

Beam Tracking for SIM-Assisted LEO Satellites: From Orbital Prediction to Wave-Domain Mobility Management

Chandan Kumar Sheemar, Giovanni Iacovelli, Wali Ullah Khan, Sourabh Solanki, and Symeon Chatzinotas,

Abstract—Low-Earth-orbit (LEO) constellations are becoming a key component of global broadband, direct-to-device, and machine-type connectivity, but their high mobility turns beam maintenance into a continuous control problem. Stacked intelligent metasurfaces (SIMs) can reduce radio-frequency chains and onboard digital processing by beamforming directly in the wave domain; however, their multilayer electromagnetic coupling makes rapid reconfiguration fundamentally different from phased-array steering. This article offers a system-level perspective on SIM beam tracking for LEO satellite communications. Building on a recently proposed one-shot geometry-aided tracker, we explain how predicted orbital geometry is turned into multilayer phase updates, present a unified predict–correct architecture, and compare four tracking paradigms—CSI-driven reactive, geometry-aided predictive, sensing-assisted, and hybrid physics/learning—in terms of the information, signaling, and computation they require and the errors they can correct; the latter two are discussed as research directions. A case study with a perturbed orbit then quantifies when prediction alone suffices: one-shot tracking updated every second retains about 85–91% of the rate of per-second full reoptimization with roughly 35–54 times less computation per update; adding a predictive full reoptimization only every 20 s raises this to about 90–95%; and one scalar SNR report per user and update limits the loss caused by a 2° attitude bias to a few percentage points, whereas open-loop prediction loses up to 21 points. A roadmap toward autonomous SIM payloads concludes the article.

Index Terms—Stacked intelligent metasurface, LEO satellite communications, beam tracking, wave-domain beamforming, orbital prediction, non-terrestrial networks.

I. INTRODUCTION

Low-Earth-orbit (LEO) satellite systems are rapidly becoming an important extension of terrestrial networks for broadband access, direct-to-device service, backhaul, maritime and aeronautical connectivity, and machine-type communications. Their relatively low propagation delay and global or regional visibility make them attractive where terrestrial infrastructure is unavailable or uneconomical. At the same time, LEO satellites move at several kilometers per second, so the satellite–terminal geometry, propagation delay, Doppler shift, visibility interval, and beam footprint evolve continuously. Beam management is therefore one of the defining challenges of LEO access networks, particularly when payload power, onboard computing, and signaling budgets are constrained [1]–[3].

The authors are with the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, L-1855 Luxembourg, Luxembourg.

Conventional multibeam satellites rely on phased arrays, hybrid analog/digital beamformers, or electronically steered antennas. These solutions are mature, but scaling them toward large apertures and dense coverage increases the number of radio-frequency (RF) chains, phase shifters, calibration paths, and baseband operations. Stacked intelligent metasurfaces (SIMs) provide an alternative architecture in which multiple programmable layers exploit inter-layer diffraction to perform beamforming directly on the electromagnetic wave [4]. Recent work has demonstrated multiuser wave-domain beamforming, double-stacked architectures, robust finite-resolution control, and satellite-mounted SIM operation [5]–[8]. Dynamic beam alignment has consequently emerged as a central challenge for practical satellite SIM deployment [9].

This creates an important paradox. SIMs are attractive for satellites because they can shift part of the beamforming complexity from digital baseband processing into the electromagnetic domain; yet LEO mobility requires the beamforming configuration to evolve continuously. An L -layer SIM with N meta-atoms per layer exposes LN coupled control variables rather than a conventional steering vector. Reoptimizing all of them from newly estimated channel state information (CSI) at every tracking instant can therefore erode the payload-level complexity advantage that motivated the SIM.

Fortunately, LEO mobility is fast but highly predictable. Satellite ephemeris, orbital dynamics, attitude information, and user location provide strong prior information about future beam geometry. Predictive beamforming has already been studied for conventional LEO architectures [10], while learning-based and dynamically scheduled beam-management methods can reduce repeated beam-search overhead [11], [12]. The open question is how these ideas should be reformulated when the front end is a multilayer wave-domain processor.

This article argues that *SIM-assisted LEO beam tracking should be viewed as predictive control of a dynamic electromagnetic configuration*. It is a system-level framework rather than the proposal of a new algorithm, as shown in Figure 1. The one-shot geometry-aided predictive tracker (GAPT) used in our case study was introduced in [13], which derives the update rule and its error bound for a single-feed SIM with a fixed one-second update interval and perfect attitude knowledge. Building on it, this article makes three contributions. First, it explains the geometry-to-SIM mapping at a level accessible to system designers, for both single-feed and multi-feed SIM transmitters, and clarifies which errors require a SIM update rather than baseband processing. Second, it organizes reactive,




Fig. 1. Concept of beam tracking for a SIM-assisted LEO satellite. Satellite motion changes the user directions and coverage footprint; orbital and user information predict the future geometry and drive the multilayer SIM reconfiguration so that the beams remain on the target users.

predictive, sensing-assisted, and learning-assisted tracking into a common predict–correct architecture and compares their requirements and failure modes (Table I); the sensing-assisted and learning-assisted paradigms are presented as research directions, not demonstrated methods. Third, it provides new results that identify where prediction alone suffices—as a function of the reconfiguration interval and of pointing error—and shows how a few scalar feedback reports per update extend that regime.

II. FROM DIGITAL BEAMFORMING TO WAVE-DOMAIN BEAM TRACKING

A. The Beamformer Has Become a Physical Processor

A conventional phased array can often be described through a steering or beamforming vector whose entries set the relative phase and amplitude across antenna elements. A SIM instead realizes the radiated field through cascaded interactions among programmable metasurface layers and fixed inter-layer propagation operators. Consequently, its output field depends jointly and nonlinearly on the complete multilayer configuration.

This distinction fundamentally changes beam tracking. Small angular motion in a conventional array can often be accommodated by updating a low-dimensional steering direction. In a SIM, a comparable displacement may require a coordinated change across many coupled meta-atoms because an update in one layer modifies the illumination of all

subsequent layers. We therefore use the term *electromagnetic configuration* to denote what is actually controlled—the LN programmable phase settings—together with the radiated field they produce. The fields inside the stack are not controlled directly; they follow from the feed excitation and the fixed propagation between layers. SIM beam tracking must predict how the radiated field should evolve, map that prediction into a feasible set of phase settings, and correct model and hardware errors with limited feedback.

B. Single-Feed and Multi-Feed SIM Transmitters

How many data streams a SIM can serve depends on how it is excited. With a single RF feed, the SIM shapes one radiation pattern that may have several lobes toward geographically separated users; it can form a steerable spot beam, a shaped coverage footprint, or a common beam for broadcast, multicast, and control channels, but it cannot spatially multiplex independent streams. With M feeds or a small active array at its input, the SIM acts as a wave-domain precoder that separates up to M streams [5], [6], so the users can be served by spatial multiplexing, possibly complemented by time and frequency multiplexing across beams. For tracking, the difference lies only in the quantity to be preserved: the per-user signal-to-noise ratio (SNR) for a single feed and the signal-to-interference-plus-noise ratio (SINR) for multiple feeds, both




Fig. 2. Unified architecture for SIM beam tracking in LEO satellite networks. Orbital physics provides the dominant open-loop prediction, while communication pilots, sensing, and feedback provide closed-loop correction.

of which depend on the same LN phases through the same cascaded propagation. The architecture and paradigms below therefore apply to both cases.

C. Why SIM Beam Tracking Is Different

Four characteristics make SIM-assisted LEO tracking distinctive. First, the control state is high dimensional: an L -layer SIM with N elements per layer exposes LN programmable phases even though the underlying satellite geometry is low dimensional. Second, the layers are electromagnetically coupled, so independent layer-by-layer steering rules are generally insufficient. Third, practical hardware imposes finite phase resolution, switching-speed limits, calibration errors, thermal drift, and possible amplitude–phase coupling. Finally, the satellite/user geometry evolves continuously, making stale SIM configurations particularly harmful for narrow beams and for interference-limited multiuser links.

Existing SIM research already addresses statistical CSI for LEO links [8] and robust finite-resolution control under imperfect CSI [7]. Beam tracking adds a further temporal dimension: the controller must determine *when* to update, *what* quantity should be tracked, and *which* errors can be handled by the SIM and which by the digital baseband.

III. LEO MOBILITY: FAST BUT PREDICTABLE

A distinguishing property of satellite mobility is that it is governed largely by orbital mechanics rather than random motion. Satellite ephemeris and attitude therefore provide accurate predictions of future user directions in the satellite body frame. For fixed or slowly moving terminals, satellite-state and, above all, attitude uncertainty dominate the tracking error; for aircraft, ships, vehicles, or UAVs, terminal mobility must also be predicted.

This predictability suggests a fundamentally different control philosophy from terrestrial beam tracking. Rather than repeatedly estimate the complete channel after the beam has moved, the satellite can first predict its future geometry and then use sparse communication or sensing measurements to correct the prediction. The resulting conceptual chain is *orbit and attitude* $\rightarrow$ *future geometry* $\rightarrow$ *target radiated field* $\rightarrow$ *SIM phase settings*. Beam tracking therefore becomes a fusion problem between a strong physics-based prior and lower-rate closed-loop observations.

The quantity being tracked need not be an angle. Depending on the service objective, a SIM controller may maintain user directions, received fields, SNR or SINR profiles, per-user rate targets, or entire coverage footprints. With dense user populations, users can also be grouped into service regions or beams: tracking K representative control points—cluster centers or samples of a footprint—instead of every user keeps the size of the tracking problem, and hence the per-update computation, independent of the population; the price is that users away from the control points are served by the pattern’s interpolation between them.

IV. A UNIFIED ARCHITECTURE FOR SIM BEAM TRACKING

Fig. 2 presents a unified architecture combining satellite ephemeris, attitude, user position or mobility, communication pilots, and optional sensing information. These inputs are converted into a predicted geometry and beam state from which the SIM tracking engine computes a feasible multilayer update. The updated configuration drives the wave-domain front end, while low-rate feedback—such as received SNR, beam quality, angular error, or tracking residual—provides closed-loop correction.




Fig. 3. Four complementary strategies for SIM beam tracking. Their relative value depends on pilot budget, ephemeris accuracy, sensing availability, hardware uncertainty, and user mobility.

The important architectural consequence is that full electromagnetic reoptimization does not need to operate at the fastest system time scale. Orbital prediction can account for most deterministic motion, whereas sparse feedback handles ephemeris error, attitude uncertainty, hardware drift, user motion, and imperfections of the electromagnetic model. Depending on the available onboard resources, the tracking engine may use full reoptimization, local model-based corrections, lookup tables, or learned residual updates. A simple and explicit switching rule is to apply local corrections while the predicted mismatch between the target and the achievable SNR or SINR profile remains small, and to trigger a full reoptimization when this mismatch, or the residual reported by the feedback, exceeds a threshold. The case study below shows where such a threshold becomes active.

V. TAXONOMY OF SIM BEAM-TRACKING STRATEGIES

The emerging design space can be organized into four complementary paradigms, summarized in Fig. 3 and compared in Table I.

A. CSI-Driven Reactive Tracking

The most direct strategy estimates the channel at each tracking instant and then reoptimizes the SIM configuration. It is general and naturally reacts to propagation effects not captured by geometry, but frequent channel acquisition, feedback, and iterative nonlinear optimization are costly for a high-dimensional SIM and a large user population, and the estimated CSI ages while the satellite moves. Reactive

tracking is therefore valuable when geometry is unavailable or multipath and unmodeled dynamics dominate, but inefficient as the sole tracking mechanism for predictable LEO motion.

B. Geometry-Aided Predictive Tracking

Geometry-aided tracking uses satellite ephemeris and terminal positions to predict future beam directions before the current configuration becomes stale. Predictive beamforming for conventional LEO systems has already shown that location and state information can reduce repeated feedback [10], while adaptive beam-tracking periods can avoid unnecessary updates [12]. For SIMs, the predicted geometry can be used either to regenerate the configuration from scratch or to correct the current configuration locally. The GAP of [13] follows the second route in four steps per update.

1) *Predict the geometry.* From two recent ephemeris points, a circular-orbit recurrence propagates the satellite position; together with Earth rotation and the nadir-pointing attitude, it yields the direction and range of each user in the satellite body frame at the next update instant.

2) *Predict the mismatch.* Keeping the current phases, the controller evaluates the SNR (or, with multiple feeds, the SINR) that each of the K users would experience under the predicted geometry and compares it with a reference profile, typically the one set by the last full optimization. The normalized deviations form a K -entry mismatch vector $\mathbf{b}$.

3) *Linearize.* The $K \times LN$ Jacobian $\mathbf{J}$, whose entries give the change of each user's normalized SNR or SINR per radian of each meta-atom phase, is computed analytically by

propagating the feed field forward through the layers and the K user steering vectors backward—the same operations as a single gradient evaluation. With multiple feeds, the same recursions are run per feed, and the interference terms enter the SINR derivatives.

4) *Correct*. A single regularized Gauss–Newton step updates all LN phases jointly,

$$\Delta\theta = -\mathbf{J}^T(\mathbf{J}\mathbf{J}^T + \eta\mathbf{I}_K)^{-1}\mathbf{b}, \quad (1)$$

after which each increment is clipped to the hardware limit $\Delta_{\max}$ and quantized to the available phase resolution. The regularization weight η penalizes large phase excursions and keeps the step well behaved when $\mathbf{J}$ is poorly conditioned. Only a $K \times K$ matrix is inverted, so one update costs $\mathcal{O}((K+1)LN^2 + K^2LN + K^3)$ operations, whereas full reoptimization costs roughly this amount per iteration. Because the step is local, its accuracy degrades when the change required per update is large—because of long update intervals or model errors—which is precisely what our case study probes.

C. Sensing-Assisted Tracking

Communication pilots are not the only source of tracking information, and the available sources differ markedly in cost. Terminals can report their own GNSS positions, as already foreseen for non-terrestrial 5G access, which adds uplink signaling but no satellite hardware. The satellite can passively estimate angles of arrival from uplink pilots or random-access preambles, which requires a receive aperture—the SIM operated in reception or a small auxiliary array—and onboard array processing. Active, radar-like sensing additionally consumes transmit power, time, and bandwidth. In all cases, the natural architecture is *predict–sense–correct*: orbital mechanics determines the nominal beam evolution and sparse measurements correct user motion, attitude errors, or model mismatch. The main challenge is allocating sensing resources without degrading the communication service.

D. Hybrid Physics/Learning Tracking

Learning is most valuable when it complements rather than replaces known orbital physics. We advocate learned components that estimate a *low-dimensional residual*—for example, a three-angle attitude bias, per-layer calibration offsets, or an ephemeris correction—and pass the corrected state to the physics-based SIM controller. This preserves physical interpretability, needs far less training data than a network that outputs all LN phases, and transfers across SIM sizes. Learning-based beam management has already been explored for LEO satellite networks [11]; for SIMs, hardware and model correction may be a more compelling role than black-box beam selection.

VI. BEAM TRACKING ACROSS MULTIPLE TIME SCALES

Not every component of a SIM-enabled satellite needs to operate at the same rate, and two time scales are easily conflated. The geometry itself changes continuously, on a

scale set by how long a user stays within the main lobe—tens of seconds for the apertures considered below. The orbit and attitude *prediction*, by contrast, is refreshed whenever new ephemeris or attitude data arrive, which may happen only every few seconds or minutes, because the prediction remains accurate in between. At the intermediate SIM-control time scale, the metasurface is reconfigured when the predicted mismatch becomes significant. At the fastest baseband time scale, synchronization, equalization, coding, and link adaptation operate on much shorter intervals.

This hierarchy does *not* mean that baseband processing can absorb residual pointing errors. With a single feed, the baseband scales the entire radiated pattern by one complex factor: it can compensate the amplitude and phase seen by a user, adapt the modulation and coding, and equalize, but it cannot redistribute power toward a user that a stale pattern has left in a null. With M feeds, digital precoding across the feeds adds some spatial correction, but only within the M -dimensional space spanned by the current SIM configuration, which a stale configuration may no longer align with the users. The error sources therefore split into three classes. *Spatial pointing errors* require a SIM update, possibly assisted by digital precoding across feeds. *Synchronization and channel-distortion errors* are corrected in the baseband. *Geometry and calibration errors* require both estimation, from feedback or sensing, and a subsequent SIM update. Excessively slow SIM updates therefore translate directly into lost rate, and the design problem is to balance tracking accuracy, reconfiguration rate, computational burden, and switching energy.

An attractive extension is event-triggered tracking. Rather than update the SIM at a fixed period, the controller can predict a beam-coherence interval or footprint error and reconfigure only when a service metric crosses a threshold. Such a threshold should reflect not only angular motion but also multilayer electromagnetic sensitivity and hardware switching cost.

VII. CASE STUDY: WHEN IS PREDICTION ENOUGH?

To make these tradeoffs concrete, we simulate a 4-GHz LEO downlink in which a satellite at 1000-km altitude and 53° inclination serves $K = 9$ ground users, placed uniformly at random within a 600-km-diameter region centered on the initial sub-satellite point, over a 50-MHz bandwidth. We use the single-feed configuration of [13], in which one SIM-shaped beam pattern must serve all users and the figure of merit is the *common rate*, i.e., the rate of the worst-served user, $\min_k \log_2(1 + \text{SNR}_k)$. This max–min coverage objective is the most demanding one for tracking, because every user must remain covered at all times, and it represents broadcast, multicast, and control channels; with multiple feeds, the same procedure applies with SINRs in place of SNRs. Each run covers 120 s of the pass, during which the satellite travels about 880 km. The true trajectory is obtained by numerically propagating the orbit with the dominant Earth-oblateness (J_2) perturbation [14] and accounting for Earth rotation, whereas the tracker propagates a circular-orbit model open loop from two ephemeris points. The satellite is nadir pointing and the user positions are known; attitude errors are studied separately.

TABLE I

COMPARISON OF SIM BEAM-TRACKING PARADIGMS FOR A SIM WITH L LAYERS OF N META-ATOMS SERVING K USERS OR CONTROL POINTS. THE COMPUTATION TIMES ARE THOSE MEASURED IN THE SINGLE-FEED CASE STUDY FOR $(L, N) = (4, 225)$ AND $K = 9$; THE LAST TWO ROWS ARE RESEARCH DIRECTIONS.

Paradigm	Information used	Feedback per update	Onboard computation per update	Errors corrected	Main limitation
CSI-driven reactive (full reoptimization)	Pilots and estimated CSI of all K users	K channel estimates	Iterative; each iteration $\mathcal{O}((K+1)LN^2)$ (46 ms)	Any effect captured by the CSI	Pilot overhead; CSI ages during estimation and optimization
Geometry-aided predictive, open loop (GAPT)	Ephemeris, attitude, user or footprint positions	None	One Jacobian and one $K \times K$ solve, $\mathcal{O}((K+1)LN^2 + K^2LN + K^3)$ (0.8 ms)	Deterministic orbital motion and Earth rotation	Attitude and model biases; long reconfiguration intervals
Predictive with scalar feedback (predict-correct)	As above, plus reported SNR/SINR or link quality	K scalars (72 bits for $K=9$ at 8 bits each)	As above	Static biases (attitude, calibration) and slow drifts	Feedback delay; requires a return link
Sensing-assisted	Uplink pilots/angles of arrival, GNSS reports, or radar echoes	Position reports or sensing resources	Estimation plus a local update	User mobility, attitude, unmodeled motion	Receive aperture, sensing accuracy, resource sharing
Hybrid physics/learning	Physics predictor and a trained low-dimensional residual model	Occasional labels; offline training	Residual inference plus a local update	Systematic hardware and model errors	Generalization, training data, certification

The SIM has half-wavelength element spacing and a total depth of 5λ , its feed and inter-layer couplings follow the Rayleigh–Sommerfeld diffraction model [4], and we consider $(L, N) = (2, 100)$ and $(4, 225)$.

Link budget: The feed radiates 20 dBW (100 W) of RF power with 6-dBi gain, the users have 0-dBi antennas and a 290-K noise temperature, so that the noise power is -127 dBW, and the free-space loss is about 165 dB. Without the SIM, the SNR would be about -11.7 dB. The SIM-shaped pattern adds about 12.1 and 14.3 dB of gain toward the worst-served user for the two SIM sizes, which yields worst-user SNRs of about 0.4 and 2.6 dB and per-second optimized common rates of 1.07 and 1.56 bit/s/Hz.

Benchmarks and metric: All schemes start from the same configuration, obtained by maximizing the minimum rate for the initial predicted geometry, and all use 6-bit phases. Full reoptimization is projected gradient ascent with backtracking on a smooth (soft-min) approximation of the minimum rate, warm-started from the current configuration and run to convergence (relative improvement below 10^{-5} , at most 500 iterations). *Reactive* reoptimization uses the true geometry at the update instant, whereas *predictive* reoptimization uses the predicted geometry at the middle of the coming hold interval. GAPT uses $\eta = 10^{-3}$ and $\Delta_{\max} = \pi/8$; imposing the same $\Delta_{\max}$ on full reoptimization changed its rate by less than 0.5%. Each configuration is held until the next update, the common rate is evaluated every second, and we report its time average normalized by that of full reoptimization every second with perfect geometry. Markers show means over 24 realizations and bars the 10th to 90th percentiles. A fixed SIM, not plotted, was observed to retain only about 7–11% of the reference.

How often must the SIM be updated? Fig. 4 leads to three observations. First, one-shot tracking needs short intervals: updated every second, GAPT retains 91% and 85% of the




Fig. 4. Normalized time-averaged common rate versus the update interval T_s . For the hybrid scheme, GAPT runs every second and T_s is the interval between predictive full reoptimizations. Solid lines: $(L, N) = (4, 225)$; dashed lines: $(L, N) = (2, 100)$.

reference for the two SIM sizes, but its retention falls to 80% and 64% at 5 s and to 39% and 26% at 20 s. The reason is the size of the required correction: in a representative run, the predicted mismatch at 20 s is about six times larger than at 1 s, the one-shot step leaves its linear regime, and almost half of the phase increments hit $\Delta_{\max}$, against about 5% at 1 s. The larger SIM, whose lobes are narrower, is the more sensitive to the update interval. Second, prediction also helps full reoptimization: a held configuration must serve the whole interval, so aiming it at the predicted mid-interval geometry retains 91% and 87% at 10 s, against 85% and 77% when it is aimed at the current geometry. Third, the combination is best:




Fig. 5. Normalized time-averaged common rate versus a constant attitude bias unknown to the tracker, with updates every second. Solid lines: $(L, N) = (4, 225)$; dashed lines: $(L, N) = (2, 100)$.

GAPT every second between predictive full reoptimizations every 20 s retains 95% and 90% of the reference while running only one full optimization every 20 s; with re-optimization every 10 s, this rises to 97% and 93%. This is a concrete instance of the switching rule discussed above: cheap local corrections at the fast scale and occasional re-anchoring that resets the reference profile, whose drift away from the max-min optimum is the main reason why GAPT alone stays below the reference even at $T_s = 1$ s.

How accurate must prediction be? The orbit is not the bottleneck: over 120 s, the open-loop circular prediction deviates from the J_2 -perturbed trajectory by less than 90 m, i.e., below 0.005° as seen from the satellite. The bottleneck is attitude knowledge instead. In Fig. 5, the satellite body frame is rotated by a constant bias that the tracker ignores; the bias also affects the initial configuration. Open-loop GAPT and even full reoptimization on the predicted geometry degrade beyond about 0.5° , retaining 75%/64% and 82%/73%, respectively, at 2° for the two SIM sizes, and the larger SIM, with its narrower lobes, is the more sensitive. Replacing the predicted SNRs in the mismatch vector $\mathbf{b}$ by the K SNRs measured by the users—72 bits per update—limits the loss of GAPT at 2° to 2 and 6 percentage points for the two SIM sizes, against 15 and 21 points in open loop, so that beyond about 1.5° it outperforms full reoptimization on the biased model. Once the model is biased, a few bits of feedback are worth more than a more sophisticated optimizer.

Computation: On a single core of a 2.1-GHz Intel Xeon processor, one GAPT update takes 0.2 and 0.8 ms for the two SIM sizes, whereas a warm-started full reoptimization takes 7 and 46 ms (34 and 46 iterations on average), i.e., about 35 and 54 times longer. Both figures include the Jacobian or gradient evaluations and the linear solves but exclude geometry prediction, which takes well below 0.1 ms. They are algorithmic runtimes, not reconfiguration latencies; the hardware steps that complete an update are discussed next.

VIII. PRACTICAL CHALLENGES AND ROADMAP TOWARD AUTONOMOUS TRACKING

Moving SIM beam tracking from algorithms to satellite hardware requires joint consideration of electromagnetic, control, and orbital effects.

Reconfiguration latency, control traffic, and energy: The computation times reported above are not the reconfiguration latency. A practical update also includes transferring the phase commands to the metasurface controller, programming and switching the meta-atoms, electromagnetic settling, and possibly calibration and feedback acquisition, and these steps may dominate. The control volume alone is not negligible: a $(4, 225)$ SIM with 6-bit control requires $900 \times 6 = 5.4$ kbit per full configuration, i.e., 5.4 kbit/s at one update per second, compared with 72 bits of SNR feedback from nine users. Incremental updates that touch only the elements whose quantized phase changes, together with the longer intervals made possible by prediction, reduce both control traffic and switching energy.

Finite-resolution and nonideal meta-atoms: Practical elements exhibit quantized states, phase-dependent amplitude, frequency selectivity, manufacturing tolerances, and temperature drift. Tracking algorithms must therefore use calibrated hardware models rather than ideal unit-modulus phase responses.

Inter-layer calibration: Propagation between SIM layers depends on spacing, alignment, frequency, and the local electromagnetic environment. Errors introduced in early layers propagate through the complete stack, motivating periodic or online self-calibration.

Orbit and attitude uncertainty: Accurate ephemeris alone does not guarantee accurate pointing. A small attitude error rotates all predicted directions in the satellite body frame at once, whereas user-position errors affect individual directions. As the case study shows, attitude bias is the dominant geometric error for fixed terminals, and it is precisely the type of low-dimensional, slowly varying error that sparse feedback or a learned residual estimator can remove.

Wideband operation and beam squint: A SIM configuration optimized at one frequency does not necessarily realize the same transformation across a wide band. This issue becomes increasingly important in Ku-, Ka-, and higher-frequency links, where beam tracking and wideband SIM synthesis may need to be solved jointly.

Fault tolerance and onboard control: Radiation, thermal cycling, and failed or drifting elements are inevitable space-payload concerns. Large SIMs provide spatial redundancy, but the controller must identify degraded elements and adapt accordingly. Likewise, centralizing every SIM update on a high-complexity digital processor risks reintroducing the very computational burden the wave-domain architecture seeks to remove.

These considerations suggest an evolutionary path for SIM-enabled LEO payloads. A first generation can rely primarily on ephemeris and known user locations for open-loop geometry-driven tracking, with periodic full reoptimization. A second generation can combine orbital prediction with sparse SNR reports, pilots, and sensing for closed-loop correction of model,

hardware, and attitude errors. A longer-term autonomous payload could integrate communication, sensing, positioning, self-calibration, and mobility management within the same wave-domain front end.

Important questions remain open: how should tracking accuracy be traded against switching energy; how should tracking and interference management be coupled when a multi-feed SIM multiplexes many users; what is the appropriate definition of SIM beam coherence time; when should users be tracked individually rather than as clusters or service footprints; how should sensing and communication share the same SIM; and how can robust beam quality be certified when the electromagnetic model itself is uncertain? Answering these questions is essential for moving SIM beam tracking from algorithmic demonstrations toward deployable satellite hardware.

IX. CONCLUSION

SIMs can reduce RF-chain count and onboard digital beam-forming complexity by moving part of the signal processing into the electromagnetic domain, but LEO mobility makes their configuration intrinsically dynamic. The central problem is therefore not only to form a beam, but to maintain the desired multilayer electromagnetic configuration as satellite and user geometry evolve. This article has argued that predictable orbital motion should be exploited as a first-class control resource and combined with sparse feedback, sensing, and hardware-aware SIM reconfiguration across multiple time scales. The case study adds three design rules. Local one-shot corrections should run at intervals of a few seconds at most and be re-anchored by occasional full optimizations aimed at the predicted geometry. Orbit prediction is accurate enough over a pass, so attitude knowledge, not ephemeris, limits open-loop tracking. A few scalar SNR reports per update remove this limitation. In our implementation, one-shot updates required 35–54 times less computation than full reoptimization; the end-to-end benefit also depends on the reconfiguration hardware, which remains to be characterized. The longer-term opportunity is an autonomous SIM payload that communicates, senses, calibrates itself, and manages mobility directly in the wave domain.

REFERENCES

- [1] O. Kodheli, E. Lagunas, N. Maturo, *et al.*, “Satellite communications in the new space era: A survey and future challenges,” *IEEE Commun. Surveys Tuts.*, vol. 23, no. 1, pp. 70–109, 2021, doi: 10.1109/COMST.2020.3028247.
- [2] Z. Xiao, J. Yang, T. Mao, C. Xu, R. Zhang, Z. Han, and X.-G. Xia, “LEO satellite access network (LEO-SAN) toward 6G: Challenges and approaches,” *IEEE Wireless Commun.*, vol. 31, no. 2, pp. 89–96, Apr. 2024, doi: 10.1109/MWC.011.2200310.
- [3] S. Bhandari, T. X. Vu, and S. Chatzinotas, “User-centric flexible resource management framework for LEO satellites with fully regenerative payload,” *IEEE J. Sel. Areas Commun.*, vol. 42, no. 5, pp. 1246–1261, May 2024, doi: 10.1109/JSAC.2024.3365883.
- [4] C. K. Sheemar, W. U. Khan, S. Solanki, G. C. Alexandropoulos, and S. Chatzinotas, “A survey on stacked intelligent metasurfaces: Fundamentals, recent advances, and challenges,” *arXiv preprint arXiv:2603.05633*, 2026.
- [5] J. An, M. Di Renzo, M. Debbah, H. V. Poor, and C. Yuen, “Stacked intelligent metasurfaces for multiuser downlink beamforming in the wave domain,” *IEEE Trans. Wireless Commun.*, vol. 24, no. 7, pp. 5525–5538, 2025, doi: 10.1109/TWC.2025.3547779.
- [6] A. Papazafeiropoulos, P. Kourtessis, S. Chatzinotas, D. I. Kaklamani, and I. Krikidis, “Performance of double-stacked intelligent metasurface-assisted multiuser massive MIMO communications in the wave domain,” *IEEE Trans. Wireless Commun.*, vol. 24, no. 5, pp. 4205–4218, 2025, doi: 10.1109/TWC.2025.3533913.
- [7] S. ul Hassan, J. An, F. Niaz, and Z. Jin, “Robust multiuser downlink beamforming for stacked intelligent metasurfaces with imperfect CSI,” *IEEE Wireless Commun. Lett.*, vol. 15, pp. 4847–4851, 2026, doi: 10.1109/LWC.2026.3726088.
- [8] S. Lin, J. An, L. Gan, M. Debbah, and C. Yuen, “Stacked intelligent metasurface enabled LEO satellite communications relying on statistical CSI,” *IEEE Wireless Commun. Lett.*, vol. 13, no. 5, pp. 1295–1299, May 2024, doi: 10.1109/LWC.2024.3368238.
- [9] J. Zhang, J. Li, L. Yang, J. Wang, L. Xu, F. Shu, and Q. Wu, “Stacked intelligent metasurface for satellite communication networks: Architectures, key technologies, and challenges,” *IEEE Wireless Commun.*, vol. 33, no. 4, pp. 31–38, 2026, doi: 10.1109/MWC.2026.3670078.
- [10] H. Xv, Y. Sun, and M. Peng, “Predictive beamforming for joint radar and communication transmission in LEO satellite aeronautical systems,” *IEEE Trans. Veh. Technol.*, vol. 74, no. 8, pp. 12206–12221, Aug. 2025, doi: 10.1109/TVT.2025.3550680.
- [11] Z. Geng, C. She, R. Wang, Y. Li, and B. Vucetic, “Zero-shot learning for beam management in LEO satellite communications,” *IEEE Trans. Wireless Commun.*, vol. 23, no. 9, pp. 12469–12483, Sep. 2024, doi: 10.1109/TWC.2024.3392644.
- [12] Y. Lee, H. Lee, and D. Hong, “Dynamic beam tracking for LEO satellite communication systems,” *IEEE Trans. Veh. Technol.*, vol. 75, no. 5, pp. 8741–8746, May 2026, doi: 10.1109/TVT.2025.3632293.
- [13] C. K. Sheemar, G. Iacovelli, S. Solanki, W. U. Khan, and S. Chatzinotas, “One-shot beam tracking for stacked intelligent metasurfaces-assisted LEO broadcasting,” *arXiv preprint arXiv:2610.05520*, 2026.
- [14] D. A. Vallado, *Fundamentals of Astrodynamics and Applications*, 4th ed. Hawthorne, CA, USA: Microcosm Press, 2013.