Adding Preferences and Moral Values in an Agent-Based Simulation Framework for High-Performance Computing

David Marin Gutierrez[3], Javier Vázquez-Salceda[1](✉), Sergio Alvarez-Napagao[1,2], and Dmitry Gnatyshak[2]

[1] Universitat Politècnica de Catalunya-BarcelonaTECH, Edifici Omega, C/ Jordi Girona 1–3, E-08034 Barcelona, Spain
{jvazquez,salvarez}@cs.upc.edu

[2] Barcelona Supercomputer Center (BSC), C/ Jordi Girona 1–3, E-08034 Barcelona, Spain
dmitry.gnatyshak@bsc.es

[3] PwC Spain, Av. Diagonal, 640, E-08017 Barcelona, Spain
maringutierrezdavid@gmail.com

Abstract. Agent-Based Simulation is a suitable approach used now-a-days to simulate and analyze complex societal environments and scenarios. Current Agent-Based Simulation frameworks either scale quite well in computation but implement very simple reasoning mechanisms, or employ complex reasoning systems at the expense of scalability. In this paper we present our work to extend an agent-based HPC platform, enabling goal-driven agents with HTN planning capabilities to scale and run parallelly. Our extension includes preferences over their objectives, preferences over their plans, actions, and moral values. We show the expresiveness of the extended platform with a sample scenario.

Keywords: Agent-based simulation · Goals · Preferences · Values

1 Introduction

Agent-Based Simulation (ABS) is a computational approach for simulating the activities and interactions of autonomous agents (individual or collective entities such as organizations or groups) in order to better understand how a system behaves. Furthermore, they allow for the simulation of complex environments where perception, decision-making processes and actions carried out are dispersed among several stakeholders or agents. The purpose of ABS is therefore to obtain explanatory insight into the behavior of a group of agents which share a common environment. ABS can be applied to many fields such as biology, social sciences, ecology, economics, policy-making, sociology.

Many ABS frameworks have been built focusing on large simulations to be run in High-Performance Computing (HPC) platforms. In current HPC-Based

N. Fornara et al. (Eds.): COINE 2023, LNAI 14002, pp. 95–117, 2023.
https://doi.org/10.1007/978-3-031-49133-7_6

ABS approaches (such as Repast [27], NETLOGO [24], and MASON [16]) models may be elevated to and examined at genuinely large scales at the expense of having agents with limited reasoning capabilities and/or limited interaction among them, sometimes even reducing agents to mere rule-based or functional input-to-output transformers. There are some simulation scenarios in which these simplifications of the agents' reasoning suffices (e.g., if we want to simulate the general traffic flow in a big city, it may be enough to have agents with simple behaviors who simply react to changes in the environment arround them), but there are other scenarios (analysing complex human-human social relations or sociotechnical systems with intricate human-AI interactions) which require agents capable of more complex deliberative, goal-driven reasoning to simulate more complex behaviors including the effects of interactions with others.

There are many Multi-Agent frameworks in literature (such as Jadex [5], 2APL [9], BDI4Jade [17] or GOAL [12], to name a few) offering implementations of cognitive agents with more powerful practical reasoning capabilities, making them capable of exhibiting more complex social behaviours. Many of these implementations are inspired by the Beliefs-Desires-Intention (BDI) theory [7] and the BDI abstract architecture [20], modelling rational agents that use their beliefs about the current state of the world to choose which goal or goals to pursue, to then select actions or plans to fulfill the intended goal or goals. But this comes at the expense of having very limited scalability: the need to explore multiple potential instantiations of abstract goals ("which of all my goals are feasible/reachable now?") and plans ("which plans are applicable now?") in a given state of the system is computationally expensive. Many other approaches in literature offer different levels of reasoning and scalability ([1] and [21] provide an interesting comparative analysis on many of them, showing the reasoning level vs. scalability trade-off).

There is a need for new ABS platforms that could support big populations of goal-driven agents. Those ABS have the potential of being very useful in the creation and animation of richer social simulations to analyze the social relationships between agents by means of computational models of policies, norms, moral values and social conventions. By having agents whose reasoning and behaviour is influenced by these computational social models we can analise how and when the agents adhere to the norms and moral values, how they affect and limit their actions, and how they may change over time as the agents interact with each other and their environment. A first step in this direction was presented by Gnatyshak et al. in [11]: a custom Python-based BDI-agent simulation framework capable of both hosting agents imbued with more powerful practical reasoning capacity *and* running simulations with large numbers of these agents. Scalability is tackled in this framework by parallelizing via PyCOMPSs [23] the reasoning cycle of goal-oriented agents, allowing them to run concurrently whenever possible.

In this paper we address the issue of further enhancing Gnatyshak et al.'s framework by giving agents the capability to deal with preferences over their objectives, preferences over the actions they take in order to accomplish those objectives and (moral) values, as the next step towards a powerful agent-based micro-simulation framework to analyse the impact of social values, norms and conventions in large populations. In this work we also aim to explore how *far* we can go without using numbers in our preference mechanisms. Generally, humans do not reason using hard numbers (e.g, "today I prefer to go to the beach with a weight of 86, but to go to the cinema with a weight of 91; therefore, I will go to the cinema") but in qualitative terms (e.g., "today it is raining; I would rather go to the cinema than to the beach; therefore, I will go to the cinema") However, all state-of-the-art approaches we have analysed [6,8,19,25,26] end up adding hard numbers and/or ad-hoc numerical formulae to their selection strategy. So we aim to explore how *not* using numbers limits the expressiveness of our system, how severe this limitation is, and draw some conclusions as to whether it is acceptable to use numbers to attain a desirable level of complex reasoning.

This paper is structured as follows: in Sect. 2 we briefly describe the previous works we used as reference; in Sect. 3 we describe the conceptual model and how we added goals, preferences over goals, preferences over plans and actions, and support for the expression of moral values; in Sect. 4 we show how our additions to the model work in a sample scenario; and in Sect. 5 we conclude by discussing some limitations of the current model and extensions to be explored as future work.

2 Related Work

Our model of goals for Agent-Based Simulations in HPC has been inspired by two agent frameworks with working implementations: GOAP and BDI4JADE.

GOAP [18] is the AI created for the enemies of the video game F.E.A.R, mainly formalized by Jeff Orkin. It is relevant for our work as it provides goal-oriented agents in a multiagent gaming platform with strong scalability requirements. In GOAP, goals are represented by specifying a **desired state of the world** that agents strive to achieve. This desired state is described using the same structure used for the current state of the world, an agent's beliefs, actions' effects, etc. Agents can have many independent goals, but they can only pursue one at the same time. In order to plan, an agent must have a set of available actions, a set of beliefs about the world and sensors to periodically update those beliefs, and a set of goals. Each goal has a current priority, and the agent will choose to plan for the goal with the highest current priority. GOAP uses numeric priorities (i.e., a quantitative relation rather than qualitative). A* is used to plan with a heuristic minimizing the weighted number of actions used to reach the desired state., i.e., minimize the sum of costs of the actions in the plan. We borrow such goals defined as desired world states (see Sect. 3.1).

However, we aim to provide our agents with a more expressive goal model where agents may have a great number of declared goals, but only a few of them are intended to be achieved in a given point in time. Oliveira de Nunes's BDI4JADE [17] platform provides a BDI layer on top of JADE [2]. It uses the same structure as Orkin's GOAP to represent goals (desired state of the world). It supports the declaration of different types of goals: *belief goals* (goals that deal with states of the world described by boolean variables), *beliefset value goals* (same as before, but variables are continuous or have more than two possible values), *composite goals* (used to represent goals composed of subgoals which have to be achieved sequentially or in parallel), etc. It also differentiates between desires (non-committed goals) and intentions (committed goals). Plans are an ordered set of actions and are executed to achieve a specific goal. In BDI4JADE agents do not have a set of actions that they can use to build plans, but rather, they have a library of plans that the agents can choose from. Each plan in the library has some applicability conditions (equivalent to actions' preconditions) that are used in the plan selection process. We get inspiration from BDI4JADE on its plan selection strategy.

Our main inspiration for the modelling of preferences over goals comes from CP-nets [4]. Although our actual implementation is definitely not an implementation of a CP-net, the main inspirations we have drawn from them is to establish one default and many conditional preorder relationships over goals, and building a graph to both visualize them and interpret them. We also analysed Cranefield et al.'s approach in [8] to model values (to adapt it to model preferences over goals), but upon closer inspection, we decided not to follow this approach since it uses numerical values and in this work we aim for a more qualitative approach.

In the case of preferences over plans, we drew a great deal of inspiration from Visser et al.'s work in [25]. It introduces the concepts of goals' properties, which we use extensively in our modeling of priorities over plans. We also make use of their mechanism for property propagation in our implementation. We should note that our implementation is simpler than theirs. For instance, the paper defines both properties of goals (discrete values that a property can take) and resources of goals (numerical values and intervals that represent how much of a resource -e.g., money, food- is being consumed by a goal or a sub-goal), but we chose to simplify the approach and add only discrete properties, as we want to explore a qualitative, scalar-free preference approach.

3 Conceptual Model

A **multi-agent system** for Agent-Based Simulation in HPC $\mathcal{M}$ is defined as the tuple $\mathcal{M} = \{E, \mathcal{A}^+, \mathcal{C}\}$ where:

- E is the model of a simulated **environment**, in which the agents reside, that they can perceive, gather information from, and act on;
- $\mathcal{A}^+$ is a non-empty **set of agents**;
- $\mathcal{C}$ is a **controller** (a structure that maintains the multiagents' environment model, regulates how agents access and act upon it, and handles agent-to-agent communication within the HPC execution environment), which is defined as the tuple $\mathcal{C} = \{\mathcal{I}, inAcs\}$ where $\mathcal{I}$ is the inbox for all the agents' outgoing messages (supporting agent communication), and $inAcs$ is the set

of all the actions to be exercised on the environment (regulating how agents access and act upon it).

An **agent** is defined as $\mathcal{A}_i = \{ID, msgQs, outAcs, Bh, \mathbb{B}, \mathbb{G}, g_c, \mathcal{P}_c, \mathcal{MP}, \mathbb{P}_g, \mathbb{P}_p\}$ where:

- $ID = \{AgID, AgDesc\}$ is $\mathcal{A}_i$'s identity data:
  - $AgID$ is the unique identifier of $\mathcal{A}_i$
  - $AgDesc$ is an arbitrary description of $\mathcal{A}_i$
- $msgQs = \{\mathcal{I}, \mathcal{O}\}$ is the set of $\mathcal{A}_i$'s message queues
  - $\mathcal{I} = \{\dots, msg_i, \dots\}$ is the Inbox, the set of messages sent *to* $\mathcal{A}_i$
  - $\mathcal{O} = \{\dots, msg_i, \dots\}$ is the Outbox, the set of messages sent *by* $\mathcal{A}_i$
  - $msg_i = \{AgID_s, AgID_r, performative, content, priority\}$ is a **message** sent from agent with $ID = AgID_s$ to the agent with $ID = AgID_r$, with the corresponding (FIPA-like) performative type, content, and priority.
- $outAcs$ is the set of **external actions** to be executed on the environment. It is composed of tuples of the form: $\{senderID, a^e\}$, where ID is the sender's ID, and a^e is the action that is being sent.
- $Bh = \{\mathbb{RG}, \mathbb{P}\}$ is $\mathcal{A}_i$'s **role behavior**, which is composed by:
  - $\mathbb{RG}$ is the set of **role goals** associated with the Bh which $\mathcal{A}_i$ is enacting
  - $\mathbb{P}$ is the set of plans $\mathcal{P}$ associated with the Bh
- $\mathbb{B}$ is the set of $\mathcal{A}_i$'s **beliefs**. It uses the same world state structure as E
- $\mathbb{G}$ is the set of $\mathcal{A}_i$'s **own goals** (see Sect. 3.1).
- $g_c \in (\mathbb{G} \cup \mathbb{RG})$ is the current **committed goal** (see Sect. 3.1).
- $\mathcal{P}_c = \{\dots, ab_i, \dots\}$ is $\mathcal{A}_i$'s current **plan**, which is an ordered set of action blocks. Each **action block** $ab_i = \{\dots, a_{ij}, \dots\}$ is an ordered set of actions (each a_{ij} is an action). There are three types of actions: **internal actions** (actions that are executed by the agent in order to change their beliefs), **external actions** (actions that are sent by the agent to the controller in order to be executed on the environment to alter it), **message actions** (actions that are used to generate messages intended to other agents)
- $\mathcal{MP}$ is the **metaplanner**, a library of plans for each goal (see Sect. 3.2).
- $\mathbb{P}_g$ is the set of **preferences over goals** (see Sect. 3.3).
- $\mathbb{P}_p$ is the set of **preferences over plans** (see Sect. 3.4).

Our conceptual model extends the one presented in [11]. Our extensions are described in the following sections.

3.1 Adding Goal Structure

We extend the conceptual model in [11] by providing a formal model for goals: *what* they are, *how* they are defined, and how they are *related with plans*. We have chosen to model goals as desired states of the world that agents strive to achieve. It is equivalent to the concept of **desires** in BDI Theory [7]. A goal is therefore defined by a collection of subsets of the variables that describe a state of the world (its **conditions**), and an assertion of desired value for each variable. These conditions are expressions such as 'cash==10' or 'speed>=50' to mean that having exactly 10 units of cash and that maintaining a speed of 50 or above are part of the desired state of the world, respectively. Each subset describes a conjunction of variables that describe a desired state of the world and, in order for a goal to be considered achieved, it is required that the goal condition evaluates as *true* in the eyes of the agent (that is, according to its **beliefs**).

We formally define the structure of a **set of goals** $\mathbb{G}$ as an unordered set of the form $\mathbb{G} = \{g_1, g_2, \ldots, g_n\}$ where each g_i is an individual goal among the many goals an agent has. A **goal** is defined as $g_i = \{name, descr, \mathbb{C}, status\}$ where *name* is a unique identifier of the goal, *descr* is an optional text describing the goal, $\mathbb{C}$ is the set of conditions over the state of the world for the goal to be considered achieved, and *status* is a boolean that is *True* if and only if the conditions $\mathbb{C}$ are satisfied according to the agent's current beliefs $\mathbb{B}$.

A **set of conditions over the state of the world** is defined as unordered collections of assertions over the state of the world (the *environment*) of the form $\mathbb{C} = \{a_1, a_2, \ldots, a_n\}$ where $a_i = \{n_1 \star v_1, n_2 \star v_2, \ldots, n_m \star v_m\}$ is a conjunction of statements over the values of variables of the agent's beliefs, defined by n_i, which is the *unique* name of a variable of the agent's beliefs; $\star$, which is a binary operator ($\{=, \neq, >, \geq, <, \leq\}$); and v_i, which is the value of interest that is being asserted to n_i.

The agent possesses the capabilities to check whether or not an individual goal has been achieved according to its beliefs: $check\_goal(g_i, \mathbb{B})$ outputs *True* if, according to the agent's beliefs, the conditions of the goal have been met, and false otherwise. Our agents are allowed to have multiple goals (own goals $\mathbb{G}$ and role goals $\mathbb{RG}$), but are restricted to pursuing only one at a time. This *commitment* to a goal that is intended to be pursued (g_c in the agent tuple) is equivalent to the concept of **intention** in BDI. Agents have the capability to re-consider which goal they want to pursue, and may change the goal they are committed to even if they have not achieved it, depending on their current beliefs and the state of the world they perceive.

3.2 Adding a Library of Plans

We also extend [11] to enable specifying different plans for each goal, and to pick different plans for a committed goal with an element that will act as a library of plans, The implementation of the means-ends reasoner for the platform is a Hierarchical Task Network (HTN) planner [15]. A HTN is a tree composed of three types of nodes: (i) Primitive Tasks, (ii) Methods, and (iii) Compound Tasks. The root of the HTN is an abstract compound task (e.g., *order food*).

Figure 2 provides an example. Our agents have a library of predefined HTN plans that the agent can pick from, and these plans will be related to goals by means of the structure of the **metaplanner**, which is the $\mathcal{MP}$ element of the agent tuple. Formally, it can be viewed as $\mathcal{MP} : \mathbb{G} \longrightarrow \mathbb{P}^*$, a matching relationship from goals towards plans, where $\mathbb{P}$ is the set of plans $\mathcal{P}$ associated with goal g_i and $\mathbb{P}^*$ is used to indicate that it can output tuples of plans of arbitrary cardinality (meaning one specific goal may have, for instance, three plans associated to it, while a different goal might have five, or two). We need also to add applicability conditions to plans: $\mathcal{P} = \{\mathbb{C}, ab_1, \dots ab_n\}$, where $\mathbb{C}$ is the set of conditions over the state of the world (see Sect. 3.1) that determine a plan to be applicable, and each ab_i is an action block.

Other noteworthy aspects of the metaplanner are that it incorporates appropriate functions for plan selection. Therefore, it will not simply act as a library/collection of plans, but it will also perform part of the reasoning. This reasoning includes both checking which of the associated plans are available for application, as well as ordering them based on the preferences.[1] For the first functionality, the metaplanner features a $get\_available\_plans(g_i, \mathbb{B})$ function which, taking into account the current beliefs of the agent, it outputs a subset of the set of plans associated with the goal, containing only all plans that are applicable. For the second functionality, the metaplanner has a $pick\_plans(g_i, \mathbb{B}, prefs_{\mathcal{P}})$ function, where $prefs_{\mathcal{P}}$ are the agent's preferences over plans, that will pick the plan that is more adequate to the current situation according to the agent's preferences and beliefs, from among all the applicable plans.

3.3 Adding Preferences Over Goals

The next extension we introduce in the model are preferences over goals. As we explained in Sect. 2 we drew inspiration from CP-nets and conditional preference formulas but we simplified the approach in order to be able to work without scalars, that is, having a fully qualitative approach for the specification of preferences over goals.

To define preferences over a set of goals, the approach we have taken is to establish a strict partial order relation between them to indicate which goals must be pursued before trying to achieve other goals. These binary relations between goals are reflexive, transitive and assymetric. To model the context-dependent nature of preferences, we allow the declaration of conditional preferences, which are also a strict preorder relation over goals, but they only apply when their trigger conditions are met. A nice property of strict preorders is that they have always a unique direct acyclic graph (DAG) associated to them.

In order to encode **preferences over goals** in our agents, we have added the following element, $\mathbb{P}_g$ (which stands for "$\underline{\mathbb{P}}$references over $\underline{\text{g}}$oals") to the agent tuple. We define it as $\mathbb{P}_g = \{dGP, cGP_1, cGP_2, \dots, cGP_n\}$, where dGP are the *default* preferences over goals (they apply under 'normal' circumstances), and cGP_i are *conditional* preferences over goals (they have some trigger set of conditions $\mathbb{C}_i$ over the state of the world as defined in Sect. 3.1).

The dGP and each cGP_i are defined as a DAG that corresponds directly to a **strict partial order** relationship between goals, and the only difference between them is that the dGP is the one active by default (it does not need any conditions to be met), while the various cGP_i become active and replace dGP if some associated conditions are true.

[1] We describe how we model preferences over plans in Sect. 3.4.

Once all the strict preorder relations have been established, we deduce their associated DAGs. From those DAGs, we compute a valid topological ordering of each, and these orders are the ones in which goals will be pursued by the agents (by choosing the first non-achieved goal in the topological ordering), e.g.:

– We have one agent $\mathcal{A}$, which has the goals $\mathbb{G}_0 = \{g_0, g_1, g_2\}$. g_0 is a goal to tidy the agent's bedroom, g_1 is a goal to tidy the agent's kitchen, and g_2 is a goal to store clothes that are hanging out to dry in the open.
– If we denote "goal i must be achieved before goal j" as $g_i \rightarrow g_j$, the **default preferences** over goals of agent $\mathcal{A}$, are $\{g_0 \rightarrow g_2, g_1 \rightarrow g_2\}$, that is, before storing the clothes that are outside, $\mathcal{A}$, must have cleaned both his bedroom and his kitchen. Notice how both g_0 and g_1 must be accomplished before focusing on g_2, but there is no established order between g_0 and g_1, as it is a strict *partial* order. A valid topological ordering might be: g_0, g_1, g_2, but also g_1, g_0, g_2. By default, $\mathcal{A}$, will pursue his goals in either of those orders.
– The set of **conditional preferences** over goals of agent $\mathcal{A}$, is $\{g_2 \rightarrow g_0, g_2 \rightarrow g_1\}$ with the associated trigger conditions that the variable 'raining' must be $True$. If it is raining, the agent's top priority goal will be to collect the clothes (g_2), then cleaning their kitchen or bedroom, in no specific order. Therefore, the moment it starts to rain, $\mathcal{A}$, will switch to any of the topological orderings that can be given to this set (for instance, g_2, g_1, g_0).[2]

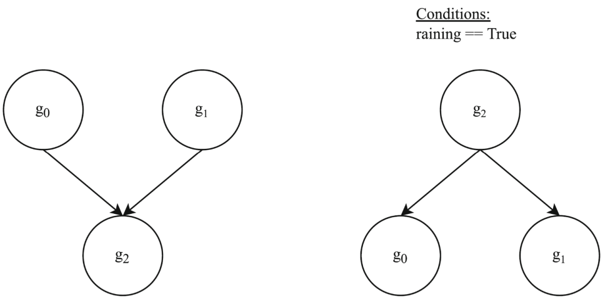

Fig. 1. Example of default and conditional preferences over goals

This example is depicted in Fig. 1. The left graph is the one deduced from the relations that defined the default preferences over goals, while the graph on the right-hand side is the one defined by the trigger condition ($raining = True$). A valid topological ordering of the left graph might be: g_0, g_1, g_2, but also g_1, g_0, g_2. By default, the agent will pursue his goals in either of those orders, but the moment it starts to rain, he will switch to any of the topological orderings that can be given to the right graph, for instance, g_2, g_1, g_0, but also g_2, g_0, g_1.

[2] In case of conflicts between preferences, the default behaviour is to choose by order of declaration in the HTN. This can be overriden by the designer. Refer to Sect. 5.

3.4 Adding Preferences Over Plans and Actions

By adding preferences over goals we provide agents with the capacity to choose *what* to pursue. But we also need to provide them with means to have preferences over *how* to achieve what they are pursuing. For example, when your goal is to eat, it is not the same to achieve that goal by eating a delicious pizza or to achieve it by eating a boring (but healthier), plain white rice, even if both actions achieve the goal all the same. We humans have preferences not only over *what* goals we want to achieve, but also over *how* we want to achieve them, and these preferences may be context-dependent. Some people might prefer to drive to their workplace, while some others would rather walk there. But the preference on walking may change in the case the weather is very cold or rainy, then prefering to commute to work by a combination of transportation modes. These examples provide us with further, key information: the preferences we have over how we achieve things are also context-dependent; we may wish achieve a specific goal by means of some actions under some circumstances, but under different circumstances we might prefer to achieve the same goal through different actions. Since the purpose of this work is to imbue agents with human-like social aspects for simulation purposes, we will need to take all these considerations into account when modeling preferences over plans and actions. In order to encode **preferences over plans and actions** in our agents, we have added element $\mathbb{P}_p$ (which stands for "$\underline{\mathbb{P}}$references over $\underline{\mathbf{p}}$lans") to the agent tuple. We define it as $\mathbb{P}_p = \{gP_1, gP_2, \ldots, gP_n\}$. We denote the preferences over plans for each goal g_i by $gP_i = \{dPP, cPP_1, cPP_2, \ldots, cPP_n\}$, where dPP are the *default* preferences over plans for goal g_i (under 'normal' circumstances), and cPP_i are *conditional* preferences over plans for goal g_i (they have some trigger set of conditions $\mathbb{C}_i$ over the state of the world).

A **property of a goal** is the name of a variable of interest that a goal has the capacity to alter. Said variable does not necessarily have to be the name of a variable in the set of beliefs of an agent. It is simply something noteworthy that achieving a goal has the capacity to give a specific set of values. For example, if a goal is to 'order dinner', some of the properties might be 'vegetarian' and 'cuisine', and their possibles values might be $\{True, False\}$ and {'French', 'Italian', 'Spanish', 'Turkish'}, respectively. In our model each goal, plan, subplan, and action may have a set of properties PS, of the form $PS = \{prop_1, prop_2, \ldots, prop_n\}$, and each property $prop_i$ is of the form $prop_i = \{v_1, v_2, \ldots, v_n\}$ where: $prop_i$ is the *unique* name/identifier of the property, and v_i is one of the possible values that the property can take. These values can be boolean, numeric, etc., depending on the nature of the property itself. The set of values that make up each property are used to indicate possible values the property can take. All properties can have the special $None$ value inside the set of their possible values. The presence of this value in a property of a plan or subplan indicates that said plan or subplan can be achieved through one or more actions that do not use or alter the property in question at all.

Propagation of properties consists in sending the properties 'upwards' from the most concrete actions, up to the root goal, passing through every subplan and subgoal in the way. The full description of the method is provided in [25]. Given two *sequential* actions that have the same parent, the parent's set of properties will be the result of computing the union between the two children's properties. Each child will not have different possible values for the same properties, since they are sequential actions, and it would not make sense to design a plan in which child action no. 1 sets 'cuisine'='Spanish' only for the child action no. 2 to set the cuisine to be 'French'. Therefore, the properties of the two (sequential) children will always be different, and the resulting properties of the parent node will simply be the joining of the children's sets of properties, and it is trivial to see that this process applies to n sequential children actions.

Given two *alternative* actions that have the same parent, the parent's set of properties will be the result of merging the properties of the children in the following manner: if both children set different values for the same property then, for the father, the values of the property will be the union of the values that the children had (e.g., if child no. 1 had 'cuisine'='Spanish' and child no. 2 had 'cuisine'='French', the parent task will have 'cuisine'={'Spanish', 'French'} to indicate that if that node is chosen, we will limit the possible values of 'cuisine' to those two values). If either child has a property that the other does not, the parent will simply take the same properties of the child that has it, and will add the special value $None$, to indicate that if that node is chosen, there is a path of the plan that accomplishes the goal without ever giving a value to that property.

Figure 2 provides an example of property propagation. It shows the set of plans associated to a goal of ordering dinner. There are three possible options: a plan to order burgers, a plan to order falafel, and a plan to order pizza. LEt us assume for this example that there is only a local burger, a local falafel, and both a local pizza restaurant and a big company that makes pizza. Other assumptions that we take are that all burgers and pizzas are non-vegan, and that all falafels are vegetarian. The designer only needs to declare properties on the actions. Then, as a result of the property propagation process, all vertices have their own set of properties that have propagated upwards, from the leaves (actions). Notice how, in general, all properties have propagated towards the upward nodes. However, most of these propagations have been very simple ones: from single child to parent, although there are two cases worth mentioning. The first one is the propagation from the subplans to order local pizza and order from big pizza company. Notice how their properties are the same in all fields except for the 'local' field, with one holding it as True, and the other as False. However, these two *alternative* subplans share a common parent, and when their properties are propagated to it, they are merged in the way we described earlier: the parent has its property 'local' with *all* of its children values, to represent that, if that subgoal (or its parent subplan) is picked, then we can still order from either a local restaurant or a big chain. The other note-worthy example is the propagation of properties to the root node, where all options have been compiled in its properties.

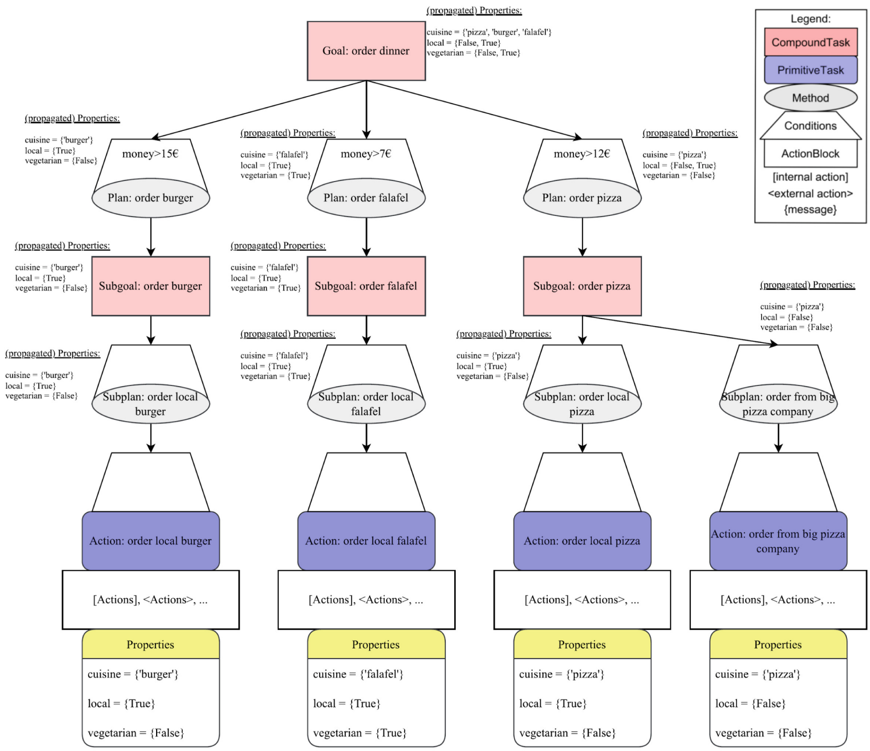

Fig. 2. Property propagation on an HTN plan assocated to the *order dinner* goal

3.5 Selection of Plans and Actions Using Properties

We will now briefly describe the process of choosing a plan taking preferences into account. An assumption we make throughout this whole example is that all plans are available, that is, our choices are not restricted by the environment in any way, shape, or form. Given a concrete goal g_i (order dinner) an agent has a set of preferences over the plans to achieve g_i. We can define this set as $gP_i = \{dPP, cPP_1, cPP_2\}$, where dPP is the default set of preferences, and cPP_1, cPP_2 are conditional sets of preferences. The dPP and each cPP_i are all an *ordered* instantiation of the values of different properties of the goal's plans. We assume that we have the following preferences over how to achieve the goal to order dinner (see Fig. 2):

1. $dPP = \{cuisine = \{falafel\}\}$: by default the metaplanner would only follow the branch with this property, and order from the falafel restaurant.
2. $cPP_1 = \{cuisine = \{burger, pizza\}, local = \{True\}\}[weather = snowy]$: in case of snow the metaplanner would follow branches that are either burger or pizza cuisine, but only those that are local (in the case of pizza this restricts it to only the local pizza place option).

3. $cPP_2 = \{local = \{False\}, vegetarian = \{False\}, cuisine = \{burger\}\}[weather = rainy]$: in case of rain the metaplanner attempts to follow branches meeting all the conditions, but even if the agent prefers to order non-vegetarian burgers, the first property prevails and leads to the only non-local option (pizza from big company).

As we can see, the agent picks from all the plans that satisfy the leftmost property, then, from those plans, it picks from those that satisfy the next leftmost property, etc. This process is for both default and conditionally triggered preferences, as they have the same structure, the only difference being that the latter need to be activated in order to take over and replace the default properties.

3.6 Adding Moral Values

Moral values can be seen as an ordering of preferences [10] that may be used by (human and artificial) agents to evaluate both individual actions and world states [14]. The main idea is that actions and world states promoting the agent's moral values are preferred over others [10].

In our framework we model the influence of moral values in the selection of actions by using the system of preferences over plans and actions described in Sects. 3.4 and 3.5. Consider the previous example of ordering food. We can ingrain moral values into each plan as extra properties in their actions. For instance, in our food ordering example (see Fig. 2) primitive tasks are associated to a *local* value (meaning the social value to favour local businesses and products over globalization-oriented trade of products coming from far away) that can be connected to *Universalism* and *Self-Transcendence* in Schwartz's theory of human values [22]. Another example is provided in Fig. 3, where bike and walk options for transportations are positively associated to the *environmentalist* value (that also can be connected to *Universalism* and *Self-Transcendence*) and the *health* value (that can be connected to *Hedonism* and *Self-Enhancement*).

As we are associating moral values to the primitive tasks, this may look as if our model pressuposes moral absolutism,[3] but actually, that is not true. As properties are defined for each plan of each agent, we can create an agent who thinks that lying is morally wrong, and an agent that thinks that it is morally right. Also, since the same action can be part of different subplans, we can also encode the fact that the morality of actions depends on their context. For example, if an agent kills an animal as part of a subplan to have fun, we can label that action as morally evil, but if the same agent kills an animal in his job as a veterinarian, then that action can be labelled as not morally evil.

[3] Moral absolutism is the position that there are universal ethical standards that apply to actions, and according to these principles, these actions are intrinsically right or wrong, regardless of what any person thinks, or context.

4 Example Scenario

We present a complex scenario to show how our agents fare with the new extensions: agents having many goals, goals decoupled from plans, preferences over goals, plans, and moral values.

In this Agent-Based Simulation, the agents' environment is a small town with some citizens living in it. These citizens are people which have their own set of daily goals (e.g., go to their workplace, have fun, eat dinner, etc.). Like real people, they have preferences over *in which order* to pursue their goals, as well as preferences over *how* to achieve them. Finally, they might have some moral inquiries into the actions we perform (e.g., being environmentalists and thinking the usage of cars is immoral, etc.).

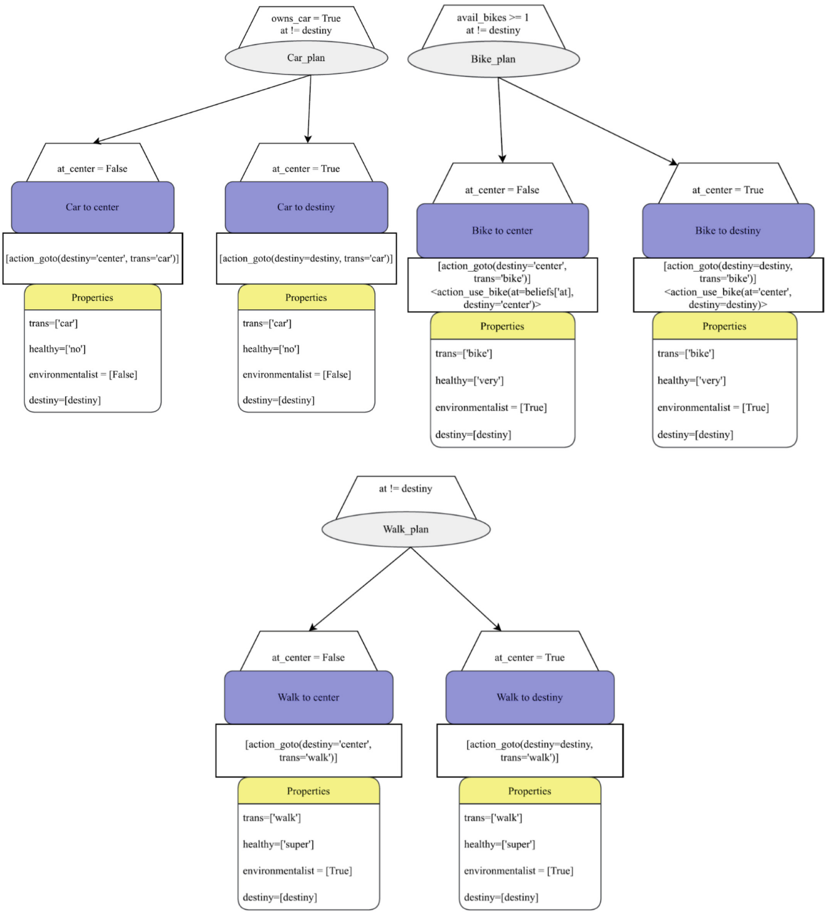

Fig. 3. Library of plans for transport goals

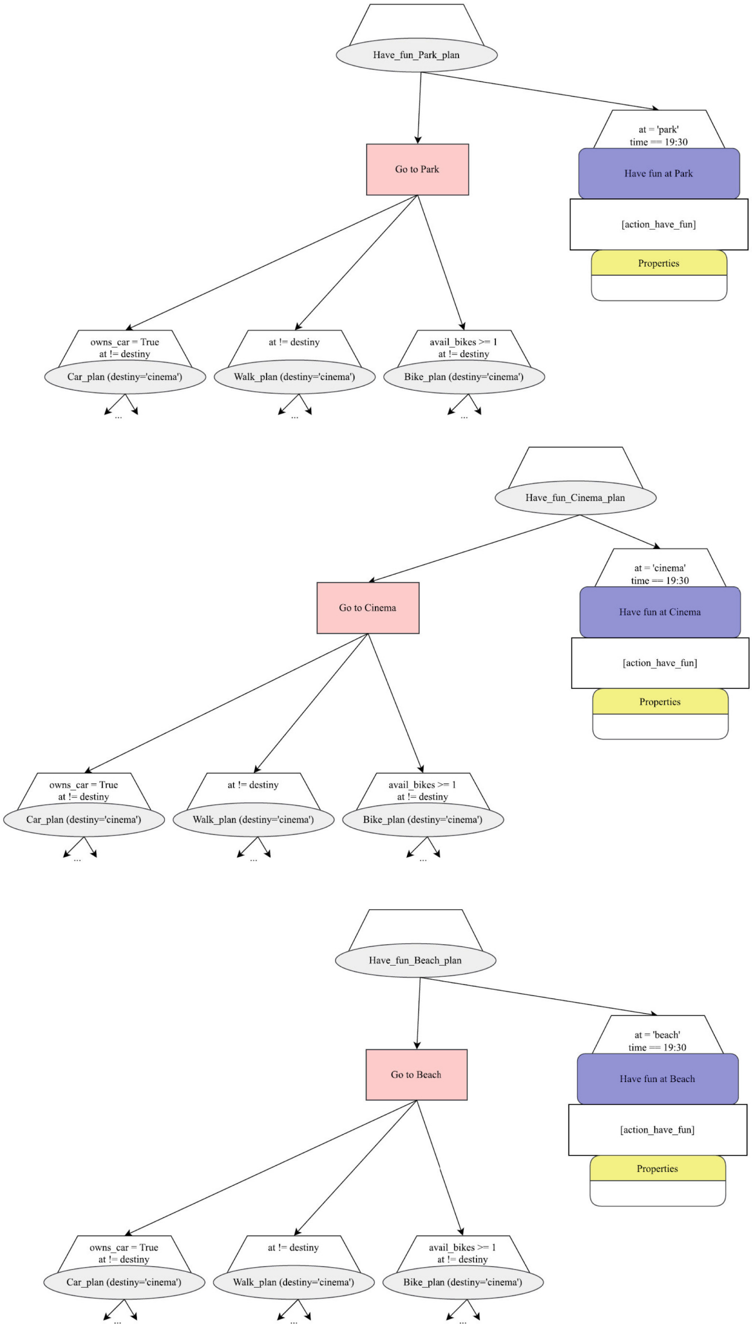

Fig. 4. Library of plans for fun-related goals

Table 1. Alice's and Bob's goals, preferences and values

<table>
<tr><td>ALICE's self goals
g_6 - Go home
g_7 - Eat dinner
g_8 - Attend any medical emergency</td><td>ALICE's role goals
g_1 - Take children to school
g_4 - Go collect her kids to school
g_5 - Have fun with her kids

g_2 - Go to work
g_3 - Work</td><td>BOB's self goals
g_3 - Have fun
g_4 - Go home
g_5 - Eat dinner
g_6 - Attend any medical emergency</td><td>BOB's role goals
g_1 - Go to work
g_2 - Work</td></tr>
<tr><td colspan="2">ALICE's preferences over goals
Default:
$[g_1 \to g_4 \to g_5 \to g_6], [g_1 \to g_2 \to g_3], [g_6 \to g_7]$
Conditional preferences:
– if (medical emergency)
$[g_8 \to g_1], [g_1 \to g_4 \to g_5 \to g_6], [g_1 \to g_2 \to g_3],$
$[g_6 \to g_7]$
– if (snowing)
$[g_2 \to g_3 \to g_6], [g_1 \to g_4 \to g_6], [g_6 \to g_7]$</td><td colspan="2">BOB's preferences over goals
Default:
$[g_1 \to g_2 \to g_4], [g_2 \to g_3], [g_4 \to g_5]$
Conditional preferences:
if (medical emergency)
$[g_6 \to g_1], [g_1 \to g_2 \to g_4], [g_2 \to g_3], [g_4 \to g_5]$</td></tr>
<tr><td colspan="2">ALICE's values
– For transport and fun-related goals:
• environmentalist = $False$
– For food-related goals:
• local = $False$ #big chains</td><td colspan="2">BOB's values
– For transport and fun-related goals:
• environmentalist = $True$
• healthy = $\{Super, Very\}$
– For food-related goals:
• local = $True$#local businesses</td></tr>
<tr><td colspan="2">ALICE's preferences over plans (transport goals)
Default:
$\{trans = \{car\}\}$</td><td colspan="2">BOB's preferences over plans (transport goals)
Default:
$\{trans = \{bike\}\}$
Conditional preferences:
$\{trans = \{walk, bike\}\}[weather = cloudy]$
$\{trans = \{car\}\}[weather = \{rainy, snowy\}]$</td></tr>
<tr><td colspan="2">ALICE's preferences over plans (fun-related goals)
Default:
$\{destiny = \{beach\}\}$
Conditional preferences:
$\{destiny = \{park\}\}[weather = cloudy]$
$\{destiny = \{cinema\}\}[weather = \{rainy, snowy\}]$</td><td colspan="2">BOB's preferences over plans (fun-related goals)
Default:
$\{destiny = \{beach\}\}$
Conditional preferences:
$\{destiny = \{cinema\}\}[weather = \{cloudy, rainy, snowy\}]$</td></tr>
<tr><td colspan="2">ALICE's preferences over plans (food-related goals)
Default:
$\{cuisine = \{pizza\}\}$
Conditional preferences:
$\{cuisine = \{chinese\}\}[weather = \{rainy\}]$</td><td colspan="2">BOB's preferences over plans (food-related goals)
Default:
$\{cuisine = \{pizza\}\}$</td></tr>
</table>

Each day of the simulated city is discretized in 64 steps. The simulated day starts at 08:00, and ends at 00:00 of the next day. Each simulation step corresponds to 15 minutes in the town. By default, the town starts with clear weather. Every iteration, there is a 10% chance of the weather changing. If that chance happens, there is a 60% chance of the weather becoming clear, 30% chance of becoming cloudy, 9% chance of raining, and 1% chance of snowing. At every iteration, there is also a 0.2% chance, for every agent, to experience a medical emergency. All these parameters are configurable by the user. The environment is randomly generated using a *seed*, and the agents will react and plan accordingly to the changes on the environment.

The town is composed of locations. Agents can move from one location to another by means of transport plans. The *town center* is the central location that connects with the others. People can live in the city center or in other city locations (residential neighborhoods). There are places where people go to have fun (a beach, a park and a cinema). There are also workplaces (a factory and corporate offices). There exist some places to go shopping: a local market (which includes Italian, Chinese and falafel restaurants), a supermarket and a big shopping center (with a big chain pizza company, a fast food burger company and a big chain of wholefood/vegetarian meals company). The city has locations wih some publicservices (a school for kids, a hospital to treat citizens).

Citizens might, by chance, experience a medical emergency, in which case, if they go to the hospital, they will be tended to and cured for free, so they can carry on with their day

The town can experience the following weather conditions: clear (sunny), cloudy, rainy and snowy. Those weather conditions may affect the citizens' choices (e.g., some may not use a bike if it rains). Some city services may be affected, too (e.g., schools are closed under snowy weather).

There are three main ways to go around the city: by car, by bike, or on foot. In order to drive a car, an agent needs to own one. In order to drive a bike, an agent needs to be at a location where a bike from the public rental system

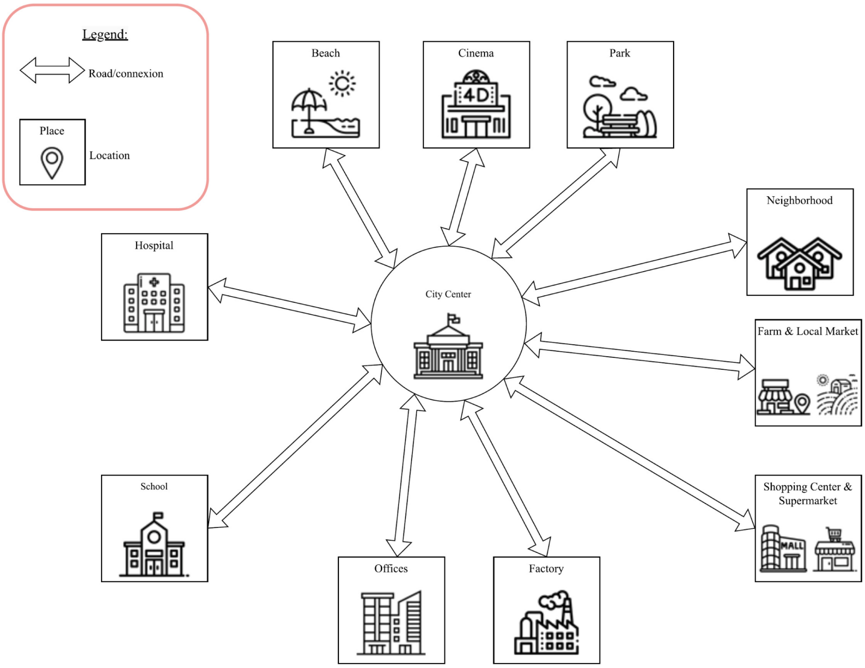

Fig. 5. Map of the town with its locations

is available, pick it, and leave it in another location. In our simulated city it is possible that some locations might not have any bikes at any given moment.

The `environment` class implements the map of city locations as well as other variables such as the current weather, the time, and extra internal variables for purposes of running the simulation. When an agent perceives the environment, they will only perceive the current time, the current weather, and the information of the location that they are currently in. For instance, if an agent is at the city center, it will not update its information about the state of the school, only about the state of the city center, the weather, and the time.

There are two main actors in our environment, **Alice** and **Bob**. They both are complex agents with numerous goals, conditional preferences over these goals, a rich library of plans, and preferences over those plans, along with moral values.

Alice is the CEO of a big company. She works at the office every day until 16:45. She has to take the children to school every morning, collect them from school at 17:00, and go have fun with them in the afternoons (until 19:45). Then, they order food at 20:00. Her initial beliefs are her current location, the current weather and time, the current location of her children, whether she owns a car, whether she has worked, if her children have gone to school, if she is at the center of the city, and whether there is a medical emergency. Table 1 shows her goals, preferences over goals and plans and her values. Alice's library of plans consists of three sets of complex plans: one set of plans for fun-related goals (see Fig. 4), one set of plans for transport goals (see Fig. 3) and one set for food plans (an extension of the one shown in Fig. 2 with an extra plan branch for Chinese food). Goals g_1, g_2, g_4, and g_6 include commuting, and therefore are mapped to transport plans by the metaplanner. g_5 and g_7 are mapped to fun and order meal plans, respectively. The other plans for other goals are trivial: they have a single plan, with a single action (e.g., in the case of the plan to work, there is only one method, with a single action).

Bob is the second agent we have created for this test scenario. Like Alice, he has his own set of beliefs, a place where he lives, a place where he goes to work, preferences over how to have fun, etc. Bob lives in the city center and is a worker in the local factory, every day until 16:45. He has no children so he goes to work directly every morning. Once he is done, he goes to have fun however he prefers. Then he goes back home and orders food at 21:00. His initial beliefs are similar to Alice's, excluding those children-related. Table 1 shows Bob's goals, preferences over goals, plans and values. Goals g_1, g_3 and g_4 include commuting and therefore are mapped to transport plans by the metaplanner. g_3 and g_5 are mapped to fun and order meal plans, respectively. Bob's goals are a subset of Alice's goals and are mapped to the same plans, but Bob will not act like Alice, as their personal preferences and moral values differ.

4.1 Tests and Results

In this section we show some execution runs to see that agents plan according to their goals, preferences and values, and that they respond to changes in the environment that might cause them to reconsider their contextual preferences and, therefore, need to replan, or even reconsider their goals.

Figure 6 shows the result of a simulation with all default parameters except for `emergencyodds = 0.2` (20%). At step 35 we can see that Alice is working in her workplace when she receives a medical emergency of one of the kids. Then, *her conditional preferences over goals activate*, she changes her current goal, and she rushes to the hospital, as we can see in the next step. Although not shown in the picture, when she goes to the hospital and is cured, her preferences over goals revert to default, and she goes back to the offices to continue working.

In Fig. 7 there is the result of a simulation with all default parameters except for `changeodds = 1`, `rainodds = clearodds = 0.5`, and `cloudoods = snowodds = 0`. At step 43, both agents were having fun at the beach. However, it suddenly started to rain, and then *their preferences over plans changed*. The goal (to have fun) does not change. What changes, however, is *how* they decide to have fun. According to their conditonal preferences for fun, in case of rain they prefer to go to the cinema, and they replan giving priority in the HTN to the branches with the `destiny={cinema}` property.

Figure 8 shows an example of the interwork of conditional preferences over plans and values. The Observer Agent tells us that it is raining. In the case of Bob, his conditional preferences over food-related goals determine that its single, permanent, default preference is always pizza (see Table 1). Therefore, Bob's HTN related to the "order food" goal (Fig. 2) will select the order pizza branches (except if Bob has less than $12, then the order falafel branch will be explored). But to choose among the two order pizza sub-branches, Bob's values ($local = True$) are used to make the choice. From the two possible options to order pizza, only "order local Pizza" has its local value True and is chosen (see Bob's mental state in Fig. 8). The rainy weather has also triggered a change

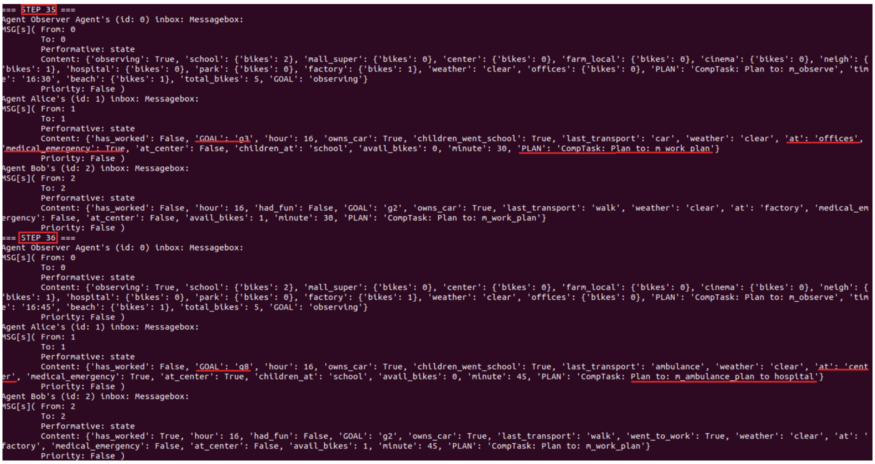

Fig. 6. Agent Alice changing preferences over goals

in his transportation means (car), which is fully mandated by his conditional preference over transportation plans. Here it is interesting to see that a conflict arises between the properties attached to the Car plan (`healthy={no}` and `environmentalist={false}`) and Bob's values ((`healthy={Super, Very}` and `environmentalist={true}`). As we have no numbers to rate the relative importance of conflicting preferences, we have to solve the conflict by explictly placing in the scenario definition file the `trans` preference before the `healthy` one.

In general, we see that our agents react to changes in their current context by changing their priorities, and always plan according to them. Additionally, by looking at the whole verbose dump of a simulation, we see that they function as expected: they pursue their default goals in the correct order, change priorities over goals whenever they should, replan according to changes in both priorities over goals and plans, and make choices based on them.

5 Conclusions

In this paper, we describe an extension to an agent-based simulation environment for High Performance Computing enabling goal-driven agents with hierarchical

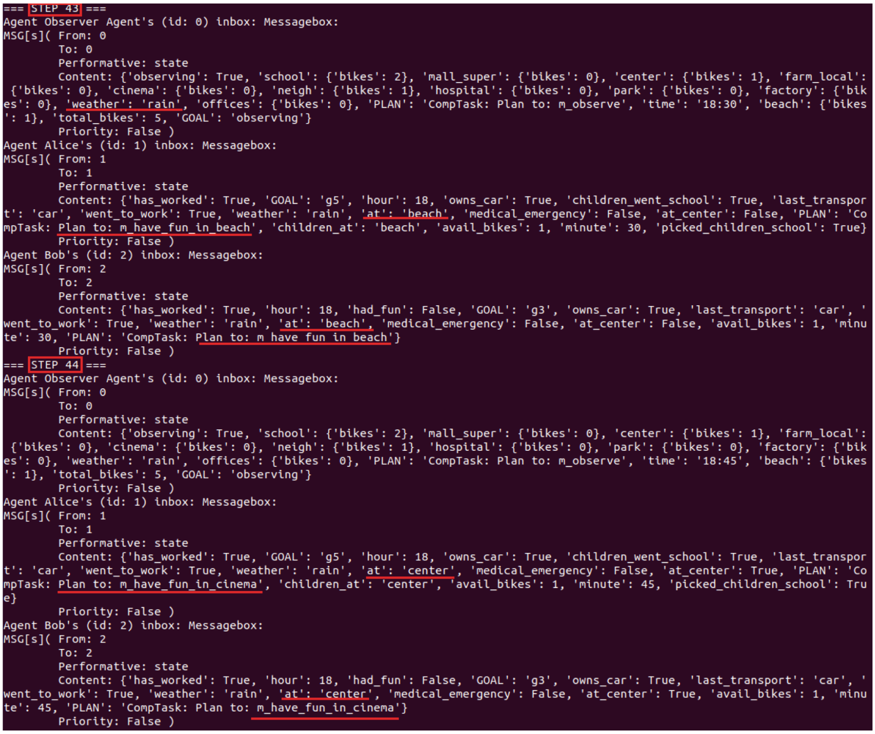

```
=== STEP 43 ===
Agent Observer Agent's (id: 0) inbox: Messagebox:
MSG[s]( From: 0
        To: 0
        Performative: state
        Content: {'observing': True, 'school': {'bikes': 2}, 'mall_super': {'bikes': 0}, 'center': {'bikes': 1}, 'farm_local':
 {'bikes': 0}, 'cinema': {'bikes': 0}, 'neigh': {'bikes': 1}, 'hospital': {'bikes': 0}, 'park': {'bikes': 0}, 'factory': {'bik
es': 0}, 'weather': 'rain', 'offices': {'bikes': 0}, 'PLAN': 'CompTask: Plan to: m_observe', 'time': '18:30', 'beach': {'bikes
': 1}, 'total_bikes': 5, 'GOAL': 'observing'}
        Priority: False )
Agent Alice's (id: 1) inbox: Messagebox:
MSG[s]( From: 1
        To: 1
        Performative: state
        Content: {'has_worked': True, 'GOAL': 'g5', 'hour': 18, 'owns_car': True, 'children_went_school': True, 'last_transpor
t': 'car', 'went_to_work': True, 'weather': 'rain', 'at': 'beach', 'medical_emergency': False, 'at_center': False, 'PLAN': 'Co
mpTask: Plan to: m_have_fun_in_beach', 'children_at': 'beach', 'avail_bikes': 1, 'minute': 30, 'picked_children_school': True}
        Priority: False )
Agent Bob's (id: 2) inbox: Messagebox:
MSG[s]( From: 2
        To: 2
        Performative: state
        Content: {'has_worked': True, 'hour': 18, 'had_fun': False, 'GOAL': 'g3', 'owns_car': True, 'last_transport': 'car', '
went_to_work': True, 'weather': 'rain', 'at': 'beach', 'medical_emergency': False, 'at_center': False, 'avail_bikes': 1, 'minu
te': 30, 'PLAN': 'CompTask: Plan to: m_have_fun_in_beach'}
        Priority: False )
=== STEP 44 ===
Agent Observer Agent's (id: 0) inbox: Messagebox:
MSG[s]( From: 0
        To: 0
        Performative: state
        Content: {'observing': True, 'school': {'bikes': 2}, 'mall_super': {'bikes': 0}, 'center': {'bikes': 1}, 'farm_local':
 {'bikes': 0}, 'cinema': {'bikes': 0}, 'neigh': {'bikes': 1}, 'hospital': {'bikes': 0}, 'park': {'bikes': 0}, 'factory': {'bik
es': 0}, 'weather': 'rain', 'offices': {'bikes': 0}, 'PLAN': 'CompTask: Plan to: m_observe', 'time': '18:45', 'beach': {'bikes
': 1}, 'total_bikes': 5, 'GOAL': 'observing'}
        Priority: False )
Agent Alice's (id: 1) inbox: Messagebox:
MSG[s]( From: 1
        To: 1
        Performative: state
        Content: {'has_worked': True, 'GOAL': 'g5', 'hour': 18, 'owns_car': True, 'children_went_school': True, 'last_transpor
t': 'car', 'went_to_work': True, 'weather': 'rain', 'at': 'center', 'medical_emergency': False, 'at_center': True, 'PLAN': 'Co
mpTask: Plan to: m_have_fun_in_cinema', 'children_at': 'center', 'avail_bikes': 1, 'minute': 45, 'picked_children_school': Tru
e}
        Priority: False )
Agent Bob's (id: 2) inbox: Messagebox:
MSG[s]( From: 2
        To: 2
        Performative: state
        Content: {'has_worked': True, 'hour': 18, 'had_fun': False, 'GOAL': 'g3', 'owns_car': True, 'last_transport': 'car', '
went_to_work': True, 'weather': 'rain', 'at': 'center', 'medical_emergency': False, 'at_center': True, 'avail_bikes': 1, 'minu
te': 45, 'PLAN': 'CompTask: Plan to: m_have_fun_in_cinema'}
        Priority: False )
```

Fig. 7. Agents Alice and Bob changing preferences over plans

task network (HTN) plans to choose among goals and among plans based on preferences and a simple moral values model. We have summarized extensions done on the agent model and how they work in a sample scenario. We have also been able to see how ‘far’ we could go without using any numbers to express preferences over goals, plans, and moral values. As we have seen, we have been able to express conditional preferences over both, have these preferences change based on context, and agents replan based on environmental changes. This work is one more step towards our goal to have a powerful agent-based micro-simulation framework to analyse the potential impact of social values, policies, norms and conventions in large populations of social-aware agents.

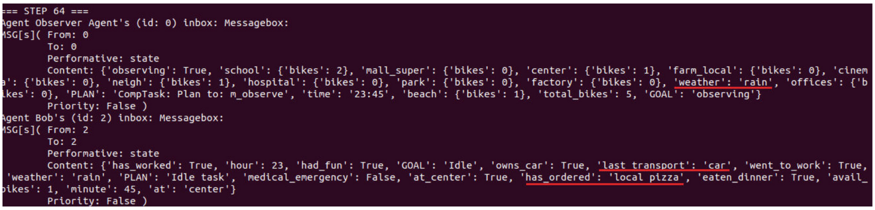

Fig. 8. Agent Bob has used his preferred means of transport for when it rains and his “local business” values to choose the local pizza option

One of the biggest limitations in how we declare goals is that, at any given moment, our agent can only pursue one goal at a time. This limitation is also common in many BDI-inspired implementations. Only few agent platforms (such as Jason [3] or 2APL [9]) allow to pursue several goals at the same time. We are already working on an extension of the model and its implementation to allow several goals at the same time, specially to allow handling combinations of achievement goals and maintenance goals. Another limitation is that our agents do not support adding (or removing) goals in runtime. Goals can be either achieved or not achieved at any given moment, but they cannot be eliminated (nor new goals can be added). This limitation was introduced for performance reasons. We plan to tackle this in future extensions.

Perhaps the biggest limitation in our declaration of preferences over goals and plans is that they are absolute, and this stems from the fact that we aimed to not use numbers in our model. Therefore, we cannot express things like ‘I prefer this *a little* more than that’, or ‘I prefer that *a lot* more than this’ that could be used to solve conflicts (such as Bob’s conflicts between the plan preference and his values). Visser’s et al.’s approach [25] provides a more complex structure that allows their agents to have more complex preferences (e.g., agents can reason about quantities, quantity optimization, limitation by quantity, etc.). Also, their agents are able to automatically extract properties of goals by looking at the actions, and then derive the relevant properties of the goals. Our model relies on the designer carefully listing (within the scenario description file) the properties and the preferences in the *right* order. In future work we will explore more flexible and expressive ways to solve this (with no numerical values, if possible).

One related issue we plan to investigate further is related to what to do when the trigger conditions of non-default preferences over goals overlap (e.g., it is snowing and a medical emergency occurs), especially in the case they define different preorders. Our current approach is to pick the first goal preorder (by declaration order in the scenario file), and to allow the designer to implement an ad-hoc, more complex solution, if their scenario requires so. It would be better to modify our model to allow for a native way to handle this issue.

Finally, our encoding of moral values also totally relies on the designer carefully listing which actions have what moral implications and, while this is good from an expressiveness point of view (it allows us to declare moral relativism as different agents having different moral convictions) and context-dependent morality (the same action carried out under different circumstances having different moral implications), it is a very exhaustive and daunting task. It would be good to have the system partly automated, perhaps employing some matching between the purpose of an action and a value-tree structure rooted in a well-founded model of values (such as Schwartz's [22], which is used in [8,13,14]).

Acknowledgements. This work has been partially supported by EU Horizon 2020 Project StairwAI (grant agreement No. 101017142).

References

1. Abar, S., Theodoropoulos, G.K., Lemarinier, P., O'Hare, G.M.: Agent based modelling and simulation tools: a review of the state-of-art software. Comput. Sci. Rev. **24**, 13–33 (2017)
2. Bellifemine, F., Poggi, A., Rimassa, G.: JADE—A FIPA-compliant agent framework, pp. 97–108. The Practical Application Company Ltd. http://jade.cselt.it/papers/PAAM.pdf (1999)
3. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-Agent Systems in AgentSpeak Using Jason (Wiley Series in Agent Technology). John Wiley & Sons Inc, Hoboken, NJ, USA (2007)
4. Boutilier, C., Brafman, R.I., Domshlak, C., Hoos, H.H., Poole, D.: CP-nets: a tool for representing and reasoning with conditional ceteris paribus preference statements. J. Artif. Intell. Res. **21**, 135–191 (2004)
5. Braubach, L., Pokahr, A., Lamersdorf, W.: Jadex: a BDI reasoning engine. In: Multi-agent programming: Languages, platforms and applications, pp. 149–174 (2005)
6. Casali, A., Godo, L., Sierra, C.: A graded BDI agent model to represent and reason about preferences. Artif. Intell. **175**(7), 1468–1478 (2011)
7. Cohen, P.R., Levesque, H.J.: Intention is choice with commitment. Artif. Intell. **42**(2–3), 213–261 (1990)
8. Cranefield, S., Winikoff, M., Dignum, V., Dignum, F.: No pizza for you: value-based plan selection in BDI agents. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence, pp. 178–184. IJCAI'17, AAAI Press (2017)
9. Dastani, M.: 2APL: a practical agent programming language. Auton. Agents Multi Agent Syst. **16**, 214–248 (2008)

10. Di Tosto, G., Dignum, F.: Simulating social behaviour implementing agents endowed with values and drives. In: Multi-Agent-Based Simulation XIII: International Workshop, MABS 2012, Valencia, Spain, June 4–8, pp. 1–12 (2012)
11. Gnatyshak, D., Oliva-Felipe, L., Álvarez Napagao, S., Padget, J., Vázquez-Salceda, J., Garcia-Gasulla, D., Cortés, U.: Towards a goal-oriented agent-based simulation framework for high-performance computing. In: Artificial Intelligence Research and Development: Proceedings of the 22nd International Conference of the Catalan Association for Artificial Intelligence, pp. 329–338. IOS Press (2019)
12. Hindriks, K.V., Roberti, T.: GOAL as a planning formalism. In: Lecture Notes in Artificial Intelligence, vol. 5774, pp. 29–40 (2009)
13. Kammler, C., Dignum, F., Wijermans, N.: Utilizing the full potential of norms for the agent's decision process. In: Proceedings of the Social Conference 2022, Milan, 12–16 September (2022)
14. Kruelen, K., de Bruin, B., Ghorbani, A., Mellema, R., Kammler, C., Vanhée, L., Dignum, V., Dignum, F.: How culture influences the management of a pandemic: a simulation of the COVID-19 crisis. J. Artif. Soc. Soc. Simul. **25**(3), 1013–1020 (2022)
15. Kutluhan, E., Hendler, J., Nau, D.S.: HTN planning: complexity and expressivity. In: Proceedings of the AAAI Conference on Artificial Intelligence, 12, pp. 1123–1128 (1994)
16. Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., Balan, G.C.: MASON: a multiagent simulation environment. SIMULATION Trans. Soc. Model. Simul. Int. **81**(7), 517–527 (2005)
17. Oliveira de Nunes, I., Lucena, C., Luck, M.: BDI4JADE: a BDI layer on top of JADE. In: Proceedings of the 9th Workshop on Programming Multiagent Systems, pp. 88–103 (2012)
18. Orkin, J.: Three states and a plan: the A.I. of F.E.A.R. In: Game Developers Conference 2006. https://alumni.media.mit.edu/jorkin/gdc2006_orkin_jeff_fear.pdf (2006)
19. Padgham, L., Singh, D.: Situational preferences for BDI Plans. In: 12th International Conference on Autonomous Agents and Multiagent Systems 2013, AAMAS 2013 2, pp. 1013–1020 (2013)
20. Rao, A.S., Georgeff, M.P.: An abstract architecture for rational agents. In: Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR-92), pp. 439–449 (1992)
21. Rousset, A., Hermann, B., Lang, C., Philippe, L.: A survey on parallel and distributed multi-agent systems for high performance computing. Comput. Sci. Rev. **22**, 27–46 (2016)
22. Schwartz, S.H.: An overview of the Schwartz theory of basic values. Online Read. Psychol. Cult. **2**(12), (2012)
23. Tejedor, E., Becerra, Y., Alomar, G., Queralt, A., Badia, R.M., Torres, J., Cortes, T., Labarta, J.: PyCOMPSs: parallel computational workflows in Python. Int. J. High Perform. Comput. Appl. **31**(1), 66–82 (2017)
24. Tisue, S., Wilensky, U.: NetLogo: a simple environment for modeling complexity. In: International Conference on Complex Systems, vol. 21, pp. 16–21 (2004)
25. Visser, S., Thangarajah, J., Harland, J., Dignum, F.: Preference-based reasoning in BDI agent systems. Auton. Agents Multi Agent Syst. **30**(3), 291–330 (2016)
26. Winikoff, M., Sidorenko, G., Dignum, V., Dignum, F.: Why bad coffee? explaining BDI agent behaviour with valuings. Artif. Intell. **300**, 103554 (2021)

27. Zia, K., Riener, A., Farrahi, K., Ferscha, A.: A new opportunity to urban evacuation analysis: very large scale simulations of social agent systems in repast HPC. In: 2012 ACM/IEEE/SCS 26th Workshop on Principles of Advanced and Distributed Simulation, pp. 233–242 (2012)