SVeDaS:
Specification and Verification of Data-aware Systems

Francesco Belardinelli

Appel projets générique 2016
Jeunes Chercheuses/Jeunes Chercheurs
Défi 7: Société de l'Information et de la Communication

Contents

Abstract

In recent years *data-aware systems* have been proposed as a comprehensive framework to model complex business workflows by considering *data* and *processes* as equally relevant tenets of the system

description [33, 45]. We claim that this setting is suited to model and verify auctions and auction-based mechanisms in electronic commerce, including for instance, English, Dutch, and Vickrey auctions [35].

The SVEDAS project is designed to advance the state-of-the-art in the modelling, analysis and deployment of data-aware systems by using a novel, compositional, agent-based approach to their specification and verification, and to apply these results to the verification of auctions against strategic behaviours of agents, such as collusion and manipulation.

The main objectives of the SVEDAS project can be summarized as follows:

1. to **introduce** agent-based, computationally-grounded models for data-aware systems, that are capable of expressing rich business workflows, including auction-based mechanisms in e-markets;
2. to **explore** logic-based formal languages for the specification of strategic behaviours of autonomous agents (including robustness against malicious behaviours in auctions) pertaining to business processes and agents operating on them;
3. to **analyse** the formal properties of these data-aware models, particularly the issues concerning formal verification by model checking in contexts of imperfect information;
4. to **find** classes of data-aware systems and expressive language fragments relevant for auction-driven applications, which have a decidable model checking problem and possibly are also amenable to practical verification;
5. to **develop** the model checker SVEDAS for the verification and validation of data-aware systems in multi-agent scenarios;
6. to **evaluate** the performance of the SVEDAS tool in popular auctioning mechanisms, including real-time bidding, and to release SVEDAS as open-source software.

We anticipate that the results of the SVEDAS project will contribute significantly to our understanding of data-aware systems, thus improving the design and management of business processes by formal verification through model checking, including through the model checker SVEDAS. In turn, these contributions will help building more secure and reliable auction-based mechanisms for e-commerce and e-business.

Table of Partners

Partner	Surname	First Name	Position	Person.months	Role
Laboratoire IBISC UEVE	Belardinelli	Francesco	MCF	27.0	PI
Laboratoire IBISC UEVE	suitable candidate		PhD student	36.0	RA

Changes with respect to the Pre-proposition

In editing the final version of the SVeDASprojet we took care to address the reviewers' remarks on the pre-proposition, whenever possible, including strengthening the collaborations for developing the scientific programme. We observe that the ANR JCJC scheme envisages a contribution only for the young researcher, collaborators will not be able to manage any funding directly. However, Dr Belardinelli will be in a position to develop his current collaborations further through regular meetings and shared MSc projects, all funded by the SVeDaS project.

As regards changes to the original budget, the initial contribution of €160k has been raised to a total of €173,971, within the brackets of variation allowed by the ANR between the preliminary and finale phase of the JCJC call. The increment is motivated by the addition of €17,280 as *décharges d'enseignement*, i.e., a contribution to relieve Dr Belardinelli from half of his annual teaching duties during the project lifetime (96h out of 192h per year), in order to focus on the development of the scientific programme of the SVeDaS project, its administrative management, and on tutoring the research assistant employed by the project. Furthermore, an extra €5k have been added to the project's budget for the organisation of a workshop to present the project's results at the end of its lifetime. This initiative is meant to increase the project's impact through dissemination. The remaining difference between the preliminary and final budget is due to a more precise calculation of the project's administrative costs.

1 Context, Positioning and Objectives

In this section we present the main problem tackled in the SVeDaS project, that is, the formal verification of data-aware systems, and outline the proposed solution leveraging on modelling and abstraction techniques from the area of logics for multi-agent system. We position the research programme within the current state-of-the-art. In particular, we discuss the novelty of the project's objectives in relation with the EU STREP project ACSI [1], which represents the inspiration and original motivation for the current submission.

1.1 The Problem and Proposed Solution

Data-aware systems (DaS) are a novel paradigm for the design, implementation and integration of business processes in service-oriented computing [61]. The originality of this approach consists in "combin[ing] data and processes in a holistic manner as the basic building block[s]" of the system's description [29]. Typically data-aware systems include a *data model*, to account for the relational structure of data, as well as the *business processes* manipulating data. Both the data model and business processes are seen as equally important tenets of the system description. This setting is in marked contrast with most of the tradition on service architectures and composition, which usually abstracts data away to reduce the complexity of the system description and thus making the verification task amenable to standard model checking techniques [61]. Recently, this data-driven approach has been successfully applied to the analysis of elaborate business scenarios, including procurement use cases by IBM [45]. However, the enhanced expressivity provided by data-aware systems comes at a computational price. In particular, we identify two main shortcomings in the present state-of-the-art.

1. Most of the literature on DaS [33, 45, 53, 34] focuses almost exclusively on the data model of business processes, while neglecting the software agents implementing the service infrastructure. These software agents might have only a partial view of the relevant data, or, in the terminology of multi-agent systems, they have *imperfect information* of the global state of the system [63]. This (lack of) information shapes the capabilities of the software agents to interact and bring about change, which in turn has an impact on the overall behaviour and performance of the data-aware system. Thus, modelling both the information state and the strategic abilities of agents operating on DaS are key to describe and predict the evolution of the system.

2. The actual deployment of DaS in concrete, safety- and security-critical scenarios demands for the development of automated verification techniques, including by model checking [28, 9]. However, while formal

methods are relatively well-understood in plain process-based modelling, the presence of data makes the typical verification questions much harder to answer and only partially explored. Notably, the common assumption of a possibly infinite data domain in the underlying system leads to an infinite state-space and undecidability of the corresponding model checking problem in the general case. Hence, the verification of DaS is highly non-trivial and it cannot be immediately handled by standard techniques for finite-state systems. Data-driven, tailored solutions need to be developed and deployed in an up-to-date model checking tool.

These are among the main challenges faced by the application of the data-centric paradigm in business process modelling.

The solution we propose in the SVEDAS project consists in

(*i*) **developing** an agent-oriented approach for modelling DaS, to account for the imperfect information in the system;

(*ii*) **investigating** verification techniques based on (bi)simulations and abstraction for contexts of imperfect information;

(*iii*) **implementing** the relevant techniques in a model checking tool – SVEDAS – to certify auctioning mechanisms represented as data-aware systems.

The proposed solution is enbodied in a series of (partial) objectives that can be summarized as follows:

1. to **introduce** agent-based, computationally-grounded models for data-aware systems, that are capable of expressing rich business workflows, including auction-based mechanisms in e-markets (e.g., English, Dutch, and Vickrey auctions, real-time bidding);

2. to **explore** logic-based formal languages for the specification of strategic behaviours of autonomous agents (including robustness against malicious behaviours, as well as manipulability and collusion in auctions) pertaining to business processes and agents operating on them;

3. to **analyse** the formal properties of these data-aware models, particularly the issues concerning formal verification by model checking in a setting with imperfect information;

4. to **find** classes of data-aware systems and expressive language fragments relevant for auction-driven applications, which have a decidable model checking problem and possibly are also amenable to practical verification;

5. to **develop** the SVeDaS model checker for the verification and validation of data-aware systems in multi-agent scenarios;

6. to **evaluate** the performance of the SVeDaS tool in the influential auctioning mechanisms mentioned above, and to release SVeDaS as open-source software.

We anticipate that the research envisaged in the SVeDaS project will impact on the application of formal methods to data-aware systems in general, and auction-based mechanisms in particular. The certification against deviant behaviours guaranteed by formal verification will contribute to design more robust and reliable auctions in e-commerce. Society as a whole will benefit from the findings of the SVeDaS project.

1.2 The State-of-the-Art

The scientific programme envisaged by the SVeDaS project draws inspiration from and is situated at the intersection of three research areas:

1. Data-aware systems, including auctioning mechanisms;

2. Multi-agent systems (MAS) and logics for strategic abilities;

3. Formal methods and verification by model checking;

In order to position the project's objectives with respect to the state-of-the-art and to illustrate the comparative advancement, here we discuss significant, recent contributions in these three areas. Given the vastity of the subject matters, the discussion will necessarily be partial and oriented towards the verification of data-aware systems by means of agent-based techniques.

1.2.1 Data-aware Systems

Data-aware systems have emerged in the last decade as the leading paradigm to analyse use cases in which data play an essential role in the system's execution [21, 33, 45]. As an example of DaS, here we briefly describe a business process inspired by a concrete IBM use case [46]: the *order-to-cash* scenario details the interactions of manufacturers, customers, and suppliers in an e-commerce situation involving the purchase and delivery of goods and services. At the start of the business process, a customer prepares and submits to some manufacturer a *purchase order* (PO), i.e., a list of products the customer requires, together with information about these products such as quantity, price, expected-by date, etc. Upon receiving a PO, the manufacturer prepares a *material order* (MO), i.e., a list of components needed to assemble the requested products, based on the information provided by the customer herself. The manufacturer then selects some suppliers and

forwards them the appropriate material orders. Upon receiving an MO, a supplier evaluates the information provided therein and either accepts or rejects the order. In the former case she then proceeds to deliver the requested components to the manufacturer, according to the relevant specs. In the latter she notifies the manufacturer of her rejection. Finally, when the manufacturer receives the components, she assembles the product and, provided that the order has been paid for, she delivers it to the customer.

Observe that, even in such a plain scenario, all key components of data-aware systems are clearly represented. The data model includes the purchase and material orders, which can be encoded in some sort of data structure, typically a relation databases; while the business processes detailing the evolution of orders from creation, through validation/rejection, to fulfilment, can be described by an appropriate set of operations on relational structures. Most importantly, the system's execution depends crucially on the data content of purchase and material orders: the supplier might chose to accept or reject a material order depending on whether she has enough resources for the requested quantity, whether the price is within a certain range of profitability, or whether she can meet the deadline for delivery. Thus, the agents' available actions and behaviour essentially depend on the information registered in the data model.

The SVeDaS project is designed to contribute to the verification and validation of data-aware systems against specifications describing the strategic behaviour of agents operating on the system. A significant contribution in this direction has come from the EU STREP project ACSI [1], to which I contributed in 2011-12. The ACSI project focused on artifact-centric systems, a particular data-driven approach to modeling and deploying business processes, and produced a stream of fundamental contributions on their verification [19, 32, 8, 38]. Among the results of the ACSI project, a key finding is represented by the notion of *uniformity*, which has been used in [18, 19] to obtain a decidable model checking problem. Intuitively, a data-aware system is uniform whenever its evolution is determined only by data values that are named explicitly in the system's description. Conversely, all data that are not exhibited can be deemed equivalent, as far as the system's execution is concerned. This allows to apply abstraction-based techniques to reduce the model checking problem to the finite case, provided that some additional constraints are met. Interestingly, the uniformity condition, which is related to the notion of *genericity* in database theory [2], is satisfied by a vast class of interesting systems, including some types of auctions.

The SVeDaS project is intended to build on these results, which are applicable to a specific class of DaS, while extending the boundaries of verification. In fact, albeit extremely relevant, we identify several criticalities regarding the methods made available by the ACSI project, as well as the current literature on DaS in general.

1. Although uniformity defines an important class of DaS, many systems of interest are not uniform. Indeed, most manipulations of data bring us outside the realm of uniformity. Even simple operations, such as increments on natural numbers, are sufficient to break uniformity [12]. Hence, a first challenge for SVeDaS is to find conditions more robust than uniformity, which still imply a decidable model checking problem.

2. According to the results of ACSI, a further assumption required to obtain decidability, besides uniformity, is *boundedness*, that is, the existence of an upper bound on the number of active elements in a data-aware system at any time in the execution [32]. However, in several scenarios assuming the existence of such a bound may appear arbitrary and artificial: databases can be expanded beyond any given size, by simply keeping on adding new entries (without removing any of the old ones). Thus, a further challenge with respect to the state-of-the-art is to identify classes of models, still general enough for representing most DaS of significance, but which can also be bounded in a natural way.

3. Related to the previous point, a third challenge is represented by unbounded systems (such as the expanding databases above). Again, in this case the decidability results of the ACSI project do not apply uncondintionally, so novel techniques need to be explored.

Thus, the SVeDaS project is solidly set within the most recent advances on the verification of data-aware systems. Yet, it is meant to question the constraints imposed on DaS, namely uniformity and boundedness, in order to develop novel verification methods suitable for a wider class of DaS.

Another element of originality of the SVeDaS project is the focus on auctioning mechanisms. Indeed, auctions can be seen as data-aware systems: the outcome depends essentially on the values of bids, base prices, and true values. As a proof of concept, in [11] a basic version of parallel and iterated English (ascending bid) auctions are formalised as DaS, then they are successfully verified against safety and liveness properties. This kind of results validates the approach proposed in the SVeDaS project. However, more elaborate cases, including real-time bidding, in which agents can modify their behaviour according to the numerical outcome of previous auctions, are not covered by ACSI, since they suffer from limitation (1) detailed above. To overcome this issue, proposals have been put forward [12] that also support arithmetic operations [34]. Yet, these contributions neglect the imperfect knowledge that agents typically have of the system's global state, which limits the applicability of these results to auctions. These considerations motivate the second tenet of the SVeDaS project.

1.2.2 Logics for Multi-agent Systems

Multi-agent systems are open, distributed systems where the processes involved, or *agents*, show highly flexible and autonomous behaviour [63]. Agents in MAS are assumed to be proactive, endowed with beliefs about the surrounding environment, as well as their own private goals and plans to achieve them [57]. Researchers in artificial intelligence have adopted multi-agent systems to model and solve problems in several areas, including economics, game theory, planning, and robotics, that are difficult, viz. impossible, for a monolithic system to tackle [60]. Most importantly for the SVeDaS project, the agent paradigm allows for a modular approach to system modelling, in which the interactions between agents are not hard-coded in the systems description, but emerge at run-time according to the agents' specification. Moreover, the description of agents in terms of intentional attitudes allows us to abstract from actual implementation details. Here the emphasis is on the local, information state of agents, as well as concepts such as knowledge and belief, that are used to describe agents. These features of MAS have been deemed extremely valuable in designing complex distributed applications, at least in the modelling stage.

We argue that data-aware systems can benefit hugely from the adoption of a multi-agent perspective. Indeed, in the other-to-cash scenario above, the clients, manufacturers, and suppliers all have their private information, that they might want to share only partially or in a controlled way. They also have different goals (e.g., profit maximisation, timeliness), and they might have various plans available to achieve them. In auction-based mechanisms bidders normally keep their true values private, as well as their bids in sealed auctions. This agent-based perspective on DaS has been explored only preliminarily [18, 19], while most of current approaches still regard DaS as monolithic systems [33, 34, 32, 8]. Moreover, the agent approach allows for the application of modular abstraction techniques to tackle the model checking problem [40].

Related to the verification of multi-agent stystems, agent-oriented specification languages have been a thriving area of research in recent years [23, 22, 39]. A diverse family of multi-modal logics has been introduced for representing and reasoning about complex strategic abilities, both individual and coalitional, including alternating-time temporal logic (ATL), strategy logic, coalition logic just to name a few [5, 26, 55]. In parallel with these developments, a well-established tradition in knowledge representation focuses on extending formalisms for reactive systems with epistemic operators, so as to reason about the systems' evolution, as well as the knowledge agents have thereof [36]. Seminal contributions on extensions of linear- and branching-time temporal logics with agent-indexed epistemic modalities date back to the '80s [42, 43]. Since then, these investigations have matured into a solid body of works, which is nowadays rightly regarded as a key contribution of

formal methods to computer science [50], particularly when combined with verification techniques [37, 48, 49].

If agent-based logics are to be applied to the specification of data-aware systems, they need to be extended with relational and first-order features to account for the role played by data. As an example, in auctions the behaviour of agents has to be checked against all admissible values for bids, asking prices, and true values, thus calling for (universal) quantification in the specification language for such properties. However, it is well-known that assuming naively unrestricted first-order quantification quickly leads to the undecidability of a number of problems, including satisfiability and model checking. Hence, more sophisticated methodologies have to be adopted to lift logics for agents to the first order. In fact, first-order logic includes some interesting fragments with nice computational properties (e.g., the monadic, guarded, and two-variable fragments [24]), which can be used to express specific behaviours of data-aware systems. For instance, quantification in DaS can be guarded by assuming that values range on appropriate subsets, suitably specified by predicates in the language. Also, whenever we want to compare two values that appear at different times of the system's execution, two variables are sufficient. The applicant Dr Belardinelli has contributed to these investigations, by proving that some sound and complete axiomatisations for multi-agent temporal epistemic logics can be lifted to the monodic fragment of first-order logic, i.e., a controlled form of quantification [15]. These results constitute a solid starting point for the research programme envisaged by the SVeDaS project.

1.2.3 Verification by Model Checking

Formal methods are widely used to represent and analyse distributed and reactive systems. In combination with verification techniques by model checking, they have become one of the success stories in computer science [9, 28]. In the model checking approach, to verify whether a system S satisfies a property P (such as a safety, liveness, or secrecy requirement), first S is modelled as (some kind of) transition system $\mathcal{M}_S$, while property P is recast as a formula φ_P in some logical language of choice. Finally, verification is reduced to check whether the formula φ_P is true in the model $\mathcal{M}_S$, or $\mathcal{M}_S \models \varphi_P$ formally. Nowadays, model checking is being successfully applied to the automated verification of real-life scenarios in safety critical systems, avionics, AUVs, robotics, and security protocols [49, 56].

Similarly, the actual deployment of data-aware systems calls for the development of verification techniques. As an example, in designing auction-based mechanisms we might require that *bidders for a particular resource bid consistently with the true value they assign to the resource (i.e., they do not exceed it), without revealing this true value publicly* ($*$). Such requirements specify the behaviour of agents with respect to a possible infinite number

of values for their bids and true values. However, verification techniques such as model checking are "mainly appropriate to control-intensive applications and less suited for data-intensive applications" [9]. Irrespectively of these difficulties, the model checking problem for auctions has received considerable attention recently [7, 41, 65, 64, 62]. Indeed, it is hard to overestimate the relevance of auctions and auction-based mechanisms in a wide range of distributed systems (e.g., task scheduling, power grid management, and resource allocation [58, 30]). However, with some notable exceptions, most of the research on this topic has focus on the design of auctioning mechanisms and the analysis of their formal properties, while the automated verification of these designs has only partially been addressed. To our knowledge [7, 41, 65, 64, 62] are among the first contributions to consider the formal verification of auctions. In [62] the authors implement a simple auction model in a BDI-based programming language, to which they apply agent verification techniques. In [41] the problem of model checking strategy-proofness of Vickrey auctions is investigated; while [65, 64] propose a formal approach to check for shilling behaviours in auctions. This list is by no means exhaustive, but it is representative of some current trends in the formal verification of auction (see also [52] for instance). Overall, [7] is among the contributions most closely related to the present proposal in spirit, but a key difference is that their models abstract from the actual data content of auctions, in order to make the problema amenable to standard model checking techniques. Moreover, these references discuss limited classes of auctions, and the solutions proposed are tailored to the cases of interest; while we advocate a general, principled account for the verification of data-aware systems that it is also capable of dealing with auction-based mechanisms. We reckon that, given the relevance that data representation and reasoning have gained in recent years, it is key for the deployment of business processes to provide data-aware systems with sound verification methodologies. In turn, this endeavour raises a number of challenges ranging from (*i*) the logic-based languages for specifying DaS behaviours, to (*ii*) the data structures to represent DaS symbolically, as well as (*iii*) efficient model checking algorithms to deal with relational and first-order features. The SVeDaS project is designed to fill this gap by developing a principled approach to DaS verification, with a specific focus on those particular business processes that are auctions.

1.3 Outline of the Methodology

Differently from traditional approaches to business process modelling, we envisage to adopt an agent-based, compositional perspective in the analysis and verification of data-aware systems. Previous work by the applicant Dr Belardinelli among others, has proved that the agent paradigm can be successfully integrated in the representation of DaS [12, 13, 14, 18, 19], in-

cluding English auctions [11], thus validating the feasibility of the project. Indeed, notions of individual knowledge, strategies and goals from artificial intelligence have already been applied to modelling and verifying distributed systems [36]. We claim that also our understanding of business processes can benefit from an analogous representation in terms of intentional attitudes of the composing services, seen as agents. Specifically, the proposed methodology can be spelled out in three directions, which correspond to the three workpackages of the SVeDaS project. Here we briefly outline them and refer to the relevant WP for further details on methodology.

WP1 We plan to explore agent-based, computationally-grounded models for data-aware systems in a context of imperfect information, and to apply the logical machinery to the formalisation of popular auctioning mechanisms.

WP2 We will investigate the formal properties of these DaS models, particularly in relation with the model checking problem. We anticipate to tackle verification by using truth-preserving (finite) abstractions that (bi)simulate the concrete, infinite-state DaS.

WP3 We will implement the verification methods developed in WP2 in the model checking tool SVeDaS for the certification of DaS. We will evaluate the performace of SVeDaS against the auction scenarios analysed in WP1, and release the tool as open-source software.

We plan to develop the three workpackages rather sequentially, by applying methods and techniques from the three areas of data-aware systems, logics for multi-agent system, and formal verification. The specific logical tools adopted are detailed in the following scientific programme.

2 Scientific Programme and Techniques, Project Organisation

The research programme is structured in three main workpackages, ranging from the more theoretical to the more applied aspect to the project, each corrisponding approximately to one year in the project lifetime. The WP are meant to be executed rather lineraly in the temporal dimension, while still allowing for feedback and adjustment according to the project objectives. The research programme will be developed in close collaboration by the Principal Investigator, Dr Belardinelli, and the PhD student who will act as Research Assistant. The commitment for PI and RA in terms of person.months will be equally distributed across the 3 years: 12 person.months per year for the RA and 9 person.months per year for the PI. This repartition is subject to change if need be. In particular, we anticipate that the commitment of the PI will be greater at the very beginning of the SVeDaS project, and in WP1 and WP2; while the RA will acquire more independence and autonomy along the project lifetime, and will be in charge for the actual coding of the SVeDaS model checker.

2.1 WP0 – Review of Background Literature (2 months)

This WP comprise a unique task and is devoted to reviewing and analysising the most recent literature on the fundamental tenets of the project, as described in the state-of-the-art: data-aware systems, logics for agents in multi-agent systems, formal verification by model checking. This task is intended to assess current methodologies in the verification of data-aware systems, with a particular focus on the application of modular, agent-based modelisation of DaS. In Section 1.2.1 we observed that current approaches typically model DaS as monolithic systems, where the components have perfect information of the system's global states. In reviewing the literature we will pay particular attention to methodologies that relax such an assumption and adopt an agent-oriented perspective.

The PI and RA will also survey the state-of-the-art in the formal verification of auctions. Specifically, we are interested in their analysis from the perspective of partecipating agents, including the agents' information state and how this changes along the auction.

This WP will be developed by the RA under the guidance of Dr Belardinelli, whose expertise lies at the intersection of the three areas mentioned above. Also, Dr Belardinelli will help the RA to become familiar with the project's background, as well as with the model checker MCMAS [49], upon which we plan to develop the SVeDaS model checker in WP3.

Outcome: At the end of the WP Dr Belardinelli and the RA will have a clear picture of the state-of-the-art in the formal verification of DaS, as

well as of the agent-oriented approaches to this problem. The RA will also become familiar with the main trends in formal verification by model checking and proficient in the use of the MCMAS model checker for multi-agent systems. This background knowledge is key for the development of subsequent work-packages.

Deliverables:

- A survey on the state-of-the-art on the verification of data-aware systems, including current approaches to formal methods for the representation of auctions and auction-based mechanisms.

2.2 WP1 – Agent-based Models for Data-aware Systems (12 months)

This workpackage is designed to lay the theoretical foundations that will be developed and applied in WP2 and WP3. The objective is to define computational models for data-aware systems based on autonomous agents. The workpackage is structured in 3 main tasks: *relational structures for DaS*, *agents with imperfect information*, and *a formal framework for auctions*. In particular, the first two tasks will develop the formal framework applied in the third, which in turn will shape and provide guidance in the developement of logical tools. These tasks will be undertake by the PI and RA in close collaboration, as the PI will be able to build upon his previous works.

2.2.1 T1.1 – Relational Structures for DaS (5 months)

In this task we develop models for representing data-aware systems statically and dynamically, also by building on previous results of the ACSI projet. The first challenge is the explicit representation of data in models: classical approaches to model checking multi-agent systems typically assumes that the local states of agents is encoded in some propositional language [36], which are indeed expressive enough for a wide range of applications [9, 49]. But, as remarked above, propositional languages are often not sufficient for expressing properties such as manipulation, collusion, and secrecy in auctions, including example $(*)$ above. Hence, more expressive formalisms, notably first-order extensions of multi-agent logics, will be explored in WP1.

The ACSI project put forward various computational models inspired to first-order formalisms: *artifact-centric multi-agent systems* [18], *relational data-centric dynamic systems* [8], *models for the situation calculus* [32]. The common feature of all these formalisms is the representation of the system's global state as some sort of relational, first-order structure, grounded in the theory of databases [2]. On the other hand, the update mechanisms applied on these relational structures to describe their temporal evolution shows varied characteristics: in [18] agents' actions are completely abstract and

system's executions are simply viewed as successions of relational structures; the research line represented by [8] make use of the well-studied formal machinery of descriptions logics to describe an update mechanisms in which T-boxes evolve according to the results of conjunctive queries; while [32] applies the update axioms of the situation calculus.

Even though these formalisms are a valuable source of inspiration, we identify two major shortcomings.

- While [18] introduces a multi-agent setting for DaS, it does not develop a computational model for partecipating agents, as protocols and actions are completely abstract. An attempt to ground artifact-centric multi-agent systems on actual computation is provided through *artifact-centric programs*, by means of actions guarded by first-order formulas, whose satisfaction entails an update on the relational, local state of the agent. Such an account, whilst of interest, assumes that each agent is working on her local copy of the system's global database, without providing guarantees on consistency. Moreover, the complexity of updates thus given is not investigated.

- In the line of relational data-centric dynamic systems and situation calculus [8, 32], systems are given monolithically, the relational structure representing the system's global state is unique, and transitions are modelled as updates on such structures. As a consequence, no notion of strategic or game-theoretic interaction is analysed in these frameworks. This feature is witnessed by the fact that the typical logics used to specify properties for these systems are LTL, CTL, and the μ-calculus, which are mostly apt at expressing temporal notions.

Furthermore, the complexity of verification in most of these settings is exponential in the size of data. So, well beyond what it is amenable in terms of practical model checking tools.

In this task we will attempt to overcome the issues outlined above by developing models for data-aware systems with an update mechanisms grounded in (possibly efficient) computations. Several directions are worth pursuing, starting with the literatures discussed in Section 1.2.1. A promising computational model is represented by *reactive modules* [4], also in the form of the *simple reactive modules* in [44]. This framework has a clear computational content as it is the basis of the programming language for the MOCHA model checker [3]. Reactive modules are basic system components, who control a number of variables whose value they can modify, but who also have access to other read-only variables controlled by other modules. To account for secrecy and privacy, a certain number of controlled variables are assumed to be not visible by other modules. Modules act on the variables they own by performing various operations, thus determining the system's transitions. The framework of reactive modules is extremely flexible and

the proponents claim that it can model pure asynchronicity (interleaving), observable asynchronicity, atomic and non-atomic synchronicity.

Our starting point will be the definition of reactive modules with infinite data types, e.g., the natural, rational, and real numbers, and operations defined thereon, which are normally considered in the literature on DaS and appears as values for auctions. To illustrate this idea, we present the definition of a bidder module in English auctions.

Bidder The bidder module m_i controls the variables in ctr_i, initialises its local state according to the guarded actions in $init_i$, and updates it following guarded actions in $update_i$, which are defined as

- set ctr_i includes variable $tvalue_i$, registering module m's true value as a real number, and bid_i to represent m's current bid, also as a real. The value of bid_i is public, while $tvalue_i$ is kept private;
- $init_i$ contains guarded actions:

$$\top \quad \rightsquigarrow \quad bid_i := \text{uu}; tvalue_i := x_3$$

  in this case the guard $\top$ is true, hence the true value is initialised with a random value x_3, while the bid is left undefined uu for the time being.

- $update_i$ contains guarded actions $skip$ and

$$\begin{aligned}
(t_out = \bot) \wedge \bigwedge_{j \in M} (bid_j = \text{uu}) \wedge (x_4 \leq tvalue_i) \quad &\rightsquigarrow \quad bid_i := x_4 \\
(t_out = \bot) \wedge \bigvee_{j \neq i} (bid_i < bid_j) \wedge \bigwedge_{j \neq i} (bid_j \neq \text{uu} \rightarrow bid_j < x_5) \wedge (x_5 \leq tvalue_i) \quad &\rightsquigarrow \quad bid_i := x_5 \\
t_out = \top \quad &\rightsquigarrow \quad bid_i := \text{uu}; \\
& \qquad\quad tvalue_i := x_6
\end{aligned}$$

  The update actions allow the bidder to bid for the item auctioned by the auctioneer module (not specified), and to raise her bid according to her true value and bid by other bidders, as long as the auctioneer does not time out.

We remark that the applicability of actions in the bidder module depends essentially on the actual values of bids and true values. Also, these variables take values from real numbers. Hence, reactive modules define data-aware systems that are infinite state in general.

Outcome: A notion of reactive module with infinite data types, including a computational update mechanism based on guarded actions. A specification language for guards in actions. A formal semantics in terms of infinite-state concurrent game structures.

2.2.2 T1.2 – Agents with Imperfect Information (5 months)

The reactive modules adopted in T1.1 partition the set of variables owned by a module into *private* and *public*, where public variables are visible (but not modificable) by any other module. Hence, we need languages and methodologies to account for the imperfect information of modules, that is, the fact that agents have access only to a partial description of the global state of the system, determined by the variables they own and can see, which may be a proper subset of all the variables appearing in the system. This modelling constraint is common in auction-based mechanisms, for instance in English auctions, where the true value of each bidder is private at the beginning of the auction, and each bidder wants to keep this piece of information secret to other bidders and the auctioneer throughout the auction. Hence, we need to model these secrecy features in our framework, particularly in the specification language.

Our starting point to specify properties of interest, including $(*)$, is a family of logics designed to express strategic properties of agents in multi-agent systems, notably alternating-time temporal logic [5] and strategy logic [26, 51], considered in their imperfect information incarnations [47]. We envisage to extend these formalism in two directions.

- We first aim at introducing first-order features, including predicates, relations and quantification, on top of these logics for strategies. To our knowledge, this step has been taken in the literature on DaS only with respect to purely temporal logics such as CTL, LTL, and the μ-calculus [34, 33]. On the other hand, first-order ATL has not been considered yet. Nonetheless, such an extension is key to express, as an example, that a bidder b can (has a strategy) to raise her bid, unless she has already hit her true value:

  $$\forall x(x = bid_b \wedge x \neq tvalue_b \rightarrow \langle\langle b \rangle\rangle F \exists y(x < y \leq tvalue_b \wedge y = bid_b))) \quad (1)$$

  In (1) we make use of quantification and equality, along with ATL operator $\langle\langle b \rangle\rangle F$. By using similar combinations of operators for strategies, we are able to express important auction-theoretic concepts, including manipulability and collusion (i.e., an agent or group of agents has a strategy to achieve a certain result in the auction).

  From a theoretical viewpoint, we aim at assessing precisely the expressive and computational power of several fragments of these first-order extensions. For instance, when specifying properties of reactive modules comparing two different values of a variable at different moments of the system's execution, including (1), it is sufficient the two-variable fragment with equality and no relation symbol, which might have better complexity [24]. Other relational features we will explore in this

task include (partial, total) orders, as well as operations on natural and real numbers (e.g., successor, min, max, etc.)

- The second extension we envisage regards epistemic logic to allow for the explicit representation of agents' information and knowledge in data-aware systems. Epistemic extensions of temporal logics have been explored since the 80's [36, 50]. Nonetheless, epistemic extensions of logics for strategies are much less investigated and a topic worth pursuing in itself. In this task we focus on the lesser endeavour of extending the first-order strategy logics developed at the previous point with epistemic operators for individual and group knowledge to express how the information agents possess affects the evolution of DaS, as well as how the information agents have about the global state of the system changes according to their actions.

  For illustrative purposes, consider the following specification stating that the true value of each bidder b is secret to all other bidders, and they cannot (have no strategy) to discover it:

$$\neg\exists x \bigvee_{j\neq b} K_j\ (x = tvalue_b) \wedge \neg\exists x \langle\langle Ag \setminus \{b\}\rangle\rangle \bigvee_{j\neq b} K_j\ (x = tvalue_b) \quad (2)$$

  where K_j is the knowledge operator for bidder j.

Outcome: A family of first-order logics for agents in data-aware systems, capable of expressing their strategic bahviour as well as their individual and group knowledge of agents. These formal languages will be interpreted on the reactive modules in T1.1, then evaluated according their expressive power and their computational properties.

2.2.3 T1.3 – A Formal Framework for Auctions (2 months)

This task will be developed in parallel with T1.1 and T1.2, and it is designed to keep the formal framework in line with the intended application to auction-based mechanisms. Specifically, this task is devoted to the formal analysis of popular models of auctions, starting with simple scenarios such has ascending bid (English) auctions, sealed auctions, Dutch and Vickrey auctions. We will assess the requirements in terms of abilities of agents and how to model them within the framework of reactive modules. For instance, we saw that in English auctions agents compare values of bids, base values, and true values. So, a notion of (partial, total) order on values is required. From these relatively simple examples we will move on to more elaborate scenarios. The aim of the task is to model real-life auctioning mechanisms such as the real-time bidding adopted in online advertisement [52]. Since real-time bidding leverages on complex machine learning algorithms, it is yet to be checked to what extent it can be captured in the framework described

in T1.1 and T1.2. It might well be that only particular cases of real-time bidding are amenable to formal verification. Hence, in this task we will also validate the effectiveness of the approach, by identifying classes of auctions that can be captured in our agent-based logical framework.

Outcome: An assessment of the proposed approach based on data-aware systems with respect to the modelling of popular auctioning mechanisms. Feedback on the formal accounts in T1.1 and T1.2

Deliverables: At the end of WP1 we will produce a technical report with the results in T1.1, T1.2, and T1.3. In particular, we aim at editing three papers with the results of WP1, to be submitted to major conferences in multi-agent systems, knowledge representation, and formal verification (e.g., AAMAS, IJCAI, AAAI, ECAI, KR,...), structured as follows:

- Infinite-state reactive modules for modelling auction-based mechanism, including use cases for the evaluation of the model checking tool SVeDaS to be developed in WP3;
- First-order extensions of logics for strategies to express manipulability and collusion in auctions;
- An epistemic extension of first-order strategy logics to express knowledge and secrecy in auctions.

2.3 WP2 – Model Checking Data-aware Systems (12 months)

This workpackage is devoted to the analysis of the structural, formal properties of the agent-based models for data-aware systems developed in WP1, particularly in relation to verificational issues and the model checking problem, in order to provide sound theoretical underpinnings to WP3. WP2 is structured in three main tasks, whose common objective is to explore ways to make the model checking problem for data-aware systems amenable to practical verification. We begin by checking to what extent the results of the ACSI project apply to infinite-state reactive modules. We anticipate that essential features of auctions call for novel methodologies and techniques, that will be developed within WP2. Results along this line will have an impact on the development of the model checker SVeDaS in WP3. Also this workpackage will be developed by the joint effort of the PI and RA; the RA is expected to work more independently in WP2 and to make personal contributions to the analysis of the model checking problem within this context.

2.3.1 T2.1 – Model Checking Reactive Modules (4 months)

In the state-of-the-art we remarked that two key assumptions underlying the results of the ACSI project are *uniformity* and *boundedness*. These features are shared by reactive modules. Indeed, in systems of reactive modules the global state is completely described by assignments of values to all agents' variables. Since agents are assumed to be in finite number and each of them controls a finite number of variables, the overall number of variables, and therefore the total number of active elements providing values to these variables, is not only finite and bounded, but actually fixed at design time. Furthermore, in many cases of interest, including English auctions, a relational language supporting only comparison $\leq$ between values, is sufficient to describe protocols and actions of the agents taking part in the auction as well as the evolution of the system (compare the specification of a bidder in T1.1). A formalism so defined satisfies the uniformity constraint as well, and therefore the ACSI results are applicable in principle, even though not naively: the specification languages in [18, 8, 32] are temporal logics ranging from CTL to the μ-calculus, where no explicit account of the agents' strategic abilities is provided. We anticipate that the model checking problem for infinite-state reactive modules against specifications in first-order ATL can be proved decidable by definining suitable alternating bisimulations [6] in a quantified context. Such preliminary results, while of interest in themselves, will help to assess to what extent these methodologies and techniques applies to the setting of the SVeDaS project.

Outcome: Alternating bisimulation relations in a quantified setting that preserve the interpretation of formulas in first-order ATL. A decidability result for particular classes of reactive module obtained by finite abstraction.

2.3.2 T2.2 – Arithmetic Operations (4 months)

A key feature to describe data-aware systems in general and auctions in particular, is the support of arithmetic operations in the language [34]: these are necessary, for instance, to describe the complex mechanisms to update prices and bids in real-time bidding, depending on previous registered values. However, by extending the logical framework with this extra expressiveness the uniformity condition fails. Specifically, it is possible to prove that model checking non-uniform systems is indeed undecidable in the general case [17]. Hence, arithmetic operations in the description and specification languages for elaborate auctions raise a number of questions as regards formal verification, particularly in relation to the existence of finite abstractions for infinite-state reactive modules. In this direction, [34] presents some valuable results. However, the specification language considered, LTL-FO, is not suitable to express subtle strategic interactions.

In this task we explore possible solutions to the issue above, to allow for arithmentic operations in DaS. Specifically, an approach based on the counterpart semantics for modal logic [31] has proved to give some interesting, preliminary results on the existence of finite simulations of data-aware systems [12]. The intuition behind counterpart semantics for first-order modal logics is that the accessibility relation between possible worlds goes together with a counterpart relation between elements in the interpretation domain that belong to different possible worlds. For our purposes, this idea allows to define finite abstractions that simulate the behaviour of the concrete, infinite state DaS, by collapsing different elements in possibly different states into a unique abstract individual, together with a counterpart relation. Notice that simulations preserve only a fragment of the specification language, typically the universal fragment. Nonetheless, since we plan to investigate various notions of simulation and abstraction (T2.3), in selected cases counterpart semantics might allow to preserve the whole of our specification language.

Outcome: An extension of the specification language for reactive modules with arithmetic operations. A diagnosis on the failure of uniformity in this setting. A counterpart semantics for reactive modules. A notion of simulation and finite abstractions for counterpart semantics. An application to the modelling of (a suitable restriction of) real-time bidding.

2.3.3 T2.3 – Bisimulations and Abstractions (4 months)

In previous works by the applicant Dr Belardinelli, among others [11, 12, 13, 14, 18, 19] the model checking problem for infinite-state DaS is proved decidable via reduction to the finite case by means of truth preserving bisimulations. More specifically, given a concrete, infinite-state DaS that satisfies the boundedness and uniformity contraints, we are able to show the existence of a finite-state abstraction that is bisimilar to the concrete DaS. Further, the bisimulation relation preserves formulas written in first-order extensions of temporal logics, such as CTL and the μ-calculus. As a result, we can verify a specification on the infinite-state DaS by model checking the same formula in the finite abstraction, by using standard techniques for finite-state systems.

In this task we aim at exploring abstraction methodologies for cases not covered by standard techniques. We remarked above that for language with arithmetic operation, the uniformity condition no longer holds. So, we plan to define more robust notions of (bi)simulation and abstraction that also work for non-uniform systems. In this direction we have some promising preliminary results [17] that yet need to be developed further into a proper methodology. Moreover, modular techniques, including agent-based simulations and data symmetry reduction, can also find application in this setting to tackle the model checking problem.

A further notion of (bi)simulation and abstraction we intend to explore is represented by many-valued simulations [10]. Also in this line the intuition is to trade precision of the verification result with efficiency. Specifically, the model checking procedure for a given specification with respect to a multi-valued simulation of a concrete system is supposed to terminate and to return an answer that applies to the concrete system as well. However, the value of this answer might be undefined, other than simply true or false. Then, we anticipate to investigate refinement techniques [27] to make the abstraction more precise and possibly obtain a definite answer.

Outcome: Result on the decidability and complexity of the model checking problems for data-aware systems based on reactive modules. General abstraction techniques for DaS, including multi-valued simulations and abstraction. Abstraction refinement techniques.

Deliverables: Also for WP2, we plan to collect the main findings in a technical report, to be made publicly available. Moreover, we aim at submitting three more papers at top conferences on MAS, KR, and formal methods. Specifically, the three contributions will deal with

- Model checking infinite-state reactive modules against specifications in first-order logics for strategies and first-order epistemic ATL, with complexity results. Applications to simple auctioning mechanisms.
- The model checking problem for first-order strategy logics supporting (fragments of) arithmetic. (Un)decidability results, (bi)simulation relations, and (possibly finite) abstractions.
- Many-valued simulations and abstractions, counterexample-guided refinement techniques.

2.4 WP3 – Tools and Techniques for DaS Verification (10 months)

In this workpackage we intend to develop practical and efficient verification techniques to model check the DaS models introduced in WP1, by making use of the theoretical results in WP2. Further, we will implement these methods and procedures into a novel verification tool – SVeDaS. Leveraging on structural properties of DaS to obtain a decidable model checking problem has been preliminarily explored by the applicant Dr Belardinelli [19]; these early results look promising. Unfortunately, mere decidability is not sufficent for real-life application, particularly complex real-time bidding scenarios. Hence, we will exploit the agent-based model checking techniques developed in WP2 in order to alleviate the *state explosion problem*, i.e., the exponential blow-up of the state space in relation to the number of modules.

In this respect, the compositional, agent-based approach to DaS modelling and verification put forward in WP1 and WP2 will contribute to obtain efficient model checking algorithms, and to scale these up to the large number of processes usually found in business use cases, including auctions. It is envisaged that the verification tool SVeDaS will be built upon the MCMAS model checker that Dr Belardinelli contributed to develop while working at Imperial College London [49]. The RA will be in charge for the actual coding, always under the guidance of the PI. The tool will be evaluated against the auction-based mechanisms analysed in WP1 and distributed as open source.

2.4.1 T3.1 – Extended ISPL and Symbolic Data Structures (3 months)

The interpreted system programming language (ISPL) is the specification language of the MCMAS model checker [49]. The first task to build our verification tool SVeDaS will be to extend ISPL to support the relational and first-order features described in WP1 and WP2.

Further, as it is customary in symbolic model checking, including MCMAS, the system's components are not represented and manipulated explicitly by the model checker; rather these are encoded symbolically. In this task we also investigate data structures to represent data-aware systems. Our starting point will be ordered binary decision diagrams (OBDDs), as this representation is used in the MCMAS model checker, as well as a number of other efficient model checking tools. OBDDs are widely used to encode finite data structures, such as the components of the ISPL modelling language. We will evaluate how OBDDs scale up in the representation of data-aware systems, including their interpretation domain. In particular, since data are exhibited explicitly in the system's description, we plan to explore other promising symbolic data structures, specifically, sentential decision diagrams (SDD) and decomposable negation normal form (DNNF) [54].

Outcome: An extended version of the ISPL programming language suitable to specify data-aware structures, including auction-based mechanisms. Compact and efficient symbolic data structures for the representation of DaS, including their data content.

2.4.2 T3.2 – Efficient Model Checking Algorithms (4 months)

In this task we will investigate efficient model checking algorithms for data-aware systems. This task is actually composed of two distinct subtasks. On one hand, we anticipate that WP2 will provide us with decidability and

complexity results on the verification of DaS. However, previous contributions in the literature seem to point to EXPSPACE-completeness in the size of data for the complexity of the model checking problem [19]. This is normally considered well beyond what is acceptable for efficient model checking. Hence, a first task is to identify relevant cases in which the complexity can be lower than the general case. For instance, in many relevant cases (including English auctions among them), agents share the same protocol and actions (e.g., all bidders share a common agent type, with a unique protocol and actions as described in T1.1). As a consequence, we plan to leverage on symmetries in data-aware systems, including simmetries on agents, to abstract even further concrete DaS, thus reducing the complexity of verification, if not in principle, at least for practical model checking.

On the other hand, we will develop the symbolic model checking algorithms to be implemented in the SVeDaS model checker. We anticipate that these will be based on the algorithms for symbolic model checking logics for strategies with imperfect information [5, 49, 25]. However, the relational and first-order features of the extended ISPL programming language for DaS call for suitable extensions capable of dealing with data. These algorithms will then be implemented in the SVeDaS model checker.

Outcome: Efficient model checking algorithms for data-aware systems. The SVeDaS model checker for the automated verification of DaS, including aution-based mechanisms.

2.4.3 T3.3 – Evaluation (3 months)

This task is dedicated to an evalution of the SVeDaS model checker against the auctioning scenarios developed in WP1. We will compare the performace of the tool against other model checkers for multi-agent systems, including MCMAS, MCK [37], and VerICS [48]. It is not yet clear whether this comparison can be extended to full data-aware systems, as these tools do not explictly support DaS either in the modelling or in the specification language. Nonetheless, we aim at using SVeDaS to certify popular auctioning mechanisms, including English, Dutch, and Vickrey auctions, real-time bidding, against malicious behaviours, such as collusion and manipulability of partecipating agents. We anticipate that in the most complex cases (e.g., real-time bidding) only partial certification is achievable, either for specific subclasses of auctions or for selected properties.

Outcome: An evaluation of the performance of the SVeDaS model checker. (Partial) certification of popular auction-based mechanisms against malicious behaviours.

Deliverables: The findings of WP3 will find their way in a technical report, and the SVeDaS model checker will be made publicly available under the GNU licence for open-source software.

Also, we plan to submit two conference papers on the results of WP3:

- a first contribution presenting the tool and its theoretical underpinnings for model checking general data-aware systems, including evaluation;
- a targeted contribution on the certification of auctioning mechanisms.

Finally, we will submit a paper reporting on the results of the SVeDaS project to a top journal in artificial intelligence and multi-agent system (AIJ, JAIR). The results of the whole project will also be presented at the workshop planned for its conclusion.

2.5 The Principal Investigator – Collaborations

The SVeDaS project originates from the research interests of the applicant, Dr Francesco Belardinelli, and benefits from his ongoing collaborations with the project partners: Prof Catalin Dima, Dr Umberto Grandi, Dr Davide Grossi, Prof Wojtek Jamroga, Prof Alessio Lomuscio, and Prof Wiebe van der Hoek, who are well-known scholars working in research institution in France, the UK, and Poland. While the project partners will not directly receive fundings from the SVeDaS project, they will collaborate with Dr Belardinelli on themes pertaining to the project's objectives and might be available to provide guidance to the RA involved in the project. The project partners have complementary expertises, given by their different approaches to formal methods in computer science. This mix is key for the success of the SVeDaS project. The personal relationship between Dr Belardinelli and the project partners is excellent, thus no issue is likely to arise during the project's lifetime.

Dr Belardinelli (Principal Investigator) is *maître de conférences* in computer science at the *Université d'Evry*. He received his *diploma di licenza* and PhD from *Scuola Normale Superiore* in Pisa, where he was awarded the highly selective SNS scholarship (€60k over four years). He has published extensively on the formal properties of modal logics (e.g., epistemic, temporal, strategy, etc.) and their applications to the specification and verification of multi-agent systems [14, 16, 18, 19, 20]. Among his achievements, in 2009-11 Dr Belardinelli was recipient of an IEF Marie Curie fellowship for the FoMMAS project (*First-order Modal Logics for the Specification and Verification of Multi-Agent Systems*), that he personally developed at the *Department of Computing*, *Imperial College London* (£160k over 2 years). In addition to FoMMAS, Dr Belardinelli worked on the EU STREP projects ACSI (Artifact-centric System Interoperation), for which he was responsible for multiple tasks including managing the project budget and liaising

with the project partners. Indeed, the ACSI project, which partly inspires the present proposal, hinged on a specific class of data-aware systems – the *artifact-centric systems* – and sanctioned the validity of the approach. The project results have been presented in major international conferences (IJCAI, ICSOC, KR) and top journals (JAIR); the project itself has been reviewed as *excellent* on completion. More recently, in 2015 the applicant has received a PEDR grant by the French Ministry of Education (4000 euros/year over 4 years), in recognition of my scientific achievements over the past 4 years. Given his track record, Dr Belardinelli has the capability to successfully develop and correctly manage the SVeDaS project.

Among the collaborators, Prof van der Hoek and Dr Grossi from the University of Liverpool both share an interest in knowledge representation and reasoning, specifically epistemic logics. Their experties will be particularly valuable in T1.2. In 2013 and 2015, Dr Belardinelli has received support from the *Université d'Evry* in form of *Fonds pour le Rayonnement de la Recherche* for collaborative projects with Prof van der Hoek and Dr Grossi, that contributed to the development of the SVeDaS project.

Further, Prof Dima and Prof Jamroga are both well-known for their work in the area of multi-agent systems, including logics for the strategic behaviour of agents and their formal verification. This collaboration will be key for both WP1 and WP2.

Dr Grandi is *maître de conférences* at IRIT and *Université Toulouse 1 Capitole*. His domain of expertise is related to computational social choice, including auctions and auction-based mechanisms analysed through game theory. So, he will be able to contribute to T1.3 among others.

Finally, Dr Belardinelli has a long established collaboration with Prof Lomuscio, Imperial College London, on themes pertaining to the MAS verification by model checking. Prof Lomuscio was also one of the principal investigators for the ACSI project, so his expertise will be valuable both in connection with data-aware systems and the development of WP3.

By building on his previous works related to the SVeDaS project, and by combining the complementary expertises of the project partners, Dr Belardinelli is well-positioned to develop successfully the proposed research programme. In particular, the justification of the required resources of €173,971k over 3 years is as follows: €110k for a PhD scholarship over 3 years; €20k to cover travel costs and fees to attend conferences and disseminate the project's results; €9k to fund 3 MSc projects on topics related to the SVeDaS project; €6k for equipment for the applicant and PhD student; €5k to organise a one-day workshop at the conclusion of the projet to disseminate the project's results. To the partial total of €167,280k it must be added 4% of overheads for project management required by the *Université d'Evry*.

3 Project Impact

The formal analysis and verification of data-aware systems is ever more relevant to ensure the efficient and reliable composition and interoperation of business processes. Current approaches to business process management are primarily centered around activity flows and often neglect the role played by data in the system's execution. In contrast, data-aware systems are focused on the combined perspective of *data models* and *business processes*. Data are visible and accessible to agents, possibly in a controlled way through some permission restrictions; they directly account for the system's evolution and can be exhibited explicitly in the system's specification. These considerations apply to auctioning processes as well: one original tenet of the SVeDaS project is to model auction-based mechanisms as data-aware systems. For the effective deployment of DaS, including auction in e-markets, verification and validation methodologies are essential. The SVeDaS project takes inspiration from the state-of-the-art in the application of formal methods to data-aware systems, and aims at developing a tailored methodology for their modelling, analysis and verification, then to apply these techniques to the formal certification of auctions.

The relevance of the SVeDaS research programme with respect to the strategic character of the ANR call cannot be overestimated. The SVeDaS project is designed to contribute to the ANR challenge *Société de l'Information et de la Communication* in the designated areas. The first objective of SVeDaS is to provide novel agent-based models for the analysis of data-aware systems, in order to improve the understanding of mechanisms related to service composition and interoperability. The proposed framework will be inspired to the state-of-the-art in multi-agent systems, and it will be evaluated against capturing complex auctioning scenarios and their dynamics in e-markets. Although preliminary results in this direction have already appeared, including some by the applicant [11, 12, 13, 14, 18, 19], an in-depth formal analysis of data-aware systems composed of agents that operate in a strategic way has not been tackled yet to our knowledge, nor efficient model checking tools have been developed. Fundamental results along this line will contribute to the topic *Fondaments du Numérique*. Moreover, the model checking tool and techniques developed in WP3 will be applied to the verification of multi-party auction-based use cases, including (possibly restricted versions of) real-time bidding. To achieve such a result, one of the typical obstacles to the large-scale deployment of automated verification techniques is the *state explosion problem*, i.e., the exponential blow-up of the state space. By adopting an agent-based, modular approach to the representation and analysis of data-aware systems we anticipate to be able to alleviate the state explosion problem considerably. Thus, the SVeDaS project will also contribute incidentally to the topic *Sciences et Technologies Logicielles*.

The dissemination strategy is designed to output two-three papers for each year of the project lifetime. These contributions will be submitted to major conferences in multi-agent systems (AAMAS, PRIMA), artificial intelligence (IJCAI, ECAI, AAAI), knowledge representation (KR, TARK), and formal verification (CAV, ATVA, CONCUR). Given the applicant's track record in publications, we deem this dissemination programme feasible. Further, the SVeDaS model checker will be made publicly available under the GNU licence for open-source software. Finally, to present the project's findings, a one-day workshop will be held at the conclusion of the project's lifetime.

The timing of the SVeDaS project is excellent for the career development of the applicant, Dr Francesco Belardinelli, as well as for *Laboratoire IBISC* of the *Université d'Evry*. The research programme originates from Dr Belardinelli's expertise on the specification and verification of multi-agent systems [16, 19, 20]. The *Jeunes Chercheuses/Jeunes Chercheurs* ANR Programme constitutes an unparalleled opportunity for the applicant to develop his academic profile and to strengthen the reputation of *Laboratoire IBISC* as a research institution. The budget of the SVeDaS project will be used mainly to fund a PhD position during the three years of the project's lifetime. This position will be available for a graduate student with a suitable profile to contribute to the main project objectives from the beginning of the studentship. He/she will work in close collaboration with Dr Belardinelli to jointly develop the research programme as outlined in Section 2. The PhD student will also contribute to strengthen the long-term, strategic positioning of *Laboratoire IBISC* in the area of formal methods.

The SVeDaS project will allow Dr Belardinelli and *Laboratoire IBISC* to acquire greater visibility within the French and international community in formal methods. The applicant already benefits from several ongoing collaborations with renown scholars, both in French research institutions and abroad (IRISA Rennes, LIP6 Paris, LACL Paris, IRIT Toulouse, Imperial College London, University of Liverpool, Università di Roma). The SVeDaS project will provide the financial support to strengthen these collaborations and to develop further research projects on related themes. In particular, by the end of the SVeDaS project Dr Belardinelli will be in a position to submit a project proposal to an appropriate scheme of the European Research Council on a topic related to the more general subject of the modelling and analysis of open dynamic systems.

References

[1] EU STREP Project ACSI. www.acsi-project.eu.

[2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-W., 1995.

[3] R. Alur et al., MOCHA: Modularity in model checking. In *Proceedings of CAV'98*, pp. 521–525. Springer-Verlag, 1998.

[4] R. Alur and T. A. Henzinger. Reactive modules. *Form. Methods Syst. Des.*, 15(1):7–48, 1999.

[5] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[6] R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In *In Proceedings CONCUR'98*, pp. 163–178. Springer-Verlag, 1998.

[7] A. Badica and C. Badica. Specification and verification of an agent-based auction service. *Information Systems Development*, pp. 239–248. Springer US, 2010.

[8] B. Bagheri Hariri et al., Verification of Relational Data-centric Dynamic Systems with External Services. In *Proceedings of PODS'13*, pp. 163–174. ACM, 2013.

[9] C. Baier and J.-P. Katoen. *Principles of Model Checking*. The MIT Press, 2008.

[10] T. Ball and O. Kupferman. An abstraction-refinement framework for multi-agent systems. In *Proceedings of LICS06*, pp. 379–388. IEEE, 2006.

[11] F. Belardinelli. Model checking auctions as artifact systems: Decidability via finite abstraction. In *Proceedings of ECAI14*, pp. 81–86. IOS Press, 2014.

[12] F. Belardinelli. Verification of non-uniform and unbounded artifact-centric systems. In *Proceedings of AAMAS'14*, pp. 717–724. IFAAMAS/ACM, 2014.

[13] F. Belardinelli and D. Grossi. On the formal verification of diffusion phenomena in open dynamic agent networks. In *Proc. of AAMAS15*, 2015.

[14] F. Belardinelli, D. Grossi, and A. Lomuscio. Finite abstractions for the verification of epistemic properties in open multi-agent systems. In *Proc. of IJCAI15*, 2015.

[15] F. Belardinelli and A. Lomuscio. Interactions between time and knowledge in a first-order logic for multi-agent systems. In *Proceedings of KR*. AAAI Press, 2010.

[16] F. Belardinelli and A. Lomuscio. Interactions between knowledge and time in a first-order logic for multi-agent systems. *JAIR*, 45:1–45, 2012.

[17] F. Belardinelli and A. Lomuscio. Decidability of model checking non-uniform artifact-centric quantified interpreted systems. In Rossi [59].

[18] F. Belardinelli, A. Lomuscio, and F. Patrizi. An Abstraction Technique for the Verification of Artifact-Centric Systems. In *Proc. of KR'12*, pp. 319 – 328, 2012.

[19] F. Belardinelli, F. Patrizi, and A. Lomuscio. Verification of agent-based artifact systems. *Journal of Artificial Intelligence Research*, 51:333–77, 2014.

[20] F. Belardinelli and W. van der Hoek. Epistemic quantified boolean logic. In *Proceedings of IJCAI15*, pp. 2748–2754. AAAI Press, 2015.

[21] K. Bhattacharya, C. E. Gerede, R. Hull, R. Liu, and J. Su. Towards Formal Analysis of Artifact-Centric Business Process Models. In *Proc. of BPM*, 2007.

[22] N. Bulling et al., Model checking logics of strategic ability: Complexity*. In *Specification and Verification of Multi-agent Systems*, pp. 125–159. Springer US, 2010.

[23] N. Bulling and W. Jamroga. Comparing variants of strategic ability. *Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.

[24] E. Borger et al., *The Classical Decision Problem*. Springer, 1997.

[25] P. Cermák et al., MCMAS-SLK: A model checker for the verification of strategy logic specifications. In *Proceedings of CAV 2014*, pp. 525–532. Springer, 2014.

[26] K. Chatterjee et al., Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.

[27] E. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement. In *Proceedings of CAV 2000*, pp. 154–169. Springer, 2000.

[28] E. M. Clarke et al., *Model Checking*. The MIT Press, 1999.

[29] D. Cohn and R. Hull. Business Artifacts: A Data-Centric Approach to Modeling Business Operations and Processes. *IEEE Data Eng. Bull.*, 32(3):3–9, 2009.

[30] J. M. Corera, I. Laresgoiti, and N. Jennings. Using archon, part 2: Electricity transportation management. *IEEE Expert*, 11(6):71–79, 1996.

[31] G. Corsi. *Counterparts and possible worlds*, CLUEB, Bologna, 2001.

[32] G. De Giacomo et al., Bounded epistemic situation calculus theories. In Rossi [59].

[33] A. Deutsch, R. Hull, F. Patrizi, and V. Vianu. Automatic verification of data-centric business

processes. In *Proceedings of ICDT*, pp. 252–267. ACM, 2009.
[34] A. Deutsch, Y. Li, and V. Vianu. Verification of hierarchical artifact systems. http://arxiv.org/abs/1604.00967, 2016.
[35] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets*. CUP, 2010.
[36] R. Fagin et al., *Reasoning about Knowledge*. MIT Press, 1995.
[37] P. Gammie and R. van der Meyden. MCK: Model checking the logic of knowledge. In *Proceedings of CAV'04*, pp. 479–483. Springer-Verlag, 2004.
[38] P. Gonzalez, A. Griesmayer, and A. Lomuscio. Verifying GSM-based business artifacts. In *Proceedings of ICWS12*, pp. 25–32. IEEE Press, 2012.
[39] V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004.
[40] O. Grumberg and D. E. Long. Model checking and modular verification. *ACM Transactions on Programming Languages and Systems*, 16(3):843–871, 1994.
[41] E. M. T. F. Guerin and W. Vasconcelos. Abstractions for model-checking game-theoretic properties of auctions. In *Proceedings of AAMAS '08*, pp. 1613–1616, 2008.
[42] J. Halpern and M. Vardi. The complexity of reasoning about knowledge and time. In *Proceedings of STOC '86)*, pp. 304–315, ACM Press.
[43] J. Halpern and M. Vardi. The complexity of reasoning about knowledge and time 1: lower bounds. *Journal of Computer and System Sciences*, 38(1):195–237, 1989.
[44] W. v. Hoek et al., On the complexity of practical atl model checking knowledge, strategies, and games in multi-agent systems. In *Proc. of AAMAS'06*, ACM Press.
[45] R. Hull. Artifact-centric business process models: Brief survey of research results and challenges. In *Proceedings of OTM Conferences (2)*, pp. 1152–1163. Springer, 2008.
[46] R. Hull et al., Business artifacts with guard-stage-milestone lifecycles. In *Proceedings of DEBS '11*, pp. 51–62, 2011. ACM.
[47] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 62:1–35, 2004.
[48] M. Kacprzak et al., Verics 2007 - a model checker for knowledge and real-time. *Fundamenta Informaticae*, 85(1-4):313–328, 2008.
[49] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: A model checker for the verification of multi-agent systems. In *Proceedings of CAV*, pp. 682–688. Springer, 2009.
[50] J.-J. C. Meyer and W. Hoek. *Epistemic Logic for AI and Computer Science*, CUP, 1995.
[51] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log.*, 15(4):34:1–34:47, 2014.
[52] N. Nisan et al., Google's auction for TV ads. In *Proceedings of ICALP 2009*, pp. 309–327. Springer, 2009.
[53] E. Nooijen et al., Automatic Discovery of Data-Centric and Artifact-Centric Processes. In *Proceedings of BPM 2012*, pp. 316–327. Springer, 2013.
[54] U. Oztok and A. Darwiche. A top-down compiler for sentential decision diagrams. In *Proceedings of IJCAI*, pages 3141–3148, 2015.
[55] M. Pauly. A modal logic for coalitional power in games. *J. Log. Comput.*, 12(1):149–166, 2002.
[56] W. Penczek and A. Lomuscio. Verifying epistemic properties of multi-agent systems via bounded model checking. *Fundamenta Informaticae*, 55(2):167–185, 2003.
[57] A. S. Rao and M. P. Georgeff. Modeling rational agents within a BDI-architecture. In *Proceedings of KR*, pp. 473–484. Morgan Kaufmann Publishers, 1991.
[58] D. M. Reeves et al., Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, 39(1):67–85, 2005.
[59] F. Rossi, editor. *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*. IJCAI/AAAI, 2013.
[60] Y. Shoham and K. Leyton-Brown. *Multiagent Systems*. CUP, 2008.
[61] M. P. Singh and M. N. Huhns. *Service-Oriented Computing*. Wiley & Sons, 2005.
[62] M. Webster, L. Dennis, and M. Fisher. Model-checking auctions, coalitions and trust. Technical report, University of Liverpool, 2009.
[63] M. Wooldridge. *An introduction to multi-agent systems*. John Wiley, England, 2002.
[64] H. Xu et al., Real-time model checking for shill detection in live online auctions. In *Software Engineering Research and Practice*, pp. 134–140. CSREA Press, 2009.
[65] H. Xu and Y. Cheng. Model checking bidding behaviors in internet concurrent auctions. *Comput. Syst. Sci. Eng.*, 22(4), 2007.