T1C

MELVIN MARK

APPLICATION NOTE
SEPTEMBER 19, 2025

Platform Goal

Neuromorphic computing combines insights from neuroscience with modern electronic hardware, enabling spiking neural networks that operate with low latency, low power, and high energy efficiency. Despite growing adoption by companies such as Intel and IBM and emerging vendors like SynSense and Innatera, there remains a gap: a practical, publicly accessible neuromorphic platform that delivers clear operational value in real-world settings. Type 1 Compute's neuromorphic platform addresses this by targeting measurable outcomes in building controls specifically, optimizing HVAC systems for comfort, stability, and energy reduction. We were engaged by Melvin Marks to improve their HVAC operations by detecting anomalies earlier, stabilizing setpoint tracking under changing conditions, and reducing overall energy consumption through real-time, on-device control

Issues

Traditional MPC/ML controllers can miss anomalies or react slowly when operating conditions drift (occupancy, weather, equipment aging). Operators need fast, robust control that cuts kWh without sacrificing comfort or requiring cloud connectivity

Architecture

We deployed a spiking SNN controller with an MPC prior on a commercial building dataset (supply air temp, airflow, static pressure, damper/fan speeds, occupancy, outdoor/discharge temps). The controller runs closed-loop on edge hardware and adapts online to changing conditions. We take into consideration components like optimization, like Cooling/Heating Capacity, Outdoor/Discharged Air Temperature, Static Pressure Points, Damper/Fan speeds, and occupancy, which have been logged by Melvin Marks' internal software, giving us real-time data to make more educated decisions

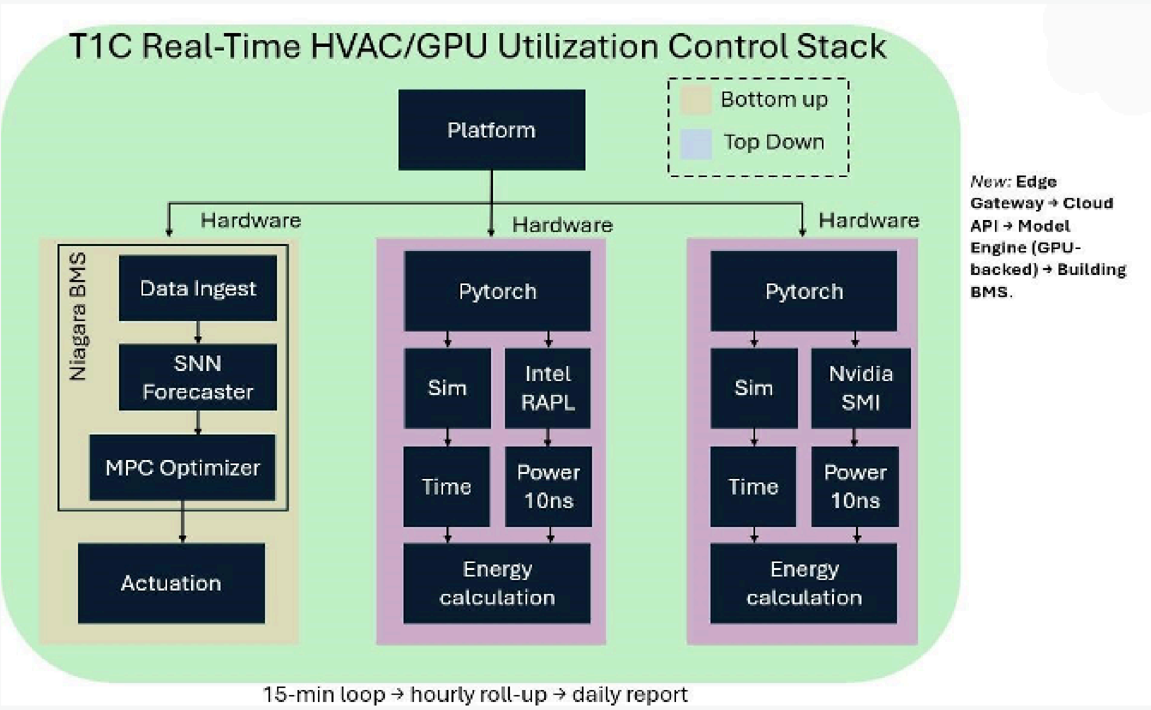

Technical Performance

Over a 300 s closed-loop run, the SNN(+MPC) HVAC controller maintained indoor temperature within the defined tolerance band after an initial ~25 s warm-up, with bounded, responsive modulation of supply-air temperature and airflow; end-to-end control-loop latency stayed well below the **100 ms target** (typically in the low tens of milliseconds, with one brief excursion), **and cumulative energy reduction performance converged to ~21–22%** versus the baseline, meeting the ≥20% goal; throughout steady state, model uncertainty remained in the high-confidence regime (≈0.8–0.9), indicating stable inference under normal operating conditions

Next Steps

Our near-term roadmap scales the platform across building automation and electrification workloads, including air-side and water-side plants, grid-interactive heat pumps, occupancy-aware ventilation, and mission-critical cooling (labs, data centers), with the same edge controller footprint.

hello@type1compute.com

503.810.8092

www.type1compute.com