FORGE KINETIC: Decentralized Swarm and Edge Autonomy with World Foundation Model Intelligence

A FORGE OS Subsystem Specification

577 Industries R&D Lab
577 Industries Incorporated
research@577industries.com

Abstract

The deployment of autonomous robotic systems in unstructured physical environments remains constrained by the sim-to-real gap, data scarcity for edge-case scenarios, the absence of integrated perception-to-action architectures, and the challenge of scaling from single-agent intelligence to coordinated multi-agent swarms under adversarial conditions. This paper presents FORGE KINETIC, the decentralized swarm and edge autonomy engine of FORGE OS—the Agent-Legible Operating System for enterprise AI. FORGE KINETIC serves as the "appendages" of FORGE OS, bridging individual agent intelligence with multi-agent coordination through six integrated capabilities: (1) **Domain-Adaptive World Model (DAWM)**, fine-tuning NVIDIA Cosmos world foundation models on operational domain data to achieve 91.3% physical consistency across generated scenarios; (2) **Hierarchical Policy Architecture (HPA)**, coupling a vision-language-model planner at 5–10 Hz with a diffusion-transformer motor controller at 200 Hz; (3) **Confidence-Gated Deployment Pipeline (CGDP)**, providing graduated autonomy through uncertainty quantification and human-in-the-loop escalation; (4) **Fractal Swarm Architecture**, organizing agents into Squad → Platoon → Command hierarchies with $O(\log n)$ communication overhead; (5) **Byzantine Fault Tolerant Collaborative SLAM (BFT-C-SLAM)**, maintaining robust shared spatial maps even when $f < n/3$ agents are compromised; (6) **Sim-to-Real Pipeline** with extreme domain randomization and Bayesian online adaptation. Experimental evaluation demonstrates 87.4% zero-shot sim-to-real transfer (a 34.2-point improvement over conventional baselines), 78.1% performance retention under degraded conditions, sub-200 ms end-to-end latency on NVIDIA Jetson AGX Orin, 1,024-agent swarm scaling with less than 15% mission degradation, and Byzantine fault tolerance for $f < n/3$ compromised agents. All capabilities integrate with FORGE OS subsystems: FORGE QBIT secures swarm communications via PQ Double Ratchet and quorum certificates, FORGE MEMORY governs mission-critical decisions through HITL gates, and FORGE CORE provides distilled intelligence for edge deployment. Case studies in warehouse logistics, defense reconnaissance, industrial inspection, and multi-agent urban search and rescue validate practical applicability.

1 Introduction

The convergence of foundation models and physical robotics represents a transformative moment for autonomous systems. World foundation models (WFMs)—neural networks that predict and generate physics-aware representations of environments—offer a fundamentally new approach to the data scarcity problem that has constrained robotic autonomy for decades [Agarwal et al., 2025]. Simultaneously, multi-agent coordination has matured from simple flocking rules to sophisticated decentralized planning [Li et al., 2025]. However, no existing system unifies both individual agent intelligence and scalable swarm coordination under enterprise governance.

1.1 The Sim-to-Real and Reality-to-Scale Challenges

Autonomous robotic deployment faces three interacting challenges.

The Sim-to-Real Gap. Policies trained in simulation frequently fail on physical hardware due to unmodeled dynamics, sensor discrepancies, and distributional shift [Salvato et al., 2021]. Domain randomization provides partial mitigation but cannot guarantee coverage of operationally critical edge cases [Tobin et al., 2017].

Data Scarcity at the Edge. Defense and industrial environments generate sparse, sensitive training data that cannot be augmented through conventional web-scale collection. A reconnaissance robot operating in a denied environment encounters scenarios absent from any public dataset [Cong and Mo, 2025].

The Scale Gap. Single-agent intelligence has advanced rapidly, but scaling from one agent to 1,024+ coordinated agents introduces communication overhead, trust challenges, and coordination complexity that monolithic architectures cannot address. Compromised agents can poison shared spatial awareness, and communication graph fragmentation can isolate subgroups.

1.2 FORGE KINETIC as the Appendages of FORGE OS

FORGE KINETIC addresses all three challenges by bridging individual agent intelligence with multi-agent swarm coordination under FORGE OS governance [577 Industries, 2025a]. The subsystem inherits and extends two lineages:

- **Individual Agent Intelligence:** The Domain-Adaptive World Model, Hierarchical Policy Architecture, and Confidence-Gated Deployment Pipeline provide each agent with world-model-grounded perception, dual-process action generation, and safety-guaranteed autonomy.
- **Multi-Agent Coordination:** The Fractal Swarm Architecture, BFT-C-SLAM, and Bayesian Sim-to-Real Pipeline enable scalable, fault-tolerant, and adaptation-capable swarm operations.

All capabilities operate under FORGE OS governance: FORGE MEMORY HITL gates govern mission-critical decisions, FORGE QBIT PQ Double Ratchet secures swarm communications, and FORGE CORE provides distilled models for edge inference [577 Industries, 2025c,b,d].

1.3 Contributions

1. **Domain-Adaptive World Model:** Cosmos WFM fine-tuning achieving 91.3% physical consistency for synthetic scenario generation (Section 4.1).
2. **Hierarchical Policy Architecture:** VLM planner (5–10 Hz) coupled with DiT controller (200 Hz) for dual-process autonomy (Section 4.2).
3. **Confidence-Gated Deployment Pipeline:** Graduated autonomy with formal uncertainty quantification and HITL escalation (Section 4.3).
4. **Fractal Swarm Architecture:** Squad $\rightarrow$ Platoon $\rightarrow$ Command hierarchy with $O(n \log n)$ communication scaling to 1,024+ agents (Section 5.1).
5. **BFT-C-SLAM:** Byzantine-tolerant decentralized mapping with quorum certificates signed by FORGE QBIT (Section 5.2).
6. **Sim-to-Real Pipeline:** Extreme domain randomization with Bayesian online adaptation achieving $<$15% sim-to-real degradation (Section 5.3).
7. **FORGE OS Telemetry:** `SWARM_UPDATE` events with KineticVector payloads for full observability (Section 6).
8. **FORGE QBIT PQ Double Ratchet:** Quantum-resistant forward-secret swarm communications (Section 6).

2 Related Work

2.1 World Foundation Models for Physical AI

World foundation models learn dynamics from massive video datasets. NVIDIA Cosmos [Agarwal et al., 2025], trained on 9,000 trillion tokens from 20 million hours of data, provides three families: Cosmos Predict (future state prediction), Cosmos Transfer (sim-to-real adaptation), and Cosmos Reason (chain-of-thought physical reasoning) [NVIDIA, 2025b]. Meta's V-JEPA 2 [Assran et al., 2025] uses self-supervised video learning for robotic planning. DeepMind's Genie 3 [DeepMind, 2025] generates interactive worlds from text. Wayve's GAIA-2 [Wayve, 2024] addresses autonomous driving. None provide the open-weight, customizable architecture necessary for sovereign deployment.

2.2 Vision-Language-Action Models

RT-2 [Brohan et al., 2023] demonstrated that vision-language pre-training transfers to robotic action. NVIDIA's GR00T N1/N1.6 [NVIDIA, 2025a,c] introduced a dual-system VLA architecture: System 2 (VLM) for deliberative reasoning and System 1 (diffusion transformer) for motor control. Skild AI [Skild AI, 2024] and π_0 [Black et al., 2024] represent general-purpose robot brains but require cloud inference, incompatible with air-gapped operations.

2.3 Sim-to-Real Transfer

Classical approaches include domain randomization [Tobin et al., 2017] and system identification [Yu et al., 2017]. Foundation models provide domain-invariant representations bridging simulated and real environments [Firoozi et al., 2023]. Cosmos Transfer converts synthetic renders to photorealistic outputs, reducing the reality gap by 40–60% versus domain randomization alone [Li et al., 2025]. Bayesian online adaptation [Yu et al., 2017] adjusts policies to real hardware during a brief calibration period, addressing residual sim-to-real discrepancies.

2.4 Edge AI for Autonomous Robotics

The NVIDIA Jetson platform—from Orin Nano (15W, 40 TOPS) to AGX Orin (60W, 275 TOPS) to Thor (100W, 800 TOPS) [NVIDIA, 2025d]—provides a hardware continuum. Model distillation, INT8/INT4 quantization, and speculative decoding enable billion-parameter deployment within thermal and power envelopes [Wan et al., 2024].

2.5 Swarm Robotics and Multi-Agent Coordination

Decentralized multi-agent reinforcement learning (MARL) enables coordination without central controllers [Lowe et al., 2017]. Graph-based communication architectures [Sukhbaatar et al., 2016] scale to large agent counts by restricting information flow to graph neighbors. Hierarchical task allocation decomposes missions across organizational levels [Gerkey and Matarić, 2004]. Reynolds flocking [Reynolds, 1987] and consensus algorithms [Olfati-Saber et al., 2007] provide theoretical foundations. However, no existing system combines foundation-model intelligence with scalable swarm coordination under enterprise governance.

2.6 Byzantine Fault Tolerance in Distributed Systems

BFT consensus protocols—PBFT [Castro and Liskov, 1999], HotStuff [Yin et al., 2019]—tolerate $f < n/3$ Byzantine faults. Robust statistics (trimmed mean, geometric median) [Yin et al., 2018] defend distributed estimation against outlier contributions. Application to multi-robot systems remains limited: standard collaborative SLAM assumes all agents are trustworthy, a dangerous assumption in adversarial environments.

2.7 Defense Robotics

DoD Directive 3000.09 [U.S. DoD, 2023] governs autonomy in weapon systems. The Replicator program [U.S. DoD, 2024] emphasizes mass-producible, software-defined autonomous platforms.

Existing commercial platforms are architecturally incompatible with air-gapped, ITAR-controlled, and classified environments.

3 Problem Formulation

3.1 Individual Agent Autonomy

Definition 1 (World Foundation Model) *A WFM $\mathcal{W} : \{v_{1:T}, \ell, s\} \rightarrow \hat{v}_{T+1:T+H}$ maps video observations $v_{1:T}$, optional language instruction ℓ, and sensor state s to predicted future frames $\hat{v}$ over horizon H, satisfying physical consistency constraints (object permanence, rigid-body dynamics, gravitational adherence).*

Definition 2 (Dual-Module Autonomy) *A robotic intelligence system $\mathcal{R} = (\mathcal{P}, \mathcal{C})$ consisting of a perception module $\mathcal{P} : \mathcal{O} \rightarrow \mathcal{Z}$ mapping observations to structured world state, and a command module $\mathcal{C} : \mathcal{Z} \times \mathcal{L} \rightarrow \mathcal{A}$ mapping world state and language instructions to actions.*

Definition 3 (Sim-to-Real Transfer Fidelity) *For policy π trained in E_{sim} and deployed in E_{real}, transfer fidelity is $F(\pi) = J_{E_{real}}(\pi)/J_{E_{sim}}(\pi)$ where $J_E(\pi) = \mathbb{E}_{\tau\sim\pi}\left[\sum_{t=0}^{T} \gamma^t R(s_t, a_t)\right]$. Reliable transfer requires $F(\pi) \geq 0.85$.*

FORGE KINETIC jointly optimizes three objectives:

$$\text{World Model Fidelity:} \quad \mathcal{L}_{\text{world}} = \mathbb{E}_{v\sim D_{\text{real}}}\left[D_{\text{KL}}(p_{\text{real}}(v_{t+1}|v_{1:t}) \| p_{\mathcal{W}}(v_{t+1}|v_{1:t}))\right] \tag{1}$$

$$\text{Policy Performance:} \quad \max_{\pi} J(\pi) \;\; \text{s.t.} \;\; P(\text{unsafe}|\pi) \leq \epsilon_{\text{safety}} \tag{2}$$

$$\text{Edge Efficiency:} \quad \min_{\theta} \mathcal{L}_{\text{latency}}(\theta) \;\; \text{s.t.} \;\; J(\pi_\theta) \geq (1-\delta) \cdot J(\pi_{\text{teacher}}) \tag{3}$$

3.2 Multi-Agent Coordination

Definition 4 (Fractal Swarm) *A hierarchically organized multi-agent system $\mathcal{S} = (L_1, L_2, L_3)$ where L_1 (squads of 3–8 agents) execute local graph Q-learning policies, L_2 (platoons of 3–5 squads) perform inter-squad coordination via auction-based allocation, and L_3 (command) decomposes global objectives, such that total communication overhead is $O(n \cdot \log_b n)$ where b is the branching factor.*

Definition 5 (Byzantine-Tolerant Mapping) *A collaborative SLAM system where n agents maintain a shared spatial map $\mathcal{M}$, tolerating up to $f < n/3$ Byzantine agents (which may report false positions, inject spurious landmarks, or corrupt map updates), such that $\mathcal{M}$ converges to within ε of ground truth with high probability.*

4 Individual Agent Intelligence

FORGE KINETIC decomposes individual agent intelligence into two tightly coupled modules—Perceive and Command—connected by a Confidence-Gated Deployment Pipeline.

4.1 FORGE KINETIC Perceive

4.1.1 Multi-Modal Sensor Fusion

Raw sensor streams (RGB, depth/LiDAR, IMU, proprioception, optional thermal/radar/acoustic) are processed through modality-specific encoders into a unified feature space:

$$z_{\text{fused}} = \text{CrossAttn}\left(\text{Enc}_{\text{rgb}}(I), \text{Enc}_{\text{depth}}(D), \text{Enc}_{\text{prop}}(q, \dot{q})\right) \tag{4}$$

producing a 768-dimensional representation updated at sensor rate (30–60 Hz). Cross-attention enables each modality to attend to complementary features.

For defense applications, graceful degradation handles sensor compromise (camera blinding, GPS denial):

$$\sigma^2_{\text{fused}} = \sum_{m\in\mathcal{M}_{\text{active}}} w_m \cdot \sigma^2_m + \sum_{m\in\mathcal{M}_{\text{denied}}} \sigma^2_{\text{prior},m} \tag{5}$$

4.1.2 Domain-Adaptive World Model (DAWM)

The DAWM fine-tunes Cosmos foundation models on operational domain data through a three-stage pipeline:

Stage 1: Domain Corpus Construction. Operational video (5,000–50,000 clips, 10–100 hours) is curated using Cosmos NeMo Curator with automatic quality filtering, deduplication, and scene classification.

Stage 2: Cosmos Fine-Tuning. Two model families are adapted:

$$\text{Cosmos Predict:} \quad \hat{v}_{t+1:t+H} = \mathcal{W}^{\text{domain}}_{\text{predict}}(v_{t-K:t}, a_t; \theta_{\text{ft}}) \tag{6}$$

$$\text{Cosmos Transfer:} \quad v^{\text{synth}}_{\text{real}} = \mathcal{W}^{\text{domain}}_{\text{transfer}}(v_{\text{sim}}; \phi_{\text{ft}}) \tag{7}$$

Fine-tuning requires 500–2,000 GPU-hours on A100/H100 hardware.

Stage 3: Scenario Generation Engine. The DAWM generates 10,000–100,000 synthetic scenarios per domain: nominal conditions for baseline training, edge cases (sensor failures, environmental extremes) via guided prompting, and adversarial conditions (EW interference, dynamic obstacles) for stress testing. Generated scenarios are validated by a learned physics verifier checking object permanence, collision plausibility, and gravitational adherence.

4.1.3 Spatial-Semantic Scene Graph

A dynamic graph $G = (V, E)$ encodes detected objects (nodes with class, pose, velocity, affordance attributes) and spatial/functional relationships (edges). Updates propagate at 10 Hz through a graph neural network:

$$h_i^{(l+1)} = \text{GRU}\left(h_i^{(l)}, \sum_{j \in \mathcal{N}(i)} \text{MLP}\left[h_i^{(l)} \| h_j^{(l)} \| e_{ij}^{(l)}\right]\right) \tag{8}$$

Defense scene graphs include threat classification nodes with Rules of Engagement (RoE) compliance attributes.

4.1.4 Predictive World Simulation

The DAWM enables "imagination"—simulating consequences of candidate actions before physical execution:

$$\hat{G}_{t+1:t+H} = \mathcal{W}^{\text{domain}}_{\text{predict}}(G_t, a_{t:t+H-1}) \tag{9}$$

supporting look-ahead planning and safety verification against constraint violations.

4.2 FORGE KINETIC Command

4.2.1 System 2: Deliberative Task Planner

The high-level planner operates at 5–10 Hz, decomposing complex instructions into physically grounded sub-goals:

$$g_{1:K} = \text{VLM}(G_t, \ell, c_{\text{mission}}; \theta_{\text{VLM}}) \tag{10}$$

The VLM backbone derives from Cosmos Reason [NVIDIA, 2025b], fine-tuned on domain-specific instruction-following data. Key properties: language-conditioned (natural language is the primary interface), context-aware (mission context maintained in structured buffer), and explainable (chain-of-thought traces logged for post-mission review).

4.2.2 System 1: Reactive Motor Controller

The low-level controller operates at 200 Hz using a diffusion transformer (DiT) [Peebles and Xie, 2023]:

$$a_{t:t+C} = \text{DiT}(z_{\text{fused}}, g_k, q_t, \dot{q}_t; \theta_{\text{DiT}}) \tag{11}$$

where C is the action chunk size (typically 16 steps at 200 Hz = 80 ms of motor commands), trained via flow matching:

$$\mathcal{L}_{\text{flow}} = \mathbb{E}_{t,a_0,\epsilon}\left[\|v_\theta(a_t, t, z) - (a_0 - \epsilon)\|^2\right] \tag{12}$$

Safety Shielding. A Lyapunov barrier function constrains actions to prevent unsafe states:

$$a_{\text{safe}} = \arg\min_{a'} \|a' - a_{\text{DiT}}\|^2 \quad \text{s.t.} \quad \dot{B}(s, a') + \alpha B(s) \leq 0 \tag{13}$$

where $B(s) \geq 0$ defines the safe set. Reflexive emergency behaviors (collision avoidance, joint limits, tipping prevention) override learned policies when triggered.

4.2.3 Cross-Embodiment Adaptation

An embodiment abstraction layer translates canonical end-effector actions to platform-specific joint commands:

$$a_{\text{robot}} = \text{ActionDecoder}_{\text{embodiment}}(a_{\text{canonical}}) \tag{14}$$

enabling deployment across manipulators, mobile platforms, quadrupeds, and humanoids with 100–500 demonstrations and 2–8 GPU-hours of fine-tuning.

4.3 Confidence-Gated Deployment Pipeline

4.3.1 Uncertainty Quantification

Three sources are independently tracked: perceptual (σ_P, Monte Carlo dropout), world model (σ_W, rollout variance), and policy (σ_C, diffusion sample variance). The composite confidence score:

$$\text{conf}(s_t) = 1 - \tanh(\lambda_P \sigma_P + \lambda_W \sigma_W + \lambda_C \sigma_C) \tag{15}$$

4.3.2 Graduated Autonomy

Four autonomy levels based on confidence:

- **Level 4** (Full Autonomy): conf ≥ 0.85. Execute without intervention.
- **Level 3** (Supervised): $0.70 \leq$ conf < 0.85. Execute with detailed logging.
- **Level 2** (Shared Control): $0.50 \leq$ conf < 0.70. Propose and wait for operator confirmation.
- **Level 1** (Teleoperation Assist): conf < 0.50. Operator maintains direct control.

Specific action categories (engagement decisions) are locked to Level 2 or below via FORGE MEMORY HITL gates, regardless of confidence score, enforcing human-in-the-loop requirements for defense operations. All autonomy transitions are logged as `ForgeEvent` records with full reasoning traces.

5 Multi-Agent Swarm Coordination

5.1 Fractal Swarm Architecture

5.1.1 Hierarchical POMDP Structure

FORGE KINETIC organizes agents into a three-level fractal hierarchy where each level operates as a Partially Observable Markov Decision Process (POMDP) with level-appropriate state abstraction and communication bandwidth.

Table 1: Fractal swarm hierarchy levels.

Level	Scope	Policy	Communication
Squad (L_1)	3–8 agents, K=2 hop radius	Local DGQL policies	Full mesh within squad; compressed to platoon leader
Platoon (L_2)	3–5 squads, abstracted state	Inter-squad task allocation (auction)	Squad leaders exchange compressed state vectors
Command (L_3)	All platoons, global objectives	Mission planning	Platoon leaders report; command broadcasts

5.1.2 Deep Graph Q-Learning (DGQL) at Squad Level

Within each squad, agents learn coordinated policies via Deep Graph Q-Learning. A graph neural network encodes the local communication topology:

$$Q_i(s_i, a_i) = f_\theta\left(s_i, \sum_{j \in \mathcal{N}_i} g_\theta(s_j, m_{j \to i})\right) \tag{16}$$

where s_i is agent i's local observation, $\mathcal{N}_i$ are communication neighbors, and $m_{j \to i}$ are messages. The GNN structure makes policy complexity independent of swarm size—each agent's Q-function conditions only on its local neighborhood.

5.1.3 Platoon-Level Coordination

Each squad is abstracted as a single entity characterized by (centroid, capability vector, task status). Inter-squad task allocation uses an auction-based mechanism where squad leaders bid on available tasks based on proximity, capability match, and resource availability. Only squad leaders exchange messages, compressing communication.

5.1.4 Swarm Command

The command level decomposes global mission objectives into platoon-level tasks, allocates resources across platoons, and provides the human operator interface for mission-level commands. Command decisions propagate downward through the hierarchy; battlefield intelligence propagates upward.

5.1.5 Communication Efficiency

Theorem 1 (Communication Scaling) *The fractal architecture achieves total message complexity $O(n \cdot \log_b n)$ compared to $O(n^2)$ for fully connected swarms, where b is the branching factor (typically $b = 3$–5).*

[Proof sketch] At L_1, each squad of k agents communicates in full mesh: $O(k^2)$ messages per squad, $O(n/k \cdot k^2) = O(nk)$ total. At L_2, n/k squad leaders form groups of g: $O(n/(kg) \cdot g^2) = O(ng/k)$ messages. At L_3, $n/(kg)$ platoon leaders communicate: $O((n/(kg))^2)$ messages. With $k, g = O(1)$, total messages are $O(n)$ per level across $O(\log_b n)$ levels: $O(n \log_b n)$.

5.1.6 Graceful Degradation Under Fragmentation

Theorem 2 (Bounded Degradation) *Under k% communication graph fragmentation, mission success rate degrades by at most $\alpha \cdot k$% where $\alpha < 0.3$, provided each squad maintains internal connectivity.*

[Proof sketch] When inter-squad links sever, each squad continues operating autonomously with its DGQL policy, which is trained to function independently. Platoons reform around available squad leaders. The bounded degradation follows because: (a) squad-local tasks continue at near-full performance, (b) coordination loss affects only inter-squad task allocation, and (c) the fractal structure ensures that fragmentation isolates groups rather than individual agents. Empirical validation (Section 8.2.2) confirms $\alpha \approx 0.24$ across deployment scenarios.

5.2 Byzantine Fault Tolerant C-SLAM (BFT-C-SLAM)

Standard collaborative SLAM assumes all agents are trustworthy—a dangerous assumption in adversarial environments where agents may be physically compromised, electronically hijacked, or spoofed.

5.2.1 Threat Model

Up to f of n agents are Byzantine ($f < n/3$). Byzantine agents may report false positions, inject spurious landmarks, corrupt map updates, or behave inconsistently across different communication partners.

5.2.2 Robust Spatial Fusion

Standard mean aggregation of agent contributions is replaced with trimmed-mean estimation:

$$\hat{x}_{\text{robust}} = \text{TrimmedMean}_{f/n}\left(\{x_1, x_2, \ldots, x_n\}\right) \tag{17}$$

where the f/n fraction of most extreme values are discarded from each dimension before averaging. This statistically discards outlier contributions from compromised agents.

5.2.3 Quorum Certificates

Map updates require a quorum certificate: at least $2f+1$ of n agents must cryptographically attest to an observation before it is incorporated into the shared map. Certificates are signed using FORGE QBIT Identity Spine keys [577 Industries, 2025b], ensuring that attestations are non-repudiable and quantum-resistant.

5.2.4 Anomaly Detection

Continuous monitoring compares each agent's contributions against swarm consensus. Agents consistently deviating beyond a configurable threshold are automatically isolated from the map fusion process. Isolation events are logged as `SWARM_UPDATE` `ForgeEvent` records in FORGE MEMORY.

5.2.5 Formal Guarantee

Theorem 3 (BFT-C-SLAM Convergence) *The decentralized map converges to within ε of ground truth with probability at least $1-\delta$, even with $f < n/3$ Byzantine agents, provided: (a) the quorum certificate threshold $2f+1$ is maintained, (b) the trimmed-mean aggregation uses trim factor f/n, and (c) at least $n-f$ honest agents observe each map region.*

[Proof sketch] With trim factor f/n, the trimmed mean excludes the f most extreme values in each dimension. Since Byzantine agents contribute at most f values, all Byzantine contributions are trimmed with high probability. The remaining $n-2f \geq f+1$ honest contributions produce a consistent estimate within ε of ground truth by the law of large numbers. The quorum certificate ensures that only observations attested by at least $f+1$ honest agents are accepted.

5.3 Sim-to-Real Pipeline

5.3.1 Extreme Domain Randomization

Training in NVIDIA Isaac Sim incorporates extreme randomization: Gaussian and salt-and-pepper sensor noise, motion blur, thermal throttling, bursty correlated packet loss, actuator delay variation (5–50 ms), and lighting/weather variation. This produces policies robust to the wide range of conditions encountered in real deployment.

5.3.2 Transfer Protocol

Trained models are exported to ONNX format and compressed through FORGE CORE's distillation pipeline for Jetson AGX Orin deployment. INT8 quantization via TensorRT maintains performance within 1.5% of full precision.

5.3.3 Bayesian Online Adaptation

On first physical deployment, a 30–60 minute adaptation phase calibrates the policy to real hardware:

$$p(\phi|D_{\text{real}}) \propto p(D_{\text{real}}|\phi) \cdot p(\phi|D_{\text{sim}}) \tag{18}$$

where ϕ represents environment dynamics parameters, $p(\phi|D_{\text{sim}})$ is the prior from simulation training, and D_{real} is a small number of real-world transitions. This Bayesian update adjusts the policy to real sensor/actuator characteristics while retaining the broad coverage from simulation training.

6 FORGE OS Integration

6.1 Telemetry

FORGE KINETIC emits `SWARM_UPDATE` events with KineticVector payloads on position change, task assignment, formation update, and engagement decision. Events conform to the `ForgeEvent` protobuf schema [577 Industries, 2025a]. During disconnected operation, events are buffered locally and batch-submitted on reconnect.

6.2 FORGE QBIT Dependency

- **PQ Double Ratchet:** Default swarm-to-swarm communication protocol, providing forward secrecy and post-quantum security [577 Industries, 2025b].
- **Quorum Certificates:** BFT-C-SLAM map updates signed by FORGE QBIT Identity Spine keys.
- **Agent Admission:** Agents must present valid FORGE QBIT certificates before swarm admission.

6.3 FORGE MEMORY Governance

Mission-critical decisions (weapons release, perimeter breach, casualty triage) are gated by FORGE MEMORY HITL gates [577 Industries, 2025c]. Communication latency to edge devices is accounted for in HITL gate timeout calculations. All swarm state changes are recorded in FORGE MEMORY's IGOM as immutable governance objects.

6.4 FORGE CORE Intelligence

Distilled models from FORGE CORE's staged post-training pipeline [577 Industries, 2025d] are deployed to Jetson hardware. Edge inference model selection uses FORGE CORE's routing engine when cloud connectivity is available. Drift feedback from FORGE KINETIC agents informs FORGE CORE's continuous distillation pipeline.

7 Implementation

7.1 Technology Stack

Table 2 summarizes the FORGE KINETIC implementation.

Table 2: FORGE KINETIC technology stack.

Component	Implementation
World Model	Cosmos Predict 2.5 + Transfer 2.5, domain fine-tuned
Reasoning VLM	Cosmos Reason 2B, domain fine-tuned
Policy DiT	Custom 16-layer DiT, flow matching
Scene Graph	PyTorch Geometric GNN, incremental update
Sensor Fusion	Custom cross-attention encoder (768-dim)
Safety Shield	Lyapunov barrier via convex QP solver
Edge Runtime	TensorRT, INT8 quantization, CUDA graphs
Sim Platform	NVIDIA Isaac Sim + Omniverse
Physics Engine	Newton (GPU-accelerated, Isaac Lab)
Hardware	Jetson AGX Orin / Jetson Thor
Orchestrator	Python asyncio, ROS 2 Humble
Swarm DGQL	PyTorch + PyTorch Geometric
BFT Consensus	Custom quorum certificate library
Swarm Comms	FORGE QBIT PQ Double Ratchet client

7.2 Training Pipeline

Phase 1: Domain World Model (500–2,000 GPU-hours). Cosmos fine-tuning on operational video using LoRA adaptation.

Phase 2: Synthetic Data Generation (100–500 GPU-hours). DAWM generates 10K–100K scenarios in Isaac Sim with Cosmos Transfer.

Phase 3: Policy Training (200–1,000 GPU-hours). Hierarchical policy trained on real demonstrations (100–1,000 episodes), synthetic scenarios, and human video data.

Phase 4: Edge Distillation (50–200 GPU-hours). VLM planner quantized to INT8 via TensorRT; DiT controller quantized with CUDA graph optimization; sensor encoders pruned and quantized.

Phase 5: Swarm Coordination (300–1,000 GPU-hours). DGQL training across simulated swarms of increasing size. BFT-C-SLAM validation with injected Byzantine agents.

Total: 1,150–4,700 GPU-hours per domain.

7.3 Deployment Modes

Mode 1: Sovereign Edge. All inference on Jetson hardware. No external data transmission. Supports ITAR, IL5/IL6, and air-gapped environments.

Mode 2: Hybrid. Edge inference for real-time control; cloud connectivity for world model updates and policy refinement during non-operational periods.

Mode 3: Cloud-Connected. Full cloud inference via DGX Cloud. Maximum model capacity; requires persistent connectivity.

8 Experimental Evaluation

8.1 Individual Agent Performance

8.1.1 Datasets and Benchmarks

Three domains: **Warehouse Logistics (WL)**—Isaac Sim, box picking (WL-P), navigation (WL-N), order fulfillment (WL-O), 500 episodes/task; **Defense Reconnaissance (DR)**—Isaac Sim + physical, route planning (DR-R), target ID (DR-I), 400 sim + 50 physical episodes/task; **Industrial Inspection (II)**—physical, defect detection, 300 cycles.

8.1.2 Zero-Shot Sim-to-Real Transfer

Table 3 presents zero-shot transfer results.

Table 3: Zero-shot sim-to-real transfer: task completion rate (%). Policies trained entirely in simulation and deployed to physical hardware without fine-tuning.

Method	WL-P	WL-N	WL-O	DR-R	DR-I	Avg
DR-Random	42.6	38.4	28.1	35.2	31.7	35.2
DR-Adaptive	51.3	46.8	34.5	42.1	38.3	42.6
BC-Real	68.2	55.1	42.3	48.7	44.2	51.7
RT-2-Adapt	71.4	62.7	48.6	53.2	49.8	57.1
GR00T-FT	82.1	78.3	68.4	72.6	67.9	73.9
FORGE KINETIC	**91.3**	**89.7**	**82.4**	**85.1**	**88.6**	**87.4**

FORGE KINETIC achieves 87.4% average task completion—a 34.2-point improvement over DR-Adaptive and 13.5 points over the cloud-inference GR00T ceiling. Improvement is most pronounced in complex multi-step tasks (WL-O: +14.0 over GR00T-FT) and defense scenarios (DR-I: +20.7).

Table 4: DAWM generation quality metrics.

Metric	Cosmos Base	Cosmos+DR	DAWM
FID ↓	42.3	31.7	**18.4**
Physics Consistency (%) ↑	74.8	79.3	**91.3**
Object Permanence (%) ↑	82.1	84.6	**94.7**
Domain Match (%) ↑	61.2	68.4	**89.2**

Table 5: Ablation study: component contribution to task completion (%).

Configuration	Avg. Completion	Δ
Full FORGE KINETIC	87.4	—
− DAWM (base Cosmos)	73.1	−14.3
− Cosmos Reason (random VLM)	76.8	−10.6
− Predictive Simulation	79.3	−8.1
− Scene Graph	81.2	−6.2
− Dual-Rate (single 50 Hz)	82.6	−4.8
− Safety Shield	85.9	−1.5
Perceive only (no Command)	0.0	−87.4
Command only (no Perceive)	34.7	−52.7

8.1.3 DAWM Generation Quality

8.1.4 Hierarchical Policy Ablation

DAWM contributes the largest single improvement (+14.3 points), confirming that domain-specific world modeling is the primary driver of sim-to-real performance. Both modules are essential: neither alone achieves acceptable performance.

8.1.5 Edge Deployment Performance

Table 6: Edge inference performance on Jetson AGX Orin (INT8).

Component	P50 (ms)	P95 (ms)	P99 (ms)
Sensor Fusion	8.2	14.1	18.7
Scene Graph Update	12.4	22.8	31.2
VLM Planner (System 2)	84.3	142.7	198.4
DiT Controller (System 1)	4.8	7.2	9.1
Safety Shield	1.2	2.4	3.8
CGDP Uncertainty Est.	3.6	6.1	8.4
End-to-End (reactive)	26.6	44.5	61.3
End-to-End (w/ planning)	114.5	192.1	261.5

The reactive loop achieves 26.6 ms P50 (37.6 Hz); the full planning loop achieves 114.5 ms P50 (8.7 Hz), both within sub-200 ms targets.

INT8 preserves 98.5% of full-precision performance at 28% power reduction. The distilled 1.1B model achieves 83.9% at 28W.

8.1.6 Robustness Under Degraded Conditions

FORGE KINETIC maintains 78.1% average under degraded conditions versus 50.9% for GR00T-FT, with the advantage most pronounced under combined degradation (68.7% vs. 31.2%).

Table 7: Model size and power consumption on Jetson AGX Orin.

Configuration	Params	Power (W)	Compl. (%)
Full Precision (FP16)	2.8B	58	87.4
INT8 Quantized	2.8B	42	86.1
INT4 Quantized	2.8B	31	82.7
Distilled (INT8)	1.1B	28	83.9
Orin Nano Target	—	≤15	74.3

Table 8: Task completion under degraded conditions (%).

Condition	FORGE KINETIC	GR00T	RT-2	DR-A
Nominal	87.4	73.9	57.1	42.6
Low Light (10 lux)	79.8	58.2	41.3	31.4
Sensor Dropout (1 cam)	82.1	54.7	38.6	28.9
GPS Denied	83.6	61.3	44.2	33.1
EMI Interference	78.4	47.8	32.1	24.7
Dust/Obscurant	76.2	52.4	36.8	27.3
Combined (3+ degrad.)	68.7	31.2	18.4	12.8
Avg. Degraded	78.1	50.9	35.2	26.4

8.1.7 Statistical Analysis

Pairwise Welch's t-tests confirm significance of all improvements ($p < 0.001$). Cohen's d effect sizes: vs. DR-Adaptive, $d = 3.42$; vs. RT-2-Adapt, $d = 2.18$; vs. GR00T-FT, $d = 1.47$ (all large). Bootstrap 95% CI for improvement over GR00T-FT: $[11.8, 15.2]$ points. Two-way ANOVA: significant main effect of method ($F(5, 2994) = 187.3$, $p < 0.001$), non-significant method$\times$domain interaction ($p = 0.142$).

8.2 Multi-Agent Swarm Performance

8.2.1 Scaling Analysis

Table 9 evaluates mission success as swarm size increases.

Table 9: Mission success rate (%) and communication overhead vs. swarm size. Flat: fully connected communication. Fractal: FORGE KINETIC hierarchical architecture.

	Mission Success (%)		Messages/step	
Agents	Flat	Fractal	Flat	Fractal
16	91.2	90.8	256	64
64	84.7	88.4	4,096	192
256	71.3	86.1	65,536	512
512	58.4	84.2	262,144	864
1,024	42.1	82.7	1,048,576	1,536

The fractal architecture scales gracefully: only 8.1-point degradation from 16 to 1,024 agents (vs. 49.1 points for flat). Communication overhead grows as $O(n \log n)$ versus $O(n^2)$, validating Theorem 1.

8.2.2 Communication Fragmentation Resilience

At 50% fragmentation, the fractal architecture retains 73.8% success (exceeding the 70% target), while flat architecture collapses to 31.4%. The measured degradation coefficient $\alpha \approx 0.24$, confirming Theorem 2.

Table 10: Mission success (%) under communication graph fragmentation.

Fragmentation	Flat (256 agents)	Fractal (256 agents)
0%	71.3	86.1
10%	62.8	83.4
25%	48.1	79.2
50%	31.4	73.8
75%	14.7	62.1

Table 11: Map accuracy (RMSE in meters) under varying numbers of Byzantine agents ($n = 24$).

Byzantine Agents (f)	Standard C-SLAM	BFT-C-SLAM
0	0.12	0.14
1	0.31	0.16
2	0.58	0.18
4	1.24	0.22
6	2.87	0.28
8 (=$n/3$)	6.41	0.34

8.2.3 BFT-C-SLAM Evaluation

Standard C-SLAM degrades catastrophically with Byzantine agents (6.41 m RMSE at $f = n/3$). BFT-C-SLAM maintains sub-0.35 m accuracy throughout, validating Theorem 3. The 0.02 m overhead at $f = 0$ is the cost of the trimmed-mean estimator and quorum certificate verification.

8.2.4 Sim-to-Real Transfer with Bayesian Adaptation

Table 12: Sim-to-real performance gap (% degradation from simulation) across adaptation phases.

Environment	No Adapt.	15 min	30 min	60 min
Indoor Warehouse	18.2%	11.4%	7.8%	5.1%
Outdoor Terrain	24.7%	16.8%	12.1%	8.9%
GPS-Denied Urban	21.3%	14.2%	10.4%	7.3%
Average	21.4%	14.1%	10.1%	7.1%

Bayesian online adaptation reduces the sim-to-real gap from 21.4% to 7.1% after 60 minutes of real-world operation, well below the 15% target. Most improvement occurs in the first 30 minutes.

8.2.5 PQ Double Ratchet Communication Overhead

The PQ Double Ratchet introduces 21% throughput reduction and 1.8 ms median latency overhead—negligible relative to the 114.5 ms planning loop—while providing post-quantum forward secrecy for all swarm communications.

8.3 Integrated FORGE OS Evaluation

Telemetry Completeness. 98.7% of swarm actions produced corresponding `ForgeEvent` records. Cross-subsystem trace completeness for defense missions: 99.2% of mission-critical decisions traceable end-to-end through FORGE KINETIC → FORGE MEMORY → FORGE QBIT.

HITL Gate Responsiveness. Median latency from HITL trigger to operator notification: 2.4 s (including satellite uplink for remote operations). Speculative execution rate for routine swarm decisions: 94.1% (i.e., 94.1% of HITL-gated decisions matched the predicted approval, enabling near-zero-latency governance).

Table 13: PQ Double Ratchet overhead for swarm communications ($n = 256$ agents).

Metric	Plaintext	PQ Double Ratchet
Message throughput (msg/s)	12,400	9,800
P50 latency overhead	—	+1.8 ms
P99 latency overhead	—	+4.2 ms
Key material per agent	—	3.2 KB
Total key material ($n = 256$)	—	819 KB

9 Deployment Case Studies

9.1 Warehouse Logistics

Setting. 3PL distribution center, 180,000 sq. ft., 4 mobile manipulators (UR10e on Clearpath Ridgeback), 6-month pilot. Multi-robot coordination via fractal swarm (1 squad of 4 agents).

Results. 93.2% task completion (vs. 71.4% scripted); 42.3 s average cycle time (vs. 68.7 s); zero safety incidents over 47,000 operations; 3.1× throughput improvement; ROI in 4.2 months. FORGE MEMORY audit trail captured all safety-relevant decisions for warehouse compliance.

9.2 Defense Reconnaissance

Setting. Controlled evaluation with U.S. Army unit, GPS-denied urban environment, 2 quadruped robots (Boston Dynamics Spot) with multi-robot coordination via BFT-C-SLAM for shared mapping, 3-month evaluation.

Results. 94.7% route completion (vs. 67.3% scripted waypoint); 91.2% threat identification accuracy (4.3% FPR); 11.4 min average mission time (vs. 18.7 min manual); CGDP escalated to Level 2 in 23.4% of missions; zero unintended autonomous actions; 4.2/5.0 operator trust. FORGE QBIT PQ Double Ratchet secured all inter-robot communications. FORGE MEMORY HITL gates enforced human approval for all engagement-relevant decisions.

9.3 Industrial Visual Inspection

Setting. Precision manufacturing, 1 articulated arm (Fanuc CRX-10iA), 4-month evaluation. DAWM generated 62,000 synthetic inspection scenarios covering 340 component types and 47 defect categories.

Results. 97.8% defect detection (vs. 89.4% human, 92.1% prior ML); 2.1% false positive rate (vs. 8.7%); 47 components/hour (vs. 22 human); 12 previously uncaught defect patterns discovered; 67% reduction in quality escapes. FORGE MEMORY audit trail provided full traceability for ISO 9001 compliance.

9.4 Multi-Agent Urban Search and Rescue

Setting. GPS-denied, partially collapsed structure. 16 agents organized as 4 squads of 4, communicating via FORGE QBIT PQ Double Ratchet. BFT-C-SLAM for shared mapping of damaged building. CGDP graduated autonomy for hazardous area exploration. FORGE MEMORY HITL for casualty triage decisions.

Results. 92.4% building coverage in 45 minutes (vs. 61.3% for uncoordinated); 0.28 m map RMSE with BFT-C-SLAM (robust to 2 simulated compromised agents); 8 simulated casualty locations identified (100% detection); mean response time to first casualty contact: 4.7 minutes. All triage recommendations routed through FORGE MEMORY HITL gates with median 3.1 s operator response time.

10 Discussion

10.1 Key Findings

1. **Domain-adaptive world modeling is the critical enabler.** The 14.3-point DAWM ablation impact exceeds all other components, confirming that the primary bottleneck is training data fidelity, not model architecture. Fine-tuning Cosmos (500–2,000 GPU-hours) yields disproportionate returns.
2. **The dual-module architecture provides genuine compositional benefit.** Neither Perceive alone (0%) nor Command alone (34.7%) achieves acceptable performance; their combination yields 87.4%. The VLM planner decomposes tasks in ways that enable the DiT controller to succeed on sub-goals intractable as monolithic policies.
3. **Fractal swarm scales to 1,024+ agents with bounded degradation.** Only 8.1% mission success degradation from 16 to 1,024 agents, versus 49.1% for flat architectures. The $O(n \log n)$ communication scaling enables real-world deployment at scales where $O(n^2)$ becomes prohibitive.
4. **BFT-C-SLAM eliminates the single-agent trust assumption.** Standard C-SLAM collapses under Byzantine attack (6.41 m RMSE at $f = n/3$); BFT-C-SLAM maintains sub-0.35 m accuracy. This is critical for defense operations where agent compromise is a realistic threat.
5. **Confidence-gated autonomy enables trust in adversarial conditions.** Zero unintended autonomous actions across all defense evaluations. The 4.2/5.0 operator trust rating demonstrates that transparent uncertainty communication is essential for defense adoption.
6. **Edge deployment is viable without catastrophic performance loss.** INT8 on Jetson AGX Orin preserves 98.5% of full-precision performance, challenging the assumption that foundation-model robotics requires cloud connectivity.

10.2 Comparison with Commercial Platforms

Table 14: FORGE KINETIC vs. commercial platforms.

Capability	FORGE KINETIC	Cloud VLA	Scripted	ROS 2 Multi
World Model	Domain-adaptive Cosmos	Generic/none	None	None
Edge Deploy	Full sovereign	Cloud only	Yes	Yes
Language Ctrl	Natural language	Natural language	Coded	Coded
Sim-to-Real	DAWM + CGDP	Basic DR	Manual	Basic DR
Safety	Formal barrier	Soft limits	Hard-coded	Soft limits
Swarm	Fractal BFT	None	None	Basic
Air-Gap	Yes	No	Yes	Yes
Explainable	CoT reasoning	Opaque	Deterministic	None

FORGE KINETIC uniquely combines foundation-model intelligence with sovereign edge deployment, formal safety guarantees, and scalable BFT swarm coordination.

10.3 Limitations

Manipulation Dexterity. Currently supports rigid-body manipulation. Deformable objects require world model extensions.

BFT Threshold. BFT-C-SLAM assumes $f < n/3$. Higher Byzantine ratios require different consensus approaches.

Fractal Command Latency. The hierarchical structure adds $O(\log n)$ latency for top-down commands—unsuitable for microsecond coordination.

Bayesian Adaptation Period. 30–60 minutes of real-world operation required; unacceptable for immediate deployment scenarios.

Adversarial Robustness. Not evaluated against deliberate adversarial attacks on perception (adversarial patches, spoofed sensors).

Long-Horizon Reasoning. Very long missions (hours, dozens of objectives) may exceed VLM context window.

10.4 FORGE OS Integration Value

FORGE KINETIC agents operating under FORGE MEMORY governance are more trustworthy than ungoverned autonomy—human oversight for critical decisions is architecturally guaranteed, not optional. FORGE QBIT PQ-secured communications enable operations in contested electromagnetic environments where classical encryption may be vulnerable. FORGE CORE distilled models provide on-edge intelligence without cloud dependency, with continuous improvement through drift feedback.

11 Conclusion

This paper presented FORGE KINETIC, the decentralized swarm and edge autonomy engine of FORGE OS. FORGE KINETIC bridges individual agent intelligence—through Domain-Adaptive World Modeling, Hierarchical Policy Architecture, and Confidence-Gated Deployment—with multi-agent coordination through Fractal Swarm Architecture, BFT-C-SLAM, and Bayesian Sim-to-Real Transfer.

Individual agents achieve 87.4% zero-shot sim-to-real transfer with sub-200 ms edge latency. Swarms scale to 1,024+ agents with less than 15% degradation, tolerate $f < n/3$ Byzantine faults, and maintain 73.8% mission success under 50% communication fragmentation. All operations are governed by FORGE OS: FORGE MEMORY HITL gates ensure human oversight, FORGE QBIT secures communications, and FORGE CORE provides distilled intelligence.

Future work includes Jetson Thor deployment for real-time deliberative control, federated learning across distributed sites, multimodal sensor expansion (hyperspectral, CBRN), and extended long-horizon mission planning with persistent memory.

References

N. Agarwal, A. Ali, M. Bala, Y. Balaji, et al. Cosmos world foundation model platform for physical AI. NVIDIA Technical Report, January 2025.

Grand View Research. Military robots market size, share & trends analysis report, 2025–2030. Market Research Report, 2025.

E. Salvato, G. Fenu, E. Medvet, and F. A. Pellegrino. Crossing the reality gap: A survey on sim-to-real transferability of robot controllers in reinforcement learning. *IEEE Access*, 9:153171–153187, 2021.

R. Firoozi, J. Tucker, S. Tian, A. Majumdar, et al. Foundation models in robotics: Applications, challenges, and the future. *arXiv preprint arXiv:2312.07843*, 2023.

Y. Cong and H. Mo. An overview of robot embodied intelligence based on multimodal models. *Int. J. Intelligent Systems*, 2025:5124400, 2025.

A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, et al. RT-2: Vision-language-action models transfer web knowledge to robotic control. In *Proc. CoRL*, 2023.

NVIDIA. Isaac GR00T N1: An open foundation model for generalist humanoid robots. NVIDIA Technical Report, March 2025.

NVIDIA. NVIDIA announces major release of Cosmos world foundation models and physical AI data tools. NVIDIA Press Release, March 2025.

NVIDIA. Isaac GR00T N1.6: A foundation model for generalist robots. NVIDIA GitHub Repository, 2025.

M. Assran, M. Bardes, D. Bouchacourt, et al. V-JEPA 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint*, 2025.

Google DeepMind. Genie 3: A new frontier for world models. DeepMind Blog, 2025.

Wayve. GAIA-2: A generative world model for autonomous driving. Wayve Technical Report, 2024.

Skild AI. A general-purpose brain for robots. Skild AI Technical Report, 2024.

K. Black, N. Brown, D. Driess, et al. π_0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *Proc. IROS*, pages 23–30, 2017.

W. Yu, J. Tan, C. K. Liu, and G. Turk. Preparing for the unknown: Learning a universal policy with online system identification. In *Proc. RSS*, 2017.

P. Li, Z. An, S. Abrar, and L. Zhou. Large language models for multi-robot systems: A survey. *arXiv preprint arXiv:2502.03814*, 2025.

NVIDIA. Jetson modules, support, ecosystem, and lineup. NVIDIA Developer, 2025.

Z. Wan, X. Wang, C. Liu, and A. Abad. Efficient large language models: A survey. *arXiv preprint arXiv:2312.03863*, 2024.

R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proc. NeurIPS*, volume 30, December 2017.

S. Sukhbaatar, A. Szlam, and R. Fergus. Learning multiagent communication with backpropagation. In *Proc. NeurIPS*, volume 29, December 2016.

B. P. Gerkey and M. J. Matarić. A formal analysis and taxonomy of task allocation in multi-robot systems. *Int. J. Robotics Research*, 23(9):939–954, 2004.

C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proc. SIGGRAPH*, pages 25–34, 1987.

R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proc. IEEE*, 95(1):215–233, 2007.

M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proc. OSDI*, pages 173–186, 1999.

M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham. HotStuff: BFT consensus with linearity and responsiveness. In *Proc. PODC*, pages 347–356, 2019.

D. Yin, Y. Chen, R. Kannan, and P. Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *Proc. ICML*, pages 5650–5659, July 2018.

U.S. Department of Defense. Autonomy in weapon systems. DoD Directive 3000.09, January 2023.

U.S. Department of Defense. Replicator initiative: Scaling autonomous systems. DoD Fact Sheet, 2024.

W. Peebles and S. Xie. Scalable diffusion models with transformers. In *Proc. ICCV*, 2023.

D. Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, 2011.

D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.

V. Makoviychuk, L. Wawrzyniak, Y. Guo, et al. Isaac Gym: High performance GPU-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.

Open X-Embodiment Collaboration. Open X-Embodiment: Robotic learning datasets and RT-X models. *arXiv preprint arXiv:2310.08864*, 2023.

577 Industries R&D Lab. FORGE OS: The Agent-Legible Operating System — Platform Spine Specification. Technical report, 577 Industries Incorporated, 2025.

577 Industries R&D Lab. FORGE QBit: Heterogeneous Post-Quantum Cryptography and Physics-Informed Computation for Sovereign AI Systems. Technical report, 577 Industries Incorporated, 2025.

577 Industries R&D Lab. FORGE Memory: Governed Multi-Agent Orchestration with Predictive Human-in-the-Loop Intelligence. Technical report, 577 Industries Incorporated, 2025.

577 Industries R&D Lab. FORGE Core: A Causal Model-Agnostic Intelligence and Routing Engine for Enterprise AI Deployment. Technical report, 577 Industries Incorporated, 2025.