Quantitative Verification and Strategy Synthesis for BDI Agents

Blair Archibald, Muffy Calder,
Michele Sevegnani, Mengwei Xu

NFM 2023, Houston

Autonomous agents

An entity

- which **perceives** its situated environment
- which **deliberates** accordingly
- which **takes actions** autonomously

in order to achieve its design objectives

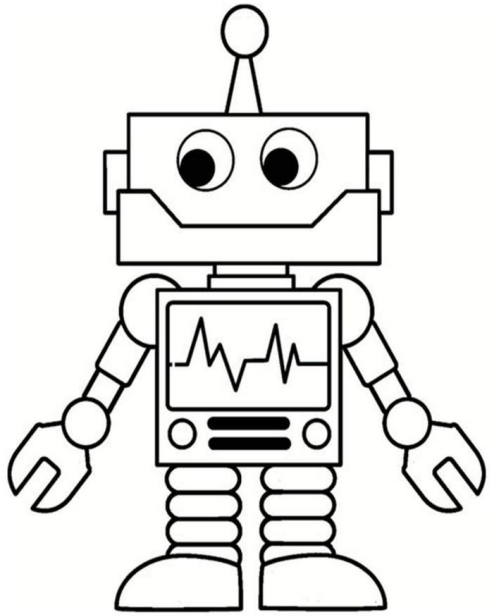

Michael Wooldridge. **An Introduction to Multiagent Systems**. John Wiley and Sons, February 2002

Autonomous agents

An entity

- which **perceives** its situated environment
- which **deliberates** accordingly
- which **takes actions** autonomously

in order to achieve its design objectives

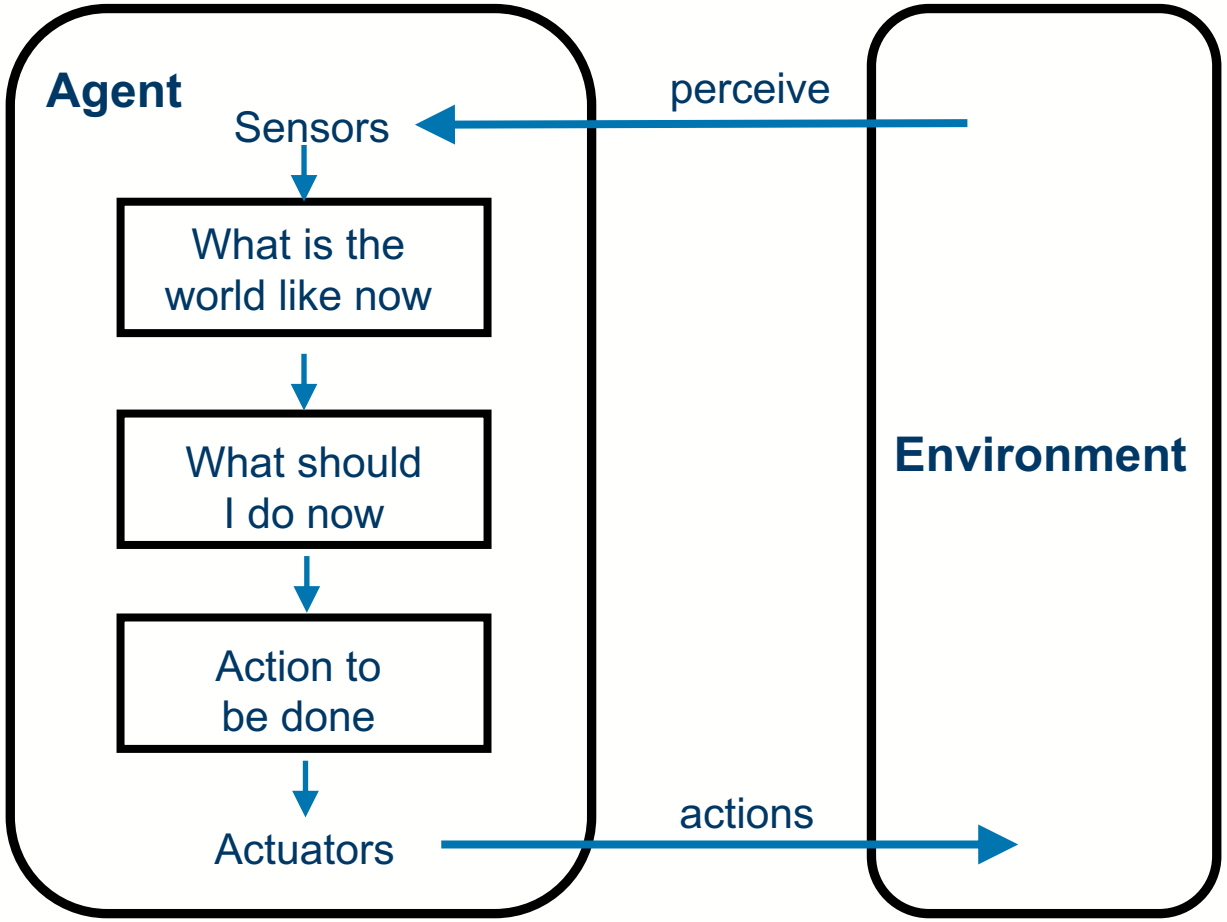

University of Glasgow

Belief-Desire-Intention (BDI) agents

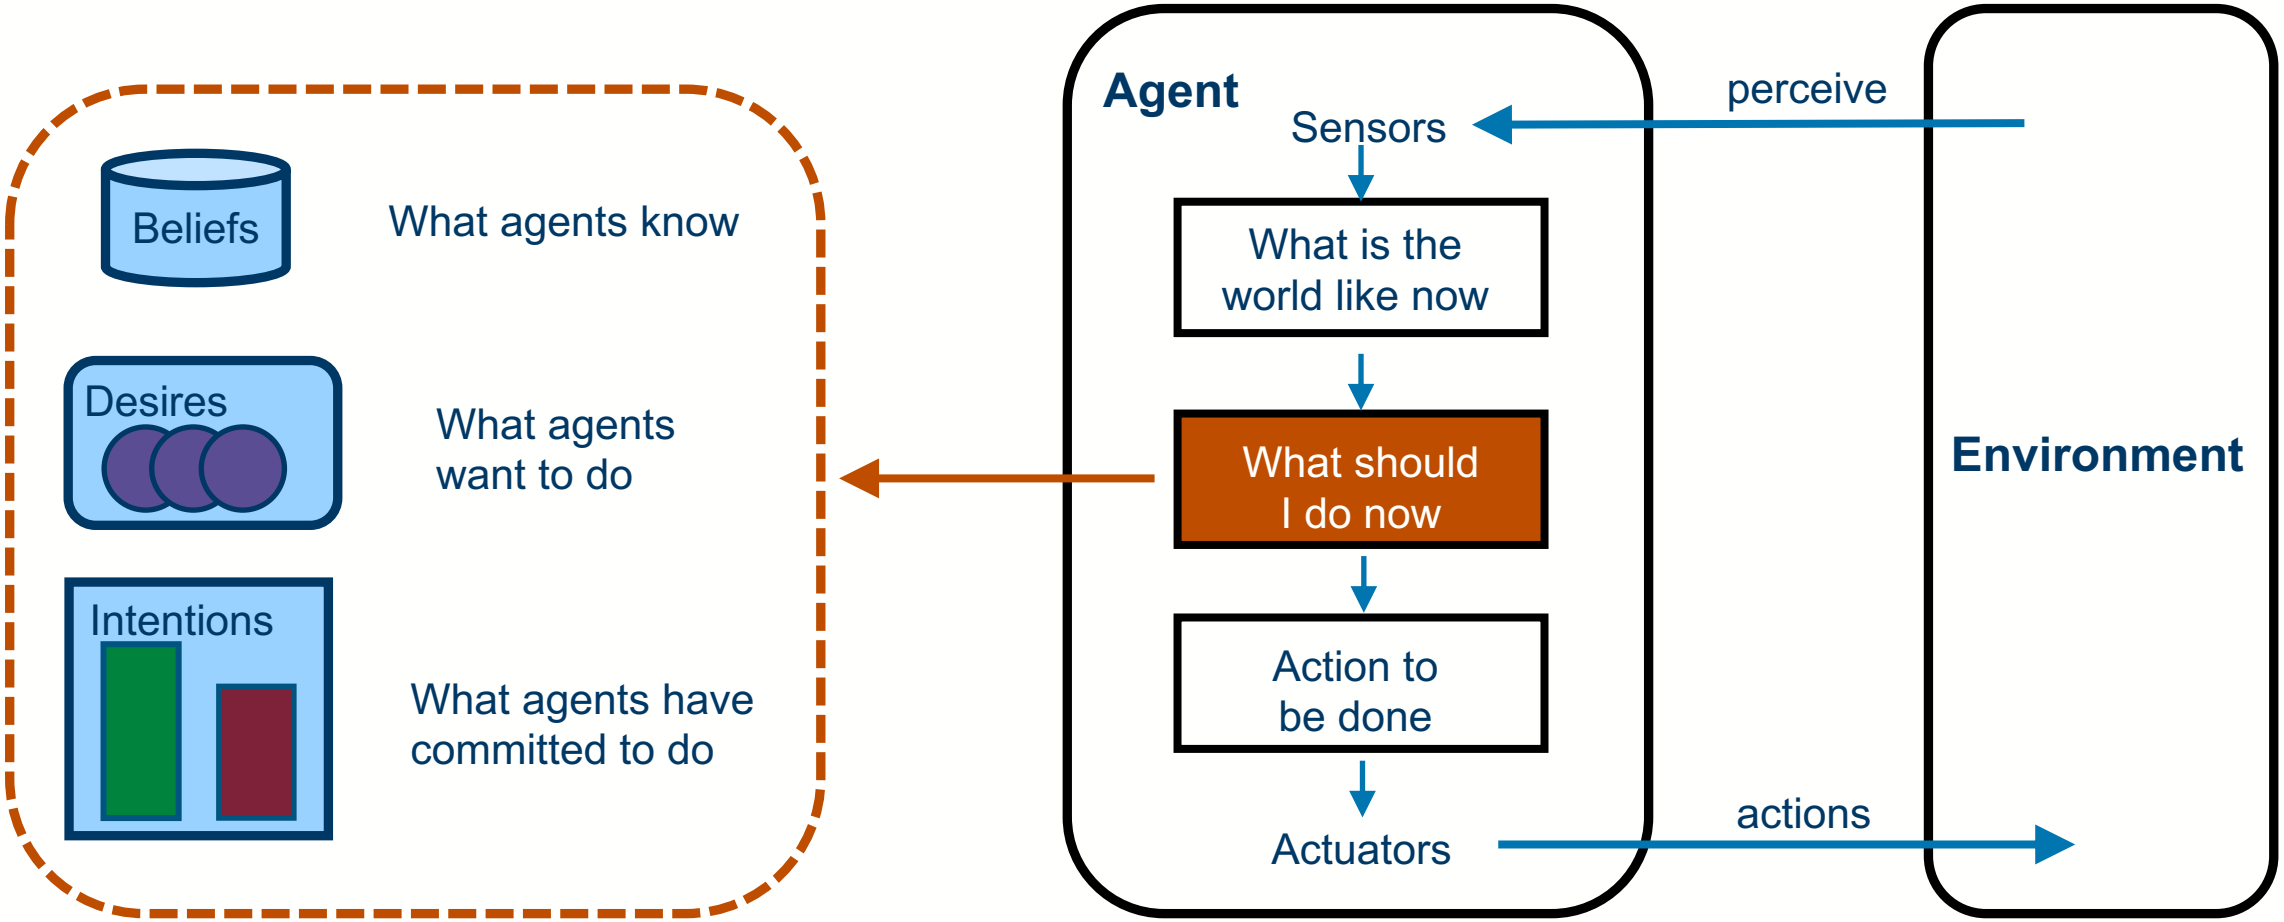

University of Glasgow

Belief-Desire-Intention (BDI) agents

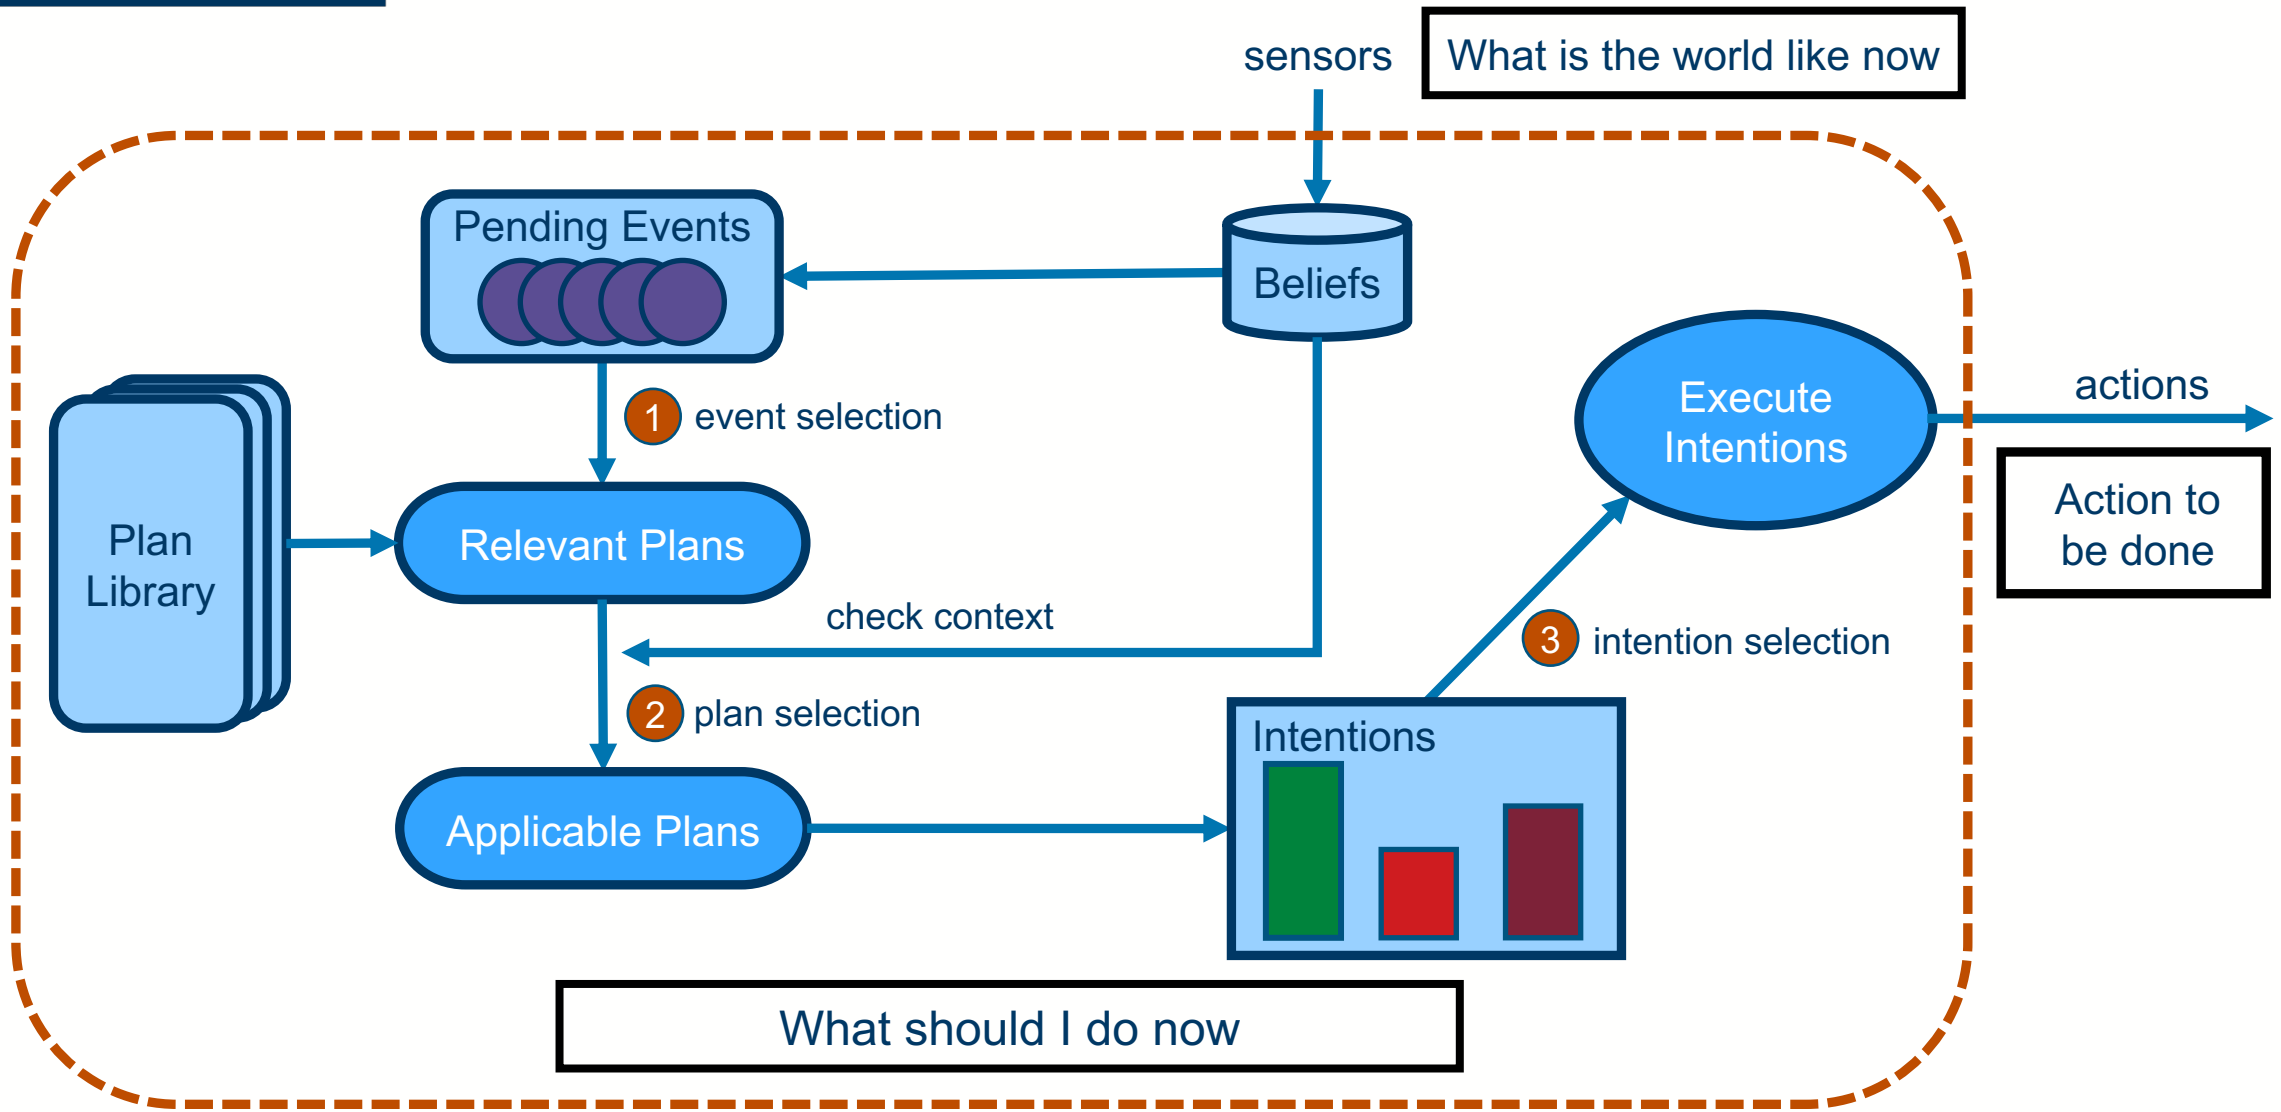

University of Glasgow

Motivation: building trustworthy agents

Writing correct BDI programs is not always easy

- Plans can include complex constructs like **declarative goals**, **failure recovery**, and **interleaved concurrency**

We need a model that allows us to

- **Verify** the probability an agent successfully completes a mission under environmental **uncertainty**
- **Synthesise** optimal strategies for internal agent decision making

University of Glasgow

Our approach: identifying non-determinism

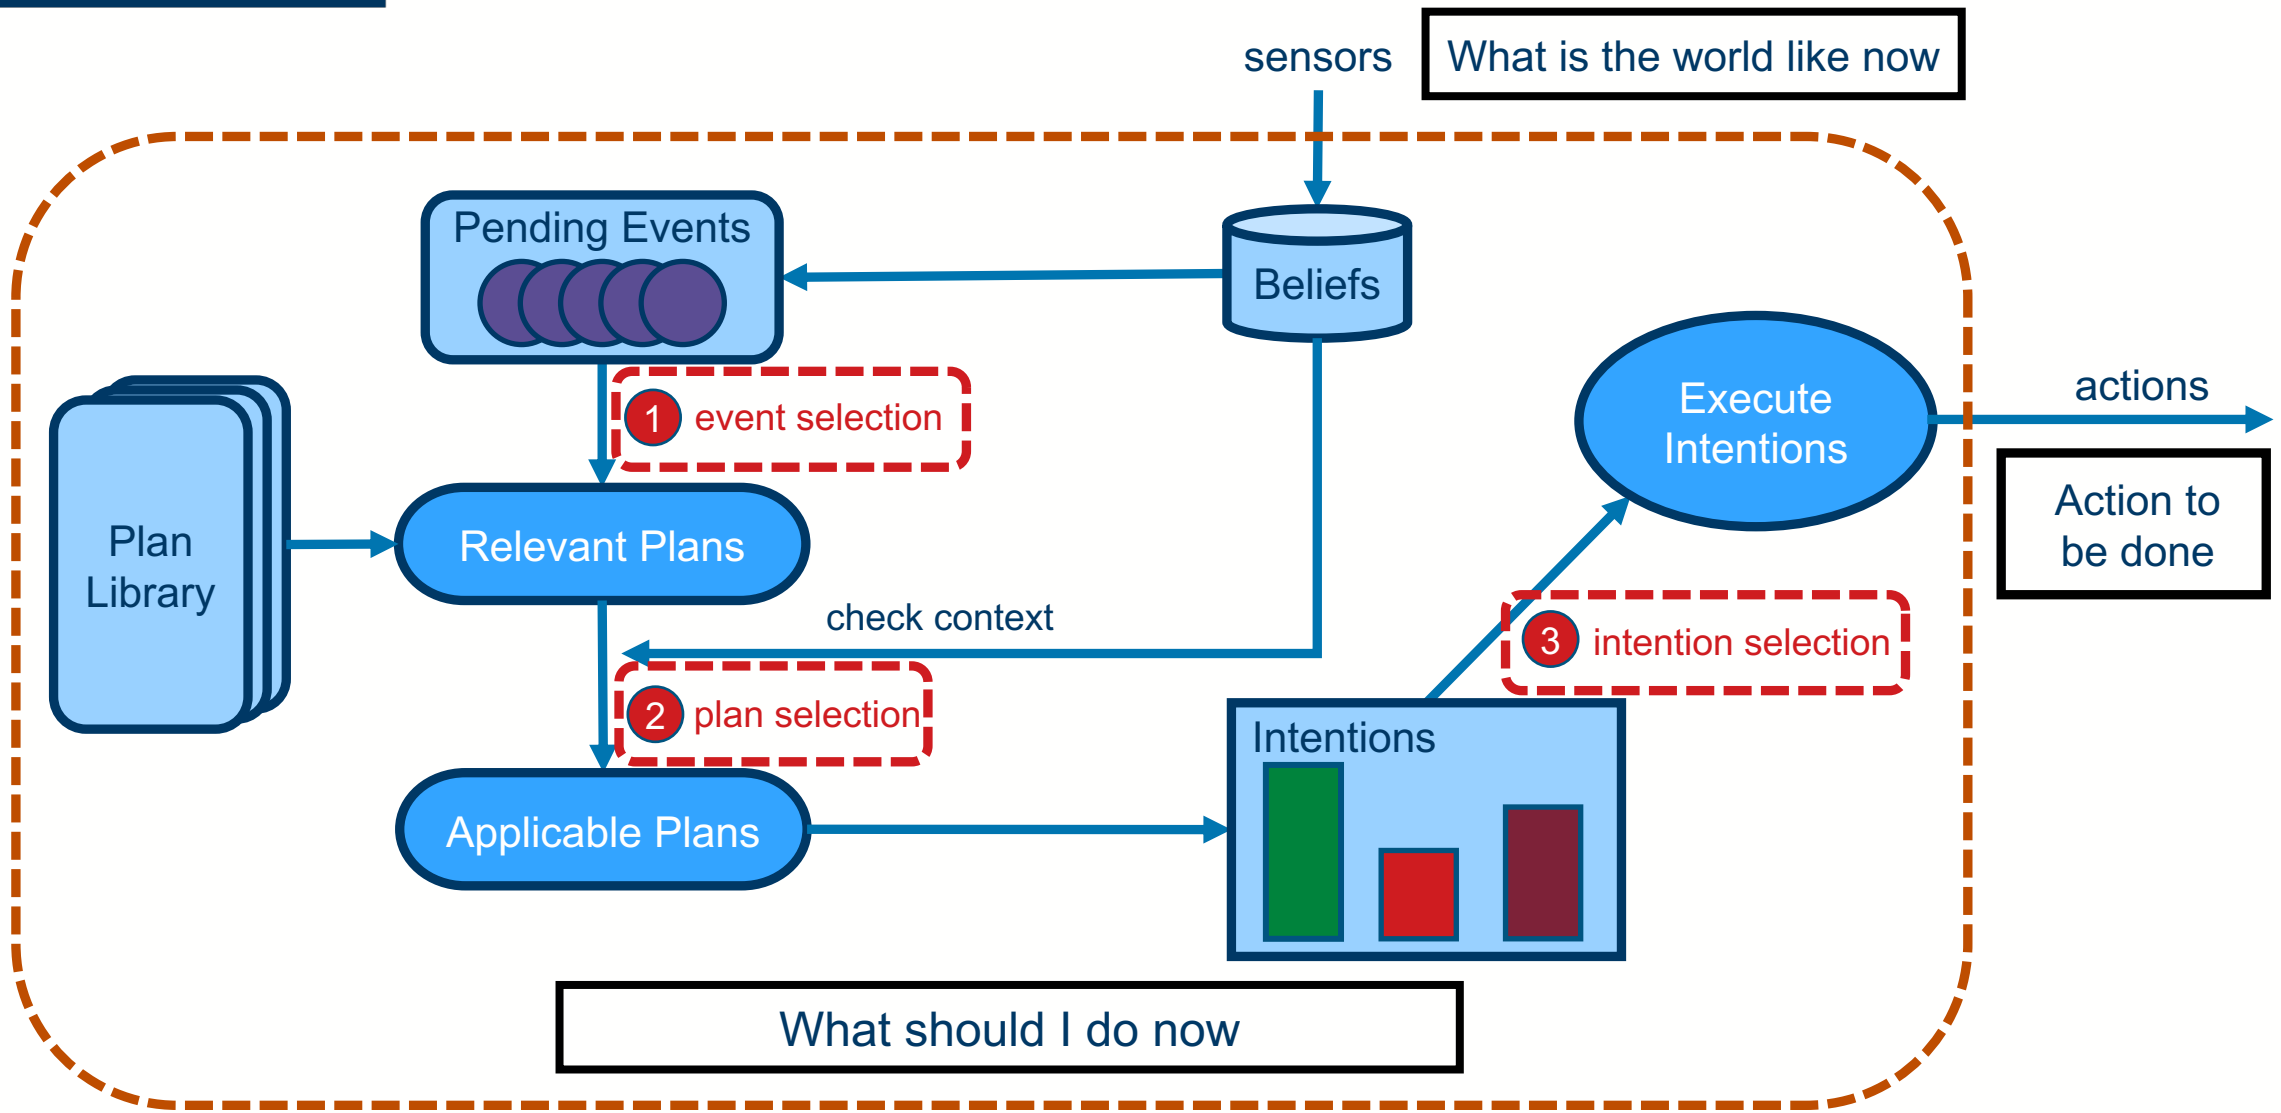

University of Glasgow

Our approach: introducing probabilistic outcomes

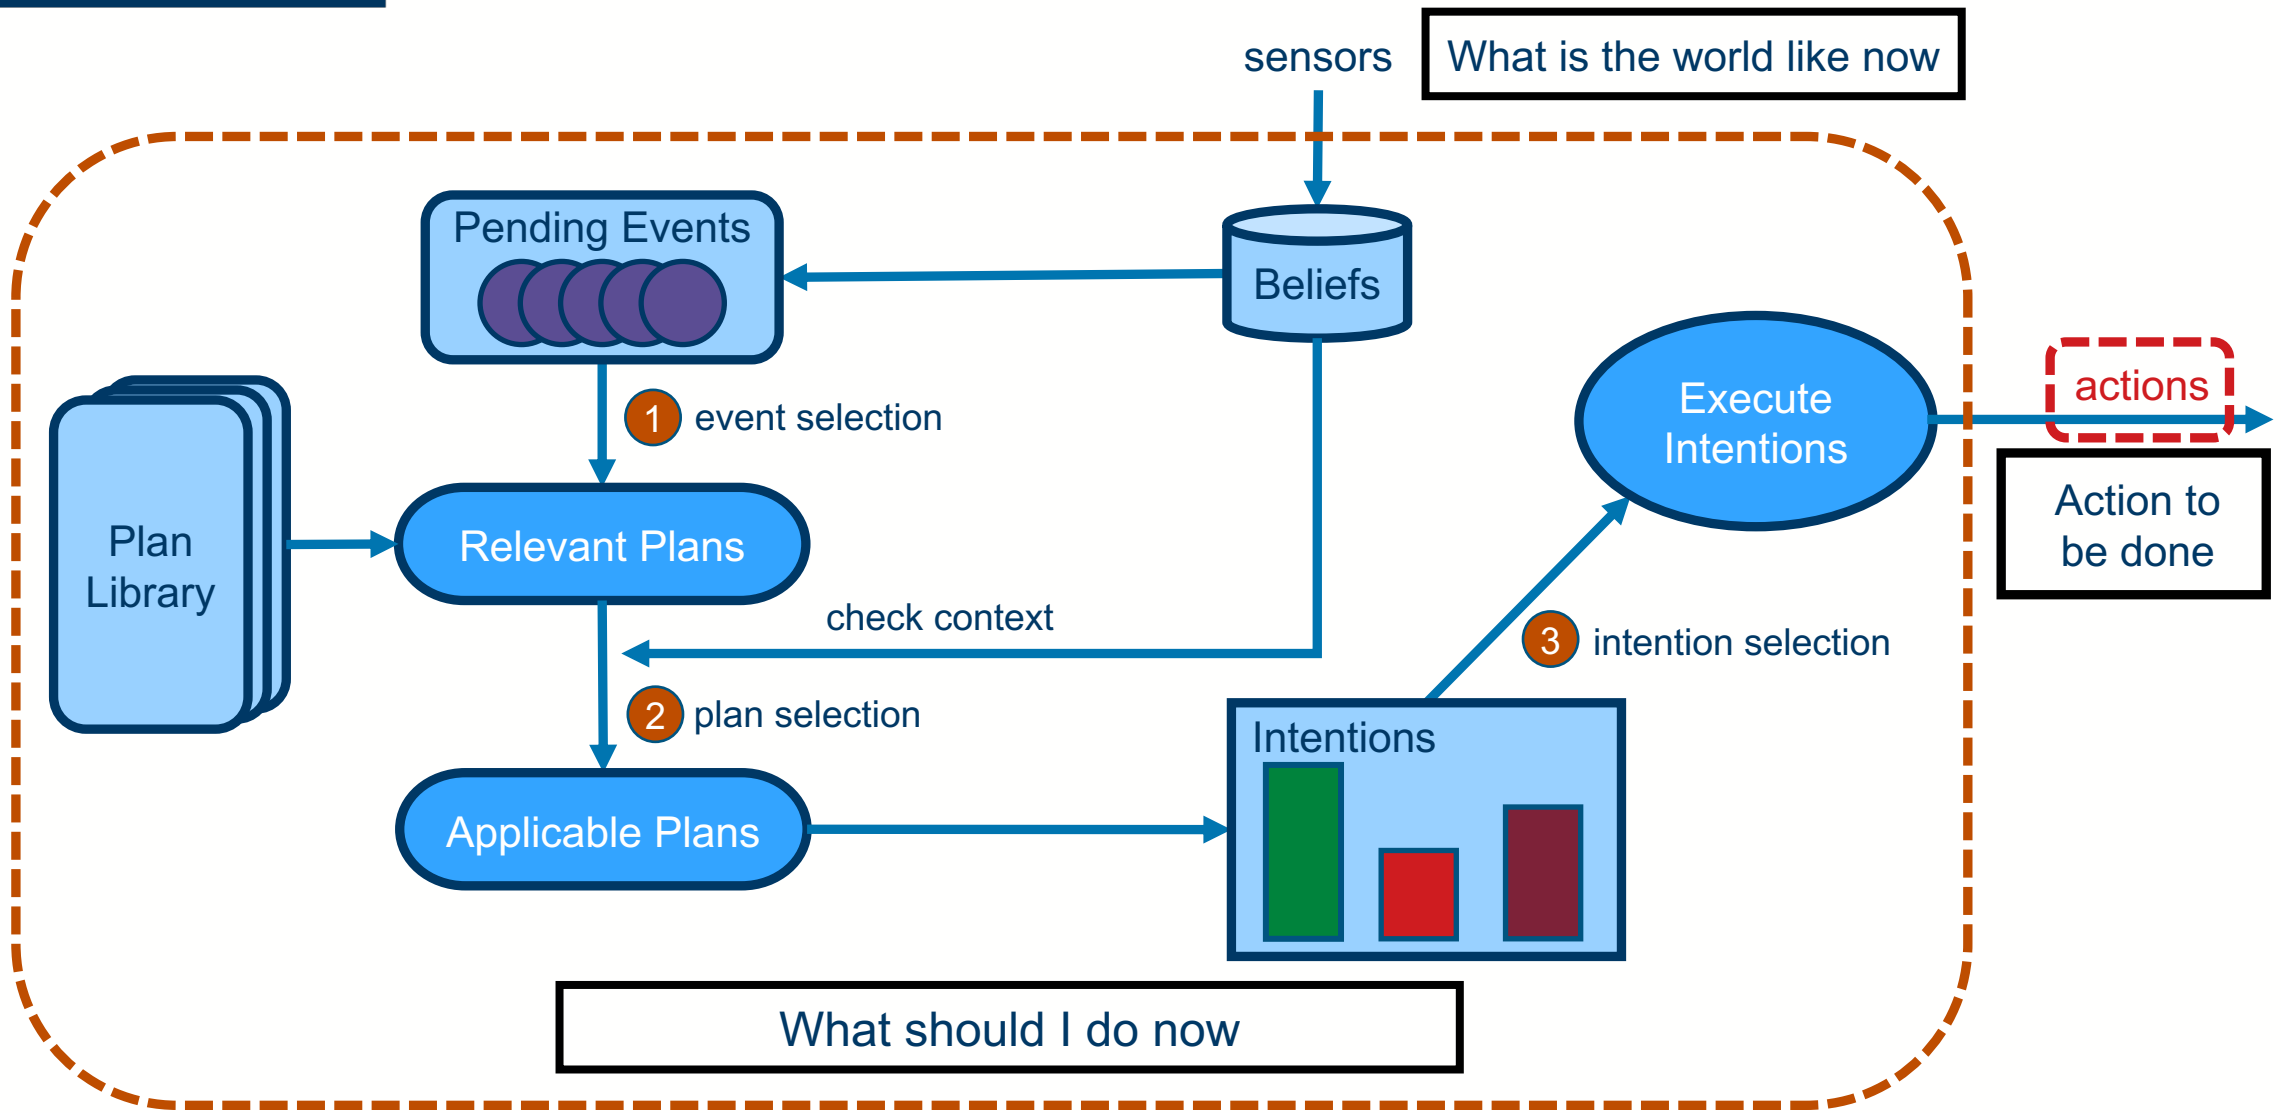

University of Glasgow

Our approach: encoding and verification pipeline

1. An MDP semantics for BDI to support non-deterministic selections and probabilistic action outcomes
2. An encoding of BDI agents into **bigraphs**

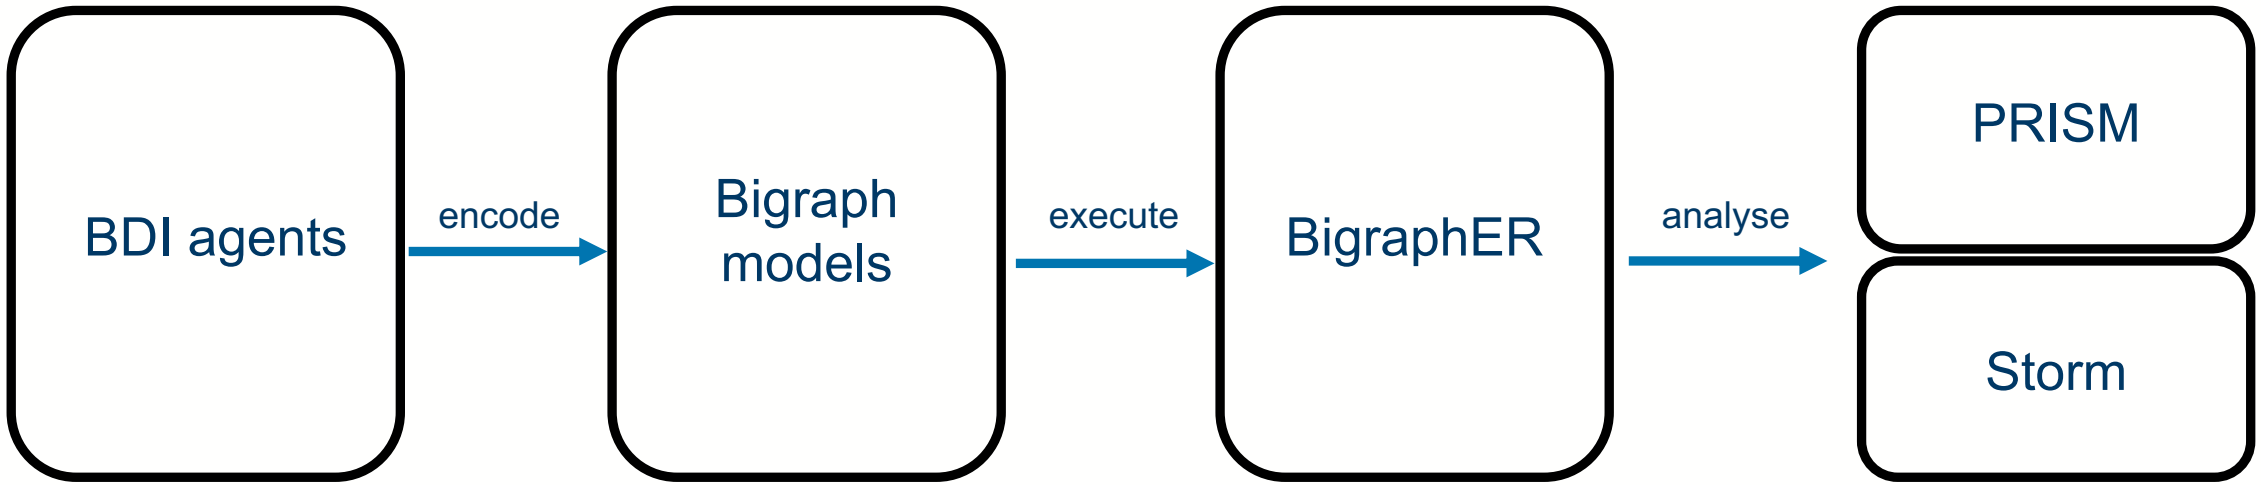

University of Glasgow

A primer on Bigraphs

1. Bigraph: superimposition of a place graph and a link graph
2. Place graph: DAG - topological space - no distances - containment relation
3. Link graph: Hypergraph - relationships between sets of entities (e.g. communication capabilities)

Robin Milner

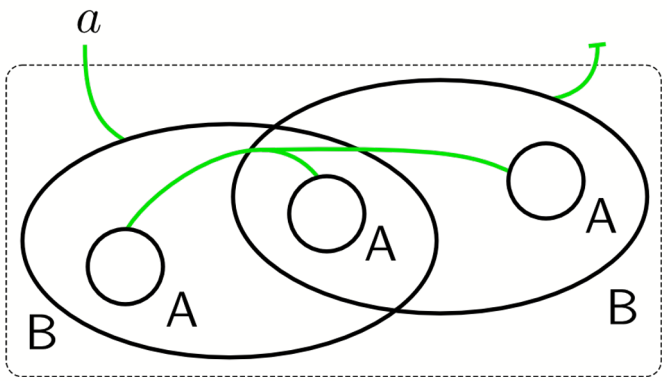

Bigraph

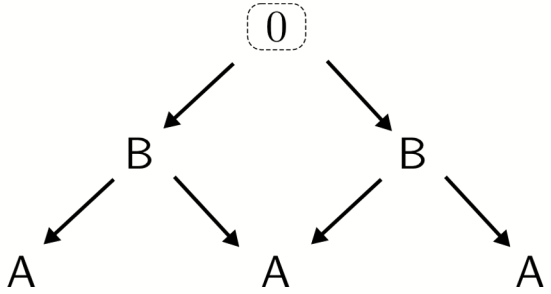

Place graph

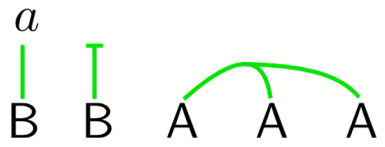

Link graph

University of Glasgow

A primer on Bigraphs (cont.)

- A set of reaction rules specify the dynamics of the system
- How to apply a rule to a bigraph (rewriting):
  1. Identify occurrences of the lhs in the bigraph
  2. Substitute each of them with the rhs

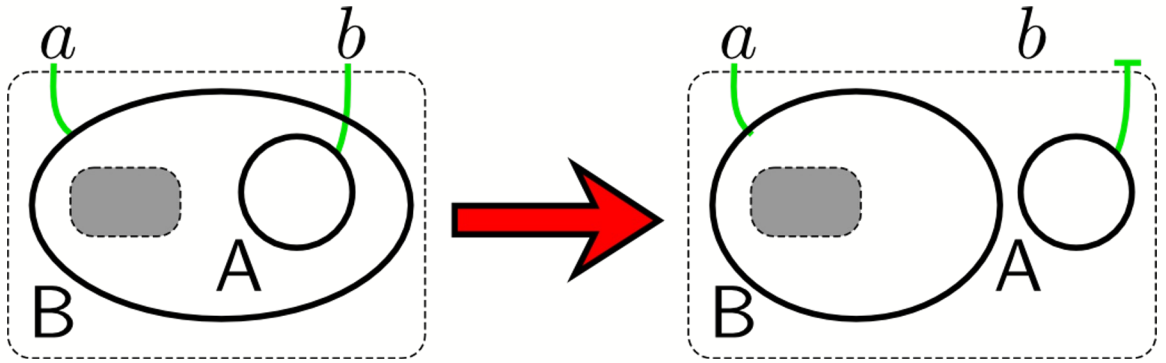

A primer on Bigraphs (cont.)

- A **Bigraphical Reactive System** consists of an initial bigraph and a set of reaction rules
- By performing all the rewriting steps we find all the reachable configurations of the system
- This can be done automatically using BigraphER

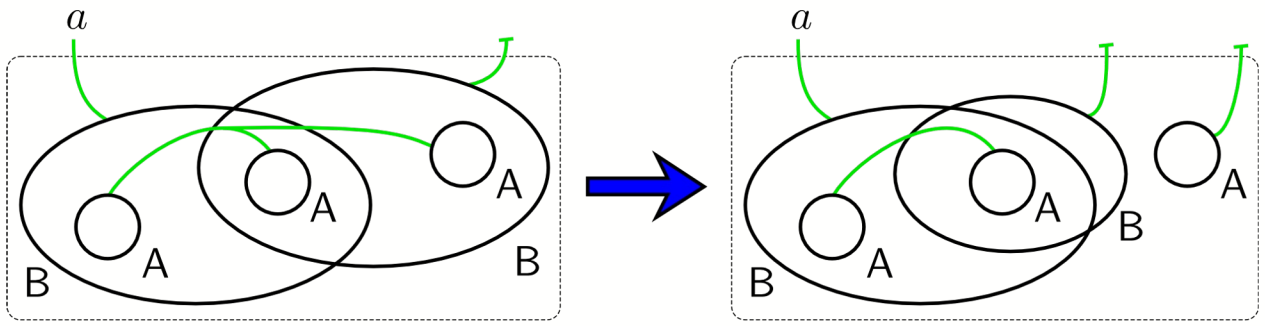

Transition

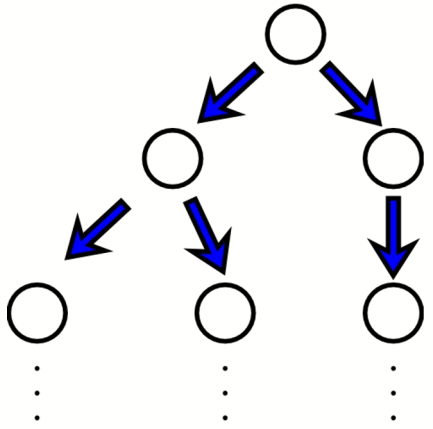

Transition system

Some technical details on our encoding

- We extend the semantics of the **CAN** language for BDI agents
- Example for probabilistic action outcomes

original BDI semantics

probabilistic BDI semantics

$$\frac{act{:}\,\varphi \leftarrow \langle \phi^{-}, \phi^{+} \rangle \quad \mathcal{B} \vDash \varphi}{\langle \mathcal{B}, act \rangle \rightarrow \langle (\mathcal{B} \backslash \phi^{-} \cup \phi^{+}), nil \rangle} \qquad act$$

$$\frac{act{:}\,\varphi \leftarrow \mu \quad \mu(\phi_i^{-}, \phi_i^{+}) = p_i \quad \mathcal{B} \vDash \varphi}{\langle \mathcal{B}, act \rangle \rightarrow_{p_i} \langle (\mathcal{B} \backslash \phi_i^{-} \cup \phi_i^{+}), nil \rangle} \qquad act^{p}$$

$$\mu = [(\phi_1^{-}, \phi_1^{+}) \mapsto p_1, \cdots, (\phi_n^{-}, \phi_n^{+}) \mapsto p_n] \qquad \sum_{i=1}^{n} p_i = 1$$

University of Glasgow

Some technical details on our encoding (cont.)

- Then we encode it as a **set** of reaction rules
- Example

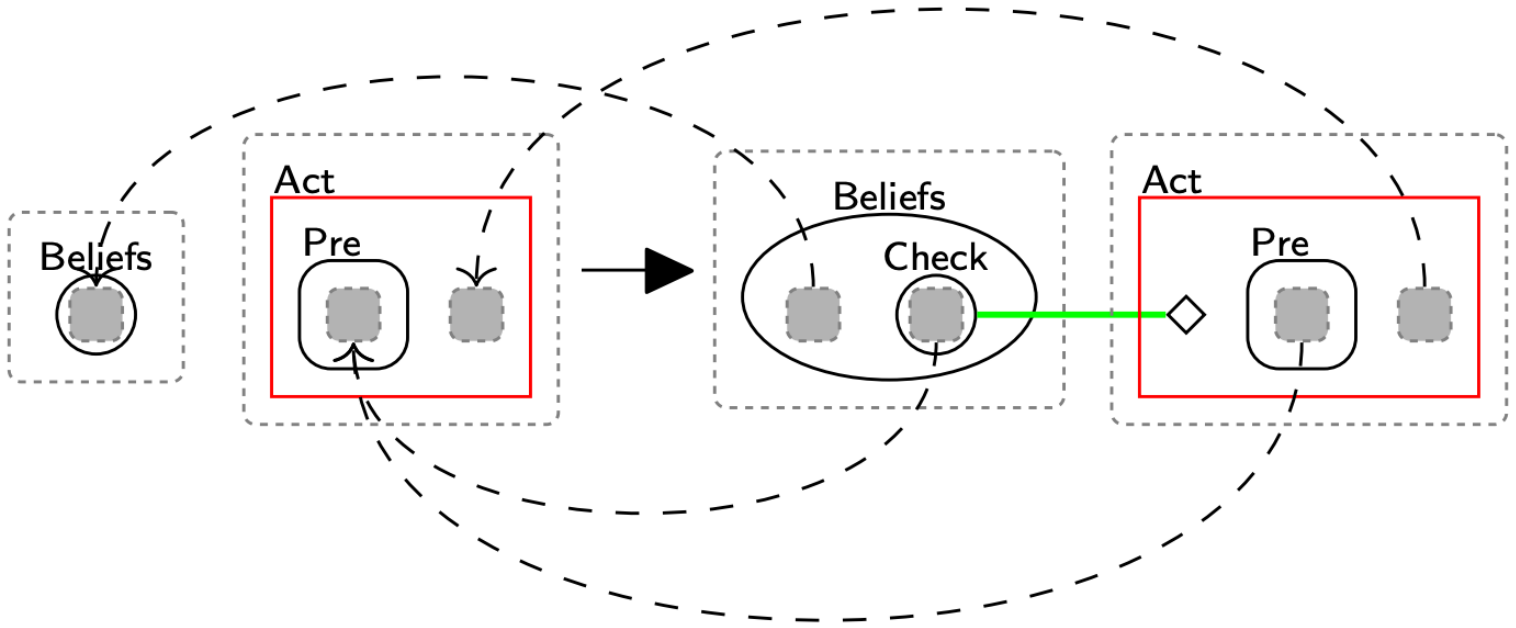

University of Glasgow

Example

- A robotic production line for packaging items with two types of wrapping
  - Cheap wrapping might break
  - Expensive wrapping never breaks
- Items decay when temperature rises over time if not wrapped

```
1  // Initial belief bases
2  deadline_1 = 10, deadline_2 = 14
3  // External events
4  product1, product2
5  // Plan library
6  product1 : true <- goal(success1,process_product1,failure1).
7  process_product1 : deadline_1 ≥ 3 <- wrap_standard1; move_product_standard1.
8  process_product1 : deadline_1 ≥ 0 <- wrap_premium1; move_product_premium1.
9  product2 : true <- goal(success2,process_product2,failure2).
10 process_product2 : deadline_2 ≥ 3 <- wrap_standard2; move_product_standard2.
11 process_product2 : deadline_1 ≥ 3 <- wrap_premium2; move_product_premium2.
```

University of Glasgow

Example: analysis and synthesis

- The max/min probability of both products being processed successfully over all possible adversaries
- Optimal strategy synthesis
  1. Wrap more urgent products first until they are packed
  2. Then switch to wrap the other products.
  3. Only after both are wrapped the robot moves them to storage
- Multi-objective analysis: obtaining high success rate while keeping the overall bag cost

University of Glasgow

Conclusions

- We have extended the CAN language BDI semantics to support non-determinism and probabilistic action outcomes
- Our extension is implemented in bigraphs, is executable with BigraphER, and verifiable with PRISM/Storm

- Future work
  - Beliefs are not probabilistic here. POMDPS? Bigraphs does not support them yet ☹
  - Runtime planning – does it scale?
  - Address some of the limitations of BDIs: actions affecting the agent instead of the environment

Thank you

Acknowledgments

- EPSRC through PETRAS: UK national centre of excellence for IoT Systems cybersecurity
- Amazon Research Award: Automated Reasoning

@michele_seve