Virtual Thrust–Based Control Formulation for Multiple Kinetic Impactor Asteroid Deflection

Moacir Fonseca Becker[(1)], Karel Hernández Bandrich[(1)], Noelle Elisabeth May[(1)], Shivani Manoj Patwardhan[(1)], Alessandro Massini[(2)], Andrea Capannolo[(1)]

[(1)] *School of Aeronautics and Astronautics, Purdue University*
West Lafayette, Indiana, USA
Email: mfonsec@purdue.edu, herna891@purdue.edu, spatwar@purdue.edu, nemay@purdue.edu, acapa@purdue.edu

[(2)] *Department of Aerospace Science and Technology, Politecnico di Milano*
Milan, Italy
Email: alessandro.massini@mail.polimi.it

Abstract **– This paper introduces *virtual thrust* (VT), a continuous low-thrust control abstraction applied directly to asteroid deflection, as a way to design multiple-kinetic-impactor (MKI) campaigns. The method first solves a continuous VT optimal-control problem with sequential convex programming, then transcribes the resulting thrust profile into a set of dense impulses, and finally compiles those impulses into a finite impact campaign. Two VT cases are studied: one maximizes deflection for a fixed control budget, and the other minimizes the asteroid velocity change needed to reach a prescribed target deviation. The method is tested on an Apophis-style case with different lead times, impact speeds, cadence limits, and control authority levels. The compiled campaigns retain almost all of the continuous VT deflection. In the cases with a fixed budget, they also closely match a separate direct impulse allocation baseline. The minimum-Δv cases require only a modest number of extra impact opportunities. These results show that VT is a useful intermediate representation: it solves the deflection problem once as a continuous control problem, while the same VT profile can be recompiled into different effective MKI campaigns as mission assumptions change.**

I. Introduction

Potentially hazardous asteroids (PHAs) are celestial bodies that are sufficiently large and whose orbits bring them sufficiently close to the Earth to pose a significant impact threat. They belong to the class of near-Earth objects (NEOs). Concretely, NASA defines PHAs as having a minimum orbit intersection distance (MOID) of less than 0.05 AU and a magnitude of 22.0 or less [1]. Even though the possibility of a civilization-ending asteroid impact is low, the catastrophic consequences of such an event have motivated extensive work on planetary defense methods, and there are many smaller asteroids that could inflict significant damage should they strike near a populated area. In 2016, NASA established the Planetary Defense Coordination Office (PDCO) to manage the agency's ongoing mission of finding, tracking, and better understanding asteroids and comets that could pose an impact hazard.

A wide range of asteroid deflection and planetary defense strategies have been explored, with abundant examples in the literature. Conway [2], for example, considered strategies to maximize deviation using an impulsive maneuver. Izzo [3] analyzed both continuous and impulsive thrust and the resulting deflection. Conway[2] and Fantino[4] focused on analysis and optimization of the trajectories of planetary defense spacecraft. Colombo [5], Santos [6], and Ketema [7] explored the continuous thrust method, while authors such as Ross [8] studied the single kinetic impactor method. Other methods such as electric propulsion systems attached to the asteroid [9], or spacecraft acting as "gravity tractors" [10] have also been explored as potential defense strategies.

A natural extension of the single kinetic impactor (SKI) concept is the multi-kinetic impactor (MKI), in which several impactors transfer momentum to the asteroid over multiple collisions. MKI can improve mission robustness, increase the probability of a successful impact, and provide greater design flexibility. It also expands the design space, since the number of impactors, impact cadence, impact directions, and impact speeds must be considered. In this paper, an MKI campaign refers to the finite set of impacts used to deflect the asteroid. At the asteroid-side level considered here, a campaign is described by the number of impacts and, for each impact, its epoch, Δv magnitude delivered to the asteroid, and direction. The spacecraft transfer arcs that produce those impacts are not included in this campaign definition.

Recent work has begun to explore these design variables. Hydrodynamic simulations by Youtao et al. [11] demonstrate that, for a fixed total system mass, MKI configurations can improve deflection performance relative to SKI, although outcomes are highly dependent on impact geometry and velocity. Multi-modal deflection studies, such as [12], further highlight the nonlinear nature of momentum transfer in multi-impact scenarios. Mission-level MKI architectures have been explored by Bolzoni [13], Qiao [14], and Massini et al. [15], examining deployment strategies, sequential impact concepts, and rapid response architectures based on impactors prepositioned in Earth–Moon parking orbits. Additionally, Chignoli [16] shows

that while increasing the number of impactors can improve deflection, diminishing returns and sensitivity to cadence and uncertainty limit performance gains.

Despite these advances, existing work often treats MKI through parametric sweeps or architecture studies. Most studies also examine only a small number of impactors, while larger or more massive targets could require campaigns with dozens of impacts or more. For that regime, a practical design procedure is still needed to go from a desired asteroid deflection to a finite campaign without first guessing the number of impacts, their timing, or how the impulse should be distributed. VT addresses that gap by solving a continuous low thrust control problem applied to the asteroid first.

VT is not intended to replace full spacecraft transfer design or to prove global optimality; rather, it provides a reference control history that can later be converted into different discrete impact campaigns under mission cadence requirements and limits on the asteroid Δv delivered by each impact. In this role, VT is useful for early mission trade studies.

This paper develops the VT design process. We first model the desired asteroid velocity changes as a bounded, fictitious continuous thrust profile applied directly to the asteroid. That profile is then transcribed into a dense set of impulses and grouped by a greedy compiler into a finite MKI campaign. The result is a bridge between continuous optimal control and discrete impact design: the optimization finds when and in what direction the asteroid should be pushed, while the compiler turns that information into impact epochs, delivered impulse magnitudes, directions, and impact count. The compiler adds flexibility by allowing different MKI campaigns to be generated under different mission assumptions while still targeting the same desired deviation.

II. Methodology

Our method follows the three-step pipeline shown in Fig. 1. In the first step, the VT optimization setup uses sequential convex programming (SCP) to produce a bounded continuous control profile. In the second step, the zero-order-hold (ZOH) VT intervals are converted into a dense sequence of impulses. In the third step, a greedy MKI compiler locally groups those dense impulses into a sparse interception campaign that respects a minimum cadence between impacts.

We use the following notation throughout.

$$\mathbf{x}_A = \begin{bmatrix} \mathbf{r}_A \\ \mathbf{v}_A \end{bmatrix} \in \mathbb{R}^6, \qquad \mathbf{r}_f = \mathbf{r}_A(t_f) - \mathbf{r}_E(t_f).$$

Here $\mathbf{x}_A$ is the heliocentric state of the asteroid, given by its position $\mathbf{r}_A$ and its velocity $\mathbf{v}_A$, and the position vector of the asteroid relative to the Earth at the final closest-approach (CA) epoch t_f is $\mathbf{r}_f$.

A. Step 1: VT Optimization Setup

Designing a full MKI campaign is a nonlinear problem with both discrete and continuous decisions: one must decide how many impactors are used, when they arrive, and what impulse magnitude and direction each impact should deliver, while also satisfying mission constraints such as minimum cadence. In the first step, the VT optimization replaces this problem with a continuous thrust control problem applied to the asteroid.

Fig. 2 illustrates the idea. We build the reference case by forcing the asteroid to strike Earth at the nominal CA point, and then back-propagate the Earth and asteroid orbits from this condition to set the state at the initial feasible control epoch t_0 (the earliest epoch at which a first interception could occur). Because the reference impact is enforced, the nominal terminal miss distance is zero. Hence, the achieved terminal deviation is the VT-induced perturbation, $\Delta\mathbf{r}_{\text{total}} = \delta\mathbf{r}$.

We model the natural asteroid motion with Cartesian heliocentric two-body dynamics,

$$\mathbf{f}(\mathbf{x}_A, t) = \begin{bmatrix} \mathbf{v}_A \\ -\mu_\odot \mathbf{r}_A / \|\mathbf{r}_A\|_2^3 \end{bmatrix}, \tag{1}$$

where $\mu_\odot$ is the solar gravitational parameter. Adding the control gives the VT model,

$$\dot{\mathbf{x}}_A = \mathbf{f}(\mathbf{x}_A, t) + \begin{bmatrix} \mathbf{0} \\ a_{\max}\mathbf{u}(t) \end{bmatrix}, \qquad \|\mathbf{u}(t)\|_2 \le 1. \tag{2}$$

Here $\mathbf{u}(t) \in \mathbb{R}^3$ is dimensionless, and the physical asteroid acceleration is $a_{\max}\mathbf{u}(t)$, which stands in for the control authority of a kinetic-impactor campaign. We demonstrate the framework using terminal Earth–asteroid deviation at the nominal CA epoch.

We discretize the control on a ZOH grid made of K intervals:

$$\begin{aligned} t_0 < t_1 < \cdots < t_K = t_f, \\ \Delta t_k = t_{k+1} - t_k, \\ \mathbf{u}(t) = \mathbf{u}_k, \quad t \in [t_k, t_{k+1}). \end{aligned} \tag{3}$$

where t_k is a ZOH grid node and $\Delta t_k = t_{k+1} - t_k$ is the duration of the k-th interval. The control $\mathbf{u}_k \in \mathbb{R}^3$ is dimensionless over interval k, and the physical asteroid acceleration is $a_{\max}\mathbf{u}_k$. The corresponding control effort (the total velocity change delivered by VT) is

$$\Delta V_A = a_{\max} \sum_{k=0}^{K-1} \|\mathbf{u}_k\|_2 \Delta t_k, \tag{4}$$

or, in continuous form,

$$\Delta V_A = a_{\max} \int_{t_0}^{t_f} \|\mathbf{u}(t)\|_2 \, dt. \tag{5}$$

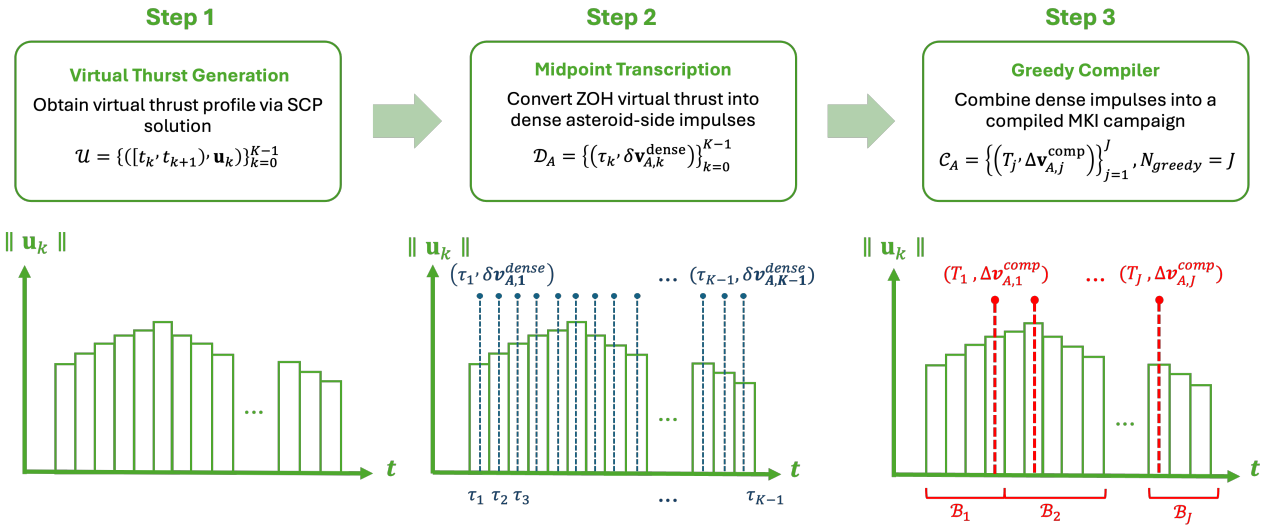

Fig. 1: Three-step VT-to-MKI workflow. Step 1 solves a continuous VT optimization problem, Step 2 transcribes the ZOH control profile into dense midpoint impulses, and Step 3 greedily compiles those impulses into a finite MKI campaign that satisfies the cadence limit.

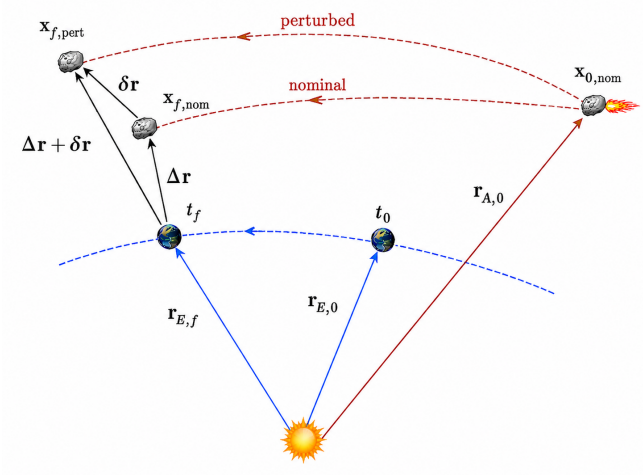

Fig. 2: VT concept schematic. The asteroid is modeled as if acted on by a fictitious continuous thrust source. Any deviation from the nominal impact condition at CA is caused by the applied VT profile.

B. *Control Authority Proxies for Kinetic Impactors*

We approximate the asteroid velocity change available as control authority from N identical kinetic impactors by assuming perfectly inelastic, aligned collisions. The scalar proxy assumes that the relative impact velocity contributes along the intended impulse direction, and the asteroid mass is treated as unchanged after each impact because $m_{\text{sc}} \ll m_A$.

$$\Delta V_{A,\max} = \frac{\beta N m_{\text{sc}} v_{\text{rel}}}{m_A}. \tag{6}$$

Here $\Delta V_{A,\max}$ is the available ΔV_A budget, while ΔV_A denotes the velocity change actually used by a VT solution. m_A is the asteroid mass, m_{sc} is the mass of one spacecraft, v_{rel} is the assumed aligned impact relative speed, β is the momentum-enhancement factor (which accounts for ejecta generated by the impact), and N is the number of full-speed impactor units used to define the available control authority. More generally, the transferred component can be written as $v_{\text{rel}} \cos\gamma$, where γ is the angle between the interceptor and the asteroid velocity direction. In this paper we assume the aligned case, $\gamma = 0$, so $\cos\gamma = 1$.

$$\Delta v_{A,1} = \frac{\beta m_{\text{sc}} v_{\text{rel}}}{m_A}. \tag{7}$$

We also define an acceleration proxy based on cadence,

$$a_{\max} \approx \frac{\beta m_{\text{sc}} v_{\text{rel}}}{m_A \Delta t_{\text{cad}}}, \tag{8}$$

where Δt_{cad} is the minimum time between successive impacts. Together, these proxies link the continuous VT solution to a control authority equivalent to an MKI campaign and a cadence limit.

In the numerical examples we set $\beta = 1$, so each impact transfers only the incident spacecraft momentum. In practice, β can exceed one and depends on impact geometry, surface properties, porosity, and ejecta behavior. For instance, DART-based estimates found $\beta = 3.6$ [17]. We treat the cadence Δt_{cad} as a mission-level constraint, not just a numerical value we can change. Therefore, enforcing a minimum time between impacts prevents unrealistically dense impact sequences from the compiler. From a mission perspective, it also leaves time for orbit determination and retargeting, and gives a way to adapt if the measured momentum transfer differs from the conservative $\beta = 1$ assumption.

The two VT optimization modes answer different design questions. In the fixed-budget maximum-deflection

mode, N, v_{rel}, and Δt_{cad} define the control authority $\Delta V_{A,\max} = N\Delta v_{A,1}$, and we ask how much final deflection that authority can produce. Here N is an authority level, not necessarily the final number of compiled impacts. In the minimum-Δv mode, no fixed authority level is imposed; we instead minimize ΔV_A needed to reach the prescribed target band. In both modes, the final impact epochs, directions, magnitudes, and count are obtained only after midpoint transcription and greedy compilation.

C. Fixed-Budget Maximum-Deflection VT Problem

The fixed-budget maximum-deflection problem asks how much final deflection can be produced by a prescribed, MKI-equivalent control budget:

$$\max_{\mathbf{x}_A(\cdot),\mathbf{u}(\cdot)} \frac{1}{2}\left\|\mathbf{r}_A(t_f) - \mathbf{r}_E(t_f)\right\|_2^2, \tag{9}$$

subject to (2), the fixed initial condition $\mathbf{x}_A(t_0) = \mathbf{x}_{A,0}$, the control bound $\|\mathbf{u}(t)\|_2 \le 1$, and the budget

$$a_{\max}\int_{t_0}^{t_f} \|\mathbf{u}(t)\|_2\, dt \le \Delta V_{A,\max}. \tag{10}$$

We solve this nonlinear problem with SCP [18, 19]. At iteration i, the current control history $\{\mathbf{u}_k^{(i)}\}_{k=0}^{K-1}$ is propagated with the nonlinear two-body dynamics to obtain the reference trajectory $\{\mathbf{x}_k^{(i)}\}_{k=0}^{K}$. We define

$$\delta\mathbf{u}_k = \mathbf{u}_k - \mathbf{u}_k^{(i)}, \qquad \mathbf{r}_f^{(i)} = \mathbf{r}_{A,K}^{(i)} - \mathbf{r}_E(t_f), \tag{11}$$

where $\mathbf{r}_{A,K}^{(i)}$ is the terminal asteroid position on this reference trajectory. Small perturbations about the reference trajectory satisfy the local linear model

$$\delta\mathbf{x}_{k+1} = \mathbf{A}_k^{(i)}\delta\mathbf{x}_k + \mathbf{B}_k^{(i)}\delta\mathbf{u}_k, \qquad \delta\mathbf{x}_0 = \mathbf{0}. \tag{12}$$

Here, $\mathbf{A}_k^{(i)}$ and $\mathbf{B}_k^{(i)}$ describe the first-order effect of state and control perturbations on interval k. Because the objective depends only on the asteroid position at t_f, we chain these local sensitivities to map each control perturbation directly to the final asteroid-position perturbation:

$$\delta\mathbf{r}_K = \sum_{k=0}^{K-1} \mathbf{S}_k^{(i)}\delta\mathbf{u}_k, \tag{13}$$

where $\mathbf{S}_k^{(i)} \in \mathbb{R}^{3\times 3}$ is the first-order sensitivity of the final asteroid position to a control change on interval k.

To build $\mathbf{S}_k^{(i)}$, we consider the incremental effect of changing only the control on interval k. Since no earlier control perturbation is applied in this isolated calculation, the state perturbation entering interval k is $\delta\mathbf{x}_k = \mathbf{0}$. The local model then gives $\delta\mathbf{x}_{k+1} = \mathbf{B}_k^{(i)}\delta\mathbf{u}_k$. This perturbation is propagated from node $k+1$ to the final node, and only the position rows are retained:

$$\mathbf{S}_k^{(i)} = \begin{bmatrix}\mathbf{I}_3 & \mathbf{0}_{3\times 3}\end{bmatrix}\mathbf{A}_{K-1}^{(i)}\mathbf{A}_{K-2}^{(i)}\cdots\mathbf{A}_{k+1}^{(i)}\mathbf{B}_k^{(i)} \in \mathbb{R}^{3\times 3}. \tag{14}$$

where $\begin{bmatrix}\mathbf{I}_3 & \mathbf{0}_{3\times 3}\end{bmatrix}$ keeps the position rows of the state. The full terminal perturbation is then obtained by summing these interval contributions over the complete control history. This simplified form lets the SCP subproblem perform faster by removing the intermediate state perturbations as decision variables while retaining the linearized effect of every control interval on the final miss vector.

The objective is linearized by writing $\mathbf{r}_f \approx \mathbf{r}_f^{(i)} + \delta\mathbf{r}_K$. Then

$$\frac{1}{2}\|\mathbf{r}_f\|_2^2 \approx \frac{1}{2}\|\mathbf{r}_f^{(i)}\|_2^2 + \left(\mathbf{r}_f^{(i)}\right)^T\delta\mathbf{r}_K.$$

The first term is constant inside the SCP subproblem, so it is dropped. Using the condensed map for $\delta\mathbf{r}_K$, the convex subproblem becomes

$$\max_{\{\delta\mathbf{u}_k\}} \left(\mathbf{r}_f^{(i)}\right)^T \sum_{k=0}^{K-1}\mathbf{S}_k^{(i)}\delta\mathbf{u}_k \tag{15}$$

subject to

$$\|\mathbf{u}_k^{(i)} + \delta\mathbf{u}_k\|_2 \le 1, \tag{16}$$

$$a_{\max}\sum_{k=0}^{K-1}\|\mathbf{u}_k^{(i)} + \delta\mathbf{u}_k\|_2\Delta t_k \le \Delta V_{A,\max}, \tag{17}$$

$$\|\delta\mathbf{u}_k\|_2 \le \Delta_u^{(i)}. \tag{18}$$

Here $\Delta_u^{(i)}$ is the trust-region radius. After each SCP step, we propagate the candidate control history with the actual nonlinear dynamics and accept it only when the nonlinear improvement agrees with the model prediction.

The maximum-deflection problem is initialized with the Conway near-optimal impulsive deflection approximation [2]. In this approximation, the two-body state transition matrix (STM) from the initial epoch to CA gives the terminal position response to an initial velocity impulse. The dominant direction is obtained from the principal eigenvector associated with $\Phi_{rv}^T\Phi_{rv}$, where Φ_{rv} maps an initial velocity perturbation into a terminal position perturbation. This direction is used only to initialize the SCP solve; the SCP iterations then optimize the full distributed VT history.

D. Minimum-Δv VT Problem

The minimum-Δv problem minimizes the asteroid velocity change needed to reach a prescribed terminal deviation,

$$\min_{\mathbf{x}_A(\cdot),\mathbf{u}(\cdot)} \Delta V_A = a_{\max}\int_{t_0}^{t_f}\|\mathbf{u}(t)\|_2\, dt, \tag{19}$$

subject to (2), the initial condition, $\|\mathbf{u}(t)\|_2 \le 1$, and the target band

$$d_{\text{req}} \le \left\|\mathbf{r}_A(t_f) - \mathbf{r}_E(t_f)\right\|_2 \le d_{\text{req}} + d_{\text{band}}. \tag{20}$$

This upper bound keeps the optimizer from meeting the minimum required deviation through unnecessarily large overshoot. Here d_{req} is the required terminal deviation and d_{band} is the allowed band width defined as a percentage of d_{req}.

Using the condensed map in (13), the SCP subproblem approximates $\mathbf{r}_f \approx \mathbf{r}_f^{(i)} + \delta\mathbf{r}_K$.

The minimum-deviation constraint is nonconvex, so we reuse the first-order squared-distance linearization from the maximum-deflection problem, but here as a constraint rather than as the objective.

$$g^{(i)}(\delta\mathbf{r}_K) = \frac{1}{2}\|\mathbf{r}_f^{(i)}\|_2^2 + \left(\mathbf{r}_f^{(i)}\right)^T \delta\mathbf{r}_K. \tag{21}$$

The resulting convex subproblem is

$$\min_{\{\delta\mathbf{u}_k\}} \sum_{k=0}^{K-1} \|\mathbf{u}_k^{(i)} + \delta\mathbf{u}_k\|_2 \Delta t_k \tag{22}$$

subject to

$$\|\mathbf{u}_k^{(i)} + \delta\mathbf{u}_k\|_2 \le 1, \qquad \|\delta\mathbf{u}_k\|_2 \le \Delta_u^{(i)}, \tag{23}$$

$$g^{(i)}(\delta\mathbf{r}_K) \ge \frac{1}{2} d_{\text{req}}^2, \tag{24}$$

$$\|\mathbf{r}_f^{(i)} + \delta\mathbf{r}_K\|_2 \le d_{\text{req}} + d_{\text{band}}. \tag{25}$$

The formulation above is the nominal SCP subproblem. In the implementation, a temporary nonnegative slack may be added to the linearized minimum-deviation constraint to avoid infeasible intermediate SCP subproblems.

The minimum-Δv target problem is initialized with a contiguous bang-arc warm start. We first build the terminal sensitivity matrices $\mathbf{S}_k^{(0)}$, where each matrix maps a control perturbation on interval k to a first-order terminal position perturbation. The most sensitive interval is selected as

$$k^\star = \arg\max_k \left\|\mathbf{S}_k^{(0)}\right\|_2, \tag{26}$$

where the matrix norm is the spectral norm. Using the Conway initial direction $\hat{\mathbf{u}}_0$, we define the terminal-deviation direction

$$\hat{\mathbf{r}}_f = \frac{\mathbf{S}_{k^\star}^{(0)} \hat{\mathbf{u}}_0}{\left\|\mathbf{S}_{k^\star}^{(0)} \hat{\mathbf{u}}_0\right\|_2}. \tag{27}$$

The sign of $\hat{\mathbf{r}}_f$ is chosen so that it increases the terminal asteroid–Earth deviation. Each interval is then scored by how strongly it can change the terminal position in this direction:

$$s_k = \left\| \left(\mathbf{S}_k^{(0)}\right)^T \hat{\mathbf{r}}_f \right\|_2. \tag{28}$$

The warm start places a full-throttle arc around one of the highest-scoring intervals,

$$\|\mathbf{u}_k^{(0)}\|_2 = \begin{cases} 1, & k_0 \le k \le k_1, \\ 0, & \text{otherwise}, \end{cases} \tag{29}$$

where k_0 and k_1 are the start and end indices of the contiguous arc. The arc length is adjusted by nonlinear propagation until the warm start is near the target band. This warm start only provides the initial control profile for the SCP optimizer; it does not force the final solution to remain bang-off-bang.

E. Step 2: Dense Midpoint Impulse Transcription

Once we have a VT solution, we replace each ZOH interval with an equivalent impulse placed at the interval midpoint,

$$\tau_k = \frac{t_k + t_{k+1}}{2}, \qquad \delta\mathbf{v}_{A,k}^{\text{dense}} = a_{\max}\mathbf{u}_k \Delta t_k. \tag{30}$$

This gives the dense impulse sequence

$$\mathcal{D}_A = \left\{ \left(\tau_k, \delta\mathbf{v}_{A,k}^{\text{dense}}\right) \right\}_{k=0}^{K-1}. \tag{31}$$

The impulsive nonlinear two-body propagation simply applies the impulses at each midpoint time

$$\mathbf{r}_A(\tau_k^+) = \mathbf{r}_A(\tau_k^-), \qquad \mathbf{v}_A(\tau_k^+) = \mathbf{v}_A(\tau_k^-) + \delta\mathbf{v}_{A,k}^{\text{dense}}. \tag{32}$$

This dense sequence is an impulsive approximation of the VT solution.

F. Step 3: Greedy MKI Campaign Compilation

The MKI campaign introduced above is written as the finite sequence of impulsive velocity changes applied to the asteroid,

$$C_A = \left\{ \left(T_j, \Delta\mathbf{v}_{A,j}^{\text{comp}}\right) \right\}_{j=1}^{N_{\text{greedy}}}, \tag{33}$$

where T_j is the epoch of the j-th compiled impact, $\Delta\mathbf{v}_{A,j}^{\text{comp}}$ is the velocity increment it delivers, and N_{greedy} is the total number of compiled impacts. Each campaign therefore fixes the number of impacts, their epochs, delivered impulse magnitudes, $\Delta v_{A,j}^{\text{comp}} = \left\|\Delta\mathbf{v}_{A,j}^{\text{comp}}\right\|_2$, and their impulse directions, $\hat{\mathbf{d}}_j = \Delta\mathbf{v}_{A,j}^{\text{comp}} / \left\|\Delta\mathbf{v}_{A,j}^{\text{comp}}\right\|_2$. These directions are ideal delivered-impulse requirements. Transfer feasibility to achieve them is left to later mission design.

Recall the dense impulse sequence $\mathcal{D}_A$ from Step 2. We keep only the nonzero candidates,

$$\mathcal{I} = \{i_1, i_2, \ldots, i_M\}, \qquad \left\|\delta\mathbf{v}_{A,i_m}^{\text{dense}}\right\|_2 > \epsilon_{\Delta v}. \tag{34}$$

Here, $\epsilon_{\Delta v}$ is a small threshold used to remove negligible dense impulses. The index i_m is the original dense grid index of the m-th retained impulse. Thus, m counts the retained impulses, while i_m points back to the original dense grid.

The greedy compiler then groups consecutive retained impulses into compiled impacts. For compiled impact j,

$$\mathcal{B}_j = \left\{ i_{m_j}, \ldots, i_{l_j} \right\}. \tag{35}$$

Here, m_j and l_j are positions in the retained list $\mathcal{I}$. The index m_j marks the first retained impulse assigned to impact j, and l_j marks the last one. Therefore, i_{m_j} and i_{l_j} are the corresponding indices on the original dense grid. For the nominal compilation, Δt_{cad} and $\Delta v_{A,1}$ match the assumptions used to build the Step 1 VT control output. The same dense sequence $\mathcal{D}_A$ can also be compiled again

with different mission assumptions, such as a different relative impact speed or cadence limit.

Each group of dense impulses assigned to one compiled impact must remain local in time,

$$\tau_{i_{l_j}} - \tau_{i_{m_j}} \le \Delta t_{\mathrm{cad}}, \tag{36}$$

where $\tau_{i_{m_j}}$ and $\tau_{i_{l_j}}$ are the midpoint times of the first and last dense impulses in group j. This condition limits the time span of the dense impulses represented by one compiled impact. It does not directly impose the spacing between the final compiled impact times T_j.

Each group must also stay within the velocity increment that one impactor can deliver,

$$\left\| \sum_{r=m_j}^{l_j} \delta \mathbf{v}_{A,i_r}^{\mathrm{dense}} \right\|_2 \le \Delta v_{A,1}. \tag{37}$$

Together, these two conditions mean that one compiled impact can replace only a nearby group of dense VT impulses, and only if their vector sum is small enough to be delivered by one impactor.

If a single dense midpoint impulse is already larger than $\Delta v_{A,1}$, then that impulse cannot be delivered by one feasible impactor. In that case, the compiler keeps it as a single impulse for reporting and nonlinear two body propagation, but flags the campaign as violating the one impactor Δv limit. No such violation occurs in the numerical cases reported below.

The compiled impulse for impact j is the vector sum of the dense impulses in its group,

$$\Delta \mathbf{v}_{A,j}^{\mathrm{comp}} = \sum_{r=m_j}^{l_j} \delta \mathbf{v}_{A,i_r}^{\mathrm{dense}}. \tag{38}$$

The compiled impact time is chosen from the dense midpoint times in the same group. We first compute the magnitude weighted average time,

$$\bar{T}_j = \frac{\sum_{r=m_j}^{l_j} \left\| \delta \mathbf{v}_{A,i_r}^{\mathrm{dense}} \right\|_2 \tau_{i_r}}{\sum_{r=m_j}^{l_j} \left\| \delta \mathbf{v}_{A,i_r}^{\mathrm{dense}} \right\|_2}. \tag{39}$$

The compiled impact time T_j is then chosen as the dense midpoint time in the group closest to $\bar{T}_j$. Larger dense impulses therefore have more influence on the selected impact time. Collecting all groups gives the compiled campaign C_A, and the number of compiled impacts N_{greedy} is the number of groups. The compiled campaign is then propagated with the same nonlinear impulsive dynamics used for the dense sequence.

This campaign provides only the impact epochs and delivered velocity increments, not the spacecraft transfer arcs. If needed, each impulse can be converted to a required relative impact velocity through the same momentum transfer relation that defines $\Delta v_{A,1}$, and then added to the asteroid velocity at the impact epoch to give the ideal pre-impact velocity for that impactor.

For the comparisons below, d_{VT}, d_{greedy}, and d_{STM} denote the final Earth–asteroid deviations obtained from the continuous VT propagation, the compiled campaign propagation, and the STM direct-allocation propagation, respectively.

G. Baseline: Direct Impulse Allocation

As a comparison, we also build a direct impulse allocation reference case based on the STM. This reference case does not use the dense VT impulse sequence or the greedy compiler. Instead, it places a prescribed number of full strength impulses directly on the same candidate midpoint grid used in Step 2. The purpose is to compare the compiled campaign with a direct STM based allocation along the Step 1 terminal deflection direction.

Let τ_ℓ be one candidate impact epoch. The STM from τ_ℓ to t_f gives the terminal position response to a velocity impulse at that epoch. We denote this velocity-to-position block by $\mathbf{\Phi}_{rv,\ell}$, so that

$$\delta \mathbf{r}_K \approx \sum_{\ell \in \mathcal{S}_N} \mathbf{\Phi}_{rv,\ell} \Delta \mathbf{v}_{A,\ell}, \tag{40}$$

where $\mathcal{S}_N$ is the selected set of N impact epochs.

The epoch selection uses the same idea as the sensitivity scoring used in the warm start. In the implementation used here, the terminal direction $\hat{\mathbf{r}}_f$ is taken as the normalized Step 1 SCP terminal deflection direction,

$$\hat{\mathbf{r}}_f = \frac{\mathbf{r}_{A,\mathrm{SCP}}(t_f) - \mathbf{r}_{A,\mathrm{nom}}(t_f)}{\left\| \mathbf{r}_{A,\mathrm{SCP}}(t_f) - \mathbf{r}_{A,\mathrm{nom}}(t_f) \right\|_2}. \tag{41}$$

Here, $\mathbf{r}_{A,\mathrm{SCP}}(t_f)$ is the final asteroid position from the Step 1 SCP solution, and $\mathbf{r}_{A,\mathrm{nom}}(t_f)$ is the nominal forced impact asteroid final position. For this terminal direction, we define

$$\mathbf{q}_\ell = \mathbf{\Phi}_{rv,\ell}^T \hat{\mathbf{r}}_f, \qquad s_\ell = \|\mathbf{q}_\ell\|_2. \tag{42}$$

The vector $\mathbf{q}_\ell$ gives the local impulse direction that most increases the terminal deviation along $\hat{\mathbf{r}}_f$, and s_ℓ measures the strength of that effect. The candidate epochs are ranked by s_ℓ, and the highest-scoring epochs are accepted until N epochs have been selected, while enforcing the minimum-cadence constraint between selected impacts.

For each selected epoch, the impulse is applied at the maximum velocity increment $\Delta v_{A,1}$ that one interceptor can deliver, in the direction of $\mathbf{q}_\ell$:

$$\Delta \mathbf{v}_{A,\ell} = \Delta v_{A,1} \frac{\mathbf{q}_\ell}{\|\mathbf{q}_\ell\|_2}, \qquad \ell \in \mathcal{S}_N. \tag{43}$$

The resulting impulse sequence is then propagated with the nonlinear dynamics to obtain the achieved final deviation. The baseline is used differently for the two VT cases. For maximum-deflection cases, the authority level N is prescribed. The STM baseline uses this same value as the

number of full-strength direct impulses, selects N epochs, and reports the nonlinear achieved deflection. This does not require the VT compiled campaign to contain exactly N impacts. For minimum-Δv cases, N is scanned as $N = 1, 2, \ldots, N_{\max}$. We then define N_{STM} as the smallest value of N for which nonlinear two-body propagation satisfies $d_{\mathrm{STM}}(N) \geq d_{\mathrm{req}}$. Thus, for minimum-$\Delta v$ cases, the STM direct allocation baseline estimates the smallest number of full-strength impactors needed to meet the target. This reference case is a sensitivity guided allocation along the Step 1 terminal deflection direction, not a globally optimal solution of the full mixed integer MKI campaign design problem.

H. *Evaluation: Nonlinear Propagation and Performance Metrics*

We compare three trajectories: the original VT solution, the dense midpoint impulse trajectory from $\mathcal{D}_A$, and the greedy compiled campaign trajectory from C_A. Every comparison uses full two-body dynamics propagation once the control or impulse sequence is fixed.

Our primary deflection metric is the terminal Earth–asteroid deviation at the nominal CA epoch,

$$d_f = \|\mathbf{r}_A(t_f) - \mathbf{r}_E(t_f)\|_2. \tag{44}$$

We also report the total ΔV_A, the compiled impact count N_{greedy}, whether the cadence and per-impact $\Delta v_{A,1}$ bounds are satisfied, and the retention ratio $d_{\mathrm{greedy}}/d_{\mathrm{VT}}$.

III. Numerical Study

We test the VT-to-MKI pipeline on a deflection case using the Apophis orbit with an assumed spherical target asteroid. The Apophis orbit provides the reference trajectory, while the assumed 100 m diameter and 2400 kg/m^3 density set the target mass used in the momentum transfer proxies. The study examines VT behavior, deflection retention after compilation, and comparison with the STM direct-allocation baseline.

A. *Test Matrix*

The batch sweep varies four mission parameters: the initial feasible control epoch t_0, the relative impact speed (affecting the per-impact $\Delta v_{A,1}$), the minimum cadence between interceptions, and the campaign authority. Table 1 lists the fixed assumptions and swept ranges, using a fixed 6-hr ZOH interval. The larger authority levels are included intentionally to test how the method behaves when many impact opportunities are available. This does not mean that realistic MKI missions always need dozens of spacecraft; smaller cases, including $N = 10$, are also included in the sweep.

The maximum-deflection sweep gives 270 cases. The minimum-Δv sweep gives 432 attempts over target deviations of $2, 4, 6, 8, 10, 20, 30$, and $40R_E$. Here R_E denotes one Earth radius. Of these, 386 reach the requested deviation within 10^{-3} km, while 46 are excluded from count comparisons because the VT solution itself does not reach the target.

Table 1. Numerical study assumptions and swept mission parameters.

Quantity	**Value / range**
Reference orbit	Apophis
Modeled asteroid diameter	100 m
Modeled asteroid density	2400 kg/m^3
Spacecraft mass	1000 kg
Momentum enhancement, β	1.0
Control lead time	36–96 months
Relative impact speed	3, 5, 7 km/s
Minimum cadence	3, 5, 7 days
Fixed-budget authority level, N	10, 25, 50, 75, 100
Target deviation (Earth radii)	$2, 4, 6, 8, 10, 20, 30, 40R_E$
ZOH interval	fixed at 6 hr

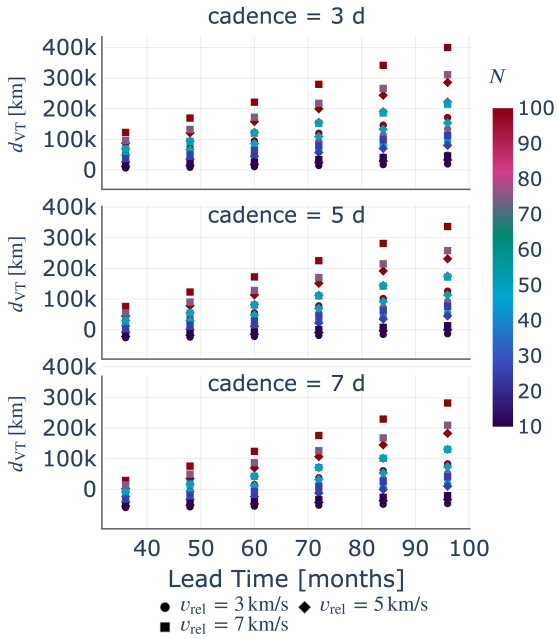

Fig. 3: Fixed-budget VT optimization design space before discrete compilation. The vertical axis gives the achieved VT deflection d_{VT}; marker shape denotes v_{rel}, and color denotes the authority level N.

B. *VT Optimization Design Space*

Fig. 3 shows the maximum-deflection design space before discrete compilation. Across the 270 cases, the achievable deflection ranges from 6,324 km to 399,697 km, with a median of 72,051 km and a median ΔV_A of 178.6 mm/s. As expected, longer lead time, larger authority, and higher impact speed increase the achievable deflection, while longer cadence reduces it by limiting placement opportunities and lowering the $a_{\max}$ proxy. Fig. 4 shows the minimum-Δv design space, where the required VT effort increases with target deviation and generally decreases with longer lead time.

For the successful minimum-Δv cases, the required VT effort has a median of 141.7 mm/s and ranges from 12.5 to 1257.7 mm/s. The lower-bound deviation constraint is active in all successful cases, so the optimizer reaches

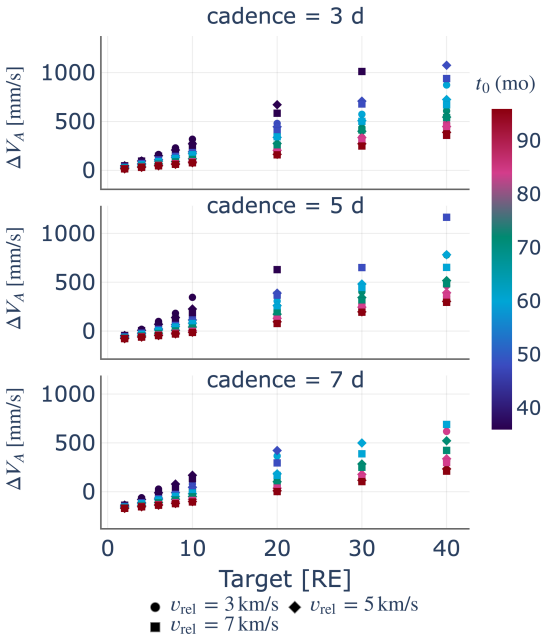

Fig. 4: Minimum-Δv VT optimization design space. Required ΔV_A is plotted against target deviation; marker shape denotes v_{rel}, and color denotes lead time t_0. Larger targets require more effort, while longer lead times generally reduce the required effort.

the deflection target with little excess margin. This is important because as we will see a minimum-effort continuous profile is not necessarily a minimum-count discrete campaign.

C. *Compiled Campaign Retention*

Fig. 5 tests whether the compiled campaign preserves the continuous VT deflection by plotting the percent difference between the compiled and continuous final deviations. The greedily compiled deviations differ from the continuous VT deviations by only a fraction of a percent: the median difference is about +0.1% for maximum deflection and +0.2% for the successful minimum-Δv cases. Thus, this indicates the VT profile retains enough timing and direction information, and is sparse enough to compile into a finite impulse campaign with essentially no loss in the final achieved deflection.

D. *Comparison with STM Direct Allocation*

A retention near one shows the three-step process is self-consistent (the same deflection achieved across all three-steps), but not necessarily that it is competitive. To test that, we compare it against the previously defined STM direct allocation, which does not use the dense impulse sequence or the greedy compiler, but it does use the Step 1 terminal deflection direction when ranking candidate impact epochs.

In the maximum-deflection cases both methods use the same assumed authority N, so they are directly comparable. Fig. 6 shows the percent difference between the

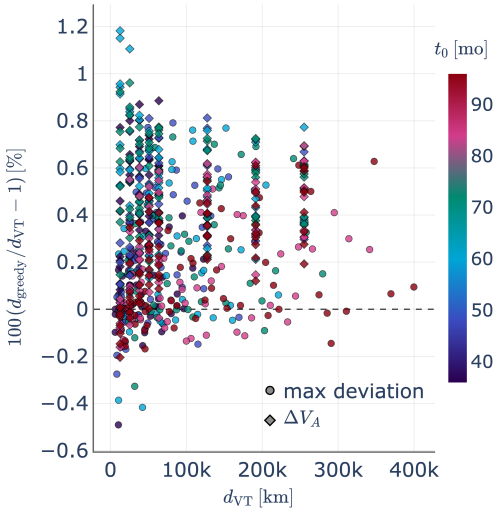

Fig. 5: Difference between the compiled campaign and the continuous VT solution. The horizontal axis gives d_{VT}, and the vertical axis gives $100(d_{\text{greedy}}/d_{\text{VT}} - 1)$. Marker shape distinguishes maximum-deflection and minimum-Δv cases, color denotes lead time t_0, and the dashed line marks zero difference.

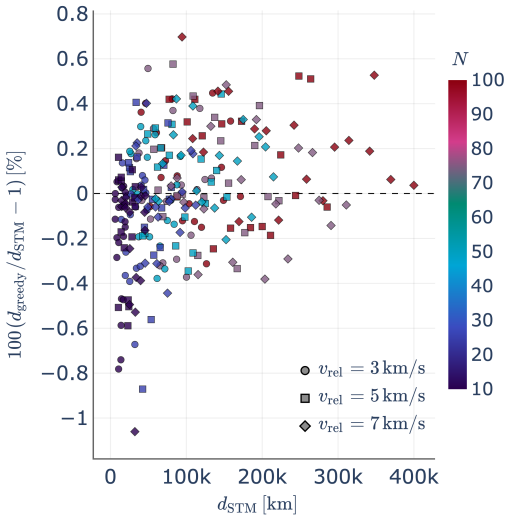

Fig. 6: Maximum-deflection comparison with the STM direct allocation baseline. The horizontal axis gives d_{STM}, and the vertical axis gives $100(d_{\text{greedy}}/d_{\text{STM}} - 1)$. Marker shape denotes v_{rel}, color denotes authority level N, and the dashed line marks zero difference.

greedy compiled VT campaign and the STM direct allocation baseline. The median value of $100(d_{\text{greedy}}/d_{\text{STM}} - 1)$ is near zero, and 269 of the 270 cases fall within 1% of the STM result. The VT-to-MKI pipeline thus recovers essentially the same design space as the baseline STM-based allocation.

The minimum-Δv comparison is different in kind, because the question is no longer the deflection obtained at a fixed count but how many impact opportunities are needed to reach the desired deviation. Fig. 7 shows the greedy compiler using more impactors than STM. The median ratio $N_{\text{greedy}}/N_{\text{STM}}$ is 1.143 over the 386 cases with valid STM count comparisons; the greedy count is larger in 274 cases, equal in 107, and smaller in 5.

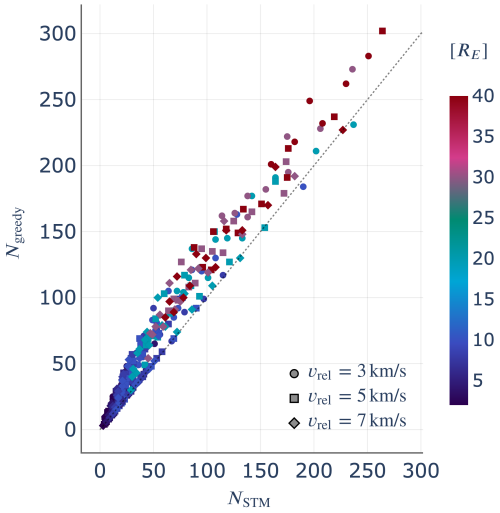

Fig. 7: Minimum-Δv comparison between greedy compiled count and STM direct-allocation count. The horizontal axis gives N_{STM}, the smallest STM count using full-size impacts that reaches the target, and the vertical axis gives N_{greedy}. Points above the 1:1 line require more compiled impacts than the STM count reference.

This gap is not a failure of the compiler, since the two methods answer related but different questions. The STM sweep searches directly for the smallest set of full-size impulses that crosses the desired deviation, and it accepts larger overshoot to get there, whereas the greedy compiler preserves the structure of the VT continuous minimum-effort profile, which more precisely tracks a desired deviation. The STM count is therefore a feasibility reference for the minimum count, while the greedy count reflects the cost of preserving the structure of the continuous VT solution as quantified next.

E. *Why Minimum-Δv Profiles Can Require Extra Impacts*

To quantify that cost, we can define an impact count based only on total effort,

$$N_{\text{LB}} = \left\lceil \frac{\Delta V_A}{\Delta v_{A,1}} \right\rceil. \tag{45}$$

This divides the total VT optimization ΔV_A by the maximum velocity increment one impactor can deliver, $\Delta v_{A,1}$. This count acts as a lower bound based only on effort

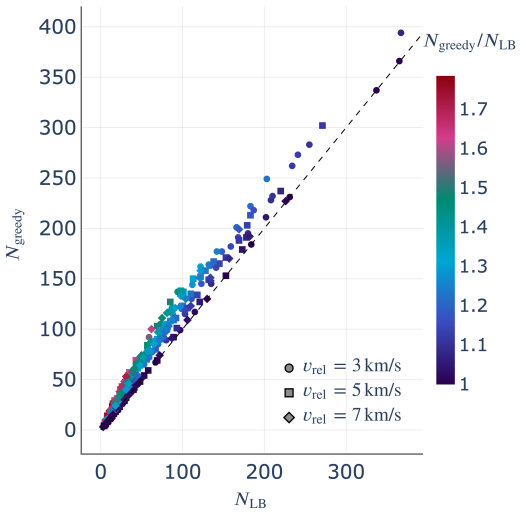

Fig. 8: Compiled count N_{greedy} compared with the lower bound based only on total effort N_{LB} for successful minimum-Δv cases. Points above the one-to-one line indicate cases where timing, cadence, or distributed small corrections require additional compiled impacts.

and it ignores timing, cadence, and grouping constraints. Fig. 8 compares N_{LB} with N_{greedy}. The greedy count is also usually larger because minimum-Δv continuous profiles whose lower deviation constraint is active often contain small distributed corrections that are useful continuously but difficult to merge into a small number of finite impacts.

Fig. 9 shows a challenging $40R_E$ target case with a 48-month lead time, $v_{\text{rel}} = 5$ km/s, and a 3-day cadence. The campaign succeeds with $N_{\text{greedy}} = 302$ interceptors and $\Delta V_A = 1075.44$ mm/s, but many corrections remain small enough to form additional lower magnitude compiled impacts on the order of $1, 2$ mm/s.

F. *Representative Campaign Examples*

We now examine two individual campaigns to show how a continuous VT profile turns into a discrete impact campaign. Unless stated otherwise, Step 3 uses the nominal compiler settings, meaning the same v_{rel} and Δt_{cad} assumptions used to set up the Step 1 case.

The first example is a maximum deflection case with an 84-month lead time, authority level $N = 75$, a 7 km/s impact speed, and a 7-day cadence. Fig. 10 shows the VT profile and the impulses selected by the compiler. If the VT profile were an ideal sequence of fully saturated arcs and zero-control arcs, the compiler would recover the original authority level $N = 75$. In this case, the profile is close to that structure but not exact. After midpoint transcription, some of the effort is split into additional small increments, so the nominal compilation gives $N_{\text{greedy}} = 77$ impact opportunities instead of 75. The total asteroid velocity change is $\Delta V_A = 417.8$ mm/s.

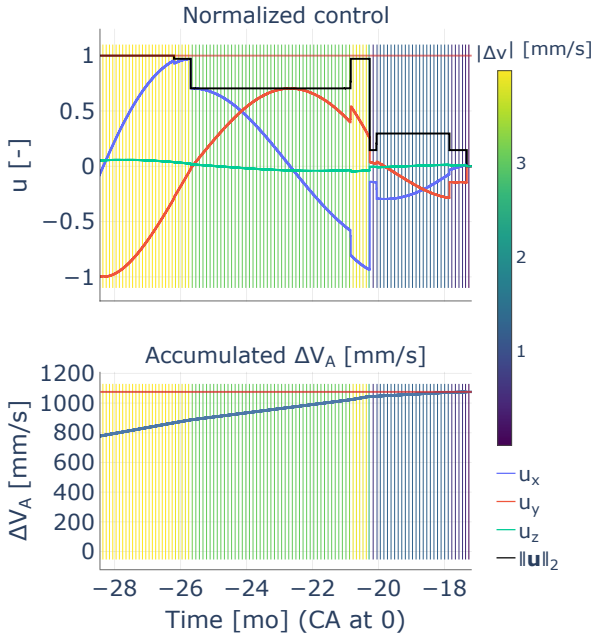

Fig. 9: Minimum-Δv VT profile and subset of the 302 compiled impulses for a $40R_E$ deflection with a 48-month lead time, $v_{\text{rel}} = 5$ km/s, and 3-day cadence. The distributed small corrections in the continuous profile produce additional small compiled impacts.

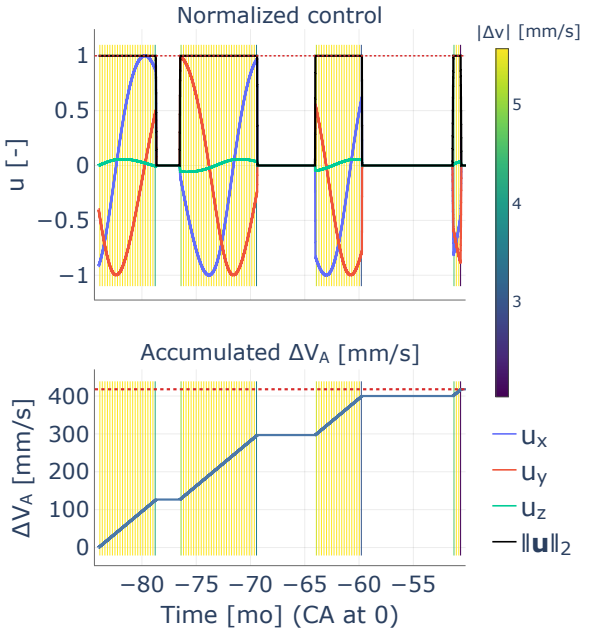

Fig. 10: Representative maximum deflection VT profile and nominal compiled impulses for the 84-month, $N = 75$, $v_{\text{rel}} = 7$ km/s, 7-day cadence case. The vertical markers show the impact opportunities selected by the greedy compiler from the continuous VT profile.

Propagating the compiled campaign with the nonlinear dynamics reproduces the continuous VT deflection almost exactly, with $d_{\text{VT}} = 233{,}399$ km and $d_{\text{greedy}} = 233{,}974$ km.

The second example is a minimum Δv case that reaches a $4R_E$ target deviation with a 36-month lead time, a 7 km/s impact speed, and a 3-day cadence. Fig. 11 shows that the active control window is compact and close to a single saturated arc, and Fig. 12 shows the matching trajectory. For minimum Δv, the number of impactors is not prescribed in Step 1. The useful reference is instead the effort equivalent count N_{LB}. Because the VT profile is already close to saturated, the nominal compiler recovers this count exactly: $N_{\text{LB}} = N_{\text{greedy}} = 18$. The achieved deviations also match closely, with $d_{\text{VT}} = 25{,}513$ km and $d_{\text{greedy}} = 25{,}550$ km, so the nominal compilation adds no extra impactors beyond the effort based estimate.

G. *Recompilation Under Different Mission Assumptions*

One advantage of the method is that the same VT solution can be recompiled under different assumed mission constraints without resolving the VT optimization. This is useful because the asteroid Δv delivered by each impact may vary with transfer feasibility. For example, a future mission design may find that some epochs allow higher impact speed than others. We demonstrate this idea with the 48-month, $N = 25$, 5 km/s, 7-day fixed-budget VT case, for which $d_{\text{VT}} = 32{,}615$ km and $\Delta V_A = 99.5$ mm/s. Table 2 compares three compiled MKI campaigns obtained from the same VT profile. The nominal campaign

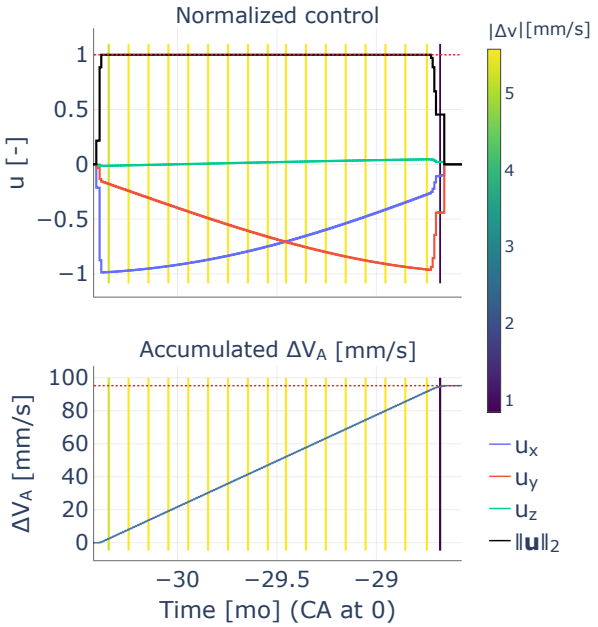

Fig. 11: Representative minimum Δv VT profile and nominal compiled impulses for the $4R_E$ target case with 36-month lead time, $v_{\text{rel}} = 7$ km/s, and 3-day cadence. The compact control arc compiles without count increase, giving $N_{\text{LB}} = N_{\text{greedy}} = 18$.

uses the original Step-1 assumptions for the fixed-budget case, namely authority level $N = 25$, constant $v_{\text{rel}} = 5$ km/s, and $\Delta t_{\text{cad}} = 7$ days. The two custom campaigns keep the same VT profile and cadence, but assign different

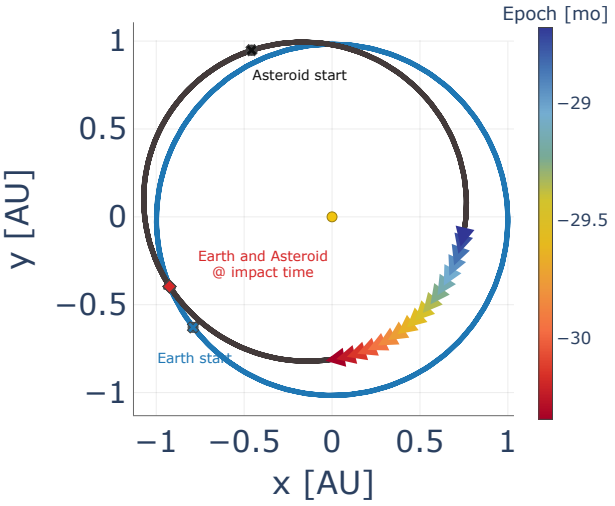

Fig. 12: Minimum Δv MKI campaign for a $4R_E$ target deviation in 36 months. The nominal compiled impacts use $v_{\text{rel}} = 7$ km/s and 3-day cadence, accrue near perihelion, and achieve the required deviation.

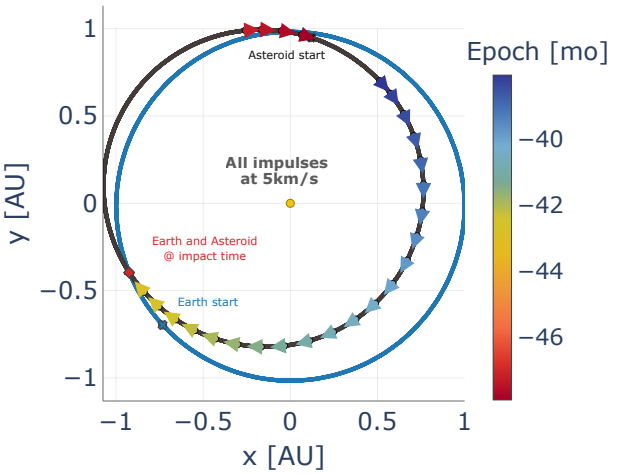

Fig. 13: 48-month, authority level $N = 25$, 7-day fixed-budget VT profile compiled with nominal constant $v_{\text{rel}} = 5$ km/s over the active window, giving $N_{\text{greedy}} = 26$ compiled impactors.

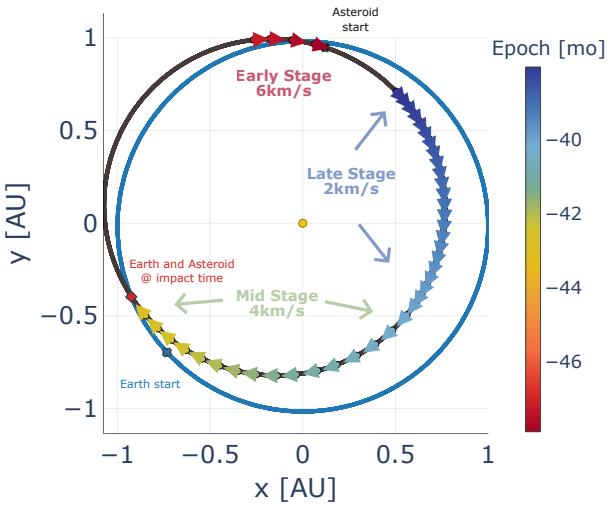

Fig. 14: 48-month, authority level $N = 25$, 7-day fixed-budget VT profile compiled with custom $v_{\text{rel}} = 6, 4$, and 2 km/s over the early, middle, and late active segments, giving $N_{\text{greedy}} = 43$ compiled impactors.

relative impact speeds to three time segments. In all three campaigns the final deflection is essentially unchanged after compilation, but the number and density of compiled impacts change with the assumed relative speed. Figs. 13–15 show the nominal and custom compiled campaigns obtained from the same VT profile.

Table 2. Recompilation of the same VT profile under different mission assumptions, with $d_{\text{VT}} = 32{,}615$ km.

Campaign	**Relative impact speeds by time segment**	N_{greedy}	d_{greedy} [km]
Nominal	Constant 5 km/s over the full active window	26	32,598
Custom 1	Time-segment speeds: 6 km/s on $[-48, -45]$ mo; 4 km/s on $[-45, -40]$ mo; 2 km/s on $[-40, -5]$ mo	43	32,604
Custom 2	Time-segment speeds: 4 km/s on $[-48, -45]$ mo; 3 km/s on $[-45, -40]$ mo; 6 km/s on $[-40, -5]$ mo	37	32,602

This example shows the practical value of separating VT optimization from campaign compilation. The continuous VT profile defines the time history of the required asteroid deflection, while the greedy compiler maps that profile into different finite MKI campaigns as mission assumptions change. If a time segment is assigned a lower relative impact speed, each impact delivers less asteroid Δv, so the compiler needs more impacts in that segment. A higher relative impact speed has the opposite effect. In the ideal case of a saturated bang-off-bang VT profile, with impacts at the same maximum relative speed used to define the authority level N, the compiled count should recover the nominal authority level, $N_{\text{greedy}} \approx N$.

IV. Discussion

The results support the main role of VT as an intermediate representation for MKI campaign design. Its value is that the impact count, timing, and distribution do not have to be fixed before solving the deflection problem. Instead, VT produces a continuous control profile that achieves the objective, and the compiler converts that profile into a finite campaign under the selected cadence limits and delivered asteroid-Δv bounds.

In the maximum-deflection cases, the continuous VT solution, the compiled campaign, and the STM direct allocation reference case all reach almost the same terminal deflection along the Step 1 deflection direction. This agreement shows that the VT profile contains enough timing and direction information to be reduced to finite impacts without losing the achieved deflection.

The minimum-Δv cases behave differently because the optimizer reaches the target with little excess margin.

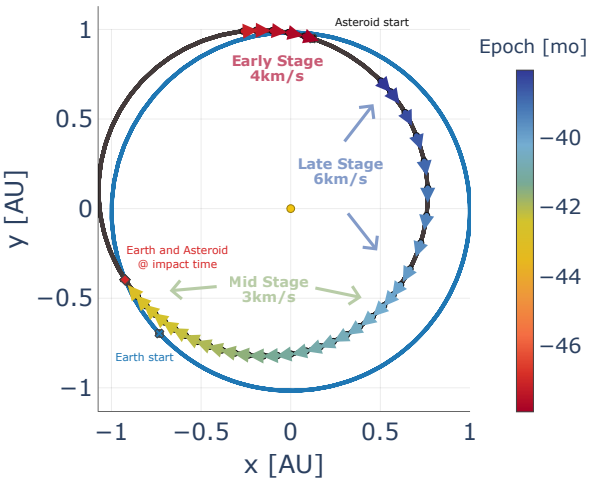

Fig. 15: 48-month, authority level $N = 25$, 7-day fixed-budget VT profile compiled with custom $v_{\text{rel}} = 4, 3$, and 6 km/s over the early, middle, and late active segments, giving $N_{\text{greedy}} = 37$ compiled impactors.

The useful part of the control profile is often close to bang-off-bang, but the active minimum-deviation constraint can introduce small corrections that keep the final deviation near the target instead of overshooting it. These corrections reduce continuous effort, but they are not always easy to merge into a small number of impact opportunities. This is why N_{greedy} can be larger than both N_{LB} and N_{STM}.

The STM direct allocation case should therefore be understood as a sensitivity based reference, not as a replacement for the VT method. STM directly places full-strength impulses and may allow more overshoot beyond the target. VT instead gives the continuous effort, timing structure, and compiled campaign associated with the minimum-Δv solution.

V. Limitations and Future Work

The present pipeline is a campaign design framework, not a full mission design tool. It prescribes the velocity increments that the asteroid must receive, but it does not compute spacecraft transfer trajectories, launch windows, or transfer feasibility. The greedy compiler is heuristic and does not prove global optimality for the full mixed-integer nonlinear MKI design problem. The terminal metric used here is Earth–asteroid deviation at the nominal CA epoch; B-plane or MOID targets could be used in the same framework but are not implemented in this study.

A useful consequence of the VT formulation is that a continuous VT profile can be recompiled under different mission assumptions, such as impact speeds that vary by epoch or different cadence limits. In this paper, those assumptions are prescribed compiler inputs rather than outputs of a transfer analysis. A mission-aware compiler would instead use interceptor reachability information, including transfers from Earth–Moon parking orbits and other escape-transfer options, to set or reject candidate impacts at each epoch. Another extension is to solve the Step-1 VT optimization with indirect optimal-control methods, which would reduce the dependence of the continuous VT profile on a prescribed time grid before the profile is transcribed into impulses. Further work also includes studying uncertainty in the momentum enhancement factor β and testing robustness to missed or off-nominal impacts.

VI. Conclusions

This paper presented virtual thrust as a way to design MKI campaigns without fixing the impact sequence in advance. The method first finds a continuous control profile that produces the desired asteroid deflection, and then compiles that profile into a finite set of impacts under the selected cadence and delivered asteroid-Δv bounds. The numerical results show that this split works well. In the maximum-deflection cases, the compiled campaigns retain essentially all of the VT deflection and closely match the direct allocation baseline based on the STM. In the minimum-Δv cases, VT reaches the requested deviation with little excess margin, while compilation requires only a modest increase in impact opportunities when small corrections cannot be merged. The results obtained show that VT can serve as a practical link between continuous optimal control and a discrete MKI campaign design.

VII. References

[1] NASA/JPL Center for Near-Earth Object Studies. NEO Basics: NEO Groups. `https://cneos.jpl.nasa.gov/about/neo_groups.html`, 2026. Accessed: 2026-05-21.

[2] B. Conway. Near-optimal deflection of earth-approaching asteroids. *AIAA*, 24, 2001.

[3] D. Izzo. Optimization of interplanetary trajectories for impulsive and continuous asteroid deflection. *Journal of Guidance, Control and Dynamics*, 30(2), 2007.

[4] E. Fantino, R. Flores, G. Donnarumma, and K. Canales, D. Howell. Direct low energy trajectories to near earth objects. *Acta Astronautica*, 229:333–344, 2024.

[5] C. Colombo, M. Vasile, and R. Gianmarco. Semi-analytical solution for the optimal low-thrust deflection of near earth objects. *Journal of Guidance, Control and Dynamics*, 32(3), 2009.

[6] R. Santos, M. Xu, Y. Zheng, X. He, and J. Ge. Conceptual design for deflecting a potentially hazardous asteroid with a space duster. *JSR*, 58, 2021.

[7] Y. Ketema. Asteroid deflection using a spacecraft in restricted keplerian motion. *Acta Astronautica*, 24, 2017.

[8] I. Ross and S. Park, S. an Porter. Gravitational effect of earth in optimizing delta v for deflecting

earth crossing asteroids. *Journal of Spacecraft and Rockets*, 38(5), 2001.

[9] J.P. Sanchez and C.R. McInnes. Synergistic approach of asteroid exploitation and planetary protection. *Advances in Space Research*, 49(4):667–685, February 2012.

[10] Edward T. Lu and Stanley G. Love. Gravitational tractor for towing asteroids. *Nature*, 438(7065):177–178, November 2005.

[11] G. Youtao, F. Yuhai, X. Bo, and W. Ying. Effectiveness analysis of multiple kinetic impacts of near-earth asteroids. *Advances in Space Research*, 69:2883–2892, 2022.

[12] A. N. Cohen, P. Lubin, D. Robertson, M. Boslough, S. Egan, and A. M. et. al. Stickle. Asteroid disruption and deflection simulations for multi-modal planetary defense. *Acta Astronautica*, 225:960–967, 2024.

[13] Ivan Bolzoni. Multiple kinetic impactor for deflection of potentially hazardous asteroids. M.s. thesis, Politecnico di Milano, Milan, Italy, 2021.

[14] P. Qiao, Y. Qian, and W. Chen. Sequential multiple kinetic impact strategy for asteroid deflection using an intermediate asteroid. *Advances in Space Research*, 77:7920–7936, 2026.

[15] Alessandro Massini, Noelle Elisabeth May, Moacir Fonseca Becker, and Andrea Capannolo. Earth-moon orbit design for a rapid-response asteroid interceptor system. In *30th International Symposium on Space Flight Dynamics*, Toulouse, France, June 2026.

[16] Lorenzo Chignoli. Multi-kinetic impactor for asteroid deflection with impact cratering physics uncertainties. M.s. thesis, Politecnico di Milano, Milan, Italy, 2019.

[17] Andrew F. Cheng, Harrison F. Agrusa, Brent W. Barbee, Alex J. Meyer, Tony L. Farnham, Sabina D. Raducan, Derek C. Richardson, et al. Momentum transfer from the DART mission kinetic impact on asteroid Dimorphos. *Nature*, 616:457–460, 2023.

[18] Yuanqi Mao, Michael Szmuk, and Behcet Acikmese. Successive convexification of non-convex optimal control problems and its convergence properties, 2016.

[19] Danylo Malyuta, Taylor P. Reynolds, Michael Szmuk, Thomas Lew, Riccardo Bonalli, Marco Pavone, and Behcet Acikmese. Convex optimization for trajectory generation, 2021.