Event-Driven Spiking Neural Networks for Emergent Robotic Swarm Behavior

M. A. Turqueti[1], C. R. daCunha[2], Rodrigo Pena[3], Om Aggarwal[1]
[1]Lawrence Berkeley National Laboratory, USA; [2]Northern Arizona University, USA; [3]Florida Atlantic University, USA

Abstract

Spiking Neural Networks promise ultra-low-power computation for edge robotics, yet their asynchronous nature makes hard real-time guarantees difficult to achieve. We present a lightweight event-driven scheduler that enforces deterministic execution on resource-constrained hardware by separating high-priority spike acquisition from downstream cascade processing. By modeling the internal processing load as a self-organized critical process constrained to a sub-critical regime, we derive a provable upper bound on spike avalanche size, guaranteeing temporal stability under real-world sensory variation.

The system is implemented on the Sea-of-Transneurons architecture developed at Lawrence Berkeley National Laboratory (LBNL) and validated on RAY (Robot for Aquatic Yield), a manta ray-inspired soft robot developed at Northern Arizona University (NAU), achieving efficient aquatic propulsion through tendon-driven fin deformation. Multiple RAY platforms are then coordinated as an autonomous swarm, where a distributed machine learning algorithm adapts collective behavior to environmental conditions, maintaining emergent heading, speed, and formation without centralized control.

Architecture

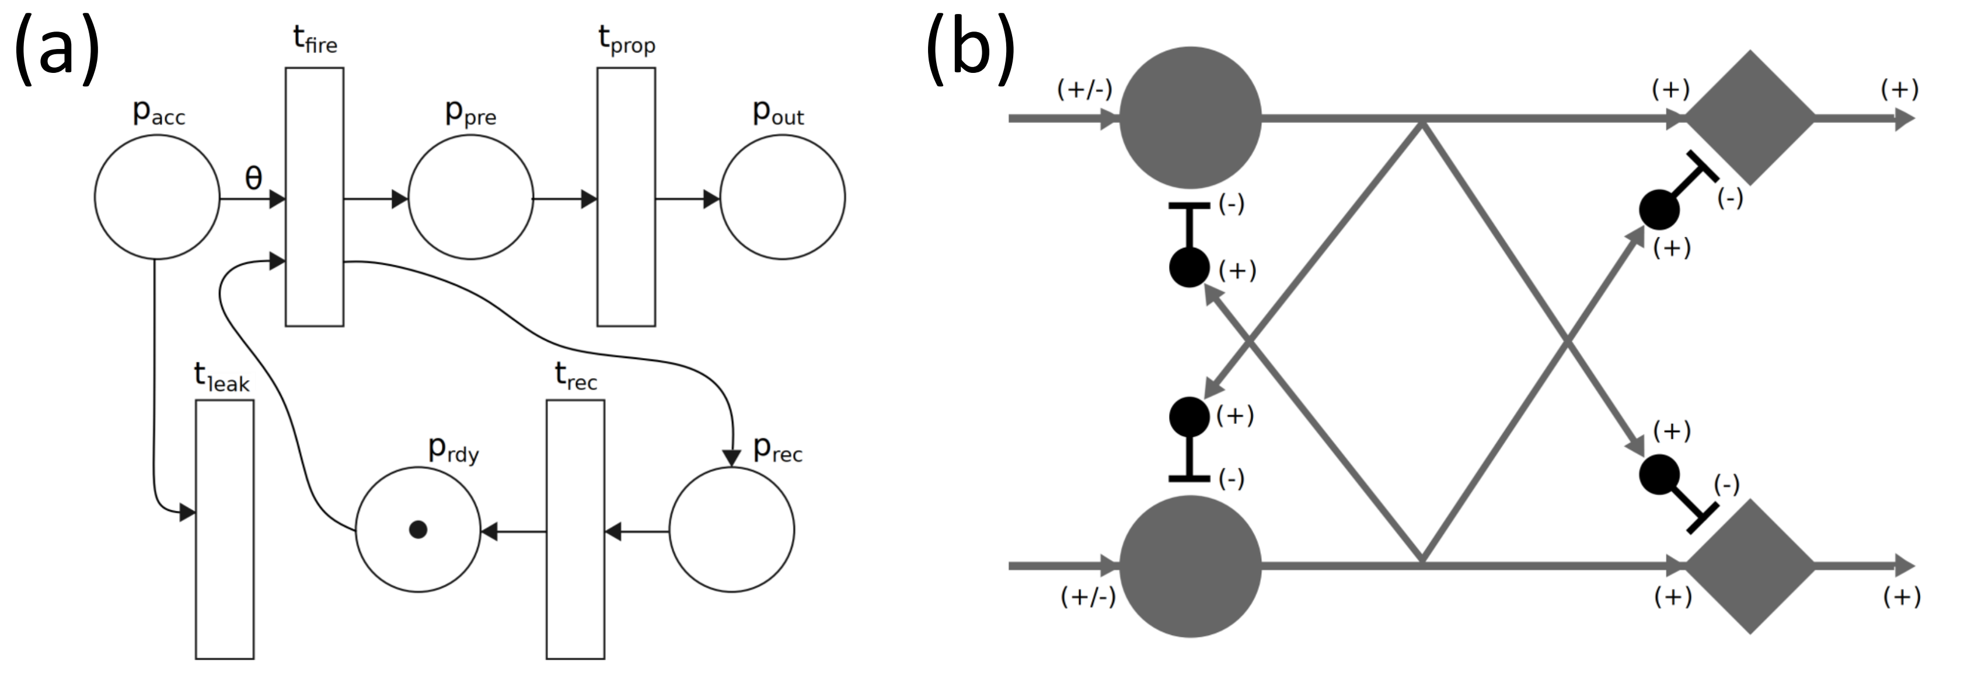

To formally characterize timing behavior, we introduce the Petri neuron (a), a leaky integrate-and-fire neuron modeled as a T-timed Petri net, and derive a closed-form Worst-Case Response Time expression computable at design time from neuron parameters. The formal framework is experimentally validated using a half-center Central Pattern Generator (CPG) implemented as two mutually inhibitory Petri neurons (b). Configured with matched parameters and driven by periodic tonic input, demonstrating that the scheduler correctly implements biologically realistic, input-driven firing rate control.

Swarm Behaviour

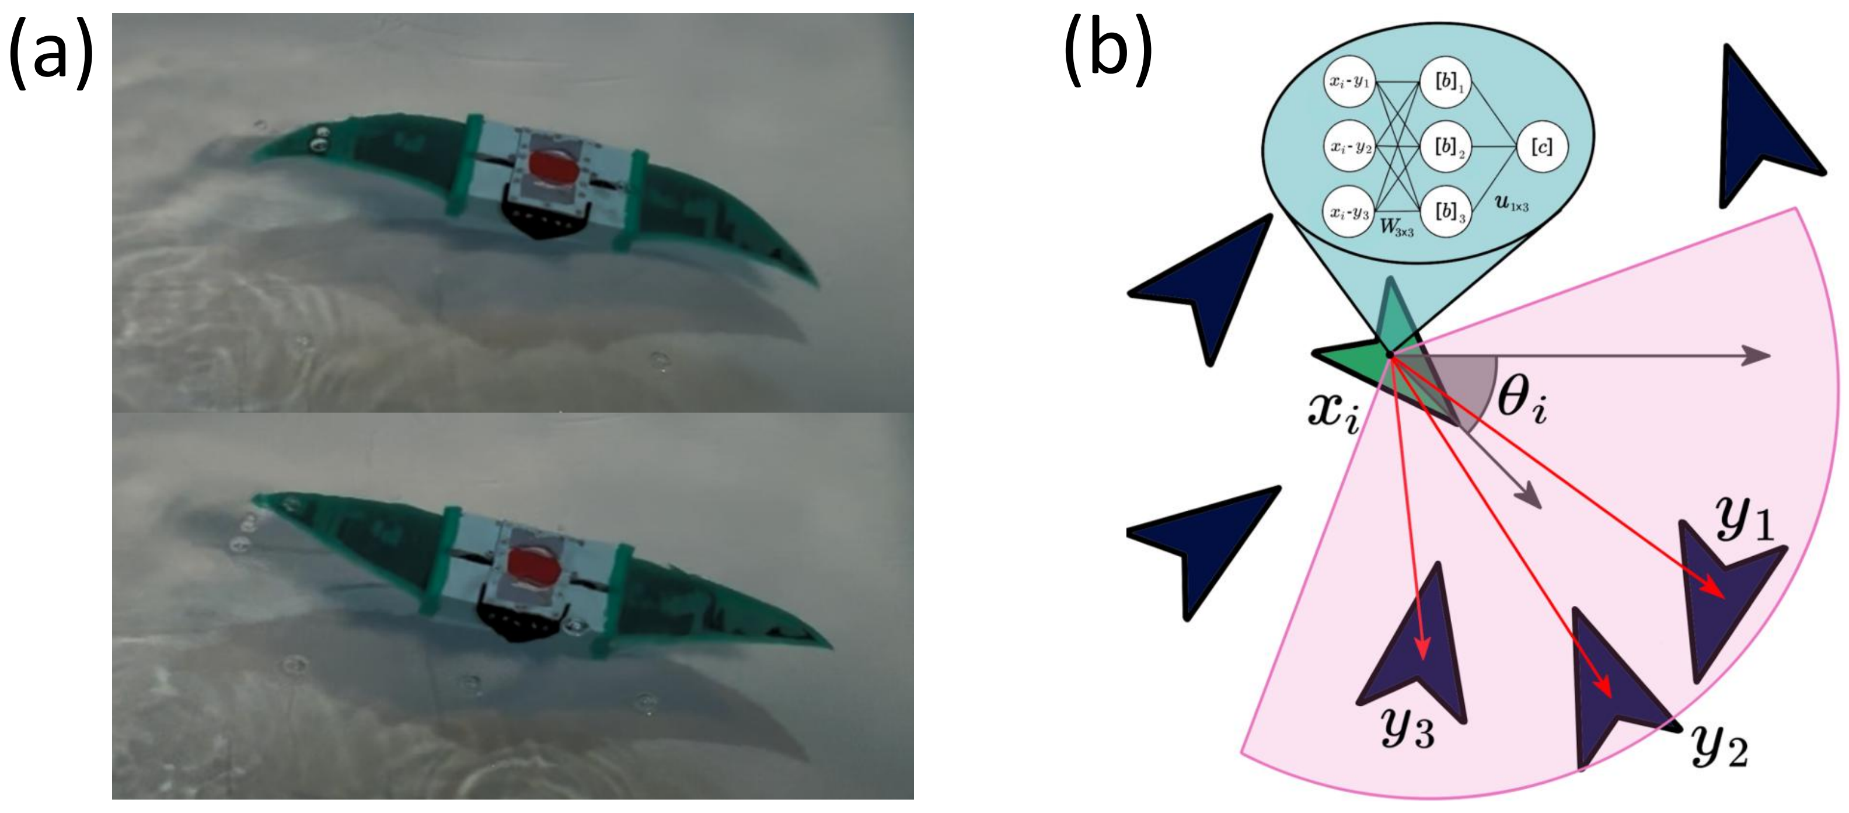

Coordinating multiple RAY platforms as a coherent autonomous swarm requires each agent to continuously share state information with its neighbors. Underwater acoustic communication provides the medium for robust, scalable inter-agent information exchange. Each RAY (a) platform is equipped with an acoustic modem that encodes spike-derived state information into sonic pulses transmitted through the water, enabling a fully decentralized swarm communication fabric without any surface relay or centralized coordinator. The bridge between the onboard SNN (b) and the acoustic communication layer is a lightweight spike encoding protocol. Rather than transmitting raw membrane potentials or continuous sensor readings, each RAY agent encodes its behavioral state as a compact spike timing pattern generated directly by the onboard CPG circuit. The rhythmic output of the CPG serves as a natural carrier signal while deviations in spike timing, and inter-firing interval relative to the nominal rhythm encode directional intent.

Hardware Implementation

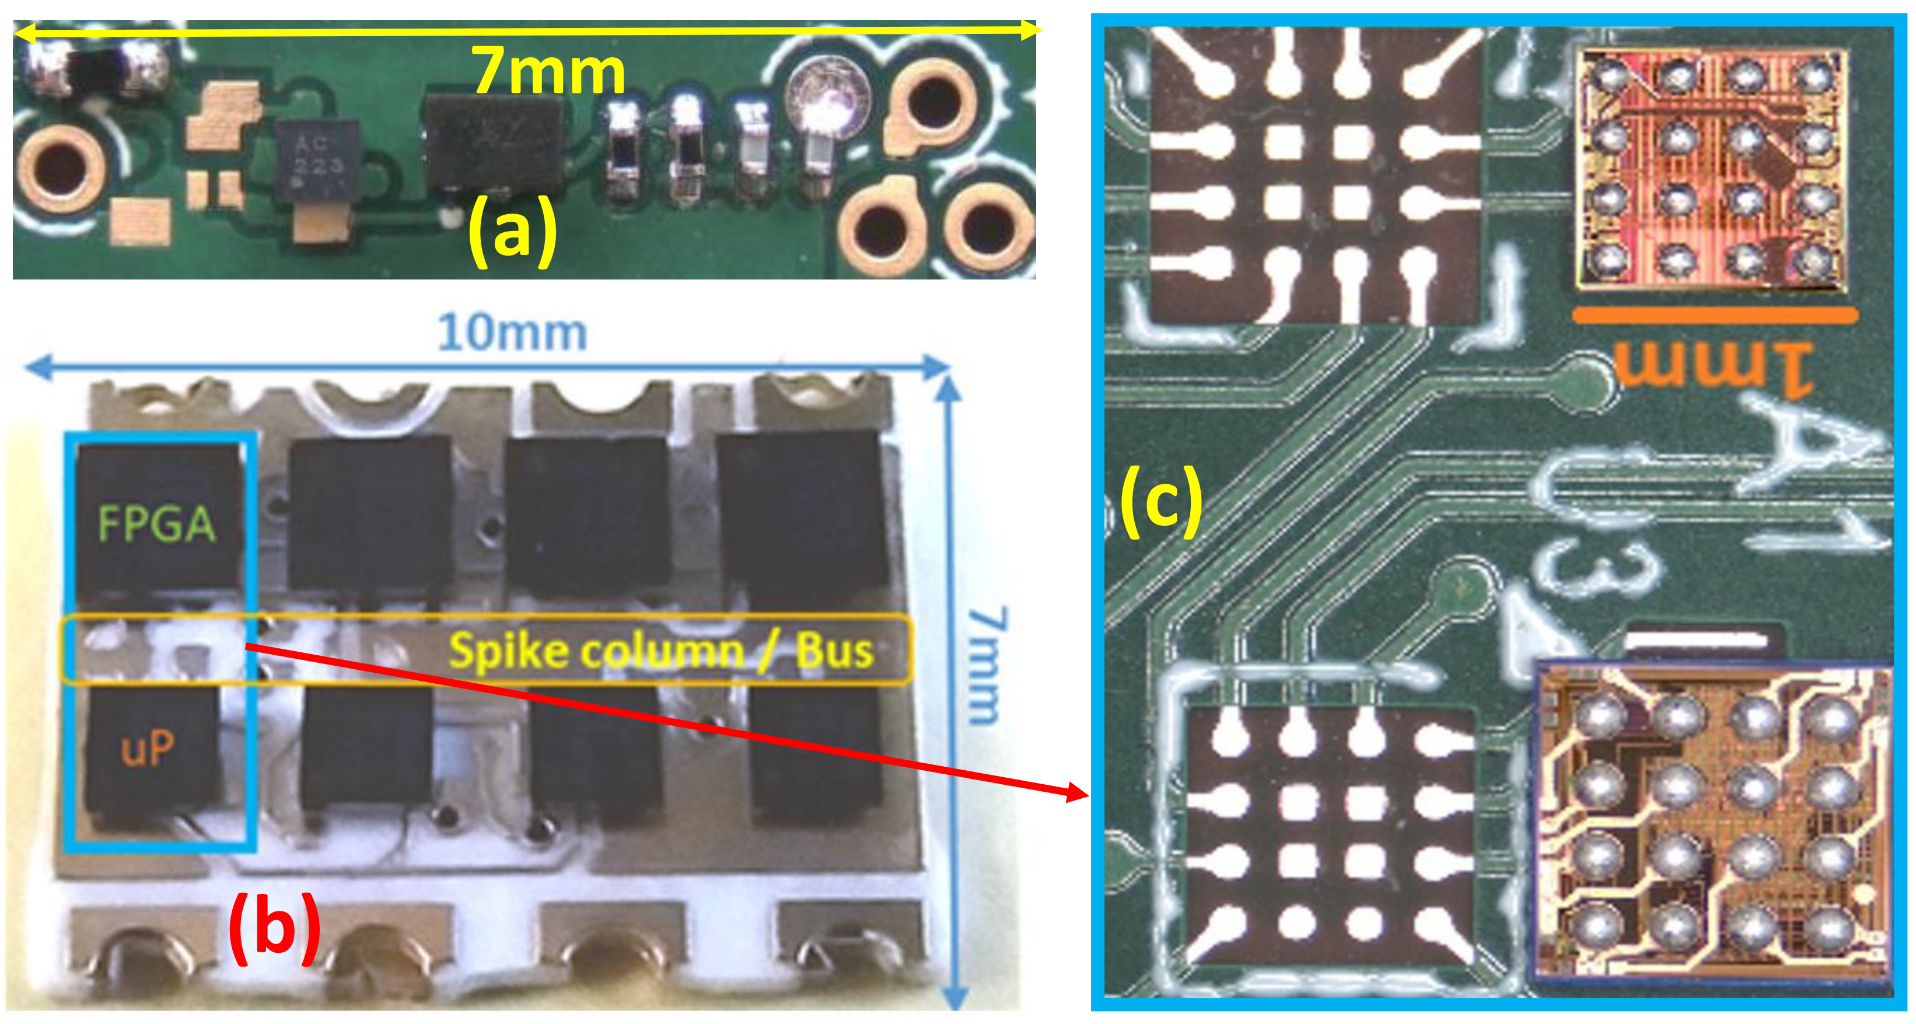

Environmental signals are conditioned by LBNL sensory transneuron (a) and then feed to the Spike Neural Network (SNN). The SNN architecture is ported to the Sea-of-Transneurons (SOT) architecture developed at LBNL. Each SOT blade (b) integrates four FPGAs with four microprocessors, supporting up to 2,000 transneurons per blade. Transneurons are distributed across tightly coupled FPGA–microprocessor nodes, enabling scalable parallel execution of large SNN populations. The event-driven communication model of the SOT architecture is naturally compatible with the Petri neuron formalism, as each transneuron implements state evolution and spike communication in a manner directly analogous to the formal model. This substrate provides the computational foundation for coordinating multiple RAY robotic platforms as an autonomous swarm.

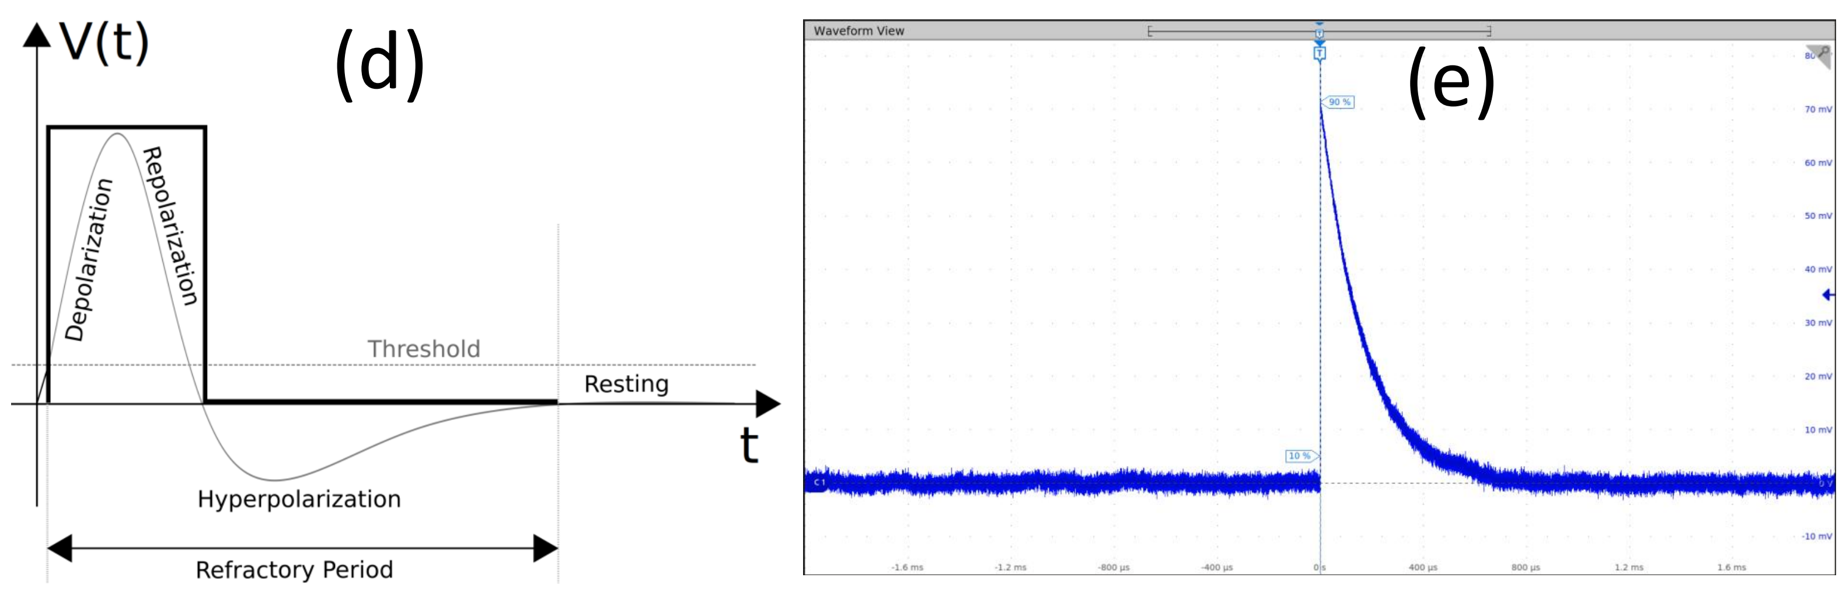

A biological neuron fires an action potential when its membrane potential crosses a defined voltage threshold, after which it undergoes a refractory period before becoming ready to fire again. In our simplified model (d), this behavior is captured as a discrete voltage pulse that resets to zero and remains there until the refractory period elapses. In the RAY , the sensory neuron layer serves as the interface between the physical aquatic environment and the onboard SNN control architecture. The sensor neuron is integrated with a spike encoder (e) that converts each sensor reading into a train of spikes.

25th IEEE Real Time Conference – La Biodola, Elba, Italy

This work was supported by the U.S. Department of Energy (DOE) under Grant DE-AC02-05CH11231