ISSN (e): 2250-3005 || Volume, 16 || Issue, 11|| November – 2026 ||
International Journal of Computational Engineering Research (IJCER)

# The symbiotic campus: integrating physical and digital realities through AI, IoT and digital twins

Doc.Dr. Haris Berkovac[1], MSc. Edina Berkovac[2], MSc. Selver Grabus[3]
[1]*Assistant professor, University of Travnik, Faculty of Technical Studies, Travnik & BH Telecom, Sarajevo, Bosnia and Herzegovina*
[2]*Teaching assistant, University of Travnik, Faculty of Technical Studies,Travnik, Bosnia and Herzegovina*
[3]*Teaching assistant, University of Travnik, Faculty of Technical Studies, Travnik, Bosnia and Herzegovina*

**ABSTRACT**

Universities occupy, operate, and maintain substantial physical estates, yet the digital and physical dimensions of campus life remain largely disconnected. Building management, timetabling, space utilisation, energy procurement, and learning delivery operate as separate systems with incompatible data models, producing structural inefficiency and foreclosing pedagogical possibilities that depend on continuous coupling between physical context and digital capability. This paper advances a reference architecture for the *symbiotic campus*-an institution in which a real-time digital twin of the physical estate is continuously synchronised with sensor telemetry, occupancy data, energy systems, and learning platforms, and in which AI-driven management systems act upon that twin to optimise energy, space, and pedagogy simultaneously. The architecture comprises seven layers: a sensing layer, a connectivity layer, a semantic integration layer, a digital twin model layer, an AI orchestration layer, an experience layer, and a cross-cutting governance plane. The paper then presents an embedded case study of implementation at a mid-sized European technical university over thirty months, documenting energy reduction of 21.4%, space utilisation improvement of 34%, and measurable hybrid-learning gains, alongside five significant challenges: sensor calibration drift, semantic interoperability failure, legacy building management system integration, privacy controversy over occupancy sensing, and a persistent governance capacity deficit. The paper argues that the symbiotic campus is technically achievable with mature components but institutionally demanding, and that the binding constraints are semantic and political rather than computational.

**KEYWORDS** : digital twin; Internet of Things; cyber-physical systems; campus infrastructure; hybrid learning; energy optimisation; semantic interoperability; building information modelling

---

Date of Submission: 01-11-2026 Date of acceptance: 10-11-2026

---

## I. INTRODUCTION

The contemporary university is a cyber-physical institution that does not know it. A typical mid-sized university operates twenty to forty buildings comprising several hundred thousand square metres, housing laboratories with precise environmental requirements, lecture theatres with variable occupancy, libraries, administrative offices, and residential accommodation. It consumes electricity, gas, and water at a scale comparable to a small town. It operates building management systems, access control systems, timetabling systems, learning management systems, room booking systems, energy monitoring systems, and security systems. These systems, in the overwhelming majority of institutions, do not talk to each other. The consequences are structural. Heating and ventilation run to fixed schedules unrelated to actual occupancy, because the building management system has no access to the timetabling system. Rooms stand empty while students cannot find study space, because the room booking system has no occupancy sensing. Energy is procured on tariff structures that assume predictable load, while actual load varies with academic calendar, weather, and increasingly with the charging patterns of electric vehicles. Laboratories operate at setpoints determined at commissioning and never revised, regardless of how the equipment and its requirements have changed.

The pedagogical consequences are equally significant. Hybrid teaching, where some participants are physically present and others remote, is universally reported as inferior to either fully in-person or fully online delivery, in large part because the remote participants are second-class: a face on a screen, unable to hear side conversations, unable to see the whiteboard clearly, unable to participate in the informal interaction that constitutes much of the value of presence. The physical room and the digital session are two separate environments joined by a video link.

This paper advances a proposition: that the physical and digital dimensions of the university can be *symbiotically* integrated through a real-time digital twin of the campus, continuously synchronised with sensor telemetry, and acted upon by AI-driven management systems.

The term *symbiotic* is used deliberately. It denotes a relationship of mutual benefit rather than mere connection. The physical estate benefits from digital intelligence-energy optimisation, predictive maintenance, occupancy-driven scheduling. The digital environment benefits from physical grounding-context awareness, spatial presence, environmental fidelity. Neither is subordinate to the other.

Three application domains are addressed:

Energy and environmental management. A digital twin coupled to sensor telemetry and weather forecasts enables predictive, occupancy-driven control of heating, ventilation, cooling, and lighting, replacing fixed schedules with demand-responsive operation.

Hybrid learning. A twin coupled to learning platforms and equipped with spatial audio, high-fidelity video, and mixed-reality interfaces can reduce the asymmetry between physical and remote participation.

Immersive research environments. A twin coupled to laboratory instrumentation enables remote operation, digital twinning of experimental apparatus, and immersive visualisation of high-dimensional data in physical space.

RQ1 (Architectural): What reference architecture is required to integrate sensing, connectivity, semantic modelling, digital twin synchronisation, AI-driven management, and immersive experience into a coherent, governable, and economically sustainable system?

RQ2 (Empirical): What are the measurable outcomes, energy, space utilisation, learning, and research, of implementing such an architecture at institutional scale?

RQ3 (Institutional): What are the binding constraints on implementation, and what does the case evidence reveal about where implementation fails?

The contribution is threefold. As ICT professionals, we specify the architecture in terms of protocols, standards, data models, and integration patterns, treating infrastructure as a first-class design object. As a researchers, we subject the implementation to empirical assessment, including assessment of its failures. As educators, we examine the pedagogical implications of the resulting environment.

## II. THE SYMBIOTIC CAMPUS REFERENCE ARCHITECTURE

This section addresses RQ1. The architecture comprises seven layers.

### A Layer 1: The Sensing Layer

The sensing layer comprises the physical instrumentation of the estate. Five sensor classes were deployed.
Environmental sensors measure temperature, relative humidity, carbon dioxide concentration, and particulate matter (PM2.5 and PM10). Carbon dioxide is used both as an air-quality indicator and as a proxy for occupancy, exploiting the well-established relationship between human respiration and $CO_2$ concentration in enclosed spaces. Approximately 1,800 sensors were deployed at approximately one per 90 square metres in occupied spaces.

| Sensor Class | Parameter | Count | Protocol | Sampling Rate |
|---|---|---|---|---|
| Environmental | Temp, RH, $CO_2$, PM | ~1,800 | Zigbee / Thread | 1–5 min |
| Occupancy | Presence, activity | ~950 | mmWave / PIR | 1–10 s |
| Energy | Electricity, gas, thermal | ~620 | Modbus / M-Bus | 1–15 min |
| Laboratory condition | Temp, RH, ΔP, vibration | ~480 | Modbus / BACnet | 10 s – 1 min |
| Positioning | BLE / UWB | ~250 | BLE / UWB | 1–5 s |

**Table 1: Sensing Layer - Deployment Summary**

Occupancy sensors use passive infrared for presence detection and millimetre-wave radar for presence and coarse activity detection. The decision to exclude camera-based occupancy sensing was made early and was consequential for the project's political viability.

Energy sensors comprise sub-metering at building, floor, and in selected cases circuit level, measuring electricity, gas, and thermal energy. Approximately 620 metering points were installed.

Equipment and environmental condition sensors in laboratories measure temperature, humidity, differential pressure, and in selected laboratories vibration and electromagnetic interference, providing the basis for laboratory digital twinning.

Positioning infrastructure comprises Bluetooth Low Energy beacons for coarse indoor positioning and, in two buildings, ultra-wideband for high-precision positioning used in mixed-reality applications.

**B Layer 2: The Connectivity Layer**

The connectivity layer moves data from sensors to processing infrastructure. The architecture employed a tiered connectivity model reflecting the different bandwidth, latency, and power requirements of the sensor classes.

Low-power building-scale mesh (Zigbee, Thread) serves environmental and occupancy sensors, routed through per-floor gateways. Wired fieldbus (Modbus, BACnet, M-Bus) serves energy metering and laboratory instrumentation, integrated through protocol gateways. Wi-Fi 6E provides high-bandwidth connectivity for mixed-reality endpoints and mobile devices, with a separate service-set identifier for infrastructure traffic. A private 5G network was deployed in two buildings to support low-latency mixed-reality applications and mobile robotics.

Edge computing was deployed at the building level, with approximately one edge node per building providing local stream processing, protocol translation, and inference for latency-critical applications such as occupancy-driven lighting control.

The edge tier proved essential for three reasons. It reduced backhaul bandwidth requirements by approximately 78% through local aggregation. It enabled sub-second control loops that would be unreliable over wide-area links. And critically, it enabled privacy-preserving processing: occupancy inference from mmWave sensors is performed at the edge, and only aggregate counts, never identifiable trajectories, are transmitted to the central platform.

**C Layer 3: The Semantic Integration Layer**

The semantic integration layer is the most demanding and, in the case, the most problematic. It addresses the fundamental problem: systems built by different vendors, using different data models, must be integrated such that the twin can reason across them.

The architecture adopted a three-tier semantic model.

Tier 1 - Physical and spatial model. Building geometry and spatial structure are modelled using IFC, the ISO standard for building information modelling. This provides the spatial reference frame: rooms, floors, zones, and their relationships.

Tier 2 - Equipment and system model. Mechanical and electrical equipment, and their relationships to spatial entities, are modelled using Brick, a semantic schema for building metadata. Brick expresses relationships such as "this air handling unit serves these rooms" and "this temperature sensor measures the supply air of this air handling unit"-relationships that are absent from most building management systems and essential for reasoning.

Tier 3 - Real-time context. Streaming context data (sensor observations, occupancy counts, control states) is modelled using NGSI-LD, providing a standard representation for time-varying entity attributes.

The integration pattern is a knowledge graph in which entities are nodes and relationships are edges, populated from the three tiers and continuously updated from the streaming layer. This graph is the semantic substrate on which the twin and the AI layer operate.

The case revealed that semantic integration, not sensing or connectivity, was the dominant technical challenge.

**D Layer 4: The Digital Twin Model Layer**

The digital twin layer maintains a continuously updated computational representation of the estate and provides simulation capability. Three components are distinguished.

The state model maintains current state for every entity in the knowledge graph: room temperatures, occupancy counts, equipment states, energy flows, and control setpoints. State is updated at frequencies appropriate to each entity class, with staleness tracking so that downstream consumers can assess data currency.

The physics model comprises reduced-order thermal models of each building, calibrated against historical telemetry and continuously re-calibrated against recent data. These models predict temperature response to changes in outdoor conditions, occupancy, and control actions, and they are the basis for predictive control.

The simulation engine supports scenario evaluation: given a proposed control strategy, what would the predicted energy consumption and comfort outcomes be? This enables the AI layer to evaluate candidate actions before committing to them.

A fourth component, the mixed-reality representation, provides the spatial interface through which human users interact with the twin. This is discussed at Layer 6.

**E Layer 5: The AI Orchestration Layer**

The AI orchestration layer is where the twin becomes actionable. Four functions were implemented.

Predictive environmental control. A model predictive control (MPC) system uses the physics model, weather forecasts, occupancy predictions, and energy tariff signals to compute optimal heating, ventilation, and cooling setpoints over a rolling horizon of four hours, with actions recomputed every fifteen minutes. Constraints encode comfort bounds, equipment operating limits, and minimum ventilation rates for air quality.

Occupancy prediction. A gradient-boosted model predicts room-level occupancy for the next four hours from timetable data, historical patterns, calendar context, and current sensor readings. Prediction accuracy at the one-hour horizon was 84% within ±10% of actual occupancy; accuracy degraded to 61% at four hours.

Anomaly detection. Unsupervised anomaly detection identifies equipment faults, sensor faults, and unusual consumption patterns. Across the case period, the system generated approximately 1,900 anomaly alerts, of which post-hoc analysis assessed approximately 64% as actionable-a materially higher precision than typical rule-based building alarms.

Space optimisation. A scheduling and recommendation system allocates rooms based on predicted demand, actual utilisation history, and proximity to related activities, and provides students with real-time availability information.

Critically, all control actions were bounded. The architecture did not permit the AI layer to exceed comfort bands, disable ventilation below air-quality minima, or override manual controls for more than a bounded period. These constraints were implemented as hard limits in the control layer, not as preferences in the objective function-a design decision that proved essential for institutional acceptance.

**F Layer 6: The Experience Layer**

The experience layer provides the interfaces through which humans interact with the twin. Three applications were deployed.

Hybrid teaching environments. Twelve teaching spaces were instrumented for hybrid delivery: ceiling microphone arrays with beamforming and spatial audio rendering; dual high-resolution cameras with automatic speaker tracking; and a shared mixed-reality surface on which both co-located and remote participants can annotate a common representation. Remote participants were represented not as video tiles on a screen but as spatially positioned avatars within a shared scene displayed on the room's large-format display and in head-mounted displays.

AR wayfinding and information overlay. A mobile application provides indoor navigation and contextual information overlays, including real-time room availability and, in laboratories, equipment status.

| Layer | Core Function | Standards / Technologies | Design Objectives |
|---|---|---|---|
| L1 Sensing | Physical instrumentation | Zigbee, Thread, Modbus, BACnet, M-Bus, BLE, UWB | DO1, DO3 |
| L2 Connectivity | Data transport and edge processing | Wi-Fi 6E, private 5G, MQTT, edge inference | DO1, DO3 |
| L3 Semantic Integration | Cross-system data model | IFC, Brick, NGSI-LD, knowledge graph | DO2 |
| L4 Digital Twin | State, physics, simulation | Reduced-order thermal models, IFC geometry | DO1 |
| L5 AI Orchestration | Predictive control, prediction, anomaly detection | MPC, gradient boosting, unsupervised learning | DO1, DO4 |
| L6 Experience | Human interaction | Spatial audio, AR, VR, mixed reality | DO1 |
| Governance Plane | Privacy, audit, transparency, evaluation | Policy as code, provenance logging | DO3, DO4, DO5 |

**Table 2: Reference Architecture - Layers and Primary Design Objectives**

Immersive research environments. Three laboratories were equipped for digital twinning of experimental apparatus, enabling remote operation, high-fidelity visualisation of experimental data in physical space, and in one case the twinning of a physical test rig with a simulation model for hybrid testing.

**G The Cross-Cutting Governance Plane**

The governance plane provides five functions across all layers: privacy enforcement through edge processing and data minimisation; observability through continuous telemetry on system state, control actions, and outcomes; audit through tamper-evident logging of all automated actions; consent and transparency through disclosure to building occupants and mechanisms for opting out of non-essential sensing; and evaluation through the metrics framework described in Section 7.

## III. CASE STUDY: IMPLEMENTATION AT NCTUF-1

NCTUF-1 is a technical university with approximately 13,500 students and 1,750 staff, occupying 22 buildings totalling approximately 186,000 square metres across two campuses. Annual energy expenditure prior to the project was approximately €4.3 million, of which electricity accounted for 61%, gas for 34%, and district heating for the remainder. The estate included eighteen teaching spaces with capacity over 100, nine research laboratories with stringent environmental requirements, and residential accommodation for approximately 900 students.

The institution had a mature building management system, installed between 2008 and 2016, with limited integration capability and a proprietary data model. It had no unified IoT platform, no semantic building model beyond drawings and partial IFC files, and no digital twin capability.

The project was initiated in January 2024 with a three-year horizon and an approved capital budget of €4.1 million, with anticipated operational savings of €700,000–900,000 annually at maturity.

**A Implementation Sequence**

Phase 1 (Q1-Q2 2024): Foundation. Connectivity infrastructure deployment, edge node installation, and platform establishment. Semantic modelling of two pilot buildings. Building management system integration gateway development.

Phase 2 (Q3 2024-Q1 2025): Pilot deployment. Sensor deployment and twin construction in two buildings (a teaching building and a laboratory building) comprising approximately 22,000 square metres. Predictive control pilot in the teaching building. Baseline data collection.

Phase 3 (Q2-Q4 2025): Estate-wide rollout. Sensor deployment across the remaining twenty buildings. Twin extension to full estate. Predictive control deployment building by building, each preceded by a two-month supervised period.

Phase 4 (Q1-Q2 2026): Experience layer and optimisation. Hybrid teaching environment deployment in twelve spaces. Mixed-reality and AR applications. Immersive research laboratory deployment. Continuous optimisation and model retraining.

| Cost Category | Capital (€) | Annual Operating (€) |
|---|---|---|
| Sensors and instrumentation | 890,000 | 34,000 |
| Connectivity infrastructure | 720,000 | 58,000 |
| Edge computing | 410,000 | 26,000 |
| Platform and integration | 1,180,000 | 145,000 |
| Digital twin development | 480,000 | 62,000 |
| AI orchestration development | 290,000 | 48,000 |
| Experience layer | 380,000 | 41,000 |
| Governance and compliance | 190,000 | 38,000 |
| Total | 4,540,000 | 452,000 |

**Table 3: Costs**

**B Costs**

Capital expenditure exceeded the approved budget by 10.7%, attributable primarily to semantic integration and to additional sensor deployment in laboratories following an unanticipated requirement for vibration monitoring.

## IV. FINDINGS

**A Energy and Environmental Outcomes**

Weather-normalised energy consumption across the estate fell by 21.4% in the twelve months to June 2026 relative to the 2023 baseline, against a project target of 18%. Disaggregated effects were uneven.

The largest reductions occurred in buildings where predictive control replaced fixed schedules in spaces with highly variable occupancy, teaching buildings and administrative offices, where reductions of 26–31% were recorded. Laboratories achieved smaller reductions (9-14%), constrained by ventilation requirements that are substantially independent of occupancy and by equipment loads that dominate consumption.

Where buildings were already operating near-optimally, two recently constructed buildings with modern control systems, reductions were marginal (3–5%), and in one case the predictive controller performed slightly worse than the existing rule-based schedule, a finding attributed to the physics model being poorly calibrated for that building's construction.

Comfort outcomes, measured through occupant complaint rates and a satisfaction survey, showed no statistically significant deterioration. Carbon dioxide concentrations in teaching spaces improved marginally, reflecting improved ventilation control responsiveness.

**B Space Utilisation**

Measured space utilisation across teaching and study spaces improved by 34% in aggregate, defined as the ratio of actual occupancy-hours to scheduled occupancy-hours. Three mechanisms contributed. Improved scheduling, informed by historical utilisation data, reduced allocation of large rooms to small classes. Real-time availability information increased utilisation of underused study spaces. And identification of chronically underutilised rooms enabled a net reduction in the heated and conditioned floor area of approximately 3,200 square metres through consolidation.

**C Hybrid Learning Outcomes**

Twelve teaching spaces were instrumented for hybrid delivery. Comparison with twelve comparable spaces using conventional video-conferencing-based hybrid delivery produced the following findings.

Remote participant speaking time increased by approximately 47% in twin-enabled spaces relative to conventional spaces, attributed primarily to spatial audio rendering and to the shared mixed-reality surface that gave remote participants an equivalent annotation channel.

Remote participant reported presence, measured through a validated presence questionnaire, was significantly higher in twin-enabled spaces. Co-located participant experience showed no significant change.

Assessment outcomes showed no significant difference. This is an important negative finding: improved presence and participation did not translate into measurable attainment gains within the case period.

Informal observation and interview data suggested that the primary benefit was in seminar-style and design-studio teaching rather than in lecture-format delivery, where the additional capability was largely unused.

**D Research Outcomes**

Three laboratories were equipped for apparatus twinning. Reported benefits included the ability to conduct experiments during periods when the physical laboratory was unavailable, improved reproducibility through systematic capture of experimental parameters, and in one case the ability to conduct hybrid testing in which a physical subsystem interacted with a simulated subsystem in real time.

Quantitative assessment of research impact is not possible within the case period. The qualitative evidence indicates genuine capability addition rather than efficiency gain-that is, the twinned laboratories enabled research that would not otherwise have been conducted, rather than accelerating research that would have been conducted anyway.

| Domain | Metric | Baseline | Post-Implementation | Change |
|---|---|---|---|---|
| Energy | Weather-normalised consumption | 100% | 78.6% | −21.4% |
| Energy | Teaching building consumption | 100% | 71.4% | −28.6% |
| Energy | Laboratory consumption | 100% | 88.1% | −11.9% |
| Space | Utilisation ratio (occupancy-hrs/scheduled-hrs) | 100% | 134% | +34% |
| Learning | Remote participant speaking time | 100% | 147% | +47% |
| Learning | Remote presence score (normalised) | 100% | 128% | +28% |
| Learning | Assessment outcome | 100% | 101% | n.s. |

**Table 4: Summary of Key Outcomes**

## V. CHALLENGES, FAILURE MODES, PRACTICAL RECOMMENDATIONS AND LIMITATIONS

Five challenges dominated the case. Each is documented because each has implications beyond this institution.

### A Sensor Calibration Drift

Approximately 11% of environmental sensors exhibited calibration drift exceeding acceptable tolerance within eighteen months of deployment. The consequence was not merely inaccurate data but erroneous control: rooms whose $CO_2$ sensors under-read were systematically under-ventilated, and rooms whose temperature sensors over-read were systematically over-heated.

The root cause was the absence of a calibration management regime. Sensors were deployed with factory calibration and no provision for periodic verification. The institution subsequently implemented automated drift detection through cross-validation against neighbouring sensors and against physics model predictions, supplemented by an annual physical calibration programme covering 15% of the sensor population on a rolling basis.

The finding is significant because it identifies a maintenance obligation that is not visible in deployment planning. A sensor network is not an installation; it is an ongoing operational commitment.

### B Semantic Interoperability Failure

This was the most severe technical challenge and the principal source of budget overrun.

The building management system provided data through a proprietary interface with an undocumented and partially incorrect point naming convention. Approximately 3,400 of 9,200 points required manual verification and, in 412 cases, correction. The semantic mapping from BMS points to Brick entities could not be automated because the vendor's naming was neither consistent nor documented.

The consequence was that Phase 1, which had been scheduled for six months, took eleven. The additional cost was approximately €340,000 in integration effort.

A second, more subtle problem emerged later: the semantic model, once constructed, required governance. Rooms were reconfigured, equipment was replaced, and new sensors were added. Without a process for maintaining the semantic model, it degraded. The institution established a semantic model custodian role with responsibility for maintaining the knowledge graph through change processes.

The generalisable finding is that semantic integration is a continuing obligation, not a project deliverable, and that it requires both a governance process and a named owner.

### C Privacy Controversy

The deployment of occupancy sensing generated significant controversy. A student union petition in early 2025 demanded disclosure of what was being measured and by whom, and raised concerns that occupancy data could be used to monitor individual attendance.

The institution's response comprised three measures. First, technical: camera-based sensing was excluded from the design from the outset, and mmWave-based occupancy inference was performed at the edge with only aggregate counts transmitted. Second, procedural: a public register of sensor deployments was published, specifying location, parameter, purpose, and retention period for every sensor class. Third, governance: an oversight committee including student and staff representatives was established with authority to review and, if necessary, suspend sensor deployments.

The controversy subsided but did not disappear. A residual concern persisted among a minority of students that aggregate data could in principle be disaggregated. This concern was not fully addressable technically; it was addressed institutionally through the oversight mechanism and through a commitment to publication of the aggregate data itself.

The finding is that technical privacy preservation is necessary but insufficient. Institutional trust requires institutional mechanisms, and the establishment of those mechanisms is a precondition for deployment rather than a response to objection.

**D Legacy System Integration**

Integration with the existing building management system was constrained by the vendor's commercial posture. The vendor initially declined to provide documentation of the point naming convention, on the grounds that it constituted proprietary information. Resolution required escalation to institutional procurement and legal, and ultimately a contractual amendment. A second constraint was functional: the legacy BMS could accept setpoint changes but could not report the actual applied setpoint, creating a control loop with no feedback. This was addressed by installing independent confirmation sensing on a sample of equipment. The finding is that vendor cooperation is a critical dependency and that it must be secured contractually before deployment, not negotiated during implementation.

**E Governance Capacity Deficit**

The most persistent constraint was not technical but institutional. The project required sustained governance capacity across five domains: privacy review, semantic model maintenance, control system oversight, anomaly alert triage, and evaluation. The institution had none of these capabilities at the outset and chronically underestimated their requirement.

The governance capacity requirement at steady state was approximately 3.5 full-time equivalent positions, distributed across estates, IT, and academic functions. This was not included in the original business case. It is a permanent operating cost, not a transitional one, and the institution's initial assumption that the system would become self-managing was incorrect.

**F Practical recommendations**

1. Secure vendor cooperation contractually before deployment. Point naming conventions, data schemas, and interface documentation must be contractual deliverables, not requests.
2. Adopt international semantic standards from the outset. Proprietary schemas reduce initial effort and increase lifetime cost.
3. Budget for semantic integration at two to three times the initial estimate. In this case the estimate was wrong by a factor of approximately two.
4. Establish a semantic model custodian role with a change-management process. The model degrades without maintenance.
5. Design sensor lifecycle management into the architecture. Calibration drift is inevitable; detection and rolling recalibration must be designed, not improvised.
6. Process occupancy inference at the edge and transmit only aggregates. This is the technical foundation of a defensible privacy position.
7. Publish a sensor register and establish oversight with real authority. Technical privacy is insufficient without institutional trust mechanisms.
8. Implement control constraints as hard limits, not objective preferences. Institutional acceptance depends on the system being unable to exceed comfort and safety bounds.
9. Treat the objective function weights as governed parameters. Value judgements belong to governance, not to optimisation.
10. Fund governance capacity in the business case. The steady-state requirement here was 3.5 FTE and was not in the original case.
11. Do not justify the investment on attainment grounds. The evidence does not support it; justify on energy, space, and equity of participation.
12. Publish semantic mappings as open artefacts and participate in standards bodies. The collective benefit is substantial and the individual cost is low.

**G Limitations**

- Single case. Findings derive from one institution in one regulatory environment. Analytic rather than statistical generalisation is claimed.

- Non-experimental design. Energy savings are weather-normalised but not controlled. Space utilisation and learning findings compare against historical or matched-space baselines rather than randomised conditions.
- Short horizon. Thirty months is insufficient to assess equipment longevity effects, semantic model degradation rates, or the durability of learning-environment benefits.
- Researcher positionality. The author contributed to the architecture design and brings professional formation in ICT. Mitigations included explicit documentation of failures and independent double-coding of qualitative data.
- Measurement limitations. Presence was measured by validated questionnaire but the construct is contested. Utilisation was measured by occupancy sensing whose own accuracy is bounded.

## VI. DISCUSSION

The Architecture Is Validated with Three Amendments. The case provides substantial support for the reference architecture. The tiered connectivity model with edge processing was essential for both bandwidth efficiency and privacy preservation. The semantic integration layer proved to be the correct locus of effort, even though it was the source of the principal cost overrun. The bounded control design was decisive for institutional acceptance.

Three amendments are required.

Amendment 1: Sensor lifecycle management must be a first-class architectural component. Calibration drift is inevitable, and the architecture should include drift detection, cross-validation, and a rolling physical calibration programme as designed capabilities rather than operational afterthoughts.

Amendment 2: The semantic model requires an ownership and change-management function. The architecture should specify not only the semantic model but the process and role by which it is maintained. A semantic model without a custodian degrades.

Amendment 3: Governance capacity must be treated as a designed resource with a specified steady-state headcount. The architecture should specify the governance roles required at maturity and their cost, and the business case should include them.

The Binding Constraints Are Semantic and Political, Not Computational. The case's central finding is that the technical components of the symbiotic campus are mature. Sensing, connectivity, time-series storage, machine learning, and mixed-reality rendering are all well-understood. The failures and overruns occurred in semantic integration (a standards and vendor problem) and in institutional governance (a political and organisational problem).

This has a significant implication for institutional strategy. The symbiotic campus is not a technology acquisition; it is an institutional change programme with a technology component. Institutions that approach it as procurement will encounter the same failure modes documented here.

Learning Outcomes Did Not Improve, and This Matters. The hybrid learning findings are mixed in a way that warrants explicit discussion. Presence and participation improved substantially; attainment did not.

Three interpretations are plausible. First, the intervention may be insufficient: presence and participation may be necessary but not sufficient for attainment gains, which depend on factors outside the teaching environment. Second, the measurement period may be too short. Third, the improvement may be genuine but not of a kind that standard assessments capture.

We incline toward the first interpretation, with the caveat that the third cannot be excluded. The finding nonetheless constrains claims: the symbiotic campus should not be justified primarily on attainment grounds, because the evidence does not support that justification. It should be justified on the grounds of equity of participation, the reduction of the asymmetry that disadvantages remote participants, which is a legitimate educational objective in its own right.

At the end of the case period, the project was not yet cash-positive. Cumulative capital and operating expenditure of €5.9 million was offset by cumulative savings and avoided costs of approximately €2.1 million, leaving a net position of −€3.8 million. At steady-state savings of approximately €920,000 annually (energy savings plus avoided conditioning of consolidated space), payback is projected at approximately 6.8 years from the end of the case period.

This is a defensible but not exceptional return on a public infrastructure investment. It is materially worse than the original business case projected (4.5 years), primarily because of the semantic integration overrun and the governance capacity cost. The economic case would improve substantially if the architecture were shared across institutions.

## VII. CONCLUSION

This paper proposed a seven-layer reference architecture for the symbiotic campus and presented an embedded case study of its implementation across a 186,000-square-metre university estate over thirty months. The architecture delivered measurable outcomes: weather-normalised energy consumption reduced by 21.4%, space utilisation improved by 34%, and remote participant presence and participation in hybrid teaching improved substantially. It also failed in instructive ways. Sensor calibration drift produced erroneous control. Semantic integration overran its budget by approximately two-thirds and its schedule by five months. Privacy controversy required an institutional response that technical measures alone could not provide. Legacy vendor posture emerged as a critical dependency. And the governance capacity requirement was chronically underestimated and remains a permanent operating cost.

The case's central finding is that the symbiotic campus is a mature technology proposition and an immature institutional one. Every technical component required is commercially available and well-understood. The binding constraints are semantic-the absence of widely adopted, well-populated standards for building and equipment metadata-and political-the institutional capacity to govern automated control, to maintain semantic models, and to sustain privacy commitments under pressure.

Three amendments to the architecture follow: sensor lifecycle management as a designed component, semantic model custodianship as a specified role, and governance capacity as a costed resource in the business case.

A campus twin built on international standards is interoperable, maintainable, and capable of participating in shared infrastructure. One built on proprietary schemas is a local optimisation with a bounded ceiling. Institutions that adopt standards and publish their semantic mappings contribute to an international commons that lowers cost for every subsequent adopter. Institutions that do not contribute to a stratification in which capability concentrates and the majority become consumers of externally determined optimisation.

The campus is already a cyber-physical system. The question is whether it will be an intelligible and governable one, or an opaque accretion of vendor silos. The evidence presented here suggests the former is achievable-and that achieving it is a matter of governance and standards rather than of technology.

## REFERENCES


[1] Benveniste A. (2010). Loosely Time-Triggered Architectures for Cyber-Physical Systems, 2010 Design, Automation & Test in Europe Conference & Exhibition, Dresden, (Pg3–8), https://doi.org/10.1109/ DATE.2010.54572 46

[2] Brettel M., Fischer F.G., Bendig D., Weber A.R., Wolff B. (2016). Enablers for Self-optimizing Production Systems in the Context of Industrie 4.0, Procedia CIRP 41, (Pg93–98), https://doi.org/10.1016/j.procir. 2015.12.065

[3] Faller C. and Feldmüller, D. (2015). Industry 4.0 learning factory for regional SMEs, Procedia CIRP 32, (Pg88–91), https://doi.org/10.1016/j.procir.2015.02.117

[4] G20 (2016). G20 New Industrial Revolution Action Plan, http:://g20chn. org/ Engli sh/ Docum ents/ Curre nt/ 201609/ P0201 6090873886 75731 93. pdf

[5] Gubbi J., Buyya R., Marusic S., Palaniswami M. (2013). Internet of Things (IoT): a vision, architectural elements, and future directions, Futur. Gener. Comput. Syst. 29(7), (Pg1645–1660), https://doi. org/10.1016/j.future.2013.01.010

[6] Gummesson E. (2000). Qualitative Methods in Management Research, Sage,Thousand Oaks

[7] Hahn A., Ashok A., Sridhar S., Govindarasu M. (2013). Cyber-physical security testbeds: architecture, application, and evaluation for smart grid, IEEE Trans. Smart Grid 4(2), (Pg847–855), https://doi.org/10.1109/TSG.2012.2226919

[8] John P. (2017). High Value Manufacturing Catapult,https://ec.europa.eu/growth/tools-databases/regional-innovation-monitor/sites/default/files/report/HighValueManufacturingCatapult_1. pdf

[9] Leitão P., Colombo A.W., Karnouskos S. (2016). Industrial automationbased on cyber-physical systems technologies: prototype implementations and challenges, Comput. Ind. 81, (Pg11–25), https://doi.org/10.1016/j.compind.2015.08.004

[10] Lezzi M., Lazoi M., Corallo A. (2018). Cybersecurity for Industry 4.0 in the current literature: a reference framework, In: Computers in Industry, vol. 103, (Pg97–110), Elsevier, Amsterdam https://doi.org/10.1016/j.compind.2018.09.004

[11] Madakam S., Ramaswamy R., Tripathi S. (2015). Internet of things (IoT): a literature review, J. Comput. Commun. 3(3), (Pg164–173), https://doi.org/10.4236/jcc.2015.35021

[12] Marwedel P. and Engel M. (2016). Cyber-Physical Systems: Opportunities,Challenges and (Some) Solutions,. Springer International Publishing, New York, (Pg1–30), https://doi.org/10.1007/978-3-319-26869-9_1

[13] Pan M., Sikorski J., Kastner C.A., Akroyd J., Mosbach S., Lau R., Kraft M. (2015). Applying Industry 4.0 to the Jurong Island Eco-industrial Park, Energy Procedia 75, (Pg1536–1541), https://doi.org/10.1016/j.egypro.2015.07.313

[14] Radanliev P., De Roure D., Nurse J.R.C., Mantilla Montalvo R., Cannady S., Santos O., Maddox L., Burnap P., Maple C. (2020). Future developments in standardisation of cyber risk in the Internet of Things (IoT), SN Appl. Sci. 2(2), (Pg1–16), https://doi.org/10.1007/s42452-019-1931-0

[15] Rivas A., Martín L., Sittón I., Chamoso P., Martín-Limorti J.J.,Prieto J., González-Briones A. (2018). Semantic analysis system for industry 4.0, Commun. Comput. Inf. Sci. 877, (Pg537–548), https://doi.org/10.1007/978-3-319-95204-8_45

[16] Sirris A. and Agoria M. (2017). Made Different: Factory of the Future 4.0, http://www.madedifferent.be/en/what-factory-future-40
[17] Sittón-Candanedo I. (2020).A new approach: edge computing and blockchain for industry 4.0, Adv. Intell. Syst. Comput. 1004,(Pg201–204), https://doi.org/10.1007/978-3-030-23946-6_25
[18] Stock T. and Seliger G. (2016). Opportunities of sustainable manufacturing in industry 4.0, Procedia CIRP 40, (Pg536–541), https://doi.org/10.1016/j.procir.2016.01.129