

Spatiotemporal Graph Neural Network for Hydrogen Concentration Prediction in Containment during Severe Accidents

Semin Joo^a, Seok Ho Song^a, Miro Seo^b, Jeong Ik Lee^{a*}

^aDept. Nuclear & Quantum Eng., KAIST, Yuseong-gu, Daejeon, Republic of Korea

^bKorea Hydro & Nuclear Power Co., Central Research Institute, Yuseong-gu, Daejeon, Korea

*Corresponding author: jeongiklee@kaist.ac.kr

***Keywords :** severe accident, SAMG, hydrogen, ST-GNN, machine learning

1. Introduction

A severe accident involves extensive core damage and may lead to reactor pressure vessel (RPV) failure, hot-leg creep rupture, or containment failure. Unlike design basis accidents (DBAs), severe accidents can alter the physical connectivity of plant systems as accident progression changes flow paths and system interactions. Effective accident management therefore requires not only technical measures but also timely and reliable decision-making under highly uncertain conditions.

Severe accident management is guided by Severe Accident Management Guidelines (SAMG), which classify plant conditions using predefined setpoints and prescribe corresponding mitigation strategies. By design, however, SAMG is primarily reactive: actions are initiated only after the setpoints are exceeded. Predictive tools capable of forecasting accident progression could support Technical Support Center (TSC) decision-making and reduce delays in implementing mitigation measures.

Recent studies have explored machine learning approaches for severe accident prediction [1-3]. These studies have focused primarily on thermal-hydraulic variables such as temperature, pressure, and water level, whereas severe challenge variables such as hydrogen concentration have received relatively little attention. Moreover, most of them rely on conventional architectures such as LSTMs and Transformers, which flatten plant variables into a single vector and thereby discard the spatial relationships among plant components [4]. Such spatial information is particularly important in severe accidents, where system connectivity itself changes as component failures occur and mitigation actions are taken.

Graph neural networks (GNNs) have been introduced to overcome this limitation by operating on non-Euclidean data, where system components are represented as nodes and their physical or functional relationships as edges. In the nuclear field, GNNs have been applied to multiple accident diagnosis by encoding the system configuration into the graph structure [5], to sensor fault detection in the primary circuit [6], and to multivariate anomaly detection using spatiotemporal graph structures [7]. These studies consistently report that embedding the plant topology into the model improves both accuracy and interpretability compared

with architectures that treat sensor signals as unstructured vectors. Their scope, however, has been largely limited to diagnosis of the current plant state, whereas the prediction of accident progression—where the connectivity itself evolves—has received little attention.

To address these limitations, this study proposes a spatiotemporal graph neural network (ST-GNN) for predicting hydrogen behavior during a station blackout (SBO) accident. By representing compartments as nodes and their flow paths as edges, the model aggregates information along the paths through which hydrogen actually diffuses, while time-varying edge weights capture the evolution of inter-compartmental interaction strengths as the accident progresses.

Two objectives are addressed. First, the contribution of each input variable group is quantified through an ablation study, and the effect of edge flexibility is examined by comparing static and dynamic graph structures. Second, the proposed models are benchmarked against conventional time-series models to quantify the benefit of incorporating spatial dependencies, with the comparison further resolved across accident phases.

2. Description of the SBO dataset

2.1 Dataset generation

The station blackout (SBO) scenario was selected as the initiating event. A SBO accident involves the long-term disablement of containment cooling and mitigation systems due to total AC power loss. The accident spans core damage, reactor vessel lower head breach, and molten core-concrete interaction (MCCI). Hydrogen generation occurs at different stages via distinct physical mechanisms—such as zirconium oxidation and MCCI—resulting in stage-dependent spatial diffusion patterns. Thus, the SBO scenario inherently provides the physical diversity suitable for investigating the temporal variability of inter-compartmental spatial dependencies.

Severe accident management actions are subsequently implemented to mitigate accident progression. Four severe accident mitigation strategies were modeled (see Table 1). Since the SBO scenario assumes a complete loss of electrical power, M6 was implemented using the

Emergency Containment Spray Backup System (ECSBS) rather than the conventional Containment Spray Pump (CSP).

Table 1. SAMG actions modeled in this study

Mitigation action number	Action
M1	RCS depressurization through POSRV opening
M2	SG feedwater external injection
M3	RCS external injection
M6	Containment spray

Each of the four mitigation strategies was assumed to be either activated or not activated at a randomly selected time between 1 hour and 24 hours after SAMG entry. Only one mitigation strategy could be initiated during a given time step, and once initiated, a strategy could not be manually terminated. However, external RCS injection and ECSBS operation were automatically terminated when the reactor cavity water level exceeded a predefined limit, while external SG injection was terminated when the SG water level exceeded its corresponding limit.

Each strategy can be either un-activated or initiated at one of the 24 hourly time steps, yielding 25 options for the each strategy. Since only one strategy can be initiated per time step, the number of available options decreases for each subsequent strategy, giving $25 \times 24 \times 23 \times 22 = 303,600$ possible sequences. Because simulating all possible cases was computationally impractical, 10,000 cases were selected for analysis. The samples were generated using the maximin sampling method, which maximizes the minimum distance between any two samples in the design space. This approach was adopted to construct a dataset with broad and uniform coverage of the possible mitigation scenarios. In addition, one unmitigated (no-action) scenario was added, yielding a total dataset of 10,001 cases.

2.2 Dataset preprocessing

All simulations were performed using MAAP 5.06. The APR1400 reactor design was selected as the reference plant, and plant behavior was simulated for a total duration of 72 h.

The total dataset of 10,001 cases was divided into training, validation, and test sets in an 8:1:1 ratio. Because the input variables comprise parameters with differing physical units and scales—including hydrogen concentration, mitigation actions, PAR/igniter counts, and thermal-hydraulic (TH) variables—standard normalization (z-score) was applied to achieve scale uniformity.

2.3 Dataset analysis

Fig. 1 presents the spatial distribution of the average hydrogen concentration across the 36 compartments at

two distinct time points ($t=3.5\text{hr}$ and $t=4.5\text{hr}$) during the accident progression. Each box reflects the spatial location of a compartment based on the MAAP nodalization, while the color indicates the average hydrogen concentration at the given time point.

At $t=3.5\text{hr}$, Compartment 4 (9.06%) and Compartment 3 (7.82%) exhibit markedly higher concentrations than surrounding compartments. This indicates that even at a given time point, significant inter-compartmental variations exist and local high-concentration regions can form. As the accident progresses to $t=4.5\text{hr}$, not only does the overall hydrogen concentration increase across all compartments, but the spatial location of the peak concentration region itself shifts. This demonstrates that the peak hydrogen concentration does not remain stationary, but rather propagates spatially over time.




Fig. 1. Average hydrogen concentration at (a) $t=3.5\text{hr}$ and (b) $t=4.5\text{hr}$ across 36 compartments

3. Methodology

3.1 Spatiotemporal Graph Neural Network

The objective of this study is to predict the future hydrogen concentration in each compartment within the containment building. The containment is discretized into 36 compartments, with each compartment defined as a node in a graph. At time step t , each node possesses state variables (hydrogen concentration) as its features, while the edges connecting the nodes correspond to the physical flow paths, or junctions, between compartments.

The prediction task is formulated as a sequence-to-sequence regression that takes the node states over the past 10 time steps (150 min) as input and outputs the compartment-wise hydrogen concentrations for the future 8 time steps (2 hr). Specifically, future hydrogen concentrations are predicted based on the past states of the 36 compartments.

The graph topology follows the APR1400 reactor containment building nodalization model in MAAP 5.06 (see Fig. 2). The physical connectivity of the containment is represented as a graph by employing the 36 compartments as nodes and the 57 junctions defined in the MAAP model as edges. The compartment corresponding to the atmospheric boundary was excluded from the prediction targets.




Fig. 2. Graph neural network structure

This configuration directly incorporates the physical mechanism—that hydrogen diffusion occurs along flow paths—into the model architecture. Unlike conventional time-series models that flatten input variables at each time step into a single vector, the graph-based representation preserves spatial adjacency information regarding compartment connectivity without losing it during the embedding phase.

The proposed model comprises two main components (see Fig. 3), capturing spatial and temporal dependencies, respectively.




Fig. 3. Architecture of the ST-GNN based hydrogen concentration predictor

- **Spatial Dependency:** At each time step, a graph convolution operation aggregates information from neighboring compartments. Each node receives information from its adjacent nodes along the edges, thereby accounting for the spatial interactions involved in inter-compartmental hydrogen diffusion.
- **Temporal Dependency:** The node representations obtained at each time step from the graph convolution are sequentially fed into a Gated Recurrent Unit (GRU) to learn the temporal influence of historical states on future conditions.

3.2 Input parameter ablation study

To analyze the contribution of different input variables to hydrogen concentration prediction, an ablation study was conducted by incrementally expanding the input set from the baseline configuration (v1) to v4 (Fig. 4):

- v1: Past hydrogen concentrations and binary indicators (0/1) for mitigation actions
- v2: Adds continuous parameters of mitigation actions (e.g., injection and spray flow rates) to v1
- v3: Adds compartment-wise counts of Passive Autocatalytic Recombiners (PARs) and igniters to v2
- v4: Adds thermal-hydraulic (TH) variables to v3




Fig. 4. Input composition of model v1~v4

Mitigation actions represent the four actions listed in Table 1. In v1, only their execution status is provided, whereas in v2, continuous parameters such as actual flow rates are provided. PARs and igniters perform recombination proportional to local hydrogen concentrations or initiate automatic combustion upon reaching flammability limits, respectively; their installed counts per compartment are assigned as static features. Although igniters remain inoperable under SBO, they were included as input variables to support future extensions of the prediction model to non-SBO accidents.

The TH variables introduced in v4 capture TH conditions associated with hydrogen generation, such as hot-leg temperature, RPV water level, and cavity water level (see Table 2). A separate LSTM model was constructed to predict these MCR-observable TH variables, and its predicted values were used as inputs to the hydrogen concentration prediction model. This framework was established to verify whether incorporating information related to hydrogen generation enhances prediction accuracy.

Table 2. MCR-observed TH variables used in v4

1	Primary system pressure
2	Cold leg 1 temperature
3	Hot leg 1 temperature
4	Cold leg 2 temperature
5	Hot leg 2 temperature
6	Reactor vessel water level
7	Cavity water level
8	IRWST water level
9	Containment upper dome pressure
10	Containment upper dome temperature
11	SG1 pressure
12	SG1 water level
13	SG2 pressure
14	SG2 water level

3.3 Static vs Dynamic Edge

Two model variants were constructed based on how the inter-compartmental interaction strengths—namely, the edge weights—are handled. Both models maintain an identical topology (57 junctions), differing solely in the treatment of the edge weights.

- **Static Model:** Assigns time-invariant, fixed values to the edge weights ($=1.0$), assuming that the interaction strength between compartments remains constant throughout the entire accident progression.
- **Dynamic Model:** Allows the edge weights to vary over time. Based on the node states at each time step, a GRU-based gating mechanism calculates individual edge weights, thereby capturing the phenomenon in which inter-compartmental interaction strengths evolve across different stages of the accident.

If high-concentration regions propagate spatially over time, the inter-compartmental interaction strengths must likewise be time-varying. Therefore, superior performance of the dynamic model over the static model would suggest that the temporal variability of spatial dependency is critical.

In addition, to evaluate the performance gain from spatial modeling, ST-GNN models were benchmarked against conventional time-series models (LSTM, Transformer) using the selected v4 input set. The LSTM processes the input sequence recurrently, using gated memory cells that regulate which information is retained

or discarded at each time step, which allows it to capture long-term temporal dependencies.

The Transformer instead attends to all time steps simultaneously through self-attention, weighting past states according to their relevance to the current prediction. Both models, however, receive the states of all compartments as a single flattened vector at each time step and therefore have no explicit representation of which compartments are physically connected; any spatial relationship must be inferred implicitly from the data. All baselines were trained on the same v4 input and constrained to a comparable parameter count, so that the performance difference against the ST-GNN models can be attributed to the representation of spatial dependencies rather than to model capacity.

All model parameters were controlled between 30k and 35k to isolate model capacity effects.

4. Results and Discussions

4.1 Input parameter ablation study

First, the effect of the input variable configuration on predictive performance was analyzed. Holding the model architecture constant as the dynamic ST-GNN, the input set was incrementally expanded from v1 to v4, with the corresponding results summarized in Fig. 5.




Fig. 5. Prediction performance of dynamic ST-GNN model with varying inputs

As the input set expanded, predictive performance improved monotonically, with the MAE decreasing by approximately 60% from v1 to v4. This demonstrates that each added group of variables provides valid information for hydrogen concentration prediction. While PARs account for hydrogen sink mechanisms, TH variables represent hydrogen source mechanisms; thus, integrating both source and sink information into the input yields optimal predictive performance.

Based on these findings, v4 was established as the final input configuration and applied to all subsequent model architecture comparisons.

4.2. Comparison among different architectures

Under the fixed v4 input configuration, the predictive performances of LSTM, Transformer, static ST-GNN, and dynamic ST-GNN were compared. The results are presented in Fig. 6.




Fig. 6. Prediction performance of various models

Given identical input conditions, the ST-GNN series models (static and dynamic) achieved lower errors than the conventional time-series models (LSTM and Transformer). This indicates that preserving spatial information via a graph structure provides an advantage in hydrogen concentration prediction. Unlike conventional models that flatten the input at each time step and thereby erase spatial information, ST-GNN enhances prediction accuracy by explicitly representing diffusion paths across compartments.

Within the ST-GNN family, the dynamic model outperformed the static model. This supports the premise that inter-compartmental interaction strengths indeed evolve across accident stages, and that accounting for this temporal variability enhances predictive performance.

4.3. Comparison across accident phases

The performance gain achieved by capturing spatial dependencies in the ST-GNN model is expected to be most prominent during phases characterized by active hydrogen diffusion. Active hydrogen diffusion primarily occurs during the early stage of an accident, where hydrogen diffuses toward the upper containment following primary pressure boundary failure or MCCI onset.

Accordingly, the accident progression was partitioned into early (0~9hr), transition (9~32hr), and late (32~72hr) phases. The early-to-transition boundary was defined by RPV failure, which marks the complete breach of the primary pressure boundary, and the transition-to-late boundary by the point at which hydrogen concentrations pass their peak and approach equilibrium.

A comparison was conducted between the LSTM model, which captures only temporal dependencies, and the dynamic ST-GNN model, which incorporates both spatiotemporal dependencies.




Fig. 7. Scenario-averaged hourly MAE of dynamic ST-GNN and LSTM model

As illustrated in Fig. 7, the dynamic ST-GNN model exhibits lower errors from the early through the transition phase. In the late phase, however, the errors of both models become comparable. This indicates that the advantage of the ST-GNN model diminishes in the late phase, where hydrogen concentrations reach equilibrium and spatial dependencies weaken. Conversely, the ST-GNN model demonstrates a clear advantage during the early accident stages, where hydrogen diffusion is active and peak concentration locations dynamically shift.




Fig. 8. Comparison of predicted values and true (MAAP-calculated) values of hydrogen concentration at nodes #3 and #17 for SBO scenario M1: 4hr, M2: 14hr, M3: 5hr, M6: 6hr.

Fig. 8 further illustrates this trend through a representative scenario at nodes 3 and 17. Node 3 is the corium chamber room, and node 17 is the containment dome. During the early phase (0~9hr), hydrogen concentrations rise sharply as diffusion becomes active following primary pressure boundary failure. Node 3 exhibits a sharp rise in hydrogen concentration driven by the breach of the primary pressure boundary, while node 17 reflects a change driven by the rapid diffusion of hydrogen from surrounding compartments. The dynamic

ST-GNN closely tracks these rapid fluctuations, whereas the LSTM exhibits a noticeable lag in capturing the timing and magnitude of the rise. This is consistent with the fixed edge-weight assumption underlying the LSTM's implicit spatial reasoning, which cannot adapt to the rapidly shifting diffusion paths characteristic of this phase.

As the scenario progresses into the late phase, both models exhibit a small deviation from the MAAP-simulated values. Because predictions are generated autoregressively, errors accumulated at each step propagate into subsequent inputs, and this rollout error compounds over the 72-hour horizon regardless of the model architecture. Consequently, the gap between the dynamic ST-GNN and LSTM narrows in the late phase, consistent with the phase-wise comparison in Fig. 7.

5. Conclusions

In this study, an ST-GNN model was proposed to predict compartment-wise hydrogen concentrations within the containment building during SBO accidents, and its validity was verified in terms of input configuration and model architecture.

The ablation study on input configurations demonstrated that incorporating information on both hydrogen sink mechanisms (PARs) and source mechanisms (TH variables) enhances the prediction accuracy. Model architecture comparisons revealed that the ST-GNN series achieved lower prediction errors than conventional time-series models (LSTM and Transformer). This confirms that preserving inter-compartmental diffusion paths via a graph structure provides an advantage in hydrogen concentration prediction. Within the ST-GNN family, the dynamic model outperformed the static model.

Phase-wise analysis further indicated that the advantage of the dynamic ST-GNN was most pronounced during the early phases, where hydrogen diffusion is active. This finding supports the premise that the performance gains of ST-GNN stem primarily from phases governed by spatial dependencies.

Future work will quantitatively evaluate the contribution of individual compartment-to-compartment connections by systematically eliminating edges within the ST-GNN architecture and observing the corresponding changes in predictive performance. Additionally, the applicability of the proposed model will be evaluated across other accident scenarios, such as loss of feedwater (LOFW) and large-break loss of coolant accident (LLOCA).

ACKNOWLEDGMENT

This work was supported by the K-Cloud project of KOREA HYDRO & NUCLEAR POWER CO., LTD (No. 2024-Tech-08).

REFERENCES

- [1] J. Song, J. Kim, J. Jeon, and S. I. Lee, A machine learning informed prediction of severe accident progressions in nuclear power plants, *Annals of Nuclear Energy*, Vol. 196, p. 110237, 2024.
- [2] S. Joo, Y. Lee, S. H. Song, K. Song, M. R. Seo, S. J. Kim, and J. I. Lee, Leveraging explainable AI for reliable prediction of nuclear power plant severe accident progression, *Reliability Engineering and System Safety*, Vol. 264, p. 111307, 2025.
- [3] Y. Tuo and X. Liu, Predicting nuclear power plant operational parameters using clustering and mutual information for feature selection and Transformer neural network optimized by TPE, *Frontiers in Energy Research*, Vol. 12, p. 1462184, 2024.
- [4] S. Liu, X. Zhou, J. Yu, Y. Wang, T. Xu, and H. Wang, Graph attention network-based model for multiple fault detection and identification of sensors in nuclear power plant, *Nuclear Engineering and Design*, Vol. 419, p. 112949, 2024.
- [5] Y. H. Chae, S. G. Kim, H. Kim, J. T. Kim, and P. H. Seong, Graph neural network based multiple accident diagnosis in nuclear power plants: Data optimization to represent the system configuration, *Nuclear Engineering and Technology*, Vol. 54, No. 8, pp. 2859-2870, 2022.
- [6] X. Wang, H. Wang, and M. Peng, Interpretability study of a typical fault diagnosis model for nuclear power plant primary circuit based on a graph neural network, *Reliability Engineering and System Safety*, Vol. 261, p. 111151, 2025.
- [7] L. Zhang, J. Cheng, Y. Zhang, and X. Zhao, Spatial-Temporal Graph Conditionalized Normalizing Flows for Nuclear Power Plant Multivariate Anomaly Detection, *IEEE Transactions on Industrial Informatics*, Vol. 20, No. 11, pp. 12945-12954, 2024.