Adaptive Digital Twin Synchronization

A Mechanism for Where, When, and What to Measure

Yang, Zhiheng; Vogt, Francisco; Belloum, Adam; Rothenberg, Christian; Grosso, Paola; Papagianni, Chrysa

DOI
10.1145/3769700.3771695

Publication date
2025

Document Version
Final published version

Published in
CoNEXT-SW '25

License
Unspecified

Adaptive Digital Twin Synchronization: A Mechanism for Where, When, and What to Measure

Zhiheng Yang
University of Amsterdam
Amsterdam, The Netherlands
z.yang@uva.nl

Francisco Vogt
Universidade Estadual de Campinas
Campinas, Brazil
franciscogermanovogt@gmail.com

Adam Belloum
University of Amsterdam
Amsterdam, The Netherlands
a.s.z.belloum@uva.nl

Christian Rothenberg
Universidade Estadual de Campinas
Campinas, Brazil
chesteve@unicamp.br

Paola Grosso
University of Amsterdam
Amsterdam, The Netherlands
p.grosso@uva.nl

Chrysa Papagianni
University of Amsterdam
Amsterdam, The Netherlands
c.papagianni@uva.nl

Abstract

Network Digital Twins (NDTs) require timely and cost-effective synchronization with physical networks. Existing telemetry and monitoring tools provide the plumbing to collect rich measurements and show potential for synchronizing the twin with its physical component. However, they lack a principled, fine-grained policy to identify and transmit only the necessary information, in order to save communication costs. In this paper, we present the idea of a *Neural Measurement Field (NMF)*, a learning-based adapter that unifies where the measurements are taken (space/location), when they are updated (time), and what information is synchronized (content) into a single continuous intensity function, trained online to maximize decision utility under the budget.

Keywords

Digital Twin, Synchronization, Network Optimization, AI-driven Network, Active Sampling

ACM Reference Format:

Zhiheng Yang, Francisco Vogt, Adam Belloum, Christian Rothenberg, Paola Grosso, and Chrysa Papagianni. 2025. Adaptive Digital Twin Synchronization: A Mechanism for Where, When, and What to Measure. In *Proceedings of the CoNEXT'25 Student workshop (CoNEXT-SW '25), December 1–4, 2025, Hong Kong, Hong Kong*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3769700.3771695

1 Introduction

The promise of NDTs lies in their ability to accurately model the behavior of physical state, so that they can be used for their control, optimization, and troubleshooting [3]. The precision of NDTs, on the other hand, depends heavily on the synchronization strategy between the twin and the physical system. A rich set of telemetry approaches has emerged, including in-band telemetry (INT), lightweight probes, and programmable headers that provide the raw mechanisms for exporting measurements. However, while these mechanisms enable the collection of measurements, they do not specify how to prioritize or select them when resources are constrained. Several policy approaches have been explored. Periodic synchronization is simple but inefficient: high sampling frequencies cause overhead, while low frequencies lead to information loss [2]. Event-triggered approaches react to anomalies, but may miss subtle yet decision-critical dynamics. More selective strategies adjust frequency or focus on specific locations, but are often heuristic, or architecture-specific, limiting their scalability [1].

We argue that synchronization for NDTs should be viewed as a *multi-dimensional decision problem* that jointly accounts for *where* to take measurements, *when* to report them, and *what* information to include. Existing solutions either ignore some of these axes or implicitly decouple them, resulting in wasteful allocations that fail to adapt to dynamic conditions. Moreover, their heuristic and architecture-specific nature limits generalization across platforms and workloads.

Our Proposal: The problem we aim to address is how to adaptively synchronize NDTs under resource constraints with intelligence. The key insight is to transform the synchronization proxy (uncertainty or decision impact) into a learning problem. We introduce the notion of a *Neural Measurement Field (NMF)*, which treats measurement allocation as a continuous decision space. Rather than separately tuning space, time, and content, NMF defines a unified intensity function $\lambda(e, t, f, p)$ that specifies the probability of triggering a measurement at entity e, time t, for feature f with precision p. This intensity function is parameterized by a neural adapter that is conditioned on the twin's predictive uncertainty, network configuration embeddings, and budget context (e.g., telemetry cost), and is trained online to maximize decision fidelity per unit cost.

2 Design Overview

The design overview is shown in Fig. 1 (❶-❺) and consists of three components together with a learning strategy:

Context encoder (CE). The CE encodes *physical state* into a compact content vector z_c. The vector can include predictive summaries (variances or proxies) as indicators of how well the twin's predictions align with the real system, short-term drift and staleness to capture temporal dynamics, topology embeddings as compact representations of the network structure, and budget/objective context to encode global constraints, etc. (❶-❷). It does not choose where/when/what; instead, z_c conditions all downstream decisions.

CoNEXT-SW '25, December 1–4, 2025, Hong Kong, Hong Kong

ACM ISBN 979-8-4007-2246-2/2025/12
https://doi.org/10.1145/3769700.3771695

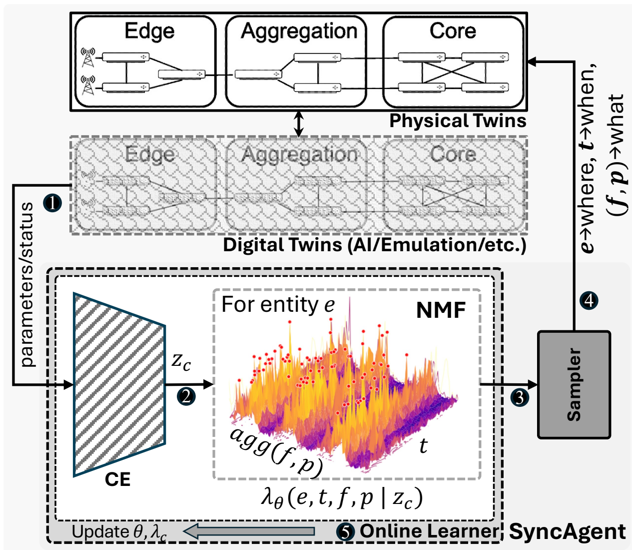

Figure 1: Overall workflow: twins exchange measurements, the encoder builds context z_c (❶-❷), the neural field outputs intensity over where/when/what (❸), the sampler selects actions (❹), and the learner updates both modules (❺).

Neural measurement field (NMF). A lightweight neural adapter parameterizes the conditional intensity function $\lambda_\theta(e, t, f, p \mid z_c)$ (❸). Here, (e, t, f, p) acts as a coordinate query: varying e reflects *where* to measure, varying t reflects *when*, and varying (f, p) reflects *what*. Intuitively, λ forms a continuous "measurement heatmap" over entities, time, and content. High-intensity regions correspond to areas of greater uncertainty or decision impact, while low-intensity regions require less synchronization. The unified formulation means that where, when, and what are not manually decoupled but arise naturally as different slices of the same field. The resulting actions can be directly instantiated by orchestrating existing telemetry capabilities, rather than inventing new data-plane functions.

Sampler. Once the intensity field $\lambda(e, t, f, p)$ is produced, it must be translated into measurement actions. A deterministic *argmax* policy is undesirable: it requires searching a high-dimensional space, often fixates on the same location, and offers no exploration to detect new dynamics. Instead, we propose a *sampling-based planner*. The field λ is treated as a probability surface: regions with higher intensity are more likely to be selected, yet other areas retain a non-zero chance of being probed (❹). This can be implemented with standard tools such as thinning for temporal point processes (to sample times), categorical sampling for entities and features, and lightweight discretization for precision levels. A simple budget guard ensures that the expected number of actions remains within resource limits. This design balances exploitation, investing more in critical regions (where misses cause high penalties), with exploration, preserving adaptivity and avoiding missed dynamics.

Learning strategy. Executed measurements are reported to the digital twins, which correct their state and evaluate both decision quality and incurred telemetry cost (❺). These outcomes form a feedback signal that updates the encoder and the measurement field jointly: parameters are reinforced when actions reduce uncertainty or improve control, and suppressed when they waste budget. We employ lightweight reinforcement-style adaptation with a budget guard. Over time, this closed loop reshapes the intensity field so that scarce measurement resources are well-allocated, while the modular encoder–field design ensures the same learning principle can transfer to other use cases or domains.

3 Discussion

Instantiation and Generalizability. Our design integrates adaptively with existing telemetry approaches: lightweight probes form a low-overhead baseline, while on-demand INT or richer headers are activated only in high-intensity regions; feature sets and bit-widths are pruned by (f, p) to stay within budget. It also generalizes: AI-based twins can estimate predictive uncertainty, while emulation-based twins derive proxies from staleness or drift. Broadly, as long as such multi-dimensional signals are available, NMF can be trained online and also applied to twins in other domains.

Stability Considerations. Since NMF evolves online, overly aggressive updates risk oscillation or overshoot. We plan to address this with a dual-variable budget guard, lightweight regularization, and smoothing, while retaining conservative periodic probes; deeper stability analysis and cost calibration remain future work.

Planned Validation. We will evaluate NMF by instantiating it with existing twin systems and comparing it against periodic and event-triggered baselines. The evaluation will measure (i) communication cost relative to digital twin fidelity, (ii) decision quality under a fixed budget B, and (iii) adaptivity to shifts and anomalies. We expect NMF to lower cost without degrading fidelity, sustain or improve decision quality under budget, and adapt more quickly to changes, while remaining stable through the budget guard.

4 Concluding Remarks

This paper outlines a preliminary vision for unifying the space, time, and content dimensions of physical twin data into a single neural intensity field. By shifting the focus from full data collection to *decision-guided measurement*, the approach strategically trades off fidelity and overhead. While the description remains conceptual, the proposed framework provides theoretical support, is compatible with existing telemetry systems without changes to data-plane primitives, and is extensible to broader digital twin domains.

Acknowledgments

This work is supported by Future Network Services (Dutch 6G flagship project, National Growth Fund) and the Sao Paulo Research Foundation (FAPESP), grant 2021/00199-8 (CPE SMARTNESS).

References

[1] Deepti Kalasapura, Jinyang Li, Shengzhong Liu, Yizhuo Chen, Ruijie Wang, Tarek Abdelzaher, Matthew Caesar, Joydeep Bhattacharyya, Jae Kim, Guijun Wang, et al. 2023. Twinsync: A digital twin synchronization protocol for bandwidth-limited IoT applications. In *Proceedings of IEEE ICCCN '23*. IEEE, 1–1.

[2] Francisco Germano Vogt, Victor HS Lopes, Fabricio Rodriguez, Marcelo Caggiani Luizelli, Chrysa Papagianni, and Christian Rothenberg. 2025. The GhostTwin: Towards live digital twins via SmartNIC network emulation. In *Proceedings of ACM SIGCOMM '25 Posters and Demos*. ACM, 187–189.

[3] Zhiheng Yang, Chrysa Papagianni, Adam Belloum, and Paola Grosso. 2024. An architectural framework for 6G network digital twins system. In *Proceedings of ACM MobiCom '24*. ACM, 2437–2441.