

Formal Verification of Storm Topologies through D-VerT

Francesco Marconi, Marcello M. Bersani, Matteo Rossi
DEIB - Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano
Piazza Leonardo da Vinci 32, 20133 Milan
{francesco.marconi, marcellomaria.bersani, matteo.rossi}@polimi.it

ABSTRACT

Data-intensive applications (DIAs) based on so-called Big Data technologies are nowadays a common solution adopted by IT companies to face their growing computational needs. The need for highly reliable applications able to handle huge amounts of data and the availability of infrastructures for distributed computing rapidly led industries to develop frameworks for streaming and big-data processing, like Apache Storm and Spark. The definition of methodologies and principles for good software design is, therefore, fundamental to support the development of DIAs. This paper presents an approach for non-functional analysis of DIAs through D-VerT, a tool for the architectural assessment of Storm applications. The verification is based on a translation of Storm topologies into the CLTL_{oc} metric temporal logic. It allows the designer of a Storm application to check for the existence of components that cannot process their workload in a timely manner, typically due to an incorrect design of the topology.

CCS Concepts

•**Theory of computation** → *Verification by model checking*; •**Software and its engineering** → *Model-driven software engineering*;

Keywords

Formal Verification; Apache Storm; MDE; Data-intensive Applications; Temporal Logic

1. INTRODUCTION

Data-intensive applications (DIAs) are computational systems that process, in a relative small amount of time, huge amounts of diversified information usually produced by data sources with high throughput. Some of the most popular companies nowadays—e. g., Twitter (www.twitter.com),

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC 2017, April 03-07, 2017, Marrakech, Morocco

Copyright 2017 ACM 978-1-4503-4486-9/17/04...\$15.00

<http://dx.doi.org/10.1145/3019612.3019769>

Groupon (www.groupon.com), Spotify (www.spotify.com), etc.—make large use of DIAs to process data gathered from millions of users.

DIAs constitute a significant asset for the production of large-scale software, and have been drawing the attention of both academia and industry. The creation of frameworks that support designers over the entire life-cycle (design, development, testing, deployment, maintenance) of DIAs is of crucial importance, and constitutes a key research challenge in this area. Topics such as techniques and tools for quality assessment, architecture enhancement, agile delivery and continuous testing of DIAs are targeted by ongoing research projects like, for instance, the DICE European project [6].

The design approach envisioned by DICE is founded on model-driven principles and can be summarized as follows. The design of an application is decomposed into three distinct and consecutive phases, each one associated with a profiled UML diagram. Each phase focuses on a specific aspect of the design and represents a refinement of the previous one that has to be validated before starting the new refinement step. If design flaws are detected, designers can either change the current model, or modify the one built in the previous step, then redo the refinement. The design process starts from a conceptual model of the application, called Platform-Independent Model (PIM); this is refined, in the second step, into the so-called Platform-Specific Model (PSM), which provides the architectural schema of the application based on a specific (data-intensive) technology; finally, in the last step, the architectural model is refined to obtain a deployment model. We approach the assessment of DIAs by applying formal verification to the architectural models described through (metric) temporal logic, according to the technique presented in [11]. The goal of the analysis is to determine, through automated techniques, whether the behavior entailed by the architecture of the application conforms to specific properties over time. The properties that an application should satisfy typically depend on the technology adopted to implement the application.

Most of the available data-intensive frameworks allow designers to specify the architecture of DIAs as a directed graph whose nodes are computational resources which carry out specific operations. The semantics underlying a graph, which reflects the runtime behavior of the application, is determined by the target technology (e. g., the same graph has two different interpretations in case we adopt a streaming

or a batch technology). In this paper, we consider Apache Storm [1], a popular technology for stream-based applications. The architecture of a Storm application is defined by means of a *topology*—i.e., a directed graph—where nodes are of two kinds: *computational nodes*, which implement the logic of the application by elaborating information and producing an outcome; and *input nodes*, which bring information into the application from its environment.

This paper presents the complete verification workflow of an architectural model defined through the concepts included in the DICE UML profile [8]. It focuses also on all the necessary transformations needed for translating the UML diagram of the architecture of the DIA into a formula of the CLTL_{oc} metric temporal logic [5], which is solved to validate the Storm architecture represented by the UML model. Finally, it presents the verification tool, called D-VerT [2], which is the component implementing the transformations.

The paper is structured as follow: Sect. 2 presents some background notions on Apache Storm and briefly recaps our approach to the modeling of Storm topologies with temporal logic introduced in [11]. Sect. 3 introduces the methodology for the verification of Storm topologies based on formal validation of refined UML models. Sect. 4 describes the structure of D-VerT and the transformations needed for enabling the verification of architectural models. Sect. 5 shows the application of the methodology through an example of Storm application which, at the end, undergoes verification with D-VerT. Sect. 6 briefly discusses some related works, and Sect. 7 concludes.

2. BACKGROUND

2.1 Apache Storm

Apache Storm [1] is a stream processing system, developed and open sourced by Twitter in 2012, allowing real-time processing of large-scale streaming data on horizontally scalable systems through a parallel and distributed computation.

The computational model of Storm applications is the Storm topology, i.e., a directed graph whose nodes realize the operations performed over the data flowing through the application and whose edges indicate how such operations are combined. Data are structured into streams that are infinite sequences of tuples, i.e., strings of characters, that are processed by the application.

A topology node set consists of *spouts* and *bolts* (in the following also referred to as *topology components*). Spouts are stream sources which usually get data from external systems, such as queuing brokers (e.g., Kafka, RabbitMQ, Kestrel), or from other data sources, e.g., Twitter Streaming APIs, whereas bolts transform the incoming data streams into new output streams to be processed by the connected bolts. Connections are statically defined at design time by the subscription of the bolts to other spouts or bolts. Fig. 1 shows an example of Storm topology that will be used in Sect. 5.

Storm is capable of guaranteeing the so-called “at least once” message processing. Reliable spouts keep track of all the tuples they emit, and if one of them fails to be processed by the entire topology within a certain timeout, then the spout




Figure 1: Example of Storm topology.

re-emits it into the topology. When message processing is “best-effort”, instead, (unreliable) spouts emit each tuple only once, without checking for the successful completion of the processing. Bolts usually perform operations, such as filtering, join, functions, database interaction, which are combined through the topology to perform complex transformations. The Java interfaces defining a bolt include the `execute()` method that implements its functionality; it reads the input tuples, processes the data, and emits (via the `emit()` method) the transformed tuples on the output streams.

The Storm runtime is designed to leverage the computational power of distributed clusters. A deployed topology is composed of one *master node* and several *worker nodes* running one or more *worker processes*, each of them executing different parts of the same topology. Each worker process runs a JVM where one or more *executors* (i.e. threads) are spawned. Executors can run one or more *tasks* which, in turn, can execute a spout or a bolt. The configuration of the topology defines the number of worker processes and, for each component (spout or bolt), the number of executors running it in parallel and the total number of tasks over those executors. Communication among workers and executors is managed through a multi-level queuing system. Each executor has its own input queue and output queue. Tuples are read from the input queue and processed by the thread handling the spout/bolt logic; afterwards, they are emitted on the outgoing queue and then moved to the parent worker’s transfer queue.

2.2 Modeling Storm topologies

Our verification approach is founded on a formal model of Storm topologies which provides an abstraction of their computations. To this end, we identified the relevant aspects of the executions of a generic topology, and some suitable assumptions that allow us to generate models that can be practically managed by state-of-the-art formal verification tools in a reasonable amount of time. The model of a topology captures how the timing parameters of its components—such as the delays between two consecutive spout events inputting tuples in the topology and the processing time of tuples for each bolt—affect the size of the bolts’ queue. The formal verification assessment of a Storm topology aims at checking for the existence of bolts that cannot process the incoming stream of tuples on time, thus causing a monotonic growth of the size of their queues.

Our approach facilitates the compositional definition of topology models, mimicking how code developers create topolo-

gies. The behavior of the relevant features and parameters of spouts and bolts is extracted by reverse-engineering the Java interfaces of the Storm API using the following assumptions:

- Deployment details are abstracted away; topologies are assumed to run on a single worker process and each executor runs a single task, which is the default Storm configuration of the runtime.
- Each bolt has a single receive queue for all its parallel instances and no sending queue, while the workers' queues are not represented (single-worker scenario). For generality, all queues have unbounded size.
- The contents of tuples is not modeled and, since tuples are all assumed to have the same size, the size of queues is represented by the number of tuples they contain.
- The external sources of information abstracted by the spouts are not represented, since they are outside of the perimeter of the application. So, their queues are not represented.
- For each component, the duration of each operation or the permanence in a given state has a minimum and a maximum time.

Given a Storm topology, its computation is captured through a set of formulae written in the CLTL_{oc} metric temporal logic [5] augmented with a notion of counters, which is used to represent the size of bolts' queues during the computation. The formal model of a topology consists of four parts, which represent: (i) the evolution of the state of the nodes; (ii) the behavior of the queues; (iii) timing constraints; (iv) failures. The full list of CLTL_{oc} formulae capturing the semantics of a Storm topology can be found in [4]; [11], instead, presents some of the technical details of the adopted decision procedure, which is validated through some experimental results.

Let us briefly recall the salient points of the behavior of a bolt. A bolt can alternatively be in one of the following states: *process*, *idle* or *failure*. If a bolt is idle and its queue is not empty, it reads tuples from its queue by an instantaneous *take* action, then it processes the tuples for α time units. This corresponds to the state *execute*, which is part of macro-state *process*, together with states *take* and *emit*. The value α is a parameter of the bolt, a positive real value which represents the amount of time that a bolt requires to process one tuple. Once the execution is completed, the bolt emits output tuples with an instantaneous action corresponding to the *emit* state. Bolts may fail and failures may occur at any moment; upon a bolt failure, the system goes to the *fail* state. If no failure occurs, after an *emit* a bolt goes to *idle*, where it stays until it reads new tuples.

To give a flavor of the formal model underlying our verification approach, we introduce a few examples of CLTL_{oc} formulae. Formulae (1)-(2) capture how the number of elements in the queue of bolt j (q_j) is updated whenever tuples are enqueued (**add_j**) or dequeued (**take_j**). They use N-valued discrete counters to represent the amounts of tuples exchanged in the topology. Term Xq_i represents the value of q_i in the next position of time. Every time the component j emits (**emit_j** holds), the queues of all bolts subscribing to j —i.e., those which are targets of arcs outgoing from j in the topology—receive r_{emit_j} tuples—i.e., the variables q_i




Figure 2: Iterative refinement process supported by the DICE framework.

representing the occupancy level of those queues are incremented by r_{emit_j} . When multiple components subscribed by a bolt emit tuples simultaneously, the increment on its queue is equal to the sum of all the tuples emitted (the value of r_{add_j} in Formulae (1)-(2)). Dually, when **take_j** holds, the occupancy level q_j is decremented by $r_{process_j}$ (number of tuples read by bolt j). Proposition **add_j** is true when at least one of the components subscribed by j is emitting, whereas **startFail_j** is true in the first instant of a failure state.

$$\mathbf{add}_j \wedge \neg \mathbf{take}_j \wedge \neg \mathbf{startFail}_j \Rightarrow (Xq_j = q_j + r_{add_j}) \quad (1)$$

$$\mathbf{take}_j \Rightarrow (Xq_j = q_j + r_{add_j} - r_{process_j}) \quad (2)$$

To measure the duration of each state and to impose timing constraints between events, we use a set of dense-time CLTL_{oc} clock variables [5] for each component of the topology. For example, Formula (3) imposes that when *emit* occurs, the duration of the current processing phase is between $\alpha - \epsilon$ and $\alpha + \epsilon$, where $\epsilon \ll \alpha$ is a positive constant that captures possible (small) variations in the duration of the processing.

$$\mathbf{process} \wedge \mathbf{emit} \Rightarrow (t_{\text{phase}} \geq \alpha - \epsilon) \wedge (t_{\text{phase}} \leq \alpha + \epsilon) \quad (3)$$

The formal model includes a number of parameters, such as α introduced above, capturing the features of the topology, which can be configured at design time. In addition to α , other parameters are, for bolts, a coefficient σ expressing the kind of operation performed by the bolt in terms of quantity of output tuples emitted given an input tuple, and also the minimum and maximum time to failure. Spouts, instead, are characterized by the average number of tuples emitted per time unit. Both spouts and bolts are also characterized by their level of parallelism, corresponding to the number of executors for the component.

3. ANALYSIS OF STORM TOPOLOGIES

D-VerT allows designers to validate temporal aspects of DIAs by means of a logic-based formal verification approach outlined in Section 2.2. The implementation of D-VerT currently supports the analysis of Storm topologies, but it can be easily extended to deal with other big-data technologies if and when their computational model is formalized through CLTL_{oc} formulae.

D-VerT is part of a more complex design process which conforms to the principles of model-driven software engineering




Figure 3: The main steps of D-VerT workflow.

pursued by the DICE methodology. As illustrated in Fig. 2, the designer defines the application by means of domain-specific models with an iterative approach consisting of three steps: (i) application design, (ii) design evaluation and (iii) monitoring of a running deployed application. D-VerT is situated at the second level of the design workflow and enables the refinement of the architectural design of the application depending on the outcome of the formal analysis. The input of D-VerT is an annotated UML (class) diagram which specifies the architecture, i.e., the topology, of the application. In case of an incorrect timing design, the outcome of D-VerT consists of a possible execution of the topology causing an undesired accumulation of tuples in some bolts. In such a case, the designer can refine the application architecture by changing the values of some parameters of the UML diagram and then redo the evaluation until (possibly) all flaws affecting the model are removed. A different scenario, which also entails a design refinement, might occur when some parameter values that are measured on a running application differ from the values used for verification at design time. In such a situation, monitored data obtained from the deployed application can be exploited to update the model, which can then be newly verified.

Relying on UML profiles for modeling is a common practice in model-driven development as it allows for the creation of domain-specific languages by extending or restricting UML to fit a specific domain. A UML Profile is made of a set of stereotypes, tags, constraints and meta-models that allow the designer to represent artifacts in a specific domain. A stereotype is a meta-concept which is used to categorize an element in a model (for example, *container*) with specific notions belonging to a domain.

As shown in Fig. 3, at the starting point of the workflow the user creates an annotated UML model describing the relevant technological aspects of the architecture of a DIA. The UML model includes suitable design abstractions, capturing specific data-intensive technologies—Storm in our case—that are adopted for implementing the architecture of an application. The diagram, called DICE Technology Specific Model (DTSM), is at the PSM level (see Sect. 1) in the model-driven approach pursued by DICE. Specifically, in a DTSM diagram, a stereotype classifies an element of an application with aspects related to a specific technology. A DTSM diagram includes generic concepts that fit many data-intensive frameworks, such as `ComputationNode`, `StorageNode` or `SourceNode`, and specific ones, depending on the selected technology. In the case of Storm, the relevant features and aspects defining a Storm topology constitute the meta-model for designing Storm applications. Some of them, that are used in Sect. 4, are `Topology`, `Spout`, `Bolt` and

`TopologyConfigurations`. For a comprehensive description of the concepts available in DTSM diagrams see [8].

DTSM diagrams for Storm include all the values of the parameters that are useful to carry out the analysis of a topology. As depicted in Fig. 3, verification of DTSM models is done through their automatic translation into a set of CLTL_{oc} formulae, which are then analyzed by the Zot bounded satisfiability checker [3] using the technique presented in [11]. More precisely, Zot is fed the CLTL_{oc} formulae capturing the application under design and the property to be checked concerning the unbounded growth of the queues of interest. The tool produces one of two responses: (i) a trace of the modeled Storm topology—a *counterexample*—corresponding to an execution of the application in which one of the queues grows in an unbounded manner—in this case, the set of formulae is *satisfiable* (SAT); or (ii) the notification that the set of formulae is *unsatisfiable* (UNSAT). In the latter case, since the language used to formalize Storm topologies is in general undecidable, we cannot conclude definitively that there is no execution of the application such that the queues grow indefinitely, but only that, within the bounds chosen for the search of counterexamples, none was found. Still, an UNSAT result increases our confidence that no major design flaws are present in the architecture of the Storm topology for what concerns its ability to process data in a timely manner.

4. TOOL DESCRIPTION

This section outlines the architecture of the D-VerT tool, the transformation enabling the verification process and the kind of analysis currently supported by the tool.

4.1 Tool Architecture

As shown in Fig. 4, D-VerT is structured as a client-server application. The client component is an Eclipse plug-in, and is part of the DICE IDE. It allows users to define the design of the DIA under development, then, after providing some additional configuration information, to launch verification tasks and to retrieve their outcomes. The server component consists in a RESTful web application written in Python. The D-VerT server exposes APIs to launch verification tasks and to obtain information about their status. To simplify the setup and deployment steps, the D-VerT server is available as a Docker¹ container. The client-server architecture decouples the application design phase from the rather computationally-intensive verification phase. Based on the needs of the user, the D-VerT server can be instantiated either on the local machine or on a remote server.

4.2 Topology creation

The design of a DIA is specified through the DICE IDE, which is based on the Papyrus tool. As mentioned above, Storm topologies are described as DICE-profiled UML Class diagrams. Each computational node of the application is defined by introducing a class tagged with a stereotype specifying whether the node is a spout or a bolt. Depending on the stereotype applied, the designer defines the values for all the necessary parameters described in Sect. 2. Subscriptions

¹ Repository omitted for double-blind reviewing.

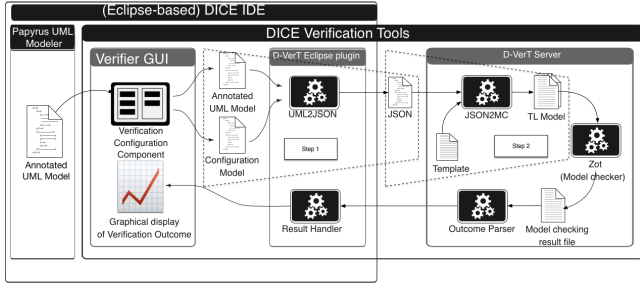


Figure 4: D-VerT workflow mapped onto the client-server architecture of the tool.

(i. e., connections) of bolts to other components are modeled as associations between the corresponding classes.

4.3 Transformations

The verification process is made possible by a two-step transformation applied on the DICE-profiled UML models to obtain the corresponding set of CLTLoc formulae.

The first step of the transformation is performed in the D-VerT client by the UML2JSON module, which extracts from the DICE-profiled UML model all parameters that are relevant for the verification. These parameters are then serialized into a JSON object, which is used to invoke the server component. The extraction of the relevant features is done by suitably navigating the UML file. DIA components and their parameters are detected thanks to their being annotated with proper stereotypes from the DICE profile.

The second step takes place in the D-VerT server, which receives the request from the client, produces the corresponding formal model and feeds it to the underlying Zot [3] tool. More precisely, the JSON2MC module, based on the contents of the JSON object included in the request, generates the temporal logic model using a templating mechanism.

4.4 Analysis

In its current stage of development, D-VerT provides support for the analysis of the boundedness of bolts’ queues. Through the run configuration dialog box of the tool (see Fig. 5) the designer can specify the bolts of interest, the depth of the search over which the verification has to be performed (the “time bound”, which corresponds to the maximum length of the trace produced) and the Zot plug-in to be used. The analysis allows for the detection of possible runs of the system leading to an unbounded growth in the size of at least one of the aforementioned bolts’ queues. This corresponds to the presence in the topology of at least one bolt that is not able to manage the incoming flow of messages in a timely manner. In this case the tool provides as output to the user the trace witnessing the existence of such a run of the system—i. e., the counterexample violating the boundedness property. The trace is returned to the user both in a textual format (i. e., the bare output of Zot) and in a graphical format, in order to provide a more user-friendly visual hint about the system execution. Figure 9 shows an example of such output trace, returned by the tool for the use case of Sect. 5. It represents the evolution of the num-

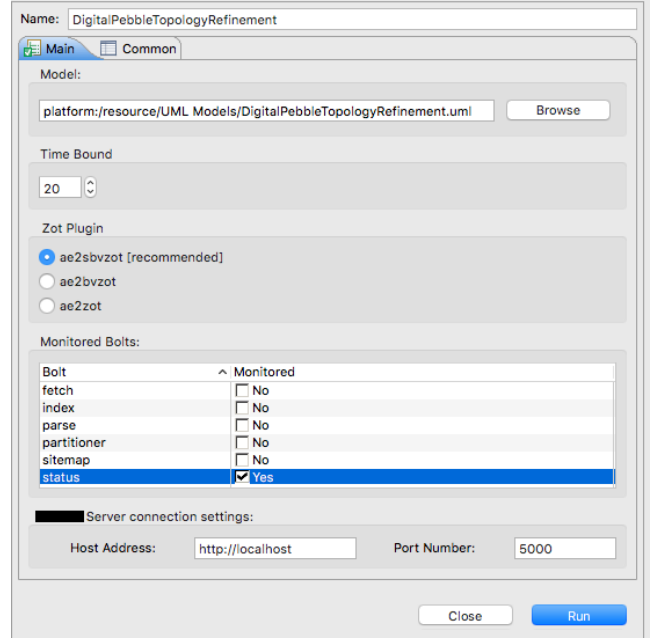


Figure 5: Run configuration view.




Figure 6: PIM of the web crawler application.

ber of tuples in the queue over time. The trace is composed by a prefix and a suffix: the latter, highlighted by the gray background, captures the growing trend of the queue size, as it corresponds to a series of operations in the system that can be repeated infinitely many times. When no trace is detected, the result is UNSAT.

5. D-VERT WORKFLOW IN ACTION

In this section we illustrate the usage flow of D-VerT for the iterative refinement of a Storm topology. The use case is taken from the open source project StormCrawler². Suppose we want to create a web crawler application to efficiently fetch, parse and index web resources of our interest. Given the dynamic nature of the web, this kind of task can be formulated as a streaming problem, where the input consists in a continuous stream of URLs that need to be processed by the streaming application with low latency, and the output is represented by the resulting indexes.

We start by modeling the application at the PIM level. In this case, the model simply includes a *source node*, a *computation node* and a *storage node*, as depicted in Fig. 6. We decide to use a Kafka queue as source node, a Storm topology as computation node and Elasticsearch as storage node.

Since we are interested in analyzing the Storm topology,

²<https://github.com/DigitalPebble/storm-crawler>




Figure 7: Initial PSM of the web crawler topology.

we focus on the computation node and consider the source node and the target storage node as “black boxes”. At the PSM level we insert more technology-related aspects, such as, in the case of Storm, the topology structure and a series of non-functional characteristics. Figure 7 shows the PSM (DICE-profiled UML diagram) of the initial design of the topology. The configuration parameters are represented as UML comments for illustrative purposes. Notice that associations between components have the opposite direction with respect to the data flow among them, since they indicate the subscription of bolts to the associated components’ streams. The diagram includes one spout in charge of fetching the input flow of URLs from Kafka and three bolts performing various steps of the web crawling process. The *partitioner* bolt partitions the incoming URLs, while the *crawler* bolt performs many operations such as resource fetching, metadata parsing and content indexing. The *status* bolt at the end of the chain indexes the URL, metadata and its status to a specific “status” Elasticsearch index. Each of these topology components can be executed by an arbitrary number of parallel threads, and is characterized by the (average) execution time (time needed to perform its task) and by the (average) number of tuples emitted with respect to the number of tuples received as input. These aspects are specified as parameters in the UML class diagram. The formal analysis on the initial topology design helped us to detect an unbounded increase in the queue of the *crawler* bolt. This outcome from the tool led us to review the topology structure, and to decide for the decomposition of the *crawler* bolt in a series of bolts, each of them performing a subtask of the original bolt (*fetch*, *sitemap*, *parse* and *index*). The refined version of the topology, shown in Fig. 8, aims at lightening the load on the core crawling phase by pipelining the main operations and by directly updating the

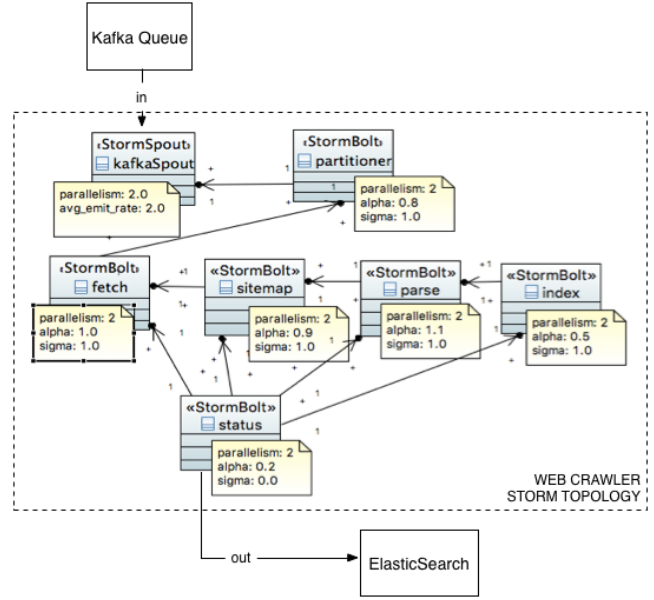


Figure 8: Refinement of the web crawler Storm topology PSM.

status bolt with partial results computed by the new bolts.

After the refactoring the tool revealed another unwanted run of the system, this time showing a growing trend in the queue of the *status* bolt (Fig. 9). This bolt, subscribing to the streams of the four newly-created bolts, needs a further refinement to avoid excessive loads in its input buffer. Increasing the parallelism level of the *status* bolt to 4 helped improving the design so that no counterexample was found by D-VerT within the selected search bounds (UNSAT result). Execution times for the verification vary significantly depending on the topology configuration, ranging from the 98 seconds of the second analysis (Fig. 8) to the 2130 seconds of the third analysis (UNSAT result)³.

6. RELATED WORKS

Many research works in recent years investigated the usage of MDE to support the design and the formal verification of software and embedded systems. [9] presents a systematic literature review on the formal verification of static software models. Most of the works make use of UML models, often enriched with OCL constraints, and only a part of them is fully supported by a tool implementing the model transformations and the verification process. A number of other works have used a model-driven approach for the formal verification of behavioral models (see, e.g., [7, 10]), without addressing the specificities of DIAs. To the best of our knowledge, few works try to address the problem of the verification of DIAs, none of them adopting the MDE approach. They mainly focus on the verification of properties that depend exclusively on the framework by building ad-hoc models; for

³Experimental analysis carried out on commodity hardware (MacBook Air running MacOSX 10.11.4. with Intel i7 1.7 GHz, 8 GB 1600 MHz DDR3 RAM; SMT solver used by Zot was z3 v.4.4.1).

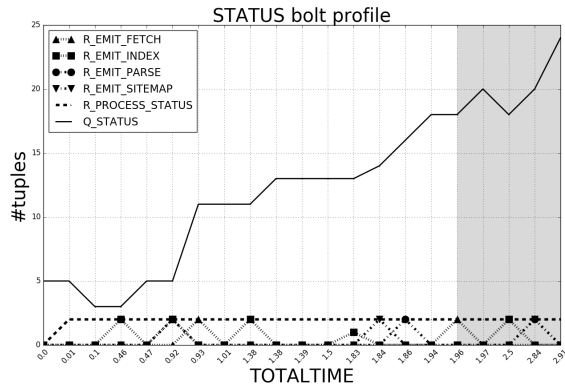


Figure 9: Graphical output trace of the *status* bolt returned by D-VerT. The black solid line represents the number of elements in the input buffer.

example, e.g., [12] verifies data locality, deadlock freedom and non-termination properties for the Hadoop parallel architecture, while [13] verifies the validity of communication data flows of Hadoop MapReduce. Our work, on the other hand, aims at allowing for the verification of properties that depend on the application design.

7. CONCLUSION

In this paper we presented the model-driven approach to the formal verification of Storm topologies supported by the D-VerT tool. It allows designers to formally check whether, given the features of the components of the topology, it is possible for the queues of some bolts to grow indefinitely, which entails that incoming tuples will not be processed in a timely manner.

Future works will focus on extending the range of big-data technologies covered by the tool (e.g., Apache Spark), on enlarging the set of properties that can be analyzed, and on improving the efficiency of the verification technique.

Acknowledgments

Work supported by Horizon 2020 project no. 644869 (DICE). We are thankful to our colleague Michele Guerriero for his precious advice and expertise in model-driven software engineering.

8. REFERENCES

- [1] Apache Storm. <http://storm.apache.org/>.
- [2] D-VerT. <http://bit.ly/2do92ao>.
- [3] Zot. <https://github.com/fm-polimi/zot>.
- [4] M. Bersani, M. Erascu, F. Marconi, and M. Rossi. DICE verification tool - initial version. Technical report, DICE Consortium, 2016. www.dice-h2020.eu.
- [5] M. M. Bersani, M. Rossi, and P. San Pietro. A tool for deciding the satisfiability of continuous-time metric temporal logic. *Acta Informatica*, 53(2):171–206, 2016.
- [6] G. Casale, D. Ardagna, M. Artac, F. Barbier, E. D. Nitto, A. Henry, G. Iuhasz, C. Joubert, J. Merseguer, V. I. Munteanu, J. Perez, D. Petcu, M. Rossi, C. Sheridan, I. Spais, and D. Vladušić. DICE: Quality-driven development of data-intensive cloud applications. In *Proc. of MiSE*, pages 78–83, 2015. www.diceh2020.eu.
- [7] Z. Daw and R. Cleaveland. Comparing model checkers for timed UML activity diagrams. *Science of Computer Programming*, 111, Part 2:277–299, 2015.
- [8] A. Gómez, M. Guerriero, J. Merseguer, E. di Nitto, and D. A. Tamburri. Design and quality abstractions - initial version. Technical report, DICE Consortium, 2016. www.dice-h2020.eu.
- [9] C. A. González and J. Cabot. Formal verification of static software models in MDE: A systematic review. *Information and Software Technology*, 56(8):821 – 838, 2014.
- [10] H. Hansen, J. Ketema, B. Luttik, M. Mousavi, and J. van de Pol. Towards model checking executable UML specifications in mCRL2. *Innovations in Systems and Software Engineering*, 6(1):83–90, 2010.
- [11] F. Marconi, M. M. Bersani, M. Erascu, and M. Rossi. Towards the formal verification of data-intensive applications through metric temporal logic. In *Formal Methods and Software Engineering - 18th International Conference on Formal Engineering Methods, ICFEM 2016*, pages 193–209, 2016.
- [12] G. S. Reddy, Y. Feng, Y. Liu, J. S. Dong, S. Jun, and R. Kanagasabai. Towards formal modeling and verification of cloud architectures: A case study on hadoop. In *2013 IEEE Ninth World Congress on Services*, pages 306–311. IEEE, 2013.
- [13] W. Su, F. Yang, H. Zhu, and Q. Li. Modeling mapreduce with CSP. In *2009 Third IEEE International Symposium on Theoretical Aspects of Software Engineering*, 2009.