Journal of Electronic Research and Application, 2026, Volume 10, Issue 5
http://ojs.bbwpublisher.com/index.php/JERA
Online ISSN: 2208-3510
Print ISSN: 2208-3502

A Bioinspired Neuromorphic Chip Based on Protein-Protein Interaction Network Topology and Its Application in Adaptive Gait Control for Legged Robots

Qiang Chen

Chongqing Ruanjiang Turing Artificial Intelligence Technology Co., Ltd., Chongqing 400020, China

Abstract: Traditional footstep controller chips of the legged robot have problems such as low energy utilization rate, lack of flexibility, and poor environmental resistance ability. In order to solve this problem, this study proposes a neuromorphic chip design based on PPI network topology, which mimics the small-world property ,scale-free degree distribution, and modular structure observed for the PPI networks using a set of protein network processors (PNPs). In memristor crossbar arrays, synaptic weights can be stored and computed in situ to eliminate the memory-processor bottleneck. A three-level motion hierarchy deals with control on different timescales: a reflex level for millisecondscale posture stabilization, a rhythm level to generate periodic gait patterns, and a strategy level for long-horizon motion planning and terrain adaptation. Fabricated in 28 nm CMOS, the chip contains 1024 PNPUs and draws a total power of 2.5W. Experiments with a quadruped robot show that our chip is able to achieve better than 10 × energy efficiency over traditional model predictive control running at 1.5m/s walkspeed on flat ground, with terrain transitions completed within 0.8s. With a quarter of processing units disabled, we still retain more than 80% of locomotion capability, verifying that our model generalizes to unseen situations.

Keywords: Protein-protein interaction networks; Neuromorphic chips; Embodied intelligence; Adaptive gait control; Memristors; Self-healing

Online publication: Jun 29, 2026

1. Introduction

Embodied AI is the intelligence that arises from constant physical interaction of an agent with its surroundings through a corporeal body, and it becomes one of the most active research frontiers in AI [1]. A fundamental problem in this area is how to generate stable, flexible, and adaptive locomotion, including gait control of biped and quadruped robots on rough terrain. With robotics applications expanding from factory floors to field exploration, disaster rescue, and home assistance, the requirements of motion control systems toward low power, low delay, and performance resilience have increased significantly.

Conventional methods for robot locomotion control can be categorized as follows, namely model-predictive control (MPC) based on the concept of zero-moment point (ZMP), bio-inspired control based on central pattern generator (CPG), and end-to-end policy trained with DRL. ZMP-MPC learns an optimal robot dynamics model online for stable walking, but its heavy dependence on the model's accuracy renders it brittle in unknown terrains or when subject to external perturbations [2]. CPG methods mimic the rhythmic motor circuits of living nervous systems and are used to generate regular periodic gaits, but tuning their parameters is nontrivial, and they do not adapt to changing environmental conditions. DRL-based controllers learn policies by interacting with the environment via agents, and are more adaptable; however, they have poor sample efficiency, lengthy training time, and lack any guarantees for stable control [3–5]. Importantly, all these paradigms operate on standard CPU/GPU hardware with large power budgets, which is not ideal given the strict power requirements of edge computation at legged robots.

Protein-protein interaction (PPI) networks are at the heart of cellular functionality, being behind everything from signal transduction and metabolism to gene regulation. Recent systems biology studies show that PPI networks are small-world, scale-free, with a degree distribution, a high clustering coefficient, and a dynamic rewiring ability [6]. The small-world property allows for fast dissemination of information in the whole network; Scale-free topology makes it robust to random node failure; modularity facilitates functional specialization and functional integration, respectively; and dynamic reconfiguration enables coupling strengths to be adjusted in response to changing environmental demands. Thus, these properties enable biological systems to perform computations which, in the face of complex, variable circumstances.

Why PPI nets are interesting to neuromorphic engineering, because they have a very close similarity with plasticity: conformation change of proteins corresponds to change of the neuron's membrane potential;protein-protein interaction, like synaptic coupling; and formation and removal of protein complexes as synaptogenesis and synaptic pruning, respectively. All these analogies provide strong biological justification to adopt PPI organization principles in designing neuromorphic chips.

The vast majority of previous work in the area of neuromorphic systems focuses on hardware implementations of spiking neural networks (e.g., Intel's Loihi or IBM's TrueNorth), but such organizational rules on a molecular level are still missing for constructing larger-scale neuromorphic systems [1,7,8]. This paper closes that gap both by presenting a bio-inspired neuromorphic chip with a PPI-network-derived topology, as well as an application of it to adaptive gait control on a legged robot. Our contributions include

(1) A kinetic PNPU design capturing protein conformational dynamics, where the memristor crossbar provides extremely efficient in-situ computation;
(2) A layered motion planning system with a reflex layer in milliseconds, a decision layer in seconds.
(3) A self-healing system in-network based on the functionality redundancy of protein networks for healing at-chip; and
(4) Verification via experiments with an actual quadruped robot, which shows that our control has small power consumption, high dynamic performance, and is robust to disturbances.

2. Bio-inspired neuromorphic chip architecture

2.1. Protein network processing unit design

The principal computational element in our design, the protein network processing unit (PNPU), relies on the following principle drawn from the kinetics of protein conformational change: inside a living cell, a protein's function depends on its 3D structure, and that 3D structure can be affected by temperature, pH, ligand binding, etc. The PNPU emulates the same process in circuit elements, producing a physical layer simulation of biological information processing.

Each PNPU contains four sub-blocks: a conformational-state register file, an interaction-weight memory, a folding-unfolding compute engine, and an energy calculation unit. The register file is built from analog storage cells; each cell holds the energy state Ei (in [0,1]) of one protein node, where 0 denotes the inactive unfolded conformation, and 1 denotes the active folded conformation. Parallel read and write access to the register file enables nanosecond-scale state updates.

The interaction-weight memory is implemented with a memristor crossbar (RRAM crossbar) that stores inter-node coupling strengths wij. Memristors provide non-volatile, multi-level storage together with in-situ computation: their continuously tunable conductance maps directly to analog synaptic weights, and within the crossbar, the weight storage and matrix multiplication occur at the same physical location. Input voltages are applied along word lines, while output currents accumulate along bit lines, eliminating the data-movement overhead that plagues conventional von Neumann architectures [8]. For a 1024-PNPU array, a single matrix-vector multiplication finishes in tens of nanoseconds, yielding two orders of magnitude better energy efficiency than traditional digital accelerators.

The folding-unfolding engine constitutes the main part of a PNPU, and it implements the state transitions according to stochastic thermodynamics, adopting the Metropolis criterion such that the transition probability from Ei to E i' is P = min(1,exp(-DeltaE/kBT)), with DeltaE being the energy change, where kB is Boltzmann's constant and Teff is some effective temperature. A mixed signal circuit consisting of a pseudo-random number generator and an analog comparator speeds up this process in the hardware, allowing hundreds of node-state updates to be processed within the same cycle.

The energy computing unit computes total energy for the protein network according to H = -Sum (wij * Ei * Ej) – Sum (hi * Ei), where wij is the interaction intensity between node i and node j, and hi is an external input field of the sensor information. The first term captures pairwise node interaction energy, whereas the second term represents the energy of incoming stimulation. The energy function leads to an optimal solution of this system.

2.2. Network topology and on-chip interconnect

The PNPU array is logically arranged as a 2-D mesh, yet physically it implements a scale-free topology: a small set of hub PNPU nodes handles the bulk of intermediate routing, while the majority of non-hub nodes manage localized data exchange. This layout preserves high concurrency across the chip and simultaneously reduces average transmission distance. Each PNPU maintains four nearest-neighbor links plus one or two long-range connections; the long-range link targets are selected according to a scale-free degree distribution.

The chip works asynchronously, in an event driven way: computation and communication happen only if there is some dramatic change of the state at a node, cutting the dynamic power by more than 50% as compared to globally clocked, synchronous designs; and uses a hybrid routing policy in its on-chip interconnect: nearest-neighbor traffic travels over direct wires with single cycle latency; long-distance traffic is routed through hierarchical routers whose latency is logarithmic in the distance. On-die integration of temperature sensors and an adaptive frequency-voltage scaling block is also included for dynamically tuning the operating frequency and power supply voltage depending on the chip temperature and workload in order to avoid thermal runaway while not compromising performance.

2.3. Self-repair mechanism design

Based upon functional redundancy of proteins, this study proposed a dynamic reconfiguration capability for compute resources as pools within a single chip [8]. Within an organism, protein networks tolerate genetic mutations and environmental stress via redundant functions and alternative paths. This study transplants the same idea in silicon, by introducing a fault-tolerant architecture composed of a primary compute array and a redundant resource pool. PNPU: The first array (N × M) is referred to as the primary, and they perform the main control operations for locomotion, whereas the second one is called redundant, which has 10 to 20 percent more elements than the first one, and it remains idle at runtime.

A fault-detection component continuously checks each PNPU for abnormal computation results, abnormal power consumption, and timeout responses. When a PNPU failure occurs, the "chaperone" reconfig sequence is auto-triggered. The procedure comprises four steps

(1) The faulty PNPU is "denatured" and electrically separated;
(2) A spare PNPU is allocated from the redundant pool;
(3) On-chip learning circuitry transfers the failed unit's weights and states to the replacement;
(4) Online finetuning restores control performance.

This study demonstrates experimentally that the whole reconstruction finishes within less than five seconds and degrades overall locomotion performance by at most 20%.

3. Hierarchical motion planning system

To fill in the gap between low-level reflexes and high-level strategy, we developed a biologically-inspired three-level motion planning architecture to coordinate whole-body behavior at different time-scales together. We base our design on the fact that vertebrates divide their motor control into three levels: reflexes are implemented through the spinal cord, the midbrain controls rhythmic pattern generation, and the cerebral cortex plans voluntary movement. The modular connection between groups of proteins is used for information transfer between neighboring layers and thus facilitates a smooth evolution towards strategy from reflex.

3.1. Protein reflex layer

The reflex layer consists of a small, densely connected graph mapping the sensor's input directly to the actuator's output (i.e., spinal reflex arc). This module allows for postural stabilization with a millisecond resolution and a control cycle between 1 and 5 ms, controlling the effect of external disturbance and keeping the body stable. The layer is implemented using a fully connected topology where each input node connects to all output nodes; the weights are trained offline through supervised learning, and they are kept frozen during run-time to ensure deterministic response time.

The reflex layer design is based on the stretch and flexor reflex circuits that exist within biology. If the foot pressure sensors sense an uneven ground, the reflex layer compensates for the perturbation by adjusting joint torques in less than 3 ms, so that the center of mass of the body remains within its support polygon. Similarly, when the inertial measurement unit detects body tilt, the layer will generate a vestibular-like reflex that extends the limbs and restores the body's balance. Since these reactions are local and do not involve higher-level deliberation, they execute in hard real-time.

3.2. Protein rhythm layer

The rhythm layer is a ring-topology protein network whose oscillatory behavior produces a stable rhythmic signal, generating the periodic pattern of gaits. Its control time period varies between 0.5 and 2s, matching walking cadence. We model each of these phases (for a single joint) as a node on a ring, and let the coupling weights between the nodes encode the coordination relationship between the joints.

The rhythm layer could also switch between different locomotion modes (e.g., walk–run and run–jump) by changing the network parameter values. The gait switching mechanism works in terms of "phosphorylation dephosphorylation": following detection of a terrain change either visually or via foot-pressure sensors, a phosphorylation pulse is introduced at one of the PNPUs, altering their free-energy profile and pushing the system towards a novel attractor configuration. This process is analogous to the modulation of proteins by signal molecules that occur *in vivo* and allow for a seamless but fast transition from one type of walk to another. For us, the transfer of motion between flat ground and slope climbing required approximately 0.8 s, while keeping a maximum angle difference of the body during this period at 5°.

3.3. Protein strategy layer

The policy level uses an RL-trained deep protein network for discovering good locomotion policies, running at the 0.1–1s scale, which deals with the path planning, obstacle avoiding, power saving, etc., higher-level thinking processes. The neural network contains an input layer, multiple hidden layers, and an output layer. The hidden layer is decomposed into functional modules, each of which is responsible for one particular job, like terrain classification, energy estimation, or risk assessment.

Communication between the three layers is realized via "protein docking", i.e., it mimics domain-domain recognition occurring during protein-protein interaction: Output neurons of an upper level connect specifically to input neurons of the lower level, ensuring targeted information flow. If a decision is made at the strategy layer to switch locomotion mode, this occurs via adjustment of the rhythm layer's network parameters. In turn, if the rhythm layer gets stuck in some state it does not know how to handle, then it alerts the strategy layer so as to generate a new plan.

4. Experimental validation and results

4.1. Experimental platform setup

To comprehensively test our chip, we developed an entire experiment control system based on a quadruped robot platform. Our hardware consists of three parts: the robot body, the control chip, and the test-analysis equipment. The body of the robot itself is a 12-DOF quadruped with three rotary joints for each leg (hip, thigh, and shank). Total mass is 12.5 kg, rated load of 5 kg, and joint actuation consisting of a brushless DC motor with a harmonic drive reducer providing 18 Nm peak torque; position feedback resolution: 0.1 degree.

Sensor suite: A 6-axis IMU sampled at 1 kHz; twelve joint encoders (with 4096 pulses-per-revolution); four foot-pressure sensors from 0 to 100 N with 0.1 N resolution; and a depth camera (Intel RealSense D435i) operating at 30 fps. All sensor data is streamed to the controller through fast SPI buses.

The controller chip is fabricated in 28nm CMOS with a die area of 8mmx8mm and consists of 1024

PNPUs, 64KB of on-chip SRAM as a weight buffer, and a number of peripheral interfaces (SPI, I2C, UART, PWM). The supply voltage can be adjusted between 0.8 V and 1.2 V, while the total power consumption equals 2.5 W. As a reference, this study additionally executed a traditional MPC controller on an NVIDIA Jetson Xavier NX (15W) and a reduced CPG algorithm on an STM32H7 microcontroller (0.5W).

4.2. Level-ground gait control experiments

This study conducted walking trials on flat ground at three speeds: slow (0.3 m/s), medium (0.8 m/s), and fast (1.5 m/s). After each 100-meter traverse, we logged stability metrics, energy use, and joint torque profiles. Results show that the robot walks steadily at 1.5 m/s under chip control, with center-of-mass excursion kept within +/- 8 mm, pitch variation within +/- 2.5 degrees, and roll variation within +/- 1.8 degrees.

In terms of power, we achieve a total system cost (chip+ motors) of 45W/kg, i.e., 30.8% less than that of the MPC baseline(65W/kg), and 22.4% less than that of the CPG baseline (58W/kg). Our chip's own compute efficiency is 3.2TOPS/W, about a factor of 15 faster than the Jetson Xavier NX with a comparable NN. The gain is coming from the following three sources: memristor in-situ computation eliminates data-movement overhead, event-driven operation suppresses redundant switching, and, in parallel, with the dynamics of the protein network, enabling control on a timescale as short as less than one millisecond.

4.3. Complex terrain adaptation experiments

This study evaluated terrain adaptation across five representative surface types: flat ground, 15-degree inclines, 10-cm steps, gravel, and grass. The robot performed walking tests on each terrain and autonomous transitions between terrains; every test was repeated ten times, and results were averaged.

On slopes, where it climbed and descended steadily (at 1.0 m/s) without exceeding an angular deviation of less than 5 degrees from the slope normal for the robot's body longitudinal axis, suggesting successful posture control. The robot was able to estimate step height on stepped terrain independently and adjust leg lift and stride accordingly, clearing ten consecutive steps at 0.6 m/s, and achieving the task with 98% success rate. Foot-pressure feedback over gravel/grass with variable ground compliance enabled the robot to balance support forces between legs and prevent slip/sinkage:

Even more impressive were the terrain-transition experiments. When switching from flat ground to a slope, the vision system detected the terrain change in 0.5 s and relayed the information to the strategy layer. The latter planned the gait transition in 0.3 s and handed it off to the rhythm layer, which completed the actual transition in 0.8 s without a perceptible pause. On the other hand, our DRL-based comparison algorithm, which can also switch between terrains, takes tens of seconds of online learning in order to adapt to a new terrain, making these methods inappropriate to be implemented in an online environment.

4.4. System robustness testing

This study conducted two types of fault-injection experiments for measuring system robustness. One is PNPU failure, i.e., this study uses software to kill PNPUs on the primary array, simulating hardware failures. When turning off 20% PPNUs, more than 80% walking ability is still available, while average velocity decreases from 1.5 m/s to 1.2 m/s, and the amplitude of body posture fluctuations increases by approximately 25%; functional reconstruction time with a mean value equal to 4.8 s. Even if we turn off half of all PNPU's (i.e., 50%), our system can still walk slowly with a speed of 0.5m/s which proves that the multi-pathway redundancy built into the protein network topology is worth it.

The second set of tests was simulations of sensor failure. We turned off the IMU in sequence with the vision sensor and then each of the foot pressure sensors, testing graceful degradation. When the IMU was off, the robot combined joint encoder and foot-pressure information for balancing, but slowed down its gait (walking at only 0.8 m/s). With one of the foot-pressure sensors failing, the remaining three legs compensated for it almost without affecting locomotion. The results shown above indicate that hierarchy and natural redundancy do prove to be practical.

4.5. Comparative analysis

This study finally compares the presented bio-inspired neuromorphic chip to classical control strategies on four axes. Control latency: Our implementation reaches a cycle of 0.5 ms for the reflex layer vs. 5 ms (MPC) and 2 ms (CPG). Adaptability: terrain switching time is only 0.8s in ours vs. manually re-tune (MPC), or more than 30 seconds of online training (DRL). Robustness: when 20% of the compute units are disabled, our scheme loses only 20% performance, but MPC can fail catastrophically in case of model mismatch. Energy Efficiency: overall system cost is 45 W/kg, a factor of 30.8× improvement over MPC, while the chip itself consumes just 2.5W, within tight power budgets for edge computing devices.

5. Conclusion

This paper proposed a neuromorphic chip, the architecture of which stems from the topology of the protein-protein interaction networks, and applied it to adaptive gait control of a legged robot. This study has done the whole design loop: theoretical development, chip implementation, and experimentation. Highlights are:

(1) On system-level, a kinetic PNPU modeling protein conformational dynamics where memristor crossbars provide in situ storage/computation of synaptic weights 10–100 × more efficiently than traditional digital accelerators;

(2) At the algorithm level, a three-layered hierarchical motion planner coordinating control on timescales ranging between milliseconds and seconds, with "protein phosphorylation" for quick change of walking gait;

(3) A system-level, online self-healing ability based on the functionality redundancy of the protein network, which can significantly increase the reliability.

Experiments on a quadruped robot confirm our chip's benefits: fabricated in a 28 nm process, the 1024-PNPU die consumes just 2.5W and supports 1.5 m/s level walking,0.8 s terrain transitions, and more than 80% locomotion retained with the removal of 20% of processors. Such numbers are far beyond those that are currently achieved by existing CPU/GPU/DRL approaches. By putting embodied intelligent robots onto a computing platform that is at once energy efficient, adaptive, and resilient, our contribution is taking a step towards pushing Robotics outside of the lab setting to complex, unstructured real-world environments.

In the future, this study could be explored in greater detail. Firstly, the use of protein-network principles to direct 3-D stacked packaging toward higher integration density with lower communication latency; multi-robot cooperative control and swarming behavior extension; and transfer of the architecture to bipedal/humanoid platforms. Another promising direction is online learning algorithms capable of deep analysis of the dynamical properties of a protein network as they unfold in real time, towards making systems more adaptable so that a robot can automatically adapt itself to completely new surroundings.

Disclosure statement

The authors declare no conflict of interest.

References

[1] Davies M, Wild A, Orchard G, et al., 2021, Advancing Neuromorphic Computing with Loihi: A Survey of Results and Outlook. Proceedings of the IEEE, 109(5): 911–934.

[2] Barabasi A, Oltvai Z, 2004, Network Biology: Understanding the Cell's Functional Organization. Nature Reviews Genetics, 5(2): 101–113.

[3] Silver D, Huang A, Maddison C, et al., 2016, Mastering the Game of Go with Deep Neural Networks and Tree Search. Nature, 529(7587): 484–489.

[4] Arulkumaran K, Deisenroth M, Brundage M, et al., 2017, Deep Reinforcement Learning: A Brief Survey. IEEE Signal Processing Magazine, 34(6): 26–38.

[5] Henderson P, Islam R, Bachman P, et al., 2018, Deep Reinforcement Learning That Matters. Proceedings of the AAAI Conference on Artificial Intelligence, 32(1): 3207–3214.

[6] Ielmini D, Wong H, 2018, In-Memory Computing with Resistive Switching Devices. Nature Electronics, 1(6): 333–343.

[7] Ijspeert A, 2014, Biorobotics: Using Robots to Emulate and Investigate Agile Locomotion. Science, 346(6206): 196–203.

[8] Merolla P, Arthur J, Alvarez-Icaza R, et al., 2014, A Million Spiking-Neuron Integrated Circuit with a Scalable Communication Network and Interface. Science, 345(6197): 668–673.