The Multi-Agent Programming Contest: A Résumé

Comparing Agent Systems 2005–2019

Tobias Ahlbrecht[1]([✉]) , Jürgen Dix[1] , Niklas Fiekas[1] ,
and Tabajara Krausburg[1,2]

[1] Department of Informatics, Clausthal University of Technology,
Clausthal-Zellerfeld, Germany
{tobias.ahlbrecht,dix,niklas.fiekas}@tu-clausthal.de
[2] School of Technology, Pontifical Catholic University of Rio Grande do Sul,
Porto Alegre, Brazil
tabajara.rodrigues@edu.pucrs.br

Abstract. The Multi-Agent Programming Contest, MAPC, is an annual event organized since 2005 out of Clausthal University of Technology. Its aim is to investigate the potential of using decentralized, autonomously acting intelligent agents, by providing a complex scenario to be solved in a competitive environment. For this we need suitable benchmarks where agent-based systems can shine. We present previous editions of the contest and also its current scenario and results from its use in the 2019 MAPCwith a special focus on its suitability. We conclude with lessons learned over the years.

Keywords: Multi-agent systems · Decentralized computing · Cooperation · Artificial intelligence · Simulation platforms

1 Introduction

The original aim of our contest, back in the humble beginnings in 2005, was to provide a platform for comparing and evaluating systems based on *computational logic*, mainly developed for knowledge representation purposes.

We wanted to develop an interesting yet simple, but non-trivial, scenario for testing systems based on different paradigms. At that time, many knowledge-based approaches were developed as smallish PhD projects: a prototype was implemented but never seriously compared against other such systems.

Emphasis was put on the *evaluation* and *comparison* of systems, not on finding an optimal solution of a particular scenario. The creation of a scenario was always driven by the need to determine the features that a system should possess for successfully solving a complicated task. We never wanted to honor a smart idea for a solution, but the features and technology that help to tackle the problem at hand.

1.1   Structure of This Work

We start with a short introduction on agent programming and the Multi-Agent Programming Contest in general. In Sect. 2, we introduce the simulation platform, followed by the history of the Contest in Sect. 3. Afterwards, we present the newest scenario and the results of the MAPC2019 in Sect. 4. We conclude with lessons learned during the contest in general and in the latest installment in particular. Throughout the article, we focus especially on the scenario aspect of the MAPC.

1.2   Agents

During the years, the systems we compared turned more and more into those based on *agent programming languages* [9] or genuine multi-agent systems (MAS) implemented in classical programming languages. The scenarios became more complex with an increasing number of agents needed to solve the task.

In contrast to many other contests, several decisions have been taken a priori:

– not to impose any restrictions on the software used;
– not to find or compare *tricky algorithms* to solve the scenario, rather we wanted to evaluate the capabilities of the system to express and model suitable constructs for dealing with the scenario;
– not to consider the perfect implementation or high performance of a system; in particular, we never considered real-time aspects, which are important for e.g. computer games.

The last bullet above reflects the situation in agent programming for many years (still today, but to a lesser extent): agent languages are still not on par with classical programming languages in terms of their efficient implementation and their maturity concerning software engineering aspects. We therefore decided to refrain from this particular aspect.

In our MAPC, each participating team develops a group of agents (during the 5–7 months between the announcement of the scenario and the contest), which remotely connect to our MAPCserver where the scenario is being run. The MAPCserver sends the current game state in the form of percepts to each agent and expects an executable action in return. The gathered actions are executed and the game state is advanced. This cycle is repeated until a predefined number of steps is reached. The remote nature of the contest also keeps the responsibility of running the agents with the participants.

The available time for each simulation step must include the latency of the internet and is, intentionally, chosen to be quite high (4 seconds): we do not consider high performance nor real-time constraints.

In addition, we have no control of the communication within a team (e.g., shared memory or not, decentralized or not). Consequently, we could not directly enforce decentralized approaches—only by designing the scenario in a way that favors them.

We also never excluded classical (i.e., non-agent) programming languages and frameworks from being used. In fact, we almost always had non-agent entries taking part in our contest and some performed very well. Obviously, one can use a classical programming language and implement certain agent-techniques that are suitable for the scenario (or use agent technology without leveraging its potential, in effect, using it like a conventional programming approach).

1.3   Goals and Purpose of the **MAPC**

The purpose of the contest is twofold:

1. to find out for which applications agent-oriented features pay off, as opposed to features available in classical programming languages; and
2. to compare and test the versatility and suitability of agent languages or platforms.

To answer the first question, we are developing and evolving scenarios building on the experiences from previous editions of the MAPC. By improving the scenarios, we simultaneously improve our ability to answer the second question.

However, it should be clarified that *we do not want to compare problem solutions*, instead we want to compare agent languages among themselves and against classical programming languages.

The difference between agents and classical, more centralized paradigms is, to a great deal, autonomy, communication, cooperation and to strike a good balance between proactiveness and reactiveness. Clearly, any (new) feature can be implemented in any (Turing-complete) programming language, but one would hope an agent language to be more versatile and efficient or offering built-in features for elegantly programming a solution.

Therefore we always try to develop our scenarios in such a way, that no smart solution will be sufficient, but instead the *interplay of various acting entities and their emerging features* counts.

In the end, we are especially interested in:

1. which technologies the teams used;
2. to which degree they were used (i.e., how difficult (or easy) it was to use agent-based features); and
3. which aspects were especially straightforward or challenging to design and implement.

To summarize again, the contest is an attempt to shed some light on these questions: when and to what extent do agent-oriented features pay off? Is there a particular complexity of the problem that makes these approaches beneficial? Or not at all? And how are these features supported by existing agent frameworks? We refer back to all of these questions in Sect. 5.

Last but not least, it almost comes naturally that we aim to support educational efforts in the design and implementation of agent systems by providing each year a ready, off-the-shelf package: this is a ready for action tool in the

classroom and could be (and has been) used in a course on agent systems of any level. We noticed in our experience that the competition idea is especially attractive for students and results in a very engaging work atmosphere.

1.4 Related Work

Many similar competitions have been and are still being held, while most do not explicitly focus on multi-agent systems. We discuss some of them that are related to MAPCand are still active nowadays.

Directly involving agents, the *(Power) Trading Agent Competition*[1] [16] provides a trading-related scenario in the energy market. However, each team only consists of a single "broker" agent, requiring no cooperation or coordination. The goal here is to see how agents can autonomously solve supply-chain problems.

Probably the best-known are the various *RoboCup Simulation Leagues*[2]. RoboCup ranges over a variety of different domains like soccer, disaster response, and industrial logistics. Each league focuses on a specific problem that must be addressed by competitors. For instance, in RoboCupRescue two major leagues are organized on: (i) robots; and (ii) agent simulation. The first centers around (virtual) robots and less around abstract agents. For example, agents have noisy virtual sensors or may be subjected to complex physics, focusing on realism. The agent simulation league provides virtual agents placed on a map of a city that has been damaged by an earthquake event. Competitors focus on different self-isolated AI problems (e.g., task allocation) provided by the contest [18]. In addition, all teams have to give a presentation on their solution, which counts towards their final score.

There is also a number of challenges targeting specific problem domains, e.g. the *International Planning Competition* [17]. Here, of course planning is in the limelight, while in our contest it is only one possible component of an agent team. At the other end, the *General Game Playing* [12] competitions do not focus on one particular feature but on the ability of general AI systems to play an arbitrary game upon receiving its rules.

Finally, there are more than a few challenges focusing on finding (autonomous) solutions for existing commercial games, like the *Mario AI Championship*[3] [15] or the *Student StarCraft AI tournament*[4], or specifically designed games like *BattleCode*[5]. The goal here is usually to benchmark game AI techniques and algorithms.

We would also like to mention a new challenge, the Intention Progression Competition[6], which focuses on a specific issue within agent systems: the Intention Progression Problem, i.e. the decision of agents about how to proceed with

---

[1] www.powertac.org.

[2] www.robocup.org.

[3] www.marioai.org.

[4] www.sscaitournament.com.

[5] www.battlecode.org.

[6] www.intentionprogression.org.

their given intentions and plans in order to reach their goals. Thus, a solution for the MAPC(e.g. an agent) could be seen as a specific input challenge in the IPC, while solutions for the IPC could be used in agent platforms that participate in the MAPC.

Not a competition but definitely worth mentioning is the Blocks World for Teams (BW4T) [14] environment, which is not quite unlike the current MAPCscenario. There, agents have to coordinate to deliver sequences of color-coded blocks.

2   The MASSim Framework

The first edition of the MAPCin 2005 presented a simple scenario description that had to be implemented in its totality by each participant and delivered as an executable.

2.1   From 2006: The Early Days

In 2006, the *MASSim* platform was introduced: an extensible simulation server written in Java that provides the environment facilities. Agent programs can connect through the network to a *MASSim* server while agents run in the competitors' own computer infrastructure.

Since then, the format of the MAPChas been that of two teams competing against each other for performance in each simulation, and the overall winner of the contest defined by summing up the points after all participants have competed in simulations against each other, in a regular sports tournament fashion.

All simulations are run in a discrete step-by-step manner. In each step all agents execute their actions simultaneously from the point of view of the server, and there is a time limit within which agents must choose an action (otherwise they are regarded as a `no-op`). In the beginning of each step's cycle, the server sends each agent their current percepts of the environment, and waits for the response that specifies the action to execute.

When the responses from all agents are received or when the timeout limit is reached, all received actions are executed in *MASSim*. The actions (mostly) have an immediate effect on the environment and the new state of the simulation is computed which results in new agent percepts for the next simulation step.

This cycle is repeated for a fixed number of steps, and then a winner is decided according to scenario-specific criteria (usually having achieved the highest score).

MASSim is fully implemented in Java, and the information exchange with the agent programs is realized through XML messages. These messages can also be accessed as ready percept objects through the EISMASSim interface layer, which is explained later.

2.2   2017 Until Today: Simply Going Forward

In early 2017, *MASSim* was completely rewritten. XML messages were removed in favor of the more efficient JSON format.

We switched from having both a Java RMI based monitor and a web monitor to a single web-based monitor.

Also, we abandoned the former plug-in architecture in favor of an annual package, which helped in keeping the package small and freed us of having to keep *MASSim* backwards-compatible to all previous scenarios.[7]

This rewrite also allowed us to create a platform with more than two concurrent teams in mind. While we have not used this yet, it remains a tempting option for future scenarios.

Figure 1 displays the current architecture. Boxes are components, while regular arrows depict that a component uses another.
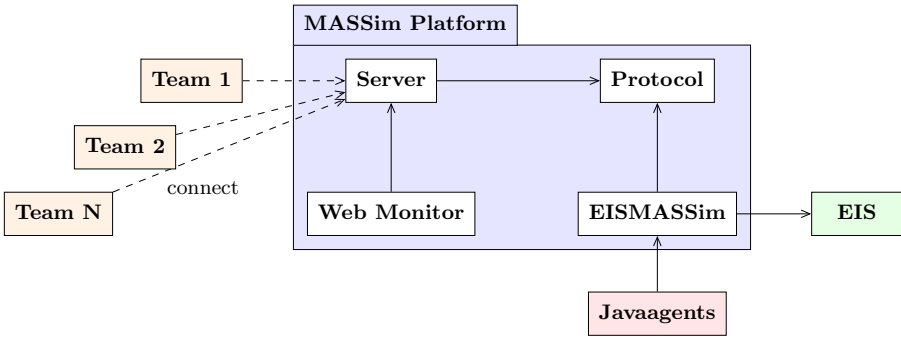

Fig. 1. The *MASSim* architecture.

The server package is responsible for running the simulations and handling connections to all agents. It only uses the facilities of the protocol package to build percept messages for agents and parse action messages it receives.

The classes used to build valid messages according to the protocol have been extracted into a self-contained protocol package that helps both with parsing JSON data into Java objects and transforming Java message objects into their JSON representation. Thus, it is e.g. used by the server to create messages for all agents. The protocol is also used by the EISMASSim component, which can be used by agent platforms to connect to the server. This component handles the whole login procedure and then translates perception and action messages into actions and percepts according to the EIS (Environment Interface Standard [7]) and vice versa. In the terms of EIS, EISMASSim is what makes *MASSim* an "EIS-enabled" environment. That is, all agent platforms that support EIS can connect to the *MASSim* server without any additional effort (though sometimes there are still some initial difficulties).

We also provide a sample implementation of agents using EISMASSim in the `Javaagents` package. Participants using Java-based platforms may connect to the server by integrating EISMASSim, using the protocol package, or, just as non-Java-teams, parse and build their own JSON messages according to the protocol.

---

[7] You can still play the old releases using their respective packages.

Finally, the web monitor is started by the server if requested and then retrieves the current game state from the server after each step.

The current *MASSim* package is fully open-source and openly available (https://multi-agentcontest.org/2019). It is not only used for the MAPC, but has also proved useful both for researchers testing their advancements in the field, and in the classroom, aiding the teaching of the multi-agent programming paradigm (https://multi-agentcontest.org/massim-in-teaching).

3   History and Evolution of the Contest

We can roughly divide the contest into two phases. In the early phase, there was not much cooperation among the agents: they acted more or less on their own. This led us to reconsider our scenario and we ended up with the *Agents on Mars* scenario, where we experienced some really interesting games. This then evolved into the *Agents in the City* (or simply *City*) scenario, which was even more realistic as it considered agents acting in a real city using actual city maps. We then adapted the *City* scenario, removing some of its complexity (regarding implementation effort for the participants) and incorporating features we think were interesting from previous scenarios, which led to the *Agents Assemble* scenario, which we will present and analyze in detail in Sect. 4.

3.1   Early Phase

The scenario used for the first edition of the MAPC(2005) consisted of a simple grid in which agents could move to empty adjacent spaces. Food units would appear randomly through the simulation, and the objective was to collect these units and carry them to a storage location.
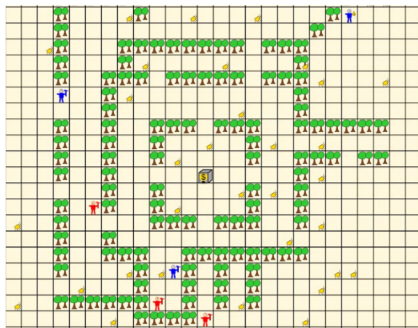

Fig. 2. The Gold Miners scenario.

The idea was refined for the second edition [6]: *Gold Miners* (Fig. 2). Now the agents were to collect gold in a competitive environment against another team,

and some obstacles were introduced to the grid to add some navigation complexity. This scenario, which was also used in the third edition of the contest, was still very simplistic, and in the proposed solutions agents acted independently of their teammates: no cooperation or coordinated behavior took place.
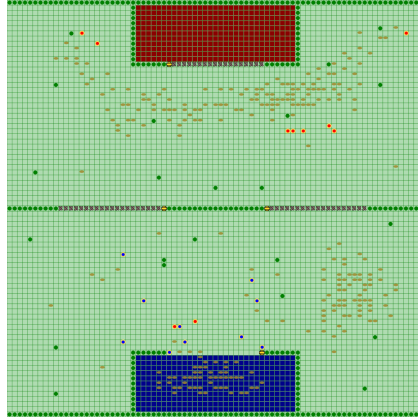

Fig. 3. The Cows and Cowboys scenario.

For the 2008–2010 editions [6], a new scenario was designed that *demands* coordination from agents: *Cows and Cowboys*, as shown in Fig. 3. Still using a grid as the underlying map, the goal for this scenario was to lead a group of cows to a particular area of the map, the team's own "corral", while preventing the opponent team from doing the same. The cows were animated entities that reacted to the agents' positions by trying to avoid them. Solving the map required agents to coordinate their positions in order to lead big groups of cows into the corrals, whereas a single agent would in most cases disperse the group of cows and fail to lead them in the desired direction.

Even in this clearly cooperative scenario, one team found a way of letting each agent work independently, always pushing a single cow. This team promptly won the contest (though out-of-competition) and we learned that *features we want to see need to be enforced rather than rewarded*, since participating teams always tend to find (and go for) the path of least resistance. Thus, a flocking algorithm for cows was introduced, which made the cows form groups and avoid agents more strongly. This allowed good teams to capture entire herds with the right agent formations, while single agents could not achieve anything anymore. In addition, fences were added as another cooperative element: agents had to stand on switches to open them and communicate to get all agents and cows safely through. In that way we achieved some cooperation among agents and saw even more interesting games.

3.2 Agents on Mars

The *Agents on Mars* scenario [5] was used from 2011–2014. It turned out to be an important step in the contest's evolution, as it introduced many innovative features and increased the game's complexity. The map took the form of a weighted graph representing the surface of the planet Mars (we always based the scenario on a fictitious story). The agents represent *All Terrain Vehicles* of different kinds, and their goal in the game is to discover the best water wells by exploring the map and then to keep control of as many wells as possible. This was done by placing themselves in specific formations that ensure a covering of an area containing the wells while keeping rival agents outside.
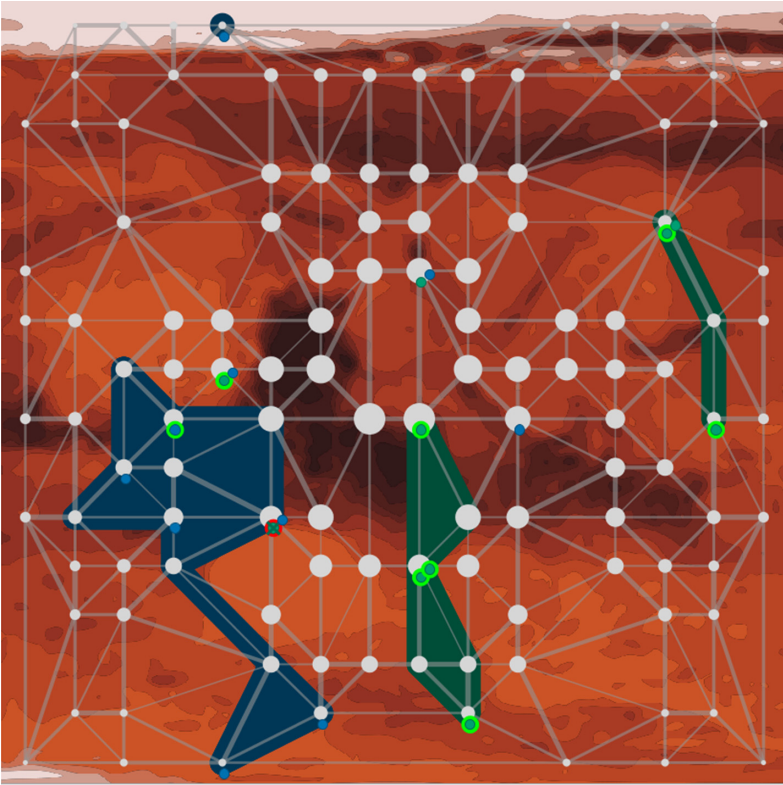

Fig. 4. The "Agents on Mars" scenario. (Color figure online)

In Fig. 4, one can see the basic graph layout, where the node sizes represent their value. The small circles at some nodes are the agents and the colored parts of the graph are currently taken by the team of the respective color.

The new agents were much more complex entities than in the previous scenarios: they had a rich set of actions to choose from, in contrast to just moving around the map. Furthermore, they dealt with a set of internal parameters that could vary through the simulation—*Energy*, *Visibility Range*, *Health* and *Strength*.

The evolution in the complexity of the scenario has remained on par with the evolution of multi-agent programming technologies used by the participating teams. A good quality of the teams has been reached and resulted in interesting games. Unlike previous scenarios, a (simple) strategy that works against each and every rival has not been discovered. However, it became difficult to further evolve the scenario. Also, it was a rather abstract problem.

3.3   The City Scenario

Our previous scenario, pictured in Fig. 5, was first used in 2016 [2] and improved two times for the editions of 2017 [3] and 2018 [4]. We started with two teams of 16 agents each moving through the streets of a digital city backed by realistic street graph data from *OpenStreetMap*[8]. The number of agents was then increased to 28 and 34 per team respectively.
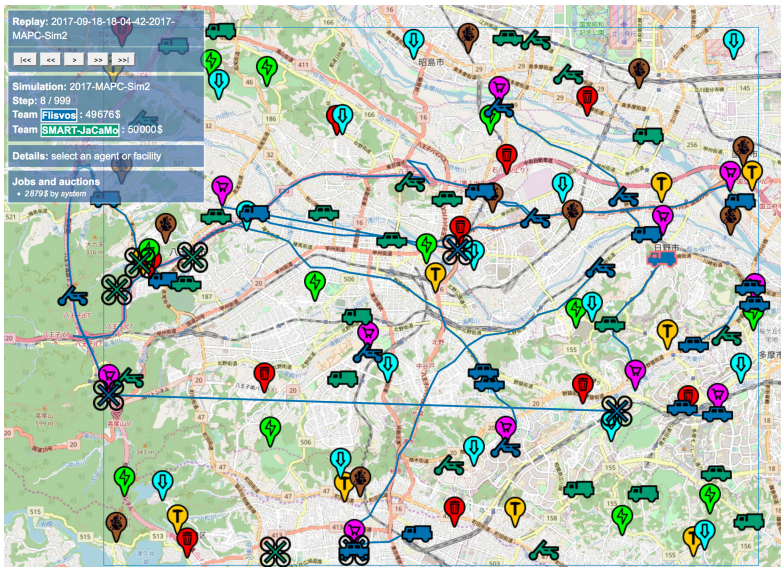

Fig. 5. The "Agents in the City" scenario.

Each team's goal was simply to earn as much money as possible by completing randomly generated jobs. These jobs required the agents to move around the

---

[8] https://www.openstreetmap.org.

city, buy certain kinds of items, cooperatively assemble these items to get new item types and finally deliver the finished products to a predefined target location. Most of these jobs were available for both agent teams simultaneously and rewarded on a *first come, first served* basis, allowing for more direct competition.

Each agent had one of four distinct roles, which characterized its movement type (air- or road-bound) and speed, as well as its maximum battery and carrying capacity. As is tradition, the number of agents was increased for each scenario to provide a greater challenge of coordination and require some more computational effort. Different agent roles were first introduced with the *Agents-on-Mars* scenario. The roles differed by certain key attributes as well as by which action was usable by which agent.

Compared to our previous scenarios, this one required more coordination and planning among agents of the same team. Some jobs are more profitable than other co-occurring jobs. Once agents are able to identify good jobs, the real challenge is the coordination of which agent secures which items from where in order to strike a good balance between time efficiency and money spent.

For the third instance of the scenario, we added a new *well* facility that teams could build and opposing teams could dismantle. To build wells, some funds had to be spent which could again be acquired by completing jobs. The wells would then generate points for as long as they existed. This change was intended to increase interaction between the teams and make the agents' actions more visible to human observers.

Lessons Learned in the City. The first run in 2016 has shown once again that participants have to be coerced into using specific features of the scenario: for example, we had to make cooperative assembly mandatory in 2017.

For the second run in 2017, we noticed a problem with the many parameters controlling the random generation of simulation instances. Finding *good* sets of parameters was not an easy task and required considerable testing. Also, for the first time we experienced that a scenario should allow for a simple naive (but far from being optimal) solution to be quickly producible. This scenario instead required considerable agent programming work before first results could be seen.

Another downside was that the visualization did not (or could not) show everything that was going on in an easily discernible way. For example, it is very impractical to display for all agents which items they are currently transporting. To amend this a little, the wells were added in 2018 to have an element that plainly shows how well a team is doing aside from the current money value.

Also, interaction between the teams was very limited and only indirectly given through the availability of shared resources (i.e. items in the shops) and the competition to get a job done first in order to receive the reward. The wells were also added to have a new entity that agents of both teams could and needed to interact with.

4   2019: Agents Assemble

After having played the *City* scenario for three consecutive years, it was once again time to come up with a fresh scenario and apply the lessons learned. We wanted to address some of the issues with the previous scenario, like visibility of agent behavior, while keeping many of the factors that made it interesting.

4.1   Scenario

In the new *Agents Assemble* scenario, as the name suggest, agents again have to construct complex structures from base objects. We switched from the map- (or graph-)based environment back to a "simple" grid structure with obstacles (see Fig. 6), comparable to the *Cow* scenario. The agents have to explore the grid to find `blocks` which also occupy one cell of the grid. Each agent has four "arms", one to each side, which can be used to pick up or connect to blocks. Blocks which are connected to an agent move in the same direction as the agent. Two adjacent blocks can also be connected to each other by two agents from the same team, when each agent is holding one of the blocks.
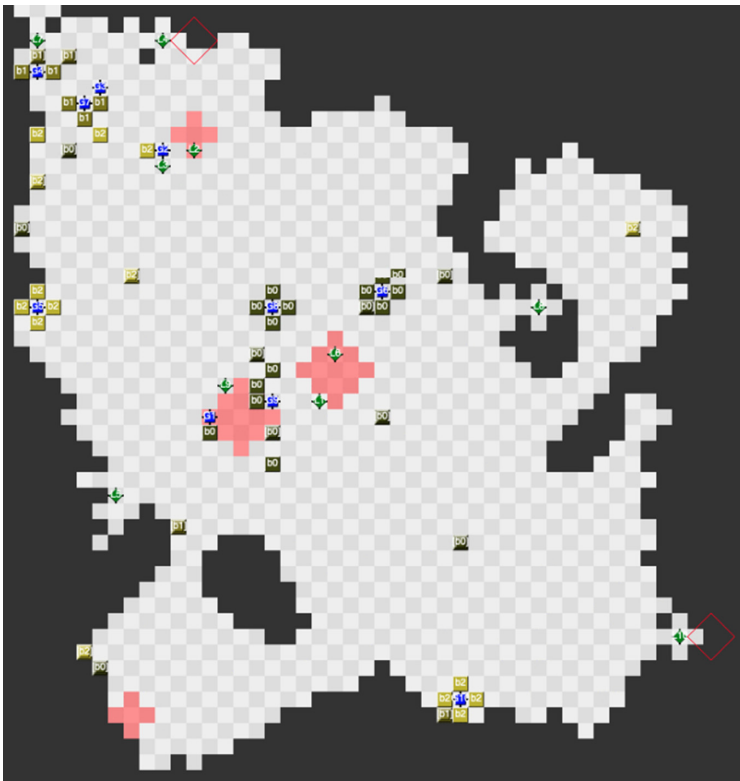

Fig. 6. MAPC2019 environment. Agents possess a local view of it and are required to assemble complex shapes to be delivered.

The system then randomly creates `tasks`, which the agents have to complete to earn reward points. The team with the most points at the end of the simulation will be the winner. Each task basically describes a structure or formation of blocks that the agents have to create. We depict an example of tasks in Fig. 7. Once the shape is assembled, the agents can deliver it to one of the goal zones to receive the points.
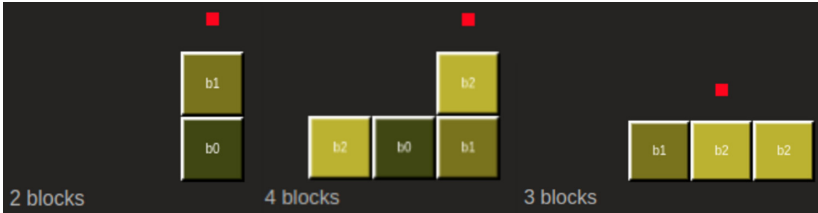

Fig. 7. Some examples of tasks in which the delivery agent should be carrying the blocks at the red dot position depicted in the figure. (Color figure online)

Actions. The agents have different actions for moving around in the grid. They can move one cell in each of the four main directions per step or rotate 90 degrees. This rotation might be handy, if the agents have blocks attached. Further, there is an action to retrieve blocks from `dispensers`, which are placed in random locations and provide one specific type of block. To work with blocks, the agents have actions for attaching and detaching things to their sides and as mentioned before, two agents can use the `connect` action to join two blocks together. An agent can also break this connection between two blocks, if the blocks are attached to the agent (directly or indirectly).

To interact with the environment and other agents, the `clear` action was added. It targets a single cell within the agent's vision radius (up to 5 cells in Manhattan distance) and has to be "charged", i.e. executed a certain number of times for the same target cell before it has an effect. Once it resolves, if the cell contained an obstacle or block, these will vanish and leave an empty cell. If instead an agent occupied that cell, it will be disabled. In that case, this agent will not be able to execute actions for a certain number of steps and also, all of its attached blocks (if any) will not be attached to the agent anymore. To give each agent a chance to avoid this, the target cells to be cleared have a perceivable marker after each `clear` action, i.e. also while charging.

As always, each action has a number of specific failure codes, indicating the reason why the action could not be executed.

Perception. One of the novelties of this scenario is that agents only perceive relative coordinates. That is, at the beginning the agents cannot know where they are. Due to their limited vision range of five cells in each direction, they

do not even know where they are relative to each other and have to find their teammates first.

This might favor solutions, where a local agent perspective is taken, rather than centralized approaches.

Dynamic Environment. To give the agents an even greater challenge, the environment dynamically changes during the matches. This makes harder for the agents to remember if they have already been at a place and requires more adaptability.

During each game, a number of `clear events` will occur. These work almost exactly like the `clear` action, only they affect a bigger region of the grid and after each event, new obstacles will appear randomly distributed around the center of the event.

Blocks and Visibility. One drawback of the *City* scenario was that it was not very interesting to watch, because most of the action did not happen in the environment. When agents bought items, these just went to their inventories. The current possessions of an agent could be displayed in a list, but it was rather difficult to keep track of multiple agents at once, not to mention all of them. Thus, in the new scenario, items (i.e., `blocks`) have received a more tangible representation, taking up considerable space in the environment. This leads to more interaction between agents and items and all of it is easily observable by human bystanders. What's more, carrying assembled shapes around becomes even more of a challenge, as the number of available routes possibly decreases.

4.2   Participants

This year, we had four teams participating in the Contest.

FIT BUT. The team from Czech Brno University of Technology consists of three people and participated in the Contest for the first time. The agents are implemented in plain Java.

GOAL-DTU. The team from Technical University of Denmark has already participated in the MAPCin one form or another for many, many years and has never missed a Contest since. As the name suggests, the agents were implemented using the GOAL [13] agent language.

LFC. The team LFC, from University of Liverpool, used JaCaMo [8] to implement its agent team. An additional fast downward planning component was developed to support the agents.

TRG. The single-person team TRG from the Canadian Carleton University also participated in the Contest for the first time. The agents were implemented with the Jason [10] framework.

An overview of the teams is listed in Table 1.

As we can see, this year, all approaches involve Java at some level. *FIT BUT* uses Java directly, while *TRG* uses Jason, which is implemented in Java. JaCaMo,

Table 1. Team overview

	FIT BUT	GOAL-DTU	LFC	TRG
System	Java	GOAL	JaCaMo	Jason
Team size	3	3	3	1
Time invested (in h)	300–400	200	200	500–600
Previously participated	No	Partly	Partly	No
Lines of Code (LOC)	5500–6300	1000	6800	9700
Started	August (29th)	August	May/September	May/mid-July
Total Score	1760	330	1790	590
Points	15	10	22	5
Win/Draw/Lose	5/-/4	3/1/5	7/1/1	1/2/6
Ranking	2	3	1	4

as used by *LFC*, in turn leverages Jason for implementing the agent reasoning. Lastly, *GOAL-DTU* uses GOAL, which is also implemented in Java.

Additionally, all teams are using an approach based on or at least somehow related to the BDI model [11], where agents' knowledge is represented in terms of **b**eliefs, agents have some **d**esires, or goals they want to achieve, and **i**ntentions, representing what an agent has elected to do. Jason, also as part of JaCaMo, is a platform for creating BDI agents. *FIT BUT* on the other hand used Java to implement their own system inspired by the BDI model. Lastly, cognitive agents implemented in GOAL also have beliefs and desires and the concept of intention also finds (informal) representation.

The teams are of similar size, except for *TRG*. Notably though, the single-person team has invested the most time. Of the four teams, *TRG* and *FIT BUT* are completely new to the Contest, while some members of *GOAL-DTU* and *LFC* had already participated before. We also note that the GOAL solution is particularly small in terms of LOC, while the Jason-based solution is a bit larger than the average.[9]

LFC and *TRG* started their initial work in May, than letting it rest until starting for real in September and July respectively. *FIT BUT* and *GOAL-DTU* both started to work in August.

4.3 Tournament and Results

In the final tournament, each team plays one match against each other team, where one match consists of three simulations with different parameters. Thus, with four teams, each team had to play 9 games. Winning a simulation is awarded

---

[9] This does not necessarily tell us anything about GOAL or Jason though.

with three tournament points, while a draw means one point for each team. The best result a team can achieve is 27 points.

We had the teams play simulations with three different sets of parameters, so that they were less likely to optimize their systems to one particular setting. Each simulation ran for 500 steps and 10 agents per team. In the second simulation, more complex tasks, with up to 5 required blocks instead of 3, were offered. In the third simulation, we increased the chance of a random `clear event` happening from 4% to 8%, leading to a more uncertain environment.

The results are also listed in Table 1. The Contest was won by the JaCaMo-based solution from Liverpool's *LFC*, with only one loss against *GOAL-DTU* and one draw against *TRG* out of 9 games, resulting in 22 tournament points. Runner-up is *FIT BUT* with 15 points, while *GOAL-DTU* achieved 10 and *TRG* 5 points. We note that each team won at least one simulation, and never only because the other team failed completely. All teams presented a workable solution.

Strategies. No team found a *strategic* advantage over the others. That is, we did not see a particular strategy being used to great effect. While the agent teams approached the problem in different ways, none of these were clearly superior to all others.

The Contest winner, *LFC*, implemented a strategy, where one agent was always waiting in a goal zone for its team members to deliver exactly the blocks needed for a particular task. We saw each agent always carry at most one block at a time. The shape required for the task was always assembled together with the agent waiting in the goal zone, who then submitted the task upon its completion. One advantage of *LFC* was clearly the capability to "dig" straight lines through obstacles with repeated `clear` actions. This technique was also used by the agents at the start of each simulation, probably to find the actual boundaries of the grid environment (which was always surrounded by a wall of obstacles). *LFC* implemented dynamic roles, where agents would start as explorers and later switch to specializations, e.g. assembling agents waiting in the goal zones.

FIT BUT in contrast had their agents meet somewhere on their routes to connect their blocks. Thus we saw *FIT BUT* agents walking around with complex shapes attached, which also worked very well.

GOAL-DTU agents could always be recognized by them proactively requesting as much as four blocks at a time and subsequently moving with four blocks of one type attached. While this ensured that they always had enough blocks at their disposal, it made it more difficult to navigate the map, especially during the late game when `clear events` could have already created narrow paths.

TRG alone tried a hybrid strategy. While some agents were coordinating to complete tasks, the other agents were trying to "defend" each goal zone by using `clear` actions on approaching opponent agents. This was an interesting decision, which unfortunately did not pay off so well, as the agents from the other teams were mostly able to circumvent these interventions. These roles were also statically assigned and did not depend on the current situation.

While we always try to build and configure the scenario in such a way, that no single best (and maybe even simple) strategy exists, we can never be sure that we succeed in this. Generally, the more features and interactions among them a scenario has, the harder it becomes to balance all of these features "correctly", so that no single feature can be used in an unforeseen way. Thus, in the new scenario, the rules governing the simulations were kept as simple as possible.

4.4 Interesting Simulations

In this section, we want to take a look at some interesting simulations[10] to see how the teams compare to each other under similar circumstances. All replays are also accessible from the contest overview page[11].

2nd Simulation of *GOAL-DTU* vs. *LFC*. Of course this simulation[12] might be interesting since it was the only one that *LFC* lost. The final score was 130 to 40 for *GOAL-DTU*. If we look at the completed tasks, we see that *GOAL-DTU* was already able to submit a task in step 78, which yielded 90 points, since the required shape consisted of three blocks. After this however, for more than 200 steps "nothing" happens. The next task is completed, again by *GOAL-DTU*, in step 317, netting 40 points for a two-block shape. *LFC* only completes one task, in step 351, receiving the 40 points. After this, no further tasks are completed. So, one question is surely what did *LFC* do before step 351. Reviewing all other simulations of *LFC*, the agents were always able to complete their first task around step 200. In step 191, we find instead the situation depicted in Fig. 8.

Fig. 8. Step 191: *LFC* clearing their blocks. (Color figure online)

---

[10] Not saying that some of the simulations were not interesting!.
[11] https://multiagentcontest.org/2019/.
[12] https://multiagentcontest.org/replays_2019/?2019-10-16-17-22-57-1571239377138_GOAL-DTU_LFC.

The green *LFC* agents (diamond-shaped, labels start with an "L", located at the top and right of this clipping), are all charging a `clear` action (red diamond markers on the grid) to remove the block they have currently attached. From their usual strategy, we conclude that their plan was to attach blocks to agent `L7`, who was already waiting in the goal zone (the red filled diamond shaped area near the center of the image; the smaller red diamond outline marks the clear action). If we go back in time, we see that the *GOAL-DTU* agents are and have been very active in this region, carrying lots of blocks, as always. First, this makes it very difficult for *LFC* to get their blocks to the `L7` agent. Secondly, the *LFC* agents decide to abandon their whole plan, even clearing all blocks they have already gathered. If we assume that *LFC* has to start anew (minus some initial discovery and exploration), we might indeed expect the next task to be completed after another 100 to 150 steps, which proved to be the case.

2nd Simulation of *LFC* vs. *TRG*. In this simulation[13] both teams were not able to score. This is quite surprising since *LFC* scored in all simulations but this one. Moreover, *LFC* was the team that scored most in the contest: 1790 points. Considering only simulations against *TRG*, it scored 180 in the first, and 210 points in the third. The question is: why did *LFC* perform much better in those other simulations?

To understand that, we need to look at *TRG*'s strategy. They always seek to position agents in the goal zone to disable any agent that enters that place. Nevertheless, it does not always work. At some times, some goal zone receives no *TRG* agents. As *LFC*'s strategy is to always choose a single goal zone to be used, at the second simulation, both strategies have collided. Every time *LFC*'s agents tried to deliver a task, a *TRG* agent was there to disable them. An example of this event is depicted in Fig. 9.

Fig. 9. The exact moment (step 174) when a *TRG* agent disables *LFC*'s agents.

---

[13] https://multiagentcontest.org/replays_2019/?2019-10-16-15-39-17-1571233157086_LFC_TRG.

This strategy seems to break the *LFC* team, because once their delivery agent is disabled, the whole team restarts to clear blocks and search for the grid's boundaries again. Whilst *LFC* is unable to deliver tasks, so is *TRG*. *TRG*'s strategy for preventing a team to score works pretty well, on the other hand, their agents were not able to coordinate themselves to form shapes required by tasks. At the end of the second simulation, no team scored a single point.

4.5   Survey Results

Traditionally, we conclude the Contest with a questionnaire that we ask each team to answer[14]. At this point, we give a brief summary of the answers from all teams. Parts of the survey results have already been used to create Table 1.

Regarding the motivation, practicing MAS development and in general learning more about agent technology were given as the main reasons for participating in the MAPC. This is aligned with MAPC's goal, in which in order to stimulate research in MAS, we need more people to learn and practice it.

We noticed that many teams mentioned that **debugging** capabilities were quite limited. Thus, they often resorted to "print(.)" as a debugging tool, i.e. adding logging statements and reading or searching the traces afterwards. Some teams stated that debugging was the most time-consuming task and some even developed their own (scenario-specific) tools that helped them to understand what was going on. Other time-consuming tasks include map navigation and merging the local views of the agents.

The most challenging aspects of the scenario according to the teams were:

– the dynamic environment;
– the local perspective of the agents; and
– coordinating agents to perform the synchronous actions.

From the survey, we also know that the teams barely added additional AI techniques to improve their solution (aside from *LFC* using a fast downward planner). This is probably due to the additional time investments required to add features to systems that are already quite complex within a limited timespan.

The main advantages of using agent technology were seen as *flexibility* and *modularity* of the system. From an agent programming perspective, agents should consider constantly the current state to select a proper action which may be a useful feature in dynamic environments. As the main drawbacks, teams cited the difficulty in debugging, a lack of portability and that it was very challenging to keep the system simple and easy to maintain.

Finally, if teams were to attend another time, they would like to improve error handling, reliability, coordination of their agents and their own debugging means.

---

[14] The reader will find these answers at the end of each team description paper.

5   Lessons Learned

The organization of our MAPCturned out to be quite work-intensive at times. Its technical implementation has been mainly done or supervised by (to this day seven) PhD students of the second author, in addition to a number of Bachelor and Master theses. However, the students also played a major role in crafting the scenarios and coming up with fresh and innovative ideas.

In the first phase of the MAPC, no real cooperation among agents was achieved. In fact *every man for himself* was a common strategy, completely against the paradigm of agent programming. Often the teams with the best working A^* path-finding algorithm won. Due to the fact that the participating agent languages were not yet mature enough, the main benefit of the contest in the early days was to serve as a debugging tool for the participating systems.

Indeed, low-level technical problems with the implementations of the agent languages often played a major role. This is in contrast to the second phase, where attention shifted to the scenario and higher-level concerns.

5.1   Agents Assemble Scenario

In the new scenario, we saw that forcing agents to work solely off their local perspectives and integrate their knowledge with other agents is a challenging task.

We once again note that it's desirable to have a problem that is easy to solve, but very difficult to solve well. In other words, it should be easy to come up with some agents that can play the game, while mastering it should require a lot of effort.

Aside from *TRG*trying to defend goal zones, we only saw limited conscious interaction between the teams. Unfortunately, our options to elicit interaction are also limited, because there is little motivation to cooperate in a zero-sum game. As such, it would only work if both teams are deceived to varying degrees. Another way would be some form of attacking, though we try to keep our scenarios as peaceful as possible. In this scenario, we had indirect interaction through presence in and modification of a shared environment, similar to the *Cow* scenario. In the *City* scenario, we had very limited interaction followed by the well-building attempt. In the *Mars* scenario, we had interaction through attacking agents, though the extent (duration and complexity) of these interactions also remained expandable. A challenge for the future is surely to design complex interactions which are interesting to realize and see in action.

5.2   General

A lesson of the early phase was the awareness that normally neglected *engineering* issues (as opposed to *scientific* ones) are of utmost importance. For example, collecting statistical data or providing visualizations turned out to be as important as the choice and the tuning of the scenarios. Without them it was extremely difficult to analyze why a team behaved as it did.

Using automatically generated statistical data, we can easily retrace a whole simulation's progress by looking at the generated charts instead of watching the whole replay. The charts mainly focused on scenario-specific data, like the development of the score or stability of dominated zones. Furthermore, we were finally able to directly and easily compare different simulation runs without having to keep a lot of details in mind. Such tools cannot only be used for debugging the teams' agents, but also for analysis of the scenario and improving it for the next round.

These insights went into the *Agents-on-Mars* scenario, where we noted an increasing number of multi-agent platforms. Since then, our scenarios have always been won by dedicated agent platforms—they seem to outperform "ad-hoc" solutions. This might be attributed to some teams taking part repeatedly, but it also points to an increasing maturity and ease of use concerning multi-agent platforms.

If you followed the years in which each scenario was used closely, you may have noticed that 2015 was missing. That particular year should have been the start of the "City Scenario". We introduced it in 2015, though we might have underestimated its complexity and readiness. As the contest date neared, the participants asked us to postpone the competition, which we did. It moved to early 2016 first and finally to the regular 2016 Contest slot. If we say complexity here, we mean the effort that was required to get a simple agent team running and dealing with all important stages of the game. We learned that, as often requested in the earlier years, we need to publish a completely new scenario as early in the year as possible. We also saw a relative core base of teams that participated in each of the three City contests. While this might tell us that teams who have made the big time investment once are likely to stick around, it is also off-putting for new teams if they have to put in a lot of work to only see some basic results. For the next scenario (Agents Assemble), we always had the concept of *easy to start, hard to master* at the back of our minds.

In order to better understand the underlying strategies of the teams, we worked out a standardized questionnaire [1] (which was further improved over the years). This did not only help to learn about the systems and the results they produced, but also to understand the whole development process. Additionally, it serves for newcomers to avoid mistakes from previous iterations.

The motivation to enter the MAPCwas for some teams simply to learn about multi-agent systems or to refine programming skills. Furthermore, most teams shared our goal of evaluating multi-agent frameworks and platforms. Regarding their structure, teams were composed of students as well as researchers with their background mostly in MAS or at least in artificial intelligence in general. This reinforced our motivation to always come up with new scenarios, rather than optimizing a particular one over the years, which only favors teams that attend each and every year (it seems this happened in the simulation league of the soccer competition).

We also asked the teams how difficult it was and how much effort had to be put into getting to a point where their system behaved as it finally did. We

got very diverse results, ranging from only a hundred to over a thousand hours and 1000 to many thousands of lines of code that had to be written, tested and debugged. This clearly hints at varying levels of usability concerning different agent platforms.

Furthermore, teams noted that they not only debugged their agents but found and fixed bugs in the agent framework or platform they used as well, which shows that the MAPCcan play a major role concerning the development and evaluation of different platforms. Nevertheless, the teams are still not satisfied with available state-of-the-art debugging tools, since it requires a lot of effort to debug even 20 agents, each with its own individual mindset.

We realized that the visualization and playability of the respective scenario is a key to reaching a broader audience, especially students, e.g. when *MASSim* is used in teaching in various courses all over the world. The competitive nature is fun for the students and this feature should never be underestimated. To this day, we cannot make out a clear correlation between the specific scenario and the number of participants. A more important factor is usually whether the interested teams are able to invest the necessary time. Similarly, the scenario doesn't seem to have a big impact on the choice of programming language or framework. Most teams either choose a framework they are already familiar with or one they want to learn, but have already heard about (from colleagues or supervisors).

Finally, coming back to our questions raised at the beginning of this paper, namely about the situations where using agent technology pays off and the strengths of agent platforms. For one, we note that the top-performing approaches are usually agent-based. Nonetheless, we have also seen conventional approaches achieve remarkable results. In situations where it is easier to take a global perspective (e.g. the City scenario), conventional approaches or even *centralized* solutions in general, usually seem to have it easier compared to situations where agents have to base their decisions mostly on their local information. Over the years, there has been no clear indication of whether agent-based solutions take less time to create or are smaller in nature. The teams that have used an agent-based approach tend to report their overall satisfaction with their chosen technology though. In the end, it may even be a surprising result that both paradigms almost see eye to eye in our test cases, as the conventional paradigms have been developed over much more time, by vastly more people and see usage that isn't even comparable to agent-oriented programming. If agent programming had a comparable maturity and similarly sophisticated tools, and if people were trained in its use as they are in traditional programming, we might see way shorter development times and in turn even better results.

6   Challenges and Outlook

While the agent paradigm plays an important role in computer science, its uptake in industry still remains small. We believe that the MAPCplays some role in determining under which conditions agent languages can be used in practice.

The ideal scenario we were (and still are) looking for, should be easily testable and not be based on difficult rules (only the solutions should be difficult), so that beginners in the area of agent programming can easily take on the challenge. A good solution should use cooperation among autonomous agents and be flexible enough so that different groups of agents evolve and work together to solve intermediate tasks.

After almost 15 years of research and experience, we still have not found such a *convincing scenario*. Nor have we yet proved that agent-based approaches are clearly superior to other, sometimes even ad-hoc, approaches using traditional programming languages. In many areas of computer science, one is often looking for a *killer application*. However, it may well be that such killer applications do not exist. In defense of MAS, there are many potential advantages that the contest is not evaluating at all, because it does not seem feasible in the context of the competition: reusability, maintenance, correctness, the possibility to model-check agents, code running on different platforms, etc.

Regardless, there are good reasons to be optimistic, because there is progress on two sides. First, multi-agent programming technologies are becoming more and more capable. Secondly, there are many lessons learned throughout the history of the contest and we are getting better at encouraging the cooperative behaviors we want to see in agents.

So what are possible ways to improve our MAPC? We are considering the idea to let *more than two* teams participate in the same simulation. The current scenario would provide for this naturally, however, we need to address better visualization (too many things happening at the same time) and evaluation (to easily find interesting situations and emergent behavior) first.

Our ultimate vision is an agent platform that allows to deploy agents written in very different agent languages, using the specific features of them. For example, it might be beneficial for BDI agents to solve very efficiently certain tasks, whereas planning agents based on some form of *hierarchical task nets* could do the planning for them. Being able to re-use agents already developed (and based on different paradigms) would certainly push the envelope for applications of multi-agent systems in general.

Agents running on a local platform (rather than participating over the Internet) would also allow more fine-grained control over communication and real-time aspects. We could then consider *many* agents, not just a few, but hundreds or thousands of *sophisticated* agents — traditional approaches do not seem to perform well in such a situation. Moreover, with many *interacting agents* we might see some interesting behavior evolve.

However, the price to pay is to standardize the communication and set up common protocols and interfaces for such agents. That would change our contest drastically.

Also, we extend our gratitude to all anonymous reviewers who helped us make improvements to this paper.

The fourth author acknowledges that his part in this study was financed by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

Last but not least, we thank all teams that entered our Contest in the last 15 years and helped to keep it alive. The second author sends a special thank you to some former PhD students in his group: Peter Novák, Tristan Behrens, Federico Schlesinger, and Michael Köster. Without their ideas and enthusiasm, the MAPCwould not have flourished.

References

1. Ahlbrecht, T., et al.: Multi-agent programming contest 2013: the teams and the design of their systems. In: Cossentino, M., El Fallah Seghrouchni, A., Winikoff, M. (eds.) EMAS 2013. LNCS (LNAI), vol. 8245, pp. 366–390. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-45343-4_22
2. Ahlbrecht, T., Dix, J., Fiekas, N.: Multi-agent programming contest 2016. Int. J. Agent-Oriented Softw. Eng. **6**(1), 58–85 (2018)
3. Ahlbrecht, T., Dix, J., Fiekas, N.: Multi-agent programming contest 2017 - the twelfth edition of the MAPC. Ann. Math. Artif. Intell. **84**(1), 1–16 (2018). https://doi.org/10.1007/s10472-018-9594-x
4. Ahlbrecht, T., Dix, J., Fiekas, N. (eds.): MAPC 2018. LNCS (LNAI), vol. 11957. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-37959-9
5. Ahlbrecht, T., Dix, J., Köster, M., Schlesinger, F.: Multi-agent programming contest 2013. In: Cossentino, M., El Fallah Seghrouchni, A., Winikoff, M. (eds.) EMAS 2013. LNCS (LNAI), vol. 8245, pp. 292–318. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-45343-4_16
6. Behrens, T., Dastani, M., Dix, J., Köster, M., Novák, P.: The multi-agent programming contest from 2005–2010. Ann. Math. Artif. Intell. **59**(3–4), 277–311 (2010). https://doi.org/10.1007/s10472-010-9219-5
7. Behrens, T.M., Hindriks, K.V., Dix, J.: Towards an environment interface standard for agent platforms. Ann. Math. Artif. Intell. **61**(4), 261–295 (2011). https://doi.org/10.1007/s10472-010-9215-9
8. Boissier, O., Bordini, R.H., Hübner, J.F., Ricci, A., Santi, A.: Multi-agent oriented programming with JaCaMo. Sci. Comput. Program. **78**(6), 747–761 (2013)
9. Bordini, R., Dastani, M., Dix, J., Segrouchni, A.E.F. (eds.): Special Issue on Multi-Agent Programming. J. Auton. Agents Multi-Agent Syst. **23** (2011). http://www.springerlink.com/content/y3621v0p2217h683/. Springer
10. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-agent Systems in AgentSpeak Using Jason, vol. 8. Wiley, Hoboken (2007)
11. Bratman, M., et al.: Intention, Plans, and Practical Reason, vol. 10. Harvard University Press, Cambridge (1987)
12. Genesereth, M., Love, N., Pell, B.: General game playing: overview of the AAAI competition. AI Mag. **26**(2), 62 (2005)
13. Hindriks, K.V.: Programming rational agents in GOAL. In: El Fallah Seghrouchni, A., Dix, J., Dastani, M., Bordini, R.H. (eds.) Multi-Agent Programming, pp. 119–157. Springer, Boston, MA (2009). https://doi.org/10.1007/978-0-387-89299-3_4

14. Johnson, M., Jonker, C., van Riemsdijk, B., Feltovich, P.J., Bradshaw, J.M.: Joint activity testbed: blocks world for teams (BW4T). In: Aldewereld, H., Dignum, V., Picard, G. (eds.) ESAW 2009. LNCS (LNAI), vol. 5881, pp. 254–256. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10203-5_26
15. Karakovskiy, S., Togelius, J.: The Mario AI benchmark and competitions. IEEE Trans. Comput. Intell. AI Games **4**(1), 55–67 (2012)
16. Ketter, W., Collins, J., Reddy, P.: Power TAC: a competitive economic simulation of the smart grid. Energy Econ. **39**, 262–270 (2013)
17. Vallati, M., Chrpa, L., Grześ, M., McCluskey, T.L., Roberts, M., Sanner, S., et al.: The 2014 international planning competition: progress and trends. AI Mag. **36**(3), 90–98 (2015)
18. Visser, A., Ito, N., Kleiner, A.: RoboCup rescue simulation innovation strategy. In: Bianchi, R.A.C., Akin, H.L., Ramamoorthy, S., Sugiura, K. (eds.) RoboCup 2014. LNCS (LNAI), vol. 8992, pp. 661–672. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-18615-3_54