

An application of Answer Set Programming in Distributed Architectures: ASP Microservices

Stefania Costantini

Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica
Università degli Studi di L'Aquila, L'Aquila, Italy
stefania.costantini@univaq.it

Lorenzo De Lauretis

Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica
Università degli Studi di L'Aquila, L'Aquila, Italy
lorenzo.delauritis@graduate.univaq.it

We propose an approach to the definition of microservices with an Answer Set Programming (ASP) 'core', where microservices are a successful abstraction for designing distributed applications as suites of independently deployable interacting components. Such ASP-based components might be employed in distributed architectures related to Cloud Computing or to the Internet of Things (IoT).

1 Introduction

An important hot topic is, in our opinion, that of defining software engineering principles for Answer Set Programming (ASP). ASP is a successfully logic programming paradigm (cf. [5] and the references therein) stemming from the Answer Set (or "Stable Model") Semantics of Gelfond and Lifschitz, and based on the programming methodology proposed by Marek, Truszczyński and Lifschitz. ASP is put into practice by means of effective inference engines, called *solvers*¹. ASP has been widely applied, e.g., to information integration, constraint satisfaction, routing, planning, diagnosis, configuration, computer-aided verification, biology/biomedicine, knowledge management, etc. .

It is important in our view to have the possibility of defining significant subprograms as *components* to be possibly distributed over different nodes of a network. Both the connections between components and the ways of exchanging information should be clearly defined. Our approach is inspired by the *microservices* architectural abstraction, which can be described as a particular way of designing distribute software applications as suites of independently deployable interacting services (cf. for instance the survey [26] or see <https://martinfowler.com/articles/microservices.html#CharacteristicsOfAMicroserviceArchitecture>). A microservice is indeed a component, as it is a unit of software that is independently replaceable and modifiable: in fact, it intended as a self-contained piece of business functionality with clear interfaces that can be accessed by the "external world". This kind of architectural abstraction enables distribution, as each microservice is meant to be executed as an independent process, and heterogeneity, allowing for different services to be written in different programming languages. Microservices are a suitable architectural abstraction for the Internet of Things (IoT): a microservice may incapsulate a physical objects, where service inputs and/or outputs can possibly be linked to sensors/actuators. Since microservices are by their very nature heterogeneous, open

¹Many performant ASP solvers are freely available, a list of them is reported at https://en.wikipedia.org/wiki/Answer_set_programming.

issues are: the way microservices communicate with each other (synchronous, asynchronous, which is the message format, etc.); and, the protocols used for the communication.

Microservices in real distributed software architectures and in cloud computing are usually deployed via *lightweight containers*. A container in software engineering is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably and can be seamlessly transferred from one computing environment to another. A widely used tool to create containers is *Docker*, available in the form of an open source Docker Engine (see <https://www.docker.com>). A Docker container image consists in a lightweight, standalone, executable package of software that includes all elements needed to run an application: code, runtime support, system tools, system libraries and settings.

Along this line, we propose ASP microservices that might be blended into heterogeneous systems, and even into Multi-Agent-System (MAS) since each such component may be seen as a reactive agent. They could in perspective be employed in cloud computing, and in IoT, including robotic applications. In this paper we discuss a first idea about how these components, that we call μ ASP-Services (μ ASPSv's), can be specified, how their interfaces to the "external world" can be defined, and how they should procedurally behave. In fact, a μ ASPSv is meant to be based upon a 'core' ASP program whose activities, however, should be triggered by external stimuli/requests coming from some source, and whose results should be returned to the requesters. Therefore, the 'core' ASP program should be included into a container, that can be possibly realized via the Docker technology, which should also include: an interface, to provide the μ ASPSv with inputs, and to deliver the outputs; solving capabilities to compute the answer sets. For ASP, standalone versions of the most important solvers are nowadays available. So, a docker deployment for a μ ASPSv should include the source program, its 'execution shell', and the solver.

A small specimen of the proposed approach is represented in the following example, which is meant to be (a fragment of) the code of a *controller* component/agent, acting in the IoT. This piece of code might be in fact the ASP 'core' of a μ ASPSv. *test_ok* is the input coming from a sensor, with value 'true' if the controlled device is working properly, (otherwise the value is set to false).

```

device_ok ← test_ok.
device_fault ← not test_ok.
wait ← not wait, not sensor_input.

```

In this example, as a programming choice, inconsistency (due to the odd cycle over *wait*) is to be interpreted as a 'no-operation' controller state, where the component is waiting for the sensor outcome. It can be assumed that the sensor provides results at a certain frequency. The outcome, i.e., *device_ok* or *device_fault*, is to be delivered to whatever other components would ask for it.

More precisely, in order to work in a standalone way within a distributed system, an interface (or 'shell') will manage the ASP program, and in particular will perform the following functions. First, manage the inputs and outputs of the μ ASPSv: i.e., be able to detect input arrival and to dispatch the outputs according to the request coming from the μ ASPSv's external environment. In the above example, inputs can be: (1) queries over the device state for which an answer has to be delivered, and (2) sensor outcomes, which are to be considered as particular inputs which activate the module. In the general case, upon the arrival of inputs, the shell will: (i) add the inputs to the ASP program as facts; (ii) evaluate the answer sets of the resulting ASP program; (iii) according to previously-received requests, extract (from the answer sets) the answers and deliver them to the external environment. Notice that the shell, after delivering the outputs, will remove the last-added program facts so as to bring back the controller to the 'no-operation' state. In a 'stateless component', all inputs will be removed, while some of the inputs can be left if instead the component is meant to have a state.

In this paper we introduce a formal definition of ASP microservices and we outline a possible logic-based semantics of an overall heterogeneous distributed system encompassing such modules. The paper is in fact structured as follows. In Section 2 we introduce basic concepts about microservices. In Section 3 we recall (for the sake of completeness) the Answer Set Programming paradigm, and in Section 4 we briefly survey and discuss existing approaches to modularity in ASP. We introduce **our contribution** in Sections 5 and 7, i.e.: (1) how to define and implement μ ASPSv's so as to be able to get inputs and extract answers, and how the inner ASP program might be structured; (2) how to provide a formal semantics to a generic microservice architecture possibly encompassing μ ASPSv's. In Section 6 we discuss a small case study, developing a specific μ ASPSv which implements an intelligent agent managing a road intersection (i.e., a “virtual traffic light”). Finally, in Section 8 we conclude.

2 Background: Microservices

To better understand Microservices, let us introduce before the concept of “Service”. A Service, as a software component, is a mechanism to enable access to one or more software capabilities [1]. It provides other applications with stable, reusable software functionality at an application-oriented, business-related level of granularity using certain standards [29]. Service-Oriented Architecture (SOA) is a software architectural style that uses services as the main building component [1]. Key features of SOA are heterogeneity, standardization and evolvability of services.

Microservices can be seen as a technique for developing software applications that, inheriting all the principles and concepts from the SOA style, permit to structure a service-based application as a collection of very small loosely coupled software services [20].

A MicroServices Architecture (MSA), is an evolution of the SOA architecture, making the communication lighter and the software parts (Microservices) smaller. It can be seen as a new paradigm for programming applications by means of the composition of small services, each running its own processes and communicating via light-weight mechanisms. Key features of MSA are bounded scope, flexibility and modularity [19]. I.e., there is a clear definition of the data a microservice service is responsible for and is “bound to.” So, the service owns this data and is responsible for its integrity and mutability.

The work in [28] shows that a distributed MSA can easily fit into an IoT system. In particular, the set of microservices can be seen as a Multi-Agent-System, cooperating to realize all system functionalities.

At the current day, microservices are still a new and emerging paradigm, having building standards not perfectly defined and communication protocols that are not well specified: in fact, following one of the definitions of microservices [20, 19], they are small loosely coupled software services that communicate, possibly exploiting service discovery to find the route of communication between any two of them. In our work, we are proposing a new approach, that is μ ASP-Services, which are based upon an inner ASP program.

3 Answer Set semantics (AS) and Answer Set Programming (ASP)

“Answer Set Programming” (ASP) (cf. [5] and the references therein) is a successful programming paradigm based on the Answer Set Semantics. In ASP, one can see an answer set program (for short, just “program”) as a set of statements that specify a problem, where each answer set represents a solution compatible with this specification. Whenever a program has no answer sets (no solution can be found), it is said to be *inconsistent*, otherwise it is said to be *consistent*.

Syntactically, an ASP program Π is a collection of *rules* of the form

$$H \leftarrow A_1, \dots, A_m, \text{not } A_{m+1}, \dots, \text{not } A_{m+n}.$$

where H is an atom, $m, n \geq 0$, and each A_i , $i \leq m+n$, is an atom. Atoms and their negations are called *literals*. Symbol \leftarrow is often indicated as $:-$ in practical programming. The left-hand side and the right-hand side of the clause are called *head* and *body*, respectively. A rule with empty body is called a *fact*. A rule with empty head is a *constraint*, where a constraint of the form ‘ $\leftarrow L_1, \dots, L_n$.’ states that literals L_1, \dots, L_n cannot be simultaneously true in any answer set. Constraints are often rephrased as ‘ $f \leftarrow \text{not } f, L_1, \dots, L_n$.’ where f is a fresh atom. To avoid the contradiction over f , some of the L_i ’s must be false thus forcing f to be false, and this, if achieved, fulfills the constraint. There are other features that for the sake of simplicity we do not consider in this paper.

The answer set (or “stable model”) semantics (AS) can be defined in several ways (cf., e.g., [31], though more recently several other definitions have appeared in the literature). However, answer sets of a program Π , if any exists, are supported minimal classical models of the program interpreted as a first-order theory in the obvious way.

The ASP approach to problem-solving consists basically in the following: (i) encoding of the given problem via an ASP program; (ii) computing the “answer sets” of the ground program via an ASP solver (as a preliminary step, solvers perform the “grounding” of the program, by substituting all variables with the constants occurring in the program); (iii) extracting the problem solutions by examining such answer sets; in fact, answer set contents can be in general reformulated in order to present the solution in terms of the given problem.

A top-down (prolog-style) query answering device has been defined in [15] for RAS, where RAS is a variation of AS where every program admits answer sets. However, RAS and AS coincide over a wide class of programs (some sufficient conditions that identify classes of programs where the two semantics coincide are reported in [14]). Queries that have been introduced are, first of all, “ $?A$ ” asking whether A is true w.r.t. some answer set of given program Π . Other queries are the following: query “ $? \text{not } A$ ” asks whether A is false w.r.t. some answer set of Π , and therefore it succeeds if $\text{not } A$ is true in some of them (this implements the operator **not** introduced in [34]); query “ $? \text{not not } A$ ” asks whether $\text{not } A$ is false in some answer set, and therefore it succeeds if A is true in some of them, which corresponds to query “ $?MA$ ”, **M** standing for ‘possibility’ in the modal logic sense; query “ $? \text{not not not } A$ ” asks whether it is not true that A is false w.r.t. some answer set of Π , i.e., that A is true in all of them, which corresponds to “ $?KA$ ”, **K** standing for ‘knowledge’ in the modal logic sense; query “ $? \text{not not not not } A$ ” asks whether A is false in every answer set, meaning **Knot** A , i.e., $\text{not } MA$ (a new operator **NOT** is a shorthand for $\text{not not not } A$).

4 Background: Modularity in ASP

Existing approaches to modularization of ASP programs have been reviewed in [21], to which the reader may refer for a complete account. Reporting from there, such approaches can be divided into two lines: “programming-in-the-large”, where programs are understood as combinations of separate and independent components, combined by means of compositional operators; “programming-in-the-small”, in which logic programming is enriched with new logical connectives for managing subprograms.

Considering the programming-in-the-small vision: in [22], program modules are viewed as generalized quantifiers; [25] proposes templates for defining subprograms; [35] developed a declarative language for modular ASP, which allows a programmer to declaratively state how one ASP module can import processed answer sets from another ASP module. The work in [13] explores how to divide an ASP program into components according to its structure in terms of cycles.

Lifschitz and Turner’s “splitting set theorem” (cf. [32]), or variants of it, is underlying many programming-in-the-large approaches. The basic idea is that a program can be divided into two parts: a “bottom” part and a “top” part, such that the former does not refer to predicates defined in the latter. Computation of the answer sets of a program can be simplified when the program is split into such parts.

[27] defines the notion of a “DLP-function” which is basically a module for which a well-defined input/output interface is provided; a suitable compositional semantics for modules is introduced. [33] provides a simple and intuitive notion of a logic programming module that interacts through an input/output interface. This is achieved by accommodating modules as proposed by [24] to the context of Answer Set Programming. Full compatibility of the module system with the stable model semantics is achieved by allowing positive recursion to occur inside modules only.

[17] focuses on modular non-monotonic logic programs (MLP) under the answer set semantics, where modules may provide input to other modules. Mutually recursive module calls are allowed.

[3] defines modules in terms of macros that can be called from a program. [2] provides modules specification with information hiding, where modules exchange information with a global state.

In [23] a technique is proposed to allow an answer set program to access the brave or cautious consequences of another answer set program. [30] proposes “modular logic programs” as a modular version of ASP. This work consider programs as modules and define modular programs as sets of modules. The authors introduce “input answer sets”, which is the key semantic object for communication between modules.

[9] proposes to adopt ASP modules in order to simulate (within reasonable complexity) possibility and necessity operators. Such operators (given the underlying modules) are meant to be usable in ASP programs, but possibly also programs written under other programming paradigms.

It can be seen that none of the above approach tackles modularization in view of using ASP modules as standalone components in distributed systems. Therefore, our approach is a novelty in the landscape of the current literature.

5 μ ASPSv’s: Specification and Implementation Guidelines

In this section we provide an abstract definition of a μ ASPSv, and some more specific indication of how such a component might be enacted and inserted into a distributed system, and how the inner ASP program might be structured.

Definition 1. *Let Π be an ASP program, and let $U = \text{Undef}(\Pi)$. A μ ASPSv based upon Π , denoted as μ ASPSv(Π), has the following specification:*

- *Inner ASP program Π ;*
- *Activation signal A (optional), with $A \in \text{Undef}(\Pi)$;*
- *Stop signal S (optional), with $S \in \text{Undef}(\Pi)$;*
- *Input set $\{I_1, \dots, I_k\} \subseteq \text{Undef}(\Pi)$;*
- *Output set $\{O_1, \dots, O_h\} \subseteq \text{Heads}(\Pi)$.*
- *Query result set $\{Q_1 = v, \dots, Q_r = v\}$ where $\{Q_1, \dots, Q_r\}$ are queries² formulated over atoms occurring in $\text{Heads}(\Pi)$ and $v = \text{true/false}$.*

The various elements listed above have the following meaning.

²c.f. previous section for possible queries.

Whenever the activation signal is specified, if A is not true in Π , then $\mu\text{ASPS}_v(\Pi)$ is in a state of no-operation.

Whenever the stop signal is specified, if S becomes true in Π , then $\mu\text{ASPS}_v(\Pi)$ will go back into a state of no-operation.

The input set is a set of atoms that, when some of them are added to Π , contribute to answer sets computation. Each of such atom corresponds to an input/request received from the μASPS_v 's surrounding environment.

The output set is a set of atoms extracted from the answer sets of Π plus the current input set. Each of these atoms corresponds to an output/answer to be delivered into the μASPS_v 's surrounding environment.

The query result set is a set of truth values elicited from the answer sets of Π . Each of these values corresponds to result of a query, to be delivered into the μASPS_v 's surrounding environment.

In order to make it possible for $\mu\text{ASPS}_v(\Pi)$ to operate dynamically, thus receiving inputs and delivering outputs and answers, a suitable *shell* program must be defined, in any programming language able to be interfaced with an answer set solver. Below we provide a schematic definition of such a shell program, to be used as a guideline for actual definition and implementation.

Definition 2. *The shell responsible to manage an ASP microservice $\mu\text{ASPS}_v(\Pi)$ can be specified by the following pseudo-code.*

```

begin
while not activation then wait endwhile;
if activation then add atom  $A$  to  $\Pi$  as a fact to bring it into operation;
while not stop do at frequency  $f$ 
    detect and annotate actual inputs  $\{I_{j_1}, \dots, I_{j_r}\} \subseteq \{I_1, \dots, I_k\}$ ;
    add  $\{I_{j_1}, \dots, I_{j_r}\}$  to  $\Pi$  as facts;
    obtain the answer sets  $\{S_1, \dots, S_n\}$  of (the augmented)  $\Pi$ ;
    elicit outputs  $\{O_1, \dots, O_v\} \subseteq \{O_1, \dots, O_h\}$ ;
    extract query results  $\{Q_1, \dots, Q_t\} \subseteq \{Q_1, \dots, Q_r\}$ ;
    deliver outputs and query results according to requests;
    remove  $\{I_{v_1}, \dots, I_{v_s}\} \subseteq \{I_{j_1}, \dots, I_{j_r}\}$  from  $\Pi$  and remove relative annotations;
endwhile;
add atom  $S$  to  $\Pi$  as a fact to bring it into no-operation;
end.

```

This shell program is able to activate and stop a μASPS_v , and to execute, until possibly a stop signal arrives, a loop where the inputs are received from the external environment and delivered to Π , and outputs and query results are extracted from the answer sets of Π (given the inputs) and delivered to the external environment. Precisely, each input will arrive from some other component, to which the output will have to be delivered. The shell program will rely upon an input-output table, where each input, expected output and related component will be annotated. Notice that at the end of each cycle some or all of the inputs will be removed from Π and the relative annotations will be eliminated (removing all inputs determines a stateless component, while omitting to remove some of the inputs, forever or for some time interval, accounts to defining a stateful component). Input detection will occur at a certain frequency, suitable for each particular kind of component, environment, and application domain. Some of the inputs may come from sensors (and therefore they do not require any answer) and some of the

outputs may go to actuators. This is also annotated in the input-output table. Notice also that the parts concerning the activation and stopping of the μ ASPSv (first and second line after the **begin**, and last line before the **end**) will be omitted if the component is running forever rather being first activated and then stopped.

Many practical aspects remain to be defined in order to obtain an implementation. For instance, if a μ ASPSv is to be situated within a multi-agent system, input-output-query exchange might happen via FIPA (cf. <http://www.fipa.org> for language specification, syntax and semantics), a widely used ACL (Agent Communication Language). The shell program can be made FIPA-compliant (i.e., able to exchange and understand FIPA messages) for instance by importing the freely available JADE library (cf. <https://jade.tilab.com/> where references to several related publications can also be found). The JADE library is an advance middleware that offers many functionalities to “agentify” imperative or object-oriented programs. In fact it provides: the agent abstraction (i.e., a given program, when running, is seen by the external environment as an agent); the ability of peer to peer inter-agent FIPA asynchronous message-passing; a yellow pages service supporting subscription of agents and a discovery mechanism, and many other facilities to support the development of distributed systems.

So for instance, an input can be sent to a μ ASPSv via a FIPA *request* message with the input as argument, to be interpreted on the μ ASPSv’s side as a request to reply with a *confirm* message containing the corresponding output. A query can be sent to the μ ASPSv via a *query-if* message whose answer will be a *confirm*, conveying the truth value of the query. Notice that, to avoid ambiguities, the FIPA syntax provides the facility to identify each message via a certain arbitrary identifier, so that the answer message can indicate that it is ‘in-reply-to’ to that identifier.

The JADE yellow pages services might be exploited by μ ASPSv’s which would want to register as agents with a name and a role, and then communicate with each other in an asynchronous way. Or, since most MAS offer such a mediator service, μ ASPSv’s might enroll in any known agent community. Finally, they might communicate peer-to-peer with other agents that they are aware of, or that they locate via the mediator.

Let us now consider how to structure ‘core’ program Π , on which a microservice μ ASPSv(Π) is based. First, activation and stopping of a module can be simply obtained by a couple of constraints, that make the program inconsistent (in no-operation state) if either activation A has not arrived, or stopping signal S has been issued:

```
:– not A.    % module activation
:– S.       % module stop
```

Then, when the module has been activated, upon arrival of new inputs, the inner program Π will in general ‘produce’ (admit) answer sets. If the answer set is unique then the outputs can be univocally identified. Otherwise the shell, in the ‘elicit outputs’ part, will have to adopt some kind of policy (e.g., preferences, utilities, costs or other) to select which answer set to consider. The queries, being by definition defined upon the whole set of answer sets, will always return an univocal result. In case, given the present input, Π should be inconsistent, then the output will consist in a failure signal (e.g., in the FIPA ACL, there is the *failure* primitive to be used in such cases).

6 Case Study

The case study that we propose here is inspired to issues raised by applications related to autonomous vehicles. Presently, machine learning mechanism have been defined to allow autonomous cars to comply with traffic lights by detecting their color, so as to pass with green and stop with red similarly to

traditional cars. Such mechanisms must be trained, are prone to errors, and are potentially subject to adversarial machine learning.

In our view, physical traffic lights might in perspective disappear, to be substituted by monitoring agents that would receive requests to pass from cars and consequently issue authorizations. This either in routes dedicated to autonomous vehicles, or in the (very reasonable) hypothesis to equip also ‘traditional’ cars with a device to interact with the monitoring agents.

Below we propose the sample design of the inner program concerning a μ ASPSv which implements the monitoring agent of a road intersection, taking the place of a physical traffic light. In the example, the traffic light agent is called *tl* and, for the sake of simplicity, behaves like a ‘real’ traffic light but just takes the colors green (*g* for short) and red (*r* for short). In fact, the yellow is no longer necessary as we assume that the involved cars (each one equipped with its own driver agent) will obey the directives. We have two lanes, one going north-south (*ns* for short) and the other one east-west (*ew* for short), crossing at the traffic light. If the traffic light is green in one direction it must be red in the other one, and vice versa. The traffic light is activated by a signal *active(t1)*, and never stopped unless there is a fault, detected by the module itself by means of a sensor. A fault is supposed to have occurred whenever *fault_tl* is true, i.e., it has been returned by the sensor.

```
tl(t1). %Traffic – Light Identifier
active(t1).
```

```
:- not active(t1). % Sensor Check activation
:- lane(L), fault_tl(t1, L, T). % Sensor Check Possible Fault
```

Each car, say here *c1*, *c2*, *c3*, *c4* and *c5*, wants to go, but it is allowed to proceed only if it gets the green traffic light. Otherwise, it remains dummy. We assume that all cars behave in the same way. Each one issues a request of format *car(C), want_go(C, t1, L, T)* where *L* is the lane, with possible values *ns* for north-south and *ew* for east-west; *T* is the time of the request. Requests by various cars may for example give rise to the addition of the following facts to the μ ASPSv’s program.

```
%INPUT : CARS
car(c1).
car(c2).
car(c3).
car(c4).
car(c5).

%INPUT : REQUESTS
want_go(c1, t1, ns, 2).
want_go(c2, t1, ns, 2).
want_go(c3, t1, ew, 2).
want_go(c4, t1, ns, 4).
want_go(c5, t1, ew, 4).
```

The following facts and rules define the lanes, and specify that this monitoring agent has a lookahead of five time instants: after that, it will have to be re-run.

```
lane(ns).
lane(ew).
time(1..5).
next(Y, X) :- time(X), time(Y), Y = X + 1.
```


Rules below define the color that the traffic light takes in a very standard way as transitions from green to red and vice versa, where the initial color is green. In reality, such a monitoring agent can employ a much more sophisticated protocol such as for instance the Contract Net Protocol (CNP). If adopting CNP, the agent might grant priority to particular kinds of vehicles, e.g., police cars, ambulances, cars transporting a disabled person, etc. More generally, any policy to grant passage according to criteria could be implemented.

$$\begin{aligned}
 tl(r, TL, L1, T1) :- & \\
 & time(T), lane(L1), lane(L2), tln(TL), L1! = L2, next(T1, T), \\
 & tl(g, TL, L1, T), tl(r, TL, L2, T). \\
 tl(g, TL, L1, T1) :- & \\
 & time(T), lane(L1), lane(L2), tln(TL), L1! = L2, next(T1, T), \\
 & tl(r, TL, L1, T), tl(g, TL, L2, T). \\
 tl(g, TL, ns, 1) :- & tln(TL). \\
 tl(r, TL, ew, T) :- & tln(TL), time(T), tl(g, TL, ns, T).
 \end{aligned}$$

In our case the implemented protocol is fair, as cars that cannot go now because it is red on their lane will be deferred to the next time instant (by delaying their request), when the color will be green (output in format $go(Car, t1, Lane, Time)$).

$$\begin{aligned}
 go(C, TL, L, T) :- & \\
 & time(T), car(C), tln(TL), lane(L), want_{go}(C, TL, L, T), tl(g, TL, L, T). \\
 wait(C, TL, L, T) :- & \\
 & time(T), car(C), tln(TL), lane(L), want_{go}(C, TL, L, T), tl(r, TL, L, T). \\
 want_{go}(C, TL, L, T1) :- & car(C), tln(TL), lane(L), wait(C, TL, L, T), next(T1, T). \\
 :- & time(T), car(C), tln(TL), lane(L), go(C, TL, L, T), tl(r, TL, L, T).
 \end{aligned}$$

Clearly, this program can be ‘cloned’ (*mutatis mutandis*) to manage any other traffic lights. For the reader’s convenience, this program is standalone and can be run exactly as it is to check its results.

7 System’s Semantics

The semantics of a single μ ASPSv is fully specified by: (i) the answer sets of the inner ASP program; (ii) the policy employed in its shell to select one single answer set; (iii) the set of queries that the shell performs over the entire set of answer sets, whose meaning is formally specified in [16, 14]. We aim however to provide a semantics for the overall distributed system composed of heterogeneous microservices, in order to provide a firm ground and a guideline for implementation.

To do so, we resort to Multi-Context Systems (MCSs), that are a well-established paradigm in Artificial Intelligence and Knowledge Representation, aimed to model information exchange among heterogeneous sources [4, 6, 7]. However, with some abuse of notation (and some slight loss of generality) we adapt and readjust the definitions to fit into our framework. To represent the heterogeneity of sources, each component in a Multi-context system, called ‘context’, is supposed to be based on its own *logic*, defined in a very general way [6]. In particular, a logic is defined out of the following.

- A set F of possible *formulas* (or *KB*-elements) under some signature;
- A set KB of knowledge bases built out of elements of F . in our framework, KB can also be a program in some programming language;

- A function ACC , where $ACC(kb, s)$ means that s is an acceptable set of consequences of knowledge base $kb \in KB$, i.e., $s \subseteq Cn$, where Cn is the set of all possible consequences that can be drawn from kb . We assume here that ACC produces a unique set of consequences. In case of a program written in a non-logical programming language, such set can be the set of legal outputs given some input (added to kb), that will be a subset of all possible outputs Cn ; for logical components, it will be (one of) the kb model(s). For instance, as we have seen the shell of a μ ASPSv will produce as consequences the elements occurring in the answer set selected according to some policy, along with query results (where queries deliver conclusions drawn in general from the whole set of answer sets)

A (Managed) *multi-context system* (MCS) $M = \{C_1, \dots, C_r\}$ is a set of $r = |M|$ contexts, each of them of the form $C_i = \langle c_i, L_i, kb_i, br_i \rangle$, where:

- c_i is the context *name* (unique for each context; if a specific name is omitted, index i can act as a name). In [8] a context's name, in the definition of a bridge rule, can be a term called *context designator* denoting the *kind* of context that should deliver that data item (for instance, *mycardiologist(c)*, *customercare(c)*, *helpdesk(h)*, etc.).
- L_i is a logic.
- $kb_i \in KB$ is a knowledge base.
- br_i is the set of *bridge rules* this context is equipped with.

Contexts in an MCS are meant to be heterogeneous distributed components, that exchange data. In fact, bridge rules are the key construct of MCSs, as it describes in a uniform way the communication/data exchange patterns between contexts. Each bridge rule $\rho \in br_i$ has the form

$$op_i(s) \leftarrow (c_1 : p_1), \dots, (c_j : p_j) \quad (1)$$

where the left-hand side s is called the *head*, and the right-hand side is called the *body*, and the comma stands for conjunction. The meaning is that, each data item p_j is supposed to come from a context c_j ³. Whenever all the c_1, \dots, c_j have delivered their data item to the *destination context* c_i , the rule becomes *applicable*. When the rule is actually applied, its conclusion s , once elaborated by operator Op , will constitute an input to c_i . Operator op_i can perform any elaboration on “raw” input s , such as format conversion, filtering, elaboration via ontologies, etc. Its operation is specified via a *management function* mng_i , which is therefore crucial for knowledge incorporation from external sources. For simplicity, here we assume mng_i to be monotonic (i.e., to produce from s one or more data items). Therefore, we can extend the previous definition of a context as $C_i = \langle c_i, L_i, kb_i, br_i, mng_i \rangle$. A bridge rule is *applicable* if its body is true, in operational terms if all the data items listed in its body have arrived. In case context designators are employed, prior to checking a bridge rule for applicability, such terms must be substituted by actual context names. For μ ASPSv's, this task will be performed by the shell, that must then be endowed with a list of contexts of each type.

A *data state* (or *belief state*) \vec{S} of an MCS M is a tuple $\vec{S} = (S_1, \dots, S_r)$ such that for $1 \leq i \leq r$, $S_i \subseteq Cn_i$. A data state can be seen as a view of the distributed system by an external “observer”. $app(\vec{S})$ is the set composed of the head of those bridge rules which are *applicable* in \vec{S} . This means, in logical terms, that their body is true w.r.t. \vec{S} . In practical terms, we may say that a bridge rule ρ associated to

³In the original formulation, there are additional literals $not(c_1 : p_{j+1}), \dots, not(c_j : p_n)$ meaning that in order for the bridge rule to be applied, the $p_{j+1} \dots p_n$ must be false in the relative contexts. We disregard this part, as non-logical components cannot use logical negation. There is no loss of generality however, as each of the p_1, \dots, p_j can state a negative fact.

context c_i is applicable in \vec{S} if all the data mentioned in the body of the bridge rule can be delivered to the destination context. This is the case whenever they are available in the contexts of origin, i.e., they occur in the present respective data state items in \vec{S} . In the original formulation of MCS, all applicable bridge rules are automatically applied, and their results, after the elaboration by the management function, are added to the destination context's knowledge base, that therefore grows via bridge-rule application.

Starting from a certain specific data state, some bridge rules will be applicable and therefore they will be applied. This will enhance the knowledge base in some of the contexts, thus determining (in these contexts) a new set of acceptable consequence, and therefore a new overall data state. In the new state other bridge rules will be applicable, and so on, until a “stable” state, called *Equilibrium*, will be reached. Technically, \vec{S} is an *equilibrium* for an mMCS M iff, for $1 \leq i \leq |M|$,

$$S_i = ACC_i(mng_i(app(\vec{S}), kb_i)) \quad (2)$$

which states that each element of the equilibrium is an acceptable set of consequences after the application of every applicable bridge rule, whose result has been incorporated into the context's knowledge base via the management function. in [8] it is proved that, in the kind of MCS that we have just described, an equilibrium will be reached in a finite number of steps. Notice however that this definition assumes the system to be isolated from any outside influence, and that an equilibrium, one reached, will last forever. Instead, in real systems there will be interactions with an external environment, and so equilibria may change over time. Moreover, each context is not necessarily a passive receiver of data sent by others.

To take these aspects into account, [8] proposes some extensions to the original formulation, among which the following, that are relevant in the present setting.

- It is noticed that contexts knowledge bases can evolve in time, not only due to bridge-rule application. In fact, contexts receive sensor inputs (passively or in consequence to active observation), or can be affected by user's modification (e.g., a context encompasses a relational database that can be modified by a user). So, each context c_i will have an associated *Update Operator* \mathcal{U}_i (that can actually consist in a tuple of operators performing each one a different kind of update). Updates and bridge rules both affect contexts' knowledge base over time. So (assuming an underlying discrete model of time) we will be able to consider, when necessary, $c_i[T]$ meaning context c_i at time T , with its knowledge base $kb_i[T]$; consequently we will have an evolution over time of contexts. Therefore, we will have a definition (not reported here) of *Timed Equilibria*. Notice only that a timed equilibrium can be reached at time $T + 1$ only if the actual elapsed time between T and $T + 1$ is sufficient for the system to “stabilize” by means of bridge-rules application on the updated knowledge bases.
- Mandatory bridge-rule application constitutes a limitation: in fact, this forces contexts to accept inputs unconditionally, and this be often inappropriate. Consider for instance a context representing a family doctor: the context may accept non-urgent patient's requests for appointments or consultation only within a certain time windows. So, [8] introduces conditional bridge-rule application, formalized via a *timed triggering function*, tr_i , which specifies which applicable bridge rules are triggered (i.e., they are practically applicable) at time T . It does so either based on certain pre-defined conditions, or by performing some reasoning over the present knowledge base contents. Therefore, the implementation of $tr_i[T]$ may require an auxiliary piece of program, that in a μ ASPSv's shell will presumably be a logic program.

So, considering contexts which are μ ASPSv's, in order to fit in the vision of the overall system as an MCS, their shell must be empowered as follows.

- Include the bridge rules associated to a μ ASPSv, and the definition/implementation of the triggering function.
- Include a facility to resolve the context designators, so as to check for applicability a triggered bridge rule after substituting context designators occurring therein with actual contexts' names.
- Include the definition of the specific management function, so as to be able to apply it on bridge-rules' results.

In the case study of previous section, each traffic light should be equipped with a bridge rule that, by means of a suitable context designator (say, *anycar(c)*) collects the cars' requests. Symmetrically, cars should be equipped with a bridge rule to collect the permission to go by the traffic light (the nearest one, whose identifier should replace a context designator of the form, e.g., *a_traffic_light(t)*). The triggering function may allow cars to enable receiving traffic-light communications only when this is deemed appropriate.

Context designators are therefore useful to write general bridge rules to be then customized to the particular situation at hand. They also allow to devise a system where components do not know or are aware of each other in advance, and where components can possibly join/leave the system at any time. A suitable middleware should be realized to allow component's shells to instantiate bridge rules. Concerning our case study, that concerns an infrastructure for autonomous vehicles, both cars and traffic lights might for instance broadcast their name and geo-localization. In this way, cars might locate the traffic light of interest, and traffic lights might become aware of nearby cars that might send them a request.

8 Concluding Remarks

We have proposed a methodology for developing microservices in Answer Set Programming, by means of the creation of a particular kind of components, that can be activated/stopped, can receive external requests and can deliver answers. We have provided a definition of a μ ASPSv and explained how it might be implemented, and we have outlined a programming methodology. We have also outlined a possible uniform semantics to specify an heterogeneous system in which μ ASPSv's can be situated. This is an absolute novelty for microservices in general, as no attempt has ever been made to provide such a uniform model for an overall system. The proposed semantics can constitute the ground for principled implementations. Overall, this work can be considered as a creative combination of existing technologies, in view of entirely new fields of application of answer set programming and logic programming in general.

Important application fields for μ ASPSv's are Cloud computing and IoT. We consider particularly important the various kinds of robotic applications and the underlying infrastructural aspects, as shown in the case study related to autonomous vehicles.

Future work includes refining the programming methodology on the one hand, devising a prototype implementation and experimenting the integration in realistic environment: we plan in particular to experiment the integration into DALI Multi-Agent-Systems, where DALI [10, 11, 12, 18] is the agent-oriented programming language defined and implemented by our research group, that has been made compatible with the Docker technology. This will allow us to experiment μ ASPSv's as components in the various applications where DALI is being applied, including cognitive robotic architectures.

References

- [1] David Ameller, Xavier Burgués, Oriol Collell, Dolors Costal, Xavier Franch & Mike P. Papazoglou (2015): *Development of service-oriented architectures using model-driven development: A mapping study*. *Information and Software Technology* 62, pp. 42 – 66, doi:10.1016/j.infsof.2015.02.006. Available at <http://www.sciencedirect.com/science/article/pii/S0950584915000361>.
- [2] Marcello Balduccini (2007): *Modules and signature declarations for a-prolog: Progress report*. *Workshop on Software Engineering for Answer Set Programming (SEA07)*, pp. 41–55.
- [3] Chitta Baral, Juraj Dzifcak & Hiro Takahashi (2006): *Macros, macro calls and use of ensembles in modular answer set programming*. In: *Logic Programming*, Springer, pp. 376–390, doi:10.1007/11799573_28.
- [4] Gerhard Brewka & Thomas Eiter (2007): *Equilibria in Heterogeneous Nonmonotonic Multi-Context Systems*. In: *Proc. of the 22nd AAAI Conf. on Artificial Intelligence*, AAAI Press, pp. 385–390. Available at <http://www.aaai.org/Library/AAAI/2007/aaai07-060.php>.
- [5] Gerhard Brewka, Thomas Eiter & Mirosław Truszczyński (eds.) (2016): *Answer Set Programming: Special Issue*. *AI Magazine* 37(3). Available at <https://dblp.org/rec/journals/aim/BrewkaEL16.bib> doi:10.1609/aimag.v37i3.2669.
- [6] Gerhard Brewka, Thomas Eiter & Michael Fink (2011): *Nonmonotonic Multi-Context Systems: A Flexible Approach for Integrating Heterogeneous Knowledge Sources*. In Marcello Balduccini & Tran Cao Son, editors: *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning - Essays Dedicated to Michael Gelfond on the Occasion of His 65th Birthday, Lecture Notes in Computer Science* 6565, Springer, pp. 233–258, doi:10.1007/978-3-642-20832-4_16.
- [7] Gerhard Brewka, Stefan Ellmauthaler & Jörg Pührer (2014): *Multi-Context Systems for Reactive Reasoning in Dynamic Environments*. In Torsten Schaub, editor: *ECAI 2014, Proc. of the 21st European Conf. on Artificial Intelligence*, IJCAI/AAAI, pp. 159–164, doi:10.1016/j.artint.2017.11.007.
- [8] Pedro Cabalar, Stefania Costantini, Giovanni De Gasperis & Andrea Formisano (2019): *Multi-context systems in dynamic environments*. *Ann. Math. Artif. Intell.* 86(1-3), pp. 87–120, doi:10.1007/s10472-019-09622-0.
- [9] S. Costantini (2011): *Answer Set Modules for Logical Agents*. In Oege de Moor, Georg Gottlob, Tim Furche & Andrew Jon Sellers, editors: *Datalog Reloaded - First Intl. Workshop, Datalog 2010, Revised Selected Papers, Lecture Notes in Computer Science* 6702, Springer, pp. 37–58, doi:10.1007/978-3-642-24206-9_3.
- [10] S. Costantini & A. Tocchio (2002): *A logic programming language for multi-agent systems*. In: *Logics in Artificial Intelligence, Proc. of the 8th Europ. Conf., JELIA 2002*, LNAI 2424, Springer-Verlag, Berlin, pp. 1–13, doi:10.1007/3-540-45757-7_1.
- [11] S. Costantini & A. Tocchio (2004): *The DALI logic programming agent-oriented language*. In: *Logics in Artificial Intelligence, Proc. of the 9th European Conf., Jelia 2004*, LNAI 3229, Springer-Verlag, Berlin, pp. 685–688, doi:10.1007/978-3-540-30227-8_57.
- [12] S. Costantini & A. Tocchio (2006): *About declarative semantics of logic-based agent languages*. In: *Declarative Agent Languages and Technologies*, LNAI 3229, Springer-Verlag, Berlin, doi:10.1007/11691792_7.
- [13] Stefania Costantini (2006): *On the existence of stable models of non-stratified logic programs*. *Theory and Practice of Logic Programming* 6(1-2), doi:10.1017/S1471068405002589.
- [14] Stefania Costantini (2019): *About Epistemic Negation and World Views in Epistemic Logic Programs*. *Theory Pract. Log. Program.* 19(5-6), pp. 790–807, doi:10.1017/S147106841900019X.
- [15] Stefania Costantini & Andrea Formisano (2013): *Negation as a Resource: A Novel View on Answer Set Semantics*. In Pedro Cabalar & Tran Cao Son, editors: *Logic Programming and Nonmonotonic Reasoning, 12th Intl. Conf., LPNMR 2013, Lecture Notes in Computer Science* 8148, Springer, pp. 257–263, doi:10.3233/FI-2015-1255.
- [16] Stefania Costantini & Andrea Formisano (2016): *Query answering in resource-based answer set semantics*. *Theory and Practice of Logic Programming* 16(5-6), pp. 619–635, doi:10.1017/S1471068416000478.

- [17] Minh Dao-Tran, Thomas Eiter, Michael Fink & Thomas Krennwallner (2009): *Modular nonmonotonic logic programming revisited*. In: *Logic Programming*, Springer, pp. 145–159, doi:10.1007/978-3-642-02846-5_16.
- [18] Giovanni De Gasperis, Stefania Costantini & Giulio Nazzicone (2014): *DALI Multi Agent Systems Framework*, DOI 10.5281/zenodo.11042. DALI GitHub Software Repository. DALI: <http://github.com/AAAI-DISIM-UnivAQ/DALI>.
- [19] Lorenzo De Lauretis (2019): *From Monolithic Architecture to Microservices Architecture*. In: *2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*, IEEE, pp. 93–96, doi:10.1109/ISSREW.2019.00050.
- [20] Nicola Dragoni, Saverio Giallorenzo, Alberto Lluch Lafuente, Manuel Mazzara, Fabrizio Montesi, Ruslan Mustafin & Larisa Safina (2017): *Microservices: yesterday, today, and tomorrow*. In: *Present and ulterior software engineering*, Springer, pp. 195–216, doi:10.1007/978-3-319-67425-4_12.
- [21] Abeer Dyoub, Stefania Costantini & Giovanni De Gasperis (2018): *Answer set programming and agents*. *Knowledge Eng. Review* 33, p. e19, doi:10.1017/S0269888918000164.
- [22] Thomas Eiter, Georg Gottlob & Helmut Veith (1997): *Modular logic programming and generalized quantifiers*. In: *Logic Programming and Nonmonotonic Reasoning*, Springer, pp. 289–308, doi:10.1007/3-540-63255-7_22.
- [23] Wolfgang Faber & Stefan Woltran (2009): *Manifold answer-set programs for meta-reasoning*. In: *Logic Programming and Nonmonotonic Reasoning*, Springer, pp. 115–128, doi:10.1007/978-3-642-04238-6_12.
- [24] Haim Gaifman & Ehud Shapiro (1989): *Fully abstract compositional semantics for logic programs*. In: *Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, ACM, pp. 134–142, doi:10.1145/75277.75289.
- [25] Giovambattista Ianni, Giuseppe Ielpa, Adriana Pietramala, Maria Carmela Santoro & Francesco Calimeri (2004): *Enhancing answer set programming with templates*. In: *10th International Workshop on Non-Monotonic Reasoning (NMR 2004), Proceedings*, pp. 233–239.
- [26] Pooyan Jamshidi, Claus Pahl, Nabor C. Mendonça, James Lewis & Stefan Tilkov (2018): *Microservices: The Journey So Far and Challenges Ahead*. *IEEE Software* 35(3), pp. 24–35, doi:10.1109/MS.2018.2141039.
- [27] Tomi Janhunen, Emilia Oikarinen, Hans Tompits & Stefan Woltran (2009): *Modularity aspects of disjunctive stable models*. *Journal of Artificial Intelligence Research*, pp. 813–857, doi:10.1613/jair.2810.
- [28] Petar Krivic, Pavle Skocir, Mario Kusek & Gordan Jezic (2017): *Microservices as agents in IoT systems*. In: *KES International Symposium on Agent and Multi-Agent Systems: Technologies and Applications*, Springer, pp. 22–31, doi:10.1007/978-3-319-59394-4_3.
- [29] Christine Legner & Roger Heutschi (2007): *SOA adoption in practice-findings from early SOA implementations*. Association for Information Systems.
- [30] Yuliya Lierler & Mirosław Truszczyński (2013): *Modular answer set solving*. *Late-Breaking Developments in the Field of Artificial Intelligence*, AAAI WS-13-17.
- [31] Vladimir Lifschitz (2008): *Twelve Definitions of a Stable Model*. In Maria Garcia de la Banda & Enrico Pontelli, editors: *Proc. of the 24th Intl. Conf. on Logic Programming, LNCS 5366*, Springer, pp. 37–51, doi:10.1007/978-3-540-89982-2_8.
- [32] Vladimir Lifschitz & Hudson Turner (1994): *Splitting a Logic Program*. In: *Logic Programming, Proceedings of the Eleventh International Conference on Logic Programming*, 94, MIT Press, pp. 23–37.
- [33] Emilia Oikarinen (2008): *Modularity in answer set programs*. Ph.D. thesis, Helsinki University of Technology, Finland.
- [34] Yi-Dong Shen & Thomas Eiter (2016): *Evaluating epistemic negation in answer set programming*. *Artificial Intelligence* 237, pp. 115–135, doi:10.1016/j.artint.2016.04.004.
- [35] Luis Tari, Chitta Baral & Saadat Anwar (2005): *A Language for Modular Answer Set Programming: Application to ACC Tournament Scheduling*. In: *Answer Set Programming, Advances in Theory and Implementation, Proceedings of the 3rd Intl. ASP'05 Workshop, CEUR Workshop Proceedings 142*, CEUR-WS.org.