accuracy of AL2 and AL3 comparable (see Fig. 2). Note that, superimposed to the general trend, there are strong fluctuations of the error (by as much as 5 orders of magnitude). This is typical of rounding uncertainties, because they can cancel out or add up unpredictably. It is remarkable that the error of the branchless scheme is equal to AL2 in the worst case, but never higher. From this point on, all test involving AL2 use $e_{thr} = 10\sqrt{\varepsilon}$.

4.2 Elliptical orbits with small inclination

In this case the eccentricity is fixed at 0.01 and the inclination varied from 10^{-2} to 10^{-16} rad.

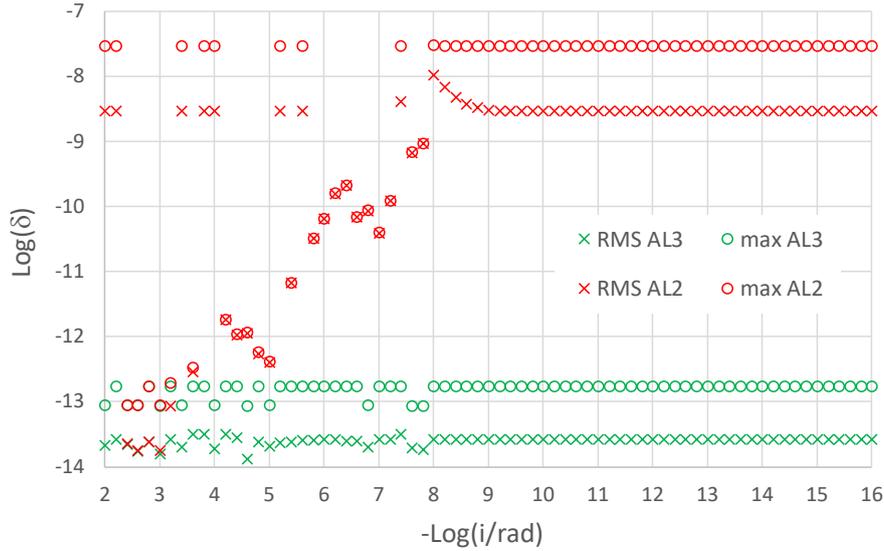


Fig. 3 - AL2 vs. AL3 error. Small-inclination elliptical orbits. $e_{thr}=10^{-7}$.

The vast superiority of the improved algorithm is evident in Fig. 3. The reason is the limitations of the arccosine function, illustrated in Eq. (3.3).

4.1 Orbits with vanishing eccentricity and inclination

Finally, we set eccentricity and inclination to the same value, which we ramp down from 10^{-2} to 10^{-16} . The behavior of both schemes is very similar (see Fig. 4). Note that, once again, AL2 is never superior to AL3 in terms of accuracy

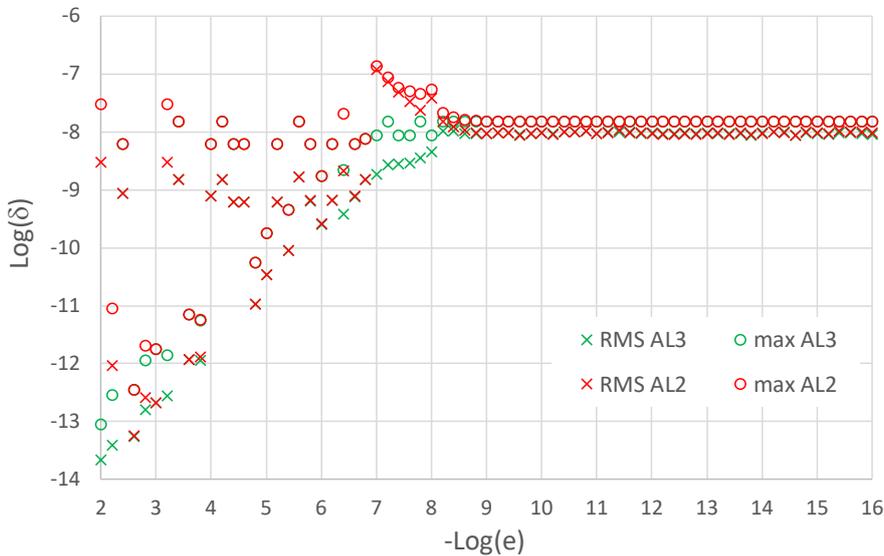


Fig. 4 - AL2 vs. AL3 error. Small eccentricity and inclination ($e=i/rad$). $e_{thr}=10^{-7}$.

5. A hybrid algorithm with reasonable performance and accuracy

The scheme AL3 is very accurate, robust and extremely simple to implement in comparison with AL2. However, the fact that it relies on calls to ATAN2 can be an issue in processors architectures with limited Floating-Point Unit (FPU) capabilities. This is often the case in microcontrollers and embedded systems. Those processors may not even support the instruction natively. Instead, the compiler (or the programmer) must

emulate it via software. Unfortunately, a suboptimal implementation can cause a performance hit. To address situations where ATAN2 is not available, is slow or does not behave as expected (i.e., does not properly handle null arguments), an alternative scheme that relies only on calls to the arccosine function (like AL2) is presented. It also removes the two transcendental function evaluations⁹ (sine and cosine) that AL3 uses to compute $\hat{\mathbf{n}}$ (step 4 in Table 3), which are expensive in systems with weak FPUs.

There are two key insights for improving the performance of AL2 and retain some of the accuracy benefits of AL3, all the while avoiding calls to ATAN2. The first one is that the loss of accuracy of AL2 when the inclination decreases is due to the presence of the arccosine function (see section 3), which has zero slope at the origin. The second one is that, as shown in section 4.1, the calculated value of the eccentricity effectively turns into numerical noise well before the orbit becomes circular (i.e., for $e < \sqrt{\varepsilon}$ it effectively becomes irrelevant if the orbit is circular or not). Thus, it is not necessary to preserve accuracy as e approaches zero. It is sufficient to guarantee that the result is reasonable.

It would be possible to remove ATAN2 from the computation of inclination and retain the full accuracy for small i by using the arcsine. Thus:

$$i = \begin{cases} \arcsin\left(\frac{\sqrt{h_x^2 + h_y^2}}{h}\right) & \text{for small inclination} \\ \arccos\left(\frac{h_z}{h}\right) & \text{otherwise} \end{cases} . \quad (5.1)$$

Note that it is not appropriate to always use the arcsine, as its accuracy degrades for polar orbits. This introduces one branch in the code, but brings several advantages to the table, as we shall see. The condition “small inclination” in Eq. (5.1) must be defined precisely and in a robust way (i.e., not linked to any particular case). We shall establish the threshold for small inclinations at the point where the accuracies of the two branches of (5.1) become comparable, with a twist to increase performance.

We can estimate the accuracy of the cosine branch assuming the evaluation of the term h_z/h is subject to an absolute error ε (small inclination is assumed, so $h_z/h \approx 1$)

$$\frac{h_z}{h} \cong \cos i + \varepsilon . \quad (5.2)$$

Using the series expansion of the cosine it is possible to estimate the computed value of inclination (\tilde{i})

⁹ For processors with capable FPUs, this performance penalty is not as large as it seems. IA64 and x86 processors support the FSINCOS instruction [27]. It computes simultaneously the sine and cosine of the argument at a lower cost than two separate evaluations.

$$\cos i + \varepsilon = 1 - \frac{\tilde{i}^2}{2} + O(\tilde{i}^4). \quad (5.3)$$

Expanding the left hand side (LHS)

$$-\frac{i^2}{2} + \varepsilon = -\frac{\tilde{i}^2}{2} + O(\tilde{i}^4), \quad (5.4)$$

where $i \approx \tilde{i}$ is assumed to lump all the high-order terms together. Retaining only second-order terms gives an estimate of the absolute error of the inclination computed with the cosine (δ^{\cos}):

$$2\varepsilon = |i^2 - \tilde{i}^2| = (i + \tilde{i})|i - \tilde{i}| \approx 2i\delta^{\cos} \rightarrow \delta^{\cos} \sim \frac{\varepsilon}{i}. \quad (5.5)$$

To estimate the error of the sine formula, we start with the series expansion

$$\frac{\sqrt{h_x^2 + h_y^2}}{h} = \sin i = i - \frac{i^3}{6} + O(i^5). \quad (5.6)$$

To maximize performance, assume only the first term of the series is retained. In this case, computing inclination is trivial. The error of this simple approximation is

$$\delta^{\sin} \sim \frac{i^3}{6}. \quad (5.7)$$

The crossover point (i_{cro}) is the inclination which makes (5.5) and (5.7) comparable:

$$\frac{\varepsilon}{i_{cro}} \sim \frac{i_{cro}^3}{6} \rightarrow i_{cro} \sim \sqrt[4]{6\varepsilon}. \quad (5.8)$$

For double precision arithmetic, $i_{cro} \sim 10^{-4}$. This is ideal, because it matches the point where the cosine approximation accuracy starts to degrade rapidly (see Fig. 3). At the cost of a branch in the scheme, the ATAN2 call for the inclination has been replaced by an arccosine (for inclined orbits) or, even better, the identity function (for small i). Moreover, it is possible to recast the calculation of $\hat{\mathbf{n}}$ in a way that avoids additional calls to trigonometric functions:

$$\hat{\mathbf{n}} = \begin{cases} \frac{-h_y}{h_{xy}} \hat{\mathbf{i}} + \frac{h_x}{h_{xy}} \hat{\mathbf{j}} & \text{if } i > i_{thr} \\ \frac{-\overline{h}_y}{\overline{h}_{xy}} \hat{\mathbf{i}} + \frac{h_x}{\overline{h}_{xy}} \hat{\mathbf{j}} & \text{otherwise} \end{cases}, \quad (5.9)$$

where $h_{xy} = \sqrt{h_x^2 + h_y^2}$ and the overline denotes “safe” values. These are computed as

$$\overline{h}_y = \max(h_y, \text{sign}(\varepsilon h, h_y)); \quad \overline{h}_{xy} = \sqrt{\overline{h}_y^2 + h_x^2}, \quad (5.10)$$

where the sign transfer function has been used:

$$\text{sign}(a, b) = \begin{cases} |a| & \text{if } b \geq 0 \\ -|a| & \text{otherwise} \end{cases} . \quad (5.11)$$

The effect of the safe values is to make $\hat{\mathbf{n}}$ parallel to the $\hat{\mathbf{i}}$ direction when $i \ll \varepsilon$ (it acts as a safeguard against division by zero). Note that

$$h_{xy} = h \sin i . \quad (5.12)$$

Therefore, the safe values (5.10) are extremely close to the original (non-overlined) variables unless the inclination is very small. They have a negligible impact on the accuracy because, by the time the correction becomes important, the orbit is almost equatorial and the errors are dominated by the XY components of position and velocity. Under these conditions, the exact location of the ascending node is not crucial. Those readers concerned with extreme accuracy can replace ε in Eq. (5.10) with a smaller value, but the effect is very limited in practice.

Moving to the efficient determination of $\{\Omega, \omega, \theta\}$, having established a branch for inclined/equatorial orbits makes it possible to completely eliminate the calls to ATAN2.

Start by computing:

$$\Omega = \text{sign}\left(\arccos(\hat{n}_x), \hat{n}_y\right), \quad (5.13)$$

$$\omega_{aux} = \arccos\left(\frac{\mathbf{e} \cdot \hat{\mathbf{n}}}{e + \varepsilon}\right), \quad (5.14)$$

$$\xi = \max\left(\min(\hat{\mathbf{r}} \cdot \hat{\mathbf{n}}, 1), -1\right) \rightarrow \theta_{aux} = \arccos(\xi). \quad (5.15)$$

The ε term in the denominator of Eq. (5.14) ensures that no exception occurs for very low eccentricity orbits ($\omega \rightarrow 0$ for quasi-circular orbits). Furthermore, it removes the need to check that the argument of the arccosine is below one. In practice, the extra term entails no loss of accuracy because, as discussed in section 4.1, the computation of eccentricity is subject to larger uncertainties due to its expression being numerically ill-conditioned.

For inclined orbits ($i > i_{thr}$), the signs of ω and θ are obtained from the vertical components of the eccentricity and position vectors:

$$\omega = \text{sign}(\omega_{aux}, e_z), \quad (5.16)$$

$$\theta = \text{sign}(\theta_{aux}, r_z) - \omega . \quad (5.17)$$

For small inclinations ($i < i_{thr}$), an auxiliary vector \mathbf{b} determines the signs:

$$\mathbf{b} = h_z (\hat{\mathbf{k}} \times \hat{\mathbf{n}}) = h_z (-\hat{n}_y \hat{\mathbf{i}} + \hat{n}_x \hat{\mathbf{j}}), \quad (5.18)$$

$$\omega = \text{sign}(\omega_{aux}, \mathbf{e} \cdot \mathbf{b}), \quad (5.19)$$

$$\theta = \text{sign}(\theta_{aux}, \mathbf{r} \cdot \mathbf{b}) - \omega. \quad (5.20)$$

The vector \mathbf{b} in Eq. (5.18) is contained in the XY plane. This simplifies the evaluation of the expressions where it appears, enhancing performance. It is orthogonal to $\hat{\mathbf{n}}$, with its direction determined by the sign of h_z . This ensures that the algorithm works also for retrograde orbits.

Note that, for the sake of clarity, the expressions above used $i > i_{thr}$ to identify inclined orbits, but the correct implementation is

$$\left| \frac{h_z}{h} \right| < \cos i_{cro}, \quad (5.21)$$

which addresses both prograde and retrograde trajectories.

To summarize, the steps of the enhanced algorithm (AL4 henceforth) are:

1. $\mathbf{h} = \mathbf{r} \times \mathbf{v}$
2. $a = \left(\frac{2}{r} - \frac{\mathbf{v} \cdot \mathbf{v}}{\mu} \right)^{-1}$
3. $\mathbf{e} = \frac{\mathbf{v} \times \mathbf{h}}{\mu} - \hat{\mathbf{r}}$
4. $e = \sqrt{\mathbf{e} \cdot \mathbf{e}}$
5. Let $\xi = \frac{h_z}{h}$, $h_{xy} = \sqrt{h_x^2 + h_y^2}$
6. If $|\xi| < \cos i_{cro}$ continue with step 7, otherwise skip to step 10
7. $i = \arccos \xi$
8. $\hat{\mathbf{n}} = \frac{-h_y}{h_{xy}} \hat{\mathbf{i}} + \frac{h_x}{h_{xy}} \hat{\mathbf{j}}$
9. Let $s_\omega = e_z$ and $s_\theta = r_z$. Skip to step 14
10. $i_{aux} = \frac{h_{xy}}{h}$. If $h_z \geq 0 \rightarrow i = i_{aux}$, otherwise $i = \pi - i_{aux}$
11. Compute safe values $\bar{h}_y = \max(h_y, \text{sign}(\varepsilon h, h_y))$ and $\bar{h}_{xy} = \sqrt{\bar{h}_y^2 + h_x^2}$
12. $\hat{\mathbf{n}} = \frac{-\bar{h}_y}{\bar{h}_{xy}} \hat{\mathbf{i}} + \frac{h_x}{\bar{h}_{xy}} \hat{\mathbf{j}}$, $\mathbf{b} = h_z (-\hat{n}_y \hat{\mathbf{i}} + \hat{n}_x \hat{\mathbf{j}})$
13. Let $s_\omega = \mathbf{e} \cdot \mathbf{b}$ and $s_\theta = \mathbf{r} \cdot \mathbf{b}$
14. $\Omega = \text{sign}(\arccos(\hat{n}_x), \hat{n}_y)$
15. $\omega_{aux} = \arccos\left(\frac{\mathbf{e} \cdot \hat{\mathbf{n}}}{e + \varepsilon}\right) \rightarrow \omega = \text{sign}(\omega_{aux}, s_\omega)$

$$16. \xi = \max(\min(\hat{\mathbf{r}} \cdot \hat{\mathbf{n}}, 1), -1) \rightarrow \theta_{aux} = \arccos(\xi) \rightarrow \theta = \text{sign}(\theta_{aux}, s_\theta) - \omega$$

Table 4 – Hybrid efficient scheme (AL4).

5.1 Accuracy of AL4 vs. AL3

The only noticeable difference between AL3 and AL4 is in the elliptical small inclination test (Fig. 1). Because the other two benchmarks yield indistinguishable results, they will not be presented to save space.

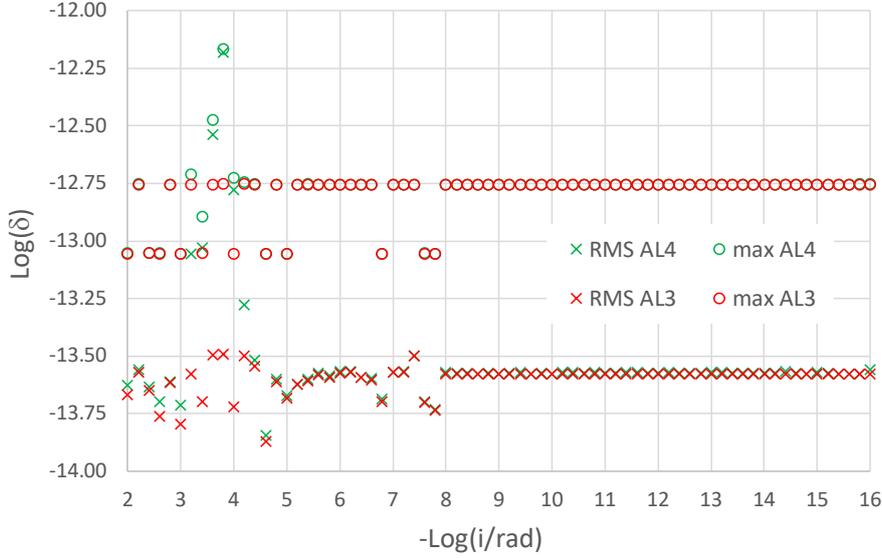


Fig. 5 – AL3 vs. AL4 error. Small-inclination elliptical orbits.

There is a moderate degradation of accuracy (up to one order of magnitude) in the neighborhood of i_{cro} , when the cosine approximation starts to lose precision and the degree-3 term of the sine series expansion is still relevant. The difference is subtle and can be ignored in most applications. If maximum accuracy is sought, AL4 can be improved by retaining one more term in the expansion of the sine. Starting from

$$\sin i = i - \frac{i^3}{3!} + \frac{i^5}{5!} + \dots, \quad (5.22)$$

and dropping terms of order 5 and higher:

$$\frac{h_{xy}}{h} = i - \frac{i^3}{6}. \quad (5.23)$$

While there is an analytical solution for Eq. (5.23), its expression is computationally costly. It is more efficient to use a single Newton-Rapson iteration to compute the approximate inclination.

Let

$$f(i) = i - \frac{i^3}{6} - \frac{h_{xy}}{h} \rightarrow f'(i) = 1 - \frac{i^2}{2}. \quad (5.24)$$

The iterative scheme is

$$i_0 = \frac{h_{xy}}{h}$$

$$i_1 = i_0 - \frac{f(i_0)}{f'(i_0)} = i_0 - \frac{-i_0^3}{6\left(1 - \frac{i_0^2}{2}\right)} = i_0 + \frac{i_0^3}{6 - 3i_0^2}. \quad (5.25)$$

The resulting algorithm (denoted AL5 hereafter) only requires a minor change in step 10 (see Table 5).

$$10. \xi = \frac{h_{xy}}{h} \rightarrow i_{aux} = \xi + \frac{\xi^3}{6 - 3\xi^2}. \text{ If } h_z \geq 0 \rightarrow i = i_{aux}, \text{ otherwise } i = \pi - i_{aux}$$

Table 5 - Improved calculation of i for small-inclination orbits (AL5).

The crossover inclination for AL5 is determined using as error estimate for the sine approximation the degree-5 term of the series (5.22)

$$\frac{\varepsilon}{i_{cro}} \sim \frac{i_{cro}^5}{5!} \rightarrow i_{cro} \sim \sqrt[5]{120\varepsilon}, \quad (5.26)$$

which yields $i_{cro} \sim 5 \cdot 10^{-3}$ for double precision. The accuracy of AL5 in the inclination test is virtually identical to AL3, so the graph is not included for the sake of brevity.

6. Additional benchmarks

This section extends the comparison to orbits with a wide range of orbital elements chosen at random. The algorithms AL2 through AL5 (AL1 is unusable in practice, so it is not included) were tested on two sets of one billion (10^9) combinations. The first set (termed “general” hereafter) uses uniformly distributed values within the following intervals:

$$a \in [10^{-3}, 10^3], e \in [0, 0.9], i \in [0, \pi], \{\Omega, \omega, \theta\} \in [0, 2\pi]. \quad (6.1)$$

For the second set (“low e/i ”), focusing on quasi-circular low-inclination orbits, the ranges of the orbital elements are:

$$a \in [10^{-3}, 10^3], \{\log e, \log i\} \in [-16, -2], \{\Omega, \omega, \theta\} \in [0, 2\pi]. \quad (6.2)$$

Note that in (6.2) it is the logarithms of the eccentricity and inclination that are uniformly distributed. This ensures that all the orders of magnitude are represented equally in the sample. The gravitational parameter is 1 in all cases.

6.1 Comparative accuracy tests

This benchmark uses the relative error because the components of the state vector are no longer of order 1. For each individual in the sample, it is computed as

$$\varphi_i = \frac{\delta_i}{\|\mathbf{x}_i^{ref}\|}, \quad (6.3)$$

where the absolute error δ_i is given by Eq. (4.1). Table 6 summarizes the results.

Orbit type	φ	AL2	AL3	AL4	AL5
General	RMS	$3.55 \cdot 10^{-11}$	$7.14 \cdot 10^{-14}$	$2.13 \cdot 10^{-12}$	$2.13 \cdot 10^{-12}$
	max	$2.76 \cdot 10^{-07}$	$1.69 \cdot 10^{-09}$	$2.73 \cdot 10^{-08}$	$2.73 \cdot 10^{-08}$
Low e/i	RMS	$1.90 \cdot 10^{-08}$	$1.80 \cdot 10^{-10}$	$1.82 \cdot 10^{-10}$	$1.82 \cdot 10^{-10}$
	max	$2.98 \cdot 10^{-07}$	$2.10 \cdot 10^{-08}$	$2.28 \cdot 10^{-08}$	$2.28 \cdot 10^{-08}$

Table 6 – Relative accuracy comparison for random orbital elements. Sample size 10^9 .

For the general set, AL3 excels at accuracy. It performs over one order of magnitude better than AL4 and AL5, whose behavior is virtually identical. AL2 produces the largest error, an order of magnitude higher than AL3/4. In the context of the low e/i set AL3 through AL5 perform very close to each other, with an average improvement of two orders of magnitude over AL2. In this case, however, the peak difference is smaller, at one order of magnitude approximately. We know from the tests in section 4.2 that the difference is larger for some specific conditions (see Fig. 3, where the accuracy of AL2 is up to 5 orders of magnitude worse). However, those are unlikely to appear in a random set.

6.2 Comparative performance tests

Finally, we assess the speed of the four algorithms in both the general and low e/i configurations. For the test, the sample sizes are reduced to 10^5 combinations of orbital elements, and the corresponding state vectors stored in memory. Then, the transformation to orbital elements is applied to each set 1000 times (i.e., a total of 10^8 transformations for each algorithm and set). To obtain a reliable timing, the entire process is repeated 10 times, averaging the duration of each pass. The set of 10^5 states gives reasonable diversity while fitting easily inside the CPU cache. This improves the estimate of raw algorithm performance by avoiding memory access bottlenecks.

The measurements were taken on an AMD Ryzen 9 7945HX laptop CPU using the Intel IFX Fortran compiler (version 2024.1) for Windows 64bit with the maximum optimization level (O3). The benchmark results are summarized in Table 7. The run times for each combination of orbit type and algorithm are normalized with respect to AL2 in the general case.

Orbit type	AL2	AL3	AL4	AL5
General	1.00	0.57	0.70	0.70
Low e/i	0.76	0.74	0.74	0.72

Table 7 - Relative time comparison for general and quasi-circular low-inclination orbits.

Code performance is affected by many implementation details. Changing the hardware, operating system, compiler version or settings can have a noticeable impact on speed. Furthermore, the effect may be different for each algorithm. The results from Table 7 should be taken as rough performance guidelines. It is always important to repeat this

test on the target system, and select the scheme that best serves the requirements of the application at hand.

Keeping in mind the caveat of the previous paragraph, the performance of AL3 for general orbits is outstanding, requiring 43% less time than AL2. AL4 and AL5 deliver a 30% reduction compared with the traditional scheme.

For the low e/i set, all schemes achieve comparable speed. Interestingly, the performance of AL3 degrades substantially relative to the general case. This is probably due to the large variation in latency of the ATAN2 function between the best and worst-case scenarios. As an example, for the processor used in the test (AMD Zen 4 family), the latency of FPATAN varies from 50 to 190 cycles [31]. It is likely that, when the protections against small values of the arguments are triggered, the latency of each ATAN2 call increases substantially. In any case, this is not a showstopper because, even in this situation, the speed is identical to AL4. While AL5 seems to fare slightly better¹⁰, the difference with respect to AL3/4 is close to the uncertainty of the time measurements (around 1%), so the advantage is marginal at best. Another noteworthy fact is that AL2 behaves much better than in the general case. This is due to the alternative paths for low eccentricity and inclination being computationally simpler (they just assign arbitrary fixed values to the ill-defined variables). This offsets the penalties due to the multiple branches, and brings the performance close to the other algorithms.

6.3 Comparison summary

The benchmarks in sections 4 and 6 evidence that, in systems with an efficient implementation of ATAN2, AL3 is the undisputed choice. It delivers the best performance, accuracy and ease of implementation. If AL3 is not feasible, AL5 should be chosen. It is as fast as AL4, equally simple to implement, and has a small accuracy advantage for low inclination orbits, as indicated in section 5.1. AL4 was a crucial stepping stone towards AL5 development, but it has no real use case. Regarding AL2, it loses even to AL4 in all categories.

7. Conclusions

This paper reviewed the standard approach to computing classical orbital elements from spacecraft state vector. Due to the singularities that affect the orbital elements in the case of circular and zero-inclination orbits, they must be addressed separately. The algorithm, if implemented in the way commonly presented in the literature (scheme AL1), suffers from accuracy and efficiency shortcomings. The accuracy can be improved to a certain extent by increasing the complexity of the implementation (AL2), at the cost of performance. An important contributor to the computational cost of AL2 is the presence of branches in the code, which is detrimental for the pipelined architectures of modern processors.

¹⁰ The small speed advantage relative to AL4 in this test is likely due to i_{cro} being larger for AL5. This allows it to take the code branch for small inclinations more frequently, avoiding an arccosine call.

The intrinsic safeguards of the ATAN2 instruction, available in many programming languages, enable a coding scheme free from branches (AL3). This approach also improves the accuracy of the transformation compared to AL2, with the errors reduced by as much as 5 orders of magnitude for some small inclination orbits. To address processor architectures with limited support for transcendental functions, a hybrid of the AL2 and AL3 schemes has been developed (AL5) that avoids ATAN2 calls while retaining a substantial part of the performance and accuracy improvements of AL3.

A test on a large set of random orbits showed that AL3 is, on average, 2 orders of magnitude more accurate than AL2. The advantage of AL5 is smaller, but still substantial at one order of magnitude. Regarding computational performance, the cost of AL3 is up to 43% less than AL2 for a random mix of all types of orbits. The advantage of AL5 is lower, at 30%. When restricted to low-inclination quasi-circular orbits, the speed gains of AL3 and AL5 are very limited, but they retain a substantial accuracy advantage.

The scheme AL3 yields important benefits in terms of simplicity (improving ease of programming and maintainability), accuracy and speed, with no obvious downsides. It is therefore an excellent alternative to the standard implementation (AL2). For those systems without good support for ATAN2, the scheme AL5 offers the same advantages, albeit to a smaller degree.

ACKNOWLEDGMENTS

The authors acknowledge Khalifa University of Science and Technology's internal grant CIRA-2021-65/8474000413 and project 8434000533 funded by Technology Innovation Institute. EF has been partially supported by the Spanish Ministry of Science and Innovation under projects PID2020-112576GB-C21 and PID2021-123968NB-I00.

REFERENCES

- [1] R. R. Bate, D. D. Mueller and J. E. White, *Fundamentals of Astrodynamics*, New York: Dover Publications, 1971, pp. 61-64.
ISBN 978-0-486-60061-1
- [2] E. Fantino, R. M. Flores, M. di Carlo, A. di Salvo and E. Cabot, "Geosynchronous inclined orbits for high-latitude communication," *Acta astronautica*, no. 140, pp. 570-582, 2017.
DOI 10.1016/j.actaastro.2017.09.014
- [3] S. Proietti, R. Flores, E. Fantino and M. Pontani, "Long-term orbit dynamics of decommissioned geostationary satellites," *Acta astronautica*, vol. 182, pp. 559-573, 2021.
DOI 10.1016/j.actaastro.2020.12.017

- [4] E. Fantino, B. Burhani, R. Flores, E. Alessi, F. Solano and M. Sanjurjo-Rivo, "End-to-end trajectory concept for close exploration of Saturn's Inner Large Moons," *Communications in Nonlinear Science and Numerical Simulation*, vol. 126, 2023, 107458.
DOI 10.1016/j.cnsns.2023.107458
- [5] J. E. Prussing and B. A. Conway, *Orbital Mechanics*, New York: Oxford University Press, 2013, pp. 50-52.
ISBN 978-0-19-983770-0
- [6] A. Tewari, *Atmospheric and Space Flight Dynamics*, Boston, MA: Birkhäuser, 2007, pp. 120-121.
ISBN 978-0-8176-4437-6
- [7] D. A. Vallado and W. D. McClain, *Fundamentals of Astrodynamics and Applications 3rd ed.*, Hawthorne, CA: Microcosm Press, 2007, pp. 119-122.
ISBN 978-1-881883-14-2
- [8] H. D. Curtis, *Orbital Mechanics for Engineering Students*, 4th ed., Elsevier, 2020, pp. 191-193.
ISBN 978-0-08-102133-0
- [9] M. Lara, J. F. San-Juan and L. M. López-Ochoa, "Efficient semi-analytic integration of GNSS orbits under tesseral effects," *Acta Astronautica*, vol. 102, pp. 355-366, 2014
DOI 10.1016/j.actaastro.2013.11.006
- [10] G. R. Hintz, "Survey of Orbit Element Sets," *Journal of Guidance, Control, and Dynamics*, vol. 31, no. 3, pp. 785-790, 2008.
DOI 10.2514/1.32237
- [11] J. L. Lagrange, *Théorie des variations séculaires des éléments des planètes. Première partie contenant les principes et les formules générales pour déterminer ces variations*, Berlin: Nouveaux mémoires de l'Académie royale des sciences et belles-lettres de Berlin, 1781.
- [12] J. L. Arsenault, K. C. Ford and P. E. Koskela, "Orbit Determination Using Analytic Partial Derivatives of Perturbed Motion," *AIAA Journal*, vol. 8, pp. 4-12, 1970.
DOI 10.2514/3.5597
- [13] F. R. Moulton, *An Introduction to Celestial Mechanics*, 2nd revised ed., New York: Dover Publications, 1970.
ISBN 978-0-486-64687-9
- [14] R. Broucke and P. Cefola, "On the equinoctial orbit elements," *Celestial Mechanics*, vol. 5, p. 303-310, 1972.
DOI 10.1007/BF01228432
- [15] V. A. Chobotov, *Orbital Mechanics*, 3rd ed., Reston, VA: AIAA, 2002, Chaps. 3 & 14.
ISBN: 978-1-56347-537-5

- [16] R. H. Battin, *An Introduction to the Mathematics and Methods of Astrodynamics*, Reston, VA: AIAA, 1999, Chap. 10.
ISBN 1-56347-342-9
- [17] G. E. O. Giacaglia, "The Equations of Motion of an Artificial Satellite in Nonsingular Variables," *Celestial Mechanics*, vol. 15, no. 2, pp. 191-215, 1977.
DOI 10.1007/BF01228462
- [18] P. E. Nacozy and S. S. Dallas, "The Geopotential in Nonsingular Orbital Elements," *Celestial Mechanics*, vol. 15, no. 4, pp. 453-466, 1977.
DOI 10.1007/BF01228611
- [19] C. J. Cohen and E. C. Hubbard, "A Nonsingular Set of Orbital Elements," *Astronomical Journal*, vol. 67, no. 1, pp. 10-15, 1962.
DOI 10.1086/108597
- [20] M. J. H. Walker, "A Set of Modified Equinoctial Orbit Elements," *Celestial Mechanics and Dynamical Astronom*, vol. 38, no. 4, pp. 391-392, 1986.
DOI 10.1007/BF01227493
- [21] M. Walker, B. Ireland and J. Owens, "A set modified equinoctial orbit elements (Erratum)," *Celestial Mechanics*, vol. 36, pp. 409-419, 1985.
DOI 10.1007/BF01238929
- [22] J. P. Shen and M. H. Lipasti, "Modern processor design: fundamentals of superscalar processors," Boston, McGraw-Hill Higher Education, 2005, pp. 39-54.
ISBN 978-0071230070
- [23] J. Smith, "A study of branch prediction strategies," in *ISCA 81: Proceedings of the 8th annual symposium on Computer Architecture*, Minneapolis, MN, 1981.
- [24] S. Eyerhan, J. E. Smith and L. Eeckhout, "Characterizing the branch misprediction penalty," in *2006 IEEE International Symposium on Performance Analysis of Systems and Software*, Austin, TX, 2006.
- [25] F. G. Pikus, "Branchless computing," in *The Art of Writing Efficient Programs*, Birmingham, UK, Packt Publishing Limite, 2021, pp. 103-110.
ISBN 9781800208117
- [26] Institute of Electrical and Electronics Engineers, "IEEE Std 754-2019: IEEE Standard for Floating-Point Arithmetic," 2019.
- [27] Intel Corp., *Intel® 64 and IA-32 Architectures Software Developer's Manual*, vol. 2: Instruction Set Reference, 2024. Available: <https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html> [Accessed 9 Apr. 2024]
- [28] J. M. A. Danby, *Fundamentals of Celestial Mechanics*, 2nd ed., Richmond, VA: William-Bell, 1992, pp. 204-206.
ISBN 0-943396-20-4
- [29] D. Golberg, "What Every Computer Scientist Should Know About Floating-Point Arithmetic," *ACM Computing Surveys*, vol. 23, no. 1, pp. 5-49, 1991.
DOI 10.1145/103162.103163

- [30] J. M. Muller et al., "Basic Properties and Algorithms," in *Handbook of Floating-Point Arithmetic*, Cham, Switzerland, Birkhäuser, 2018, p. 102.
ISBN 978-0-8176-4705-6
- [31] A. Fog, "Software optimization resources: Instruction tables," [Online]. Available: https://www.agner.org/optimize/instruction_tables.pdf [Accessed 9 Apr. 2024].