# Autonomous Spatial Networks with *Bigraphs*

**Josh Millar**, Ryan Gibb, Anil Madhavapeddy, Hamed Haddadi

# The Spatial Disconnect

**Problem:** Networks lack coherent spatial models

- Networks increasingly rely on spatial reasoning
  - e.g. policy enforcement or access control tied to physical presence
- Yet, lack a unified way to coordinate physical space.
- Spatial policies are *fragile* and *manual*.

# A Model for Spatial Coordination

**Bigraphs: a unifying model for spatial networks**

- Capture *spatial*, *social*, and *network* linkages in a single formalism
- Properties
  - Dynamic; localized *reaction rules* respond to events
  - Verifiable; safe and correct of spatial policies
  - Compositional
    - policy reuse and reasoning at different spatial scales
    - nodes update only their local graph view.

# Bigraphs for Spatial Coordination

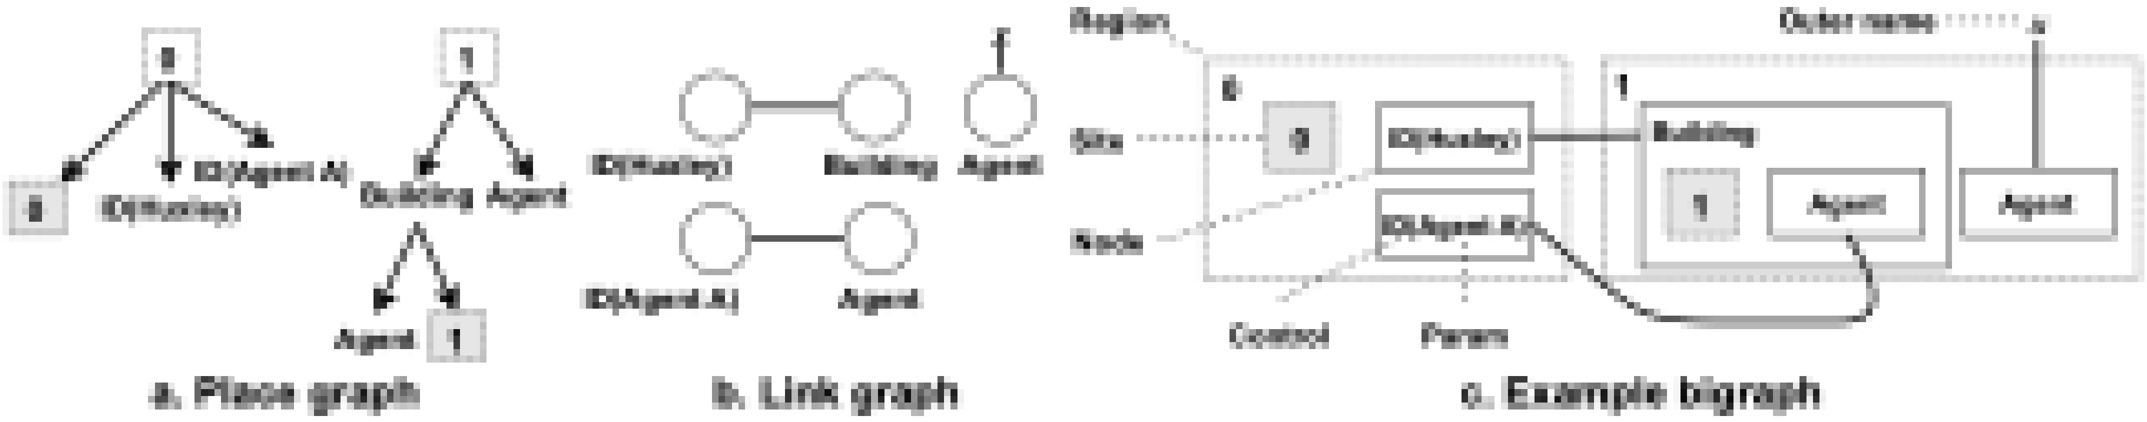


Example bigraph, built from its underlying place and link graphs.

# Bigraphs for Spatial Coordination

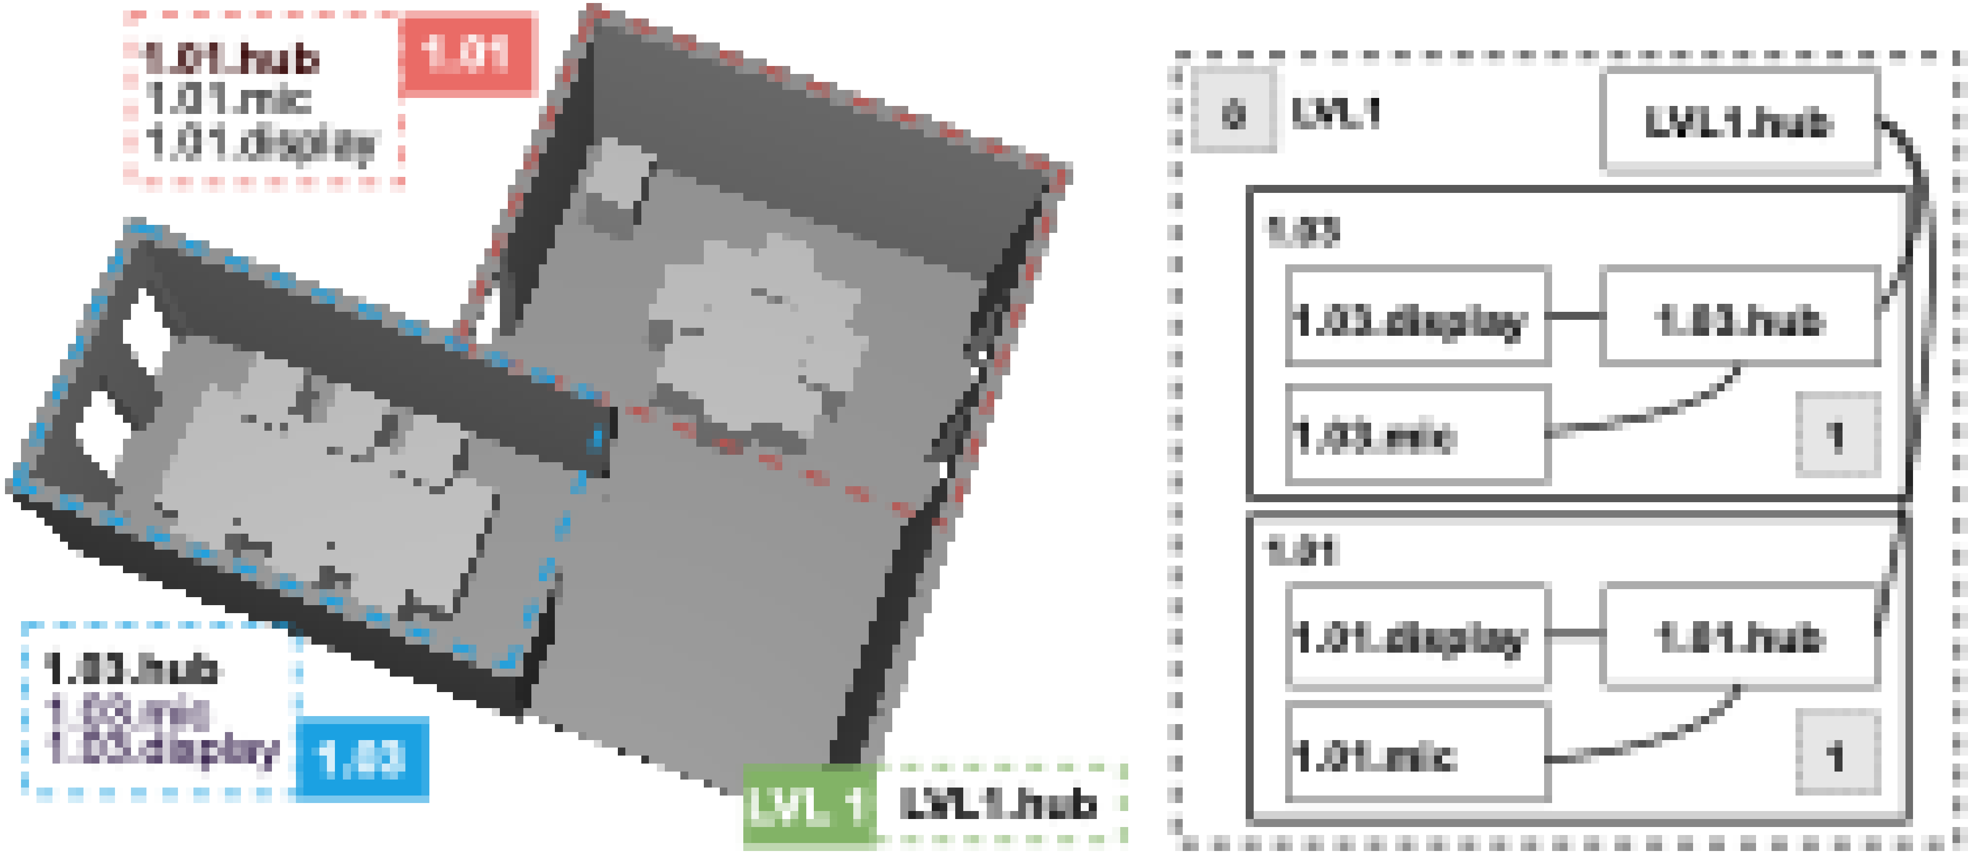


a. RoomPlan visualisation of office

b. Bigraph representation of office

# Spatially-Scoped Agents

- Growth of autonomous agents in smart environments [1, 2]
- To act reliability, agents need a structured ‘world model’
  - free-form LLM outputs are insufficient for spatial reasoning [3].
- Decisions should be *local* and *rule-based* whenever possible
  - Why do I need an LLM to turn on the lights?!
- But also support contextual reasoning over multimodal data

**[1] An et al., 2025** – IoT-LLM: Enhancing Real-World IoT Task Reasoning with Large Language Models
**[2] Shen et al., 2025** – GPIoT: Tailoring Small Language Models for IoT Program Synthesis and Development
**[3] Lu et al., 2025** – Learning to Generate Structured Output with Schema Reinforcement Learning

# Spatially-Scoped Agents

**Agent Tiers:**

- **Leaf:** tiny models for fast, local decisions
- **Delegated:** small LMs reason over sub-graphs
  **Root:** LLMs for complex reasoning & model updates

Reasoning is **scoped to the smallest sub-graph & tier.**

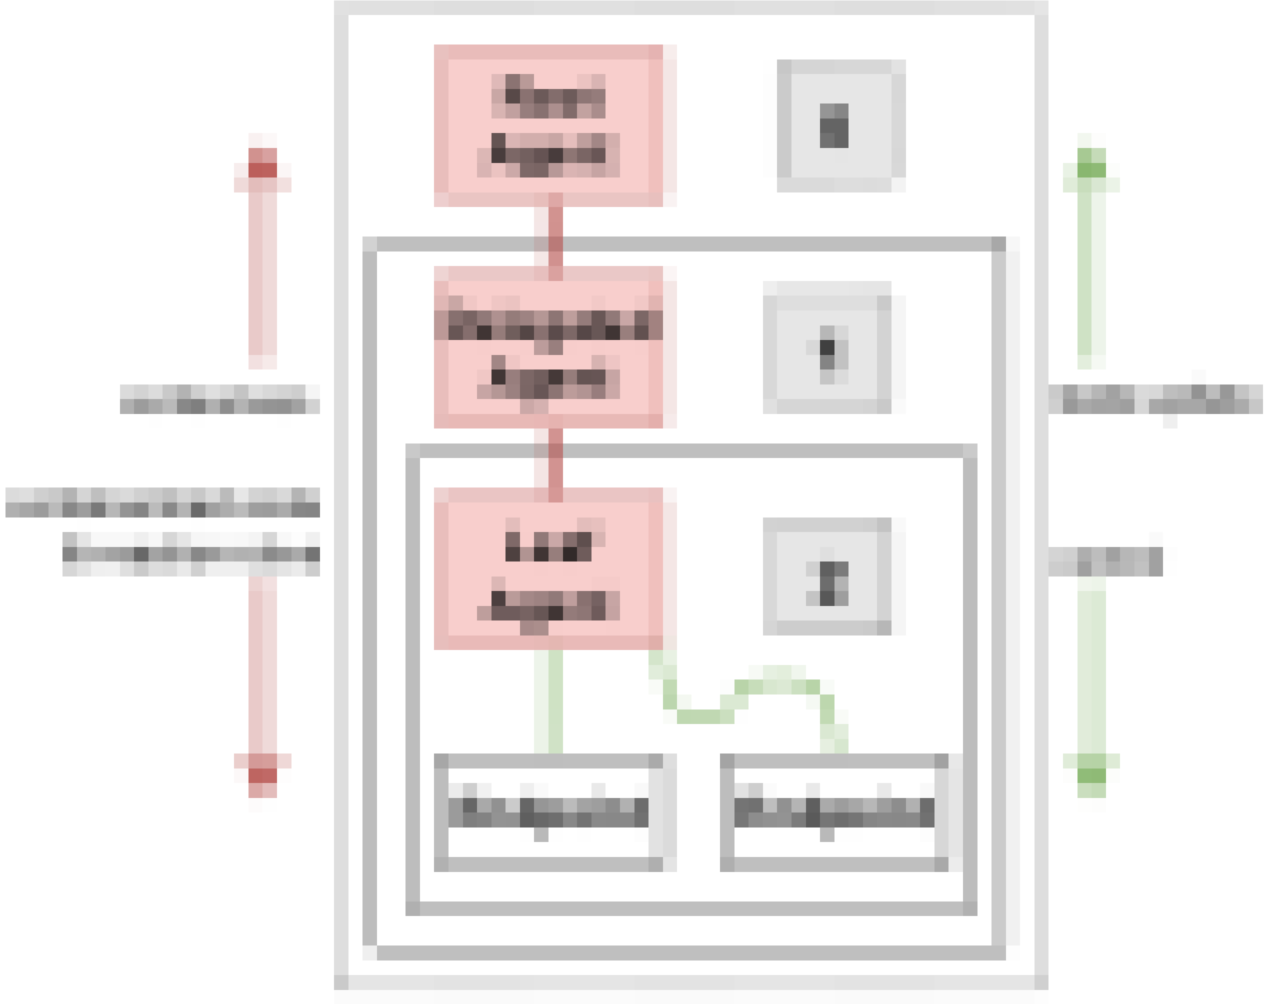

# Spatially-Scoped Agents

- **Latency**:
  - Local agents act without cloud roundtrips.
  - Immediate reactions via graph pattern matching to events.
- **Privacy**: agents only act on local graph views.
  - data captured shouldn’t leave the physical space.
- **Reliability**: distributed, no single point of failure.
- **Efficiency**: minimized data movement and computation.

# Contextual Reasoning over Bigraphs

- MCP server sits atop the graph layer
  - exposes its structure as a callable interface.
- Agents reason, update, or query over the graph via function calls
- Higher-tier LLMs push containerized programs, rules, and models to relevant nodes [1], verified and composed by the MCP server

**[1] Shen et al., 2025** – GPIoT: Tailoring Small Language Models for IoT Program Synthesis and Development

# Distributed Decision-Making

- User intents expressed as natural language policy
  - e.g. “during off-hours, dim lights in unoccupied common areas”
- LLM translates to *reaction rules* which *fire* on spatial transitions, updating local subgraphs.
- Maintains graph consistency along with availability

```
react shutdown_nodes =
/x (FN06.(Users.() || Node{x} || rest))
--> FN06.(rest);

begin brs
init ...;
rules = [{shutdown_nodes}, {...}];
end
```

**Example reaction rule:** all nodes in FN06 are shut down if no users are present.*

* Written in BigraphER, an OCaml library for the manipulation of bigraphs.
[https://uog-bigraph.bitbucket.io/](https://uog-bigraph.bitbucket.io/)

# Escalation over Tiers

- From Leaf Nodes to Agents
  - Rule-driven
    - “if more than one user present in the space, escalate with identifiers".
  - Unknown or ambiguous (i.e., no matching patterns)
    - graceful handling of novel scenarios.
- Between Agents
  - Self-Assessed Uncertainty;
  - Policy Scope Violation; manifest details the agent’s jurisdiction
- Bounded Escalation
  - Defined schema, clear scoping, auditability & hash tracing.

# What's Possible Locally?

- MCUs handle complex jobs in real-time (e.g., audio transcription) with
  - integrated $\mu$NPUs,
  - lightweight architectures,
  - and neural compression

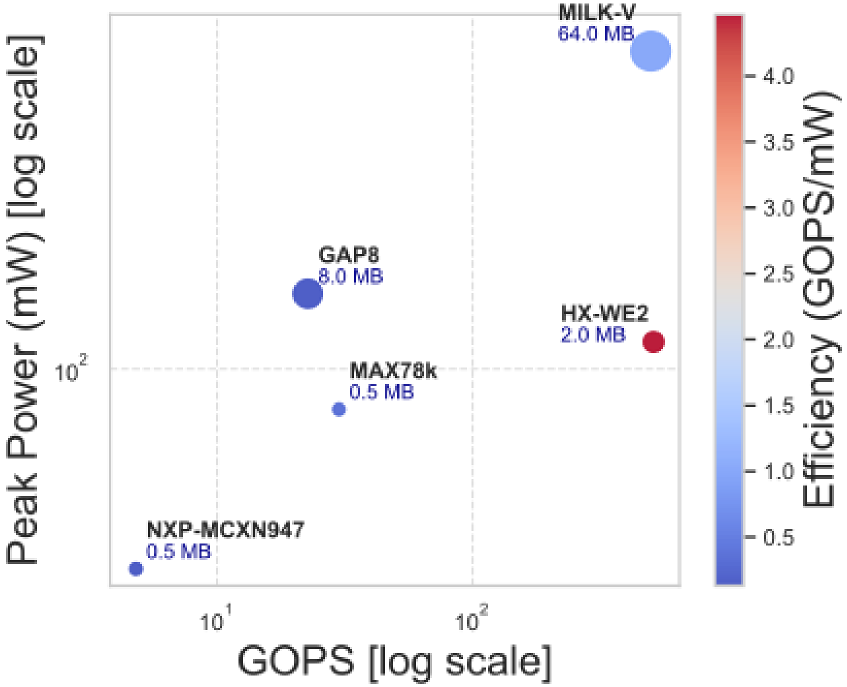


$\mu$NPUs used in our recent benchmark [1], and how they compare in terms of GOPS, peak power draw, and theoretical efficiency (GOPS/mW).


**[1] Millar et al., 2025** – Benchmarking Ultra-Low-Power μNPUs

# Challenges

- User-friendly bigraph building tools with
  - LiDAR-based indoor mapping (e.g, Apple's RoomPlan API)
  - RSSI fingerprinting
  - ML-based floor-plan annotations
- Distributed management and orchestration of logic
  - (i.e., what each agent knows, what it is allowed to act on, and how its coordination is scoped).
- Dealing with social/spatial ambiguity
  - e.g., how user's privacy settings apply when groups overlap / in shared space

# Takeaways

- **Spatial coordination is a missing layer** in modern networks.
- **Bigraphs** offer a **powerful model for connectivity**.
- Spatially-scoped agentic control improves
    - **latency**,
    - **privacy**,
    - **efficiency**,
    - and **reliability**
- Local reasoning & context escalation balance efficiency with flexibility.

# Questions?

jm4622@ic.ac.uk

**Benchmarking Ultra-Low-Power µNPUs**
- preprint, to appear at MobiCom 2025

**Acknowledgements:**

- Ryan Gibb, Dr. Sadiq Jaffer & Prof. Anil Madhavapeddy
- Adan Abu Naaj, Yushan Huang & Prof. Hamed Haddadi