

Adding Semantics to Measurements: Ontology-Guided, Systematic Performance Analysis

Attila Klenik^{ab} and András Pataricza^{ac}

Abstract

The design and operation of modern software systems exhibit a shift towards virtualization, containerization and service-based orchestration. Performance capacity engineering and resource utilization tuning become priority requirements in such environments.

Measurement-based performance evaluation is the cornerstone of capacity engineering and designing for performance. Moreover, the increasing complexity of systems necessitates rigorous performance analysis approaches. However, empirical performance analysis lacks sophisticated model-based support similar to the functional design of the system.

The paper proposes an ontology-based approach for facilitating and guiding the empirical evaluation throughout its various steps. Hyperledger Fabric (HLF), an open-source blockchain platform by the Linux Foundation, is modelled and evaluated as a pilot example of the approach, using the standard TPC-C performance benchmark workload.

Keywords: performance, measurement, bottleneck identification, EDA, ontology, blockchain, Hyperledger Fabric, TPC-C

1 Introduction

The rapidly increasing number of IT service customers made the performance of such systems a high priority. Performant systems are not just a question of powerful hardware anymore, they also require the system-wide careful design of the software stack. The systematic detection and diagnosis of performance bottlenecks by analysing multi-dimensional measurement data becomes an integrated part of both the development and operational (DevOps) parts of the system life-cycle.

The industrialization of general-purpose data analysis resulted in typical standard workflows, like CRISP-DM [59], or ASUM-DM [11]. Such workflows are typically centered around the following high-level, domain-agnostic steps [6, 20, 52]:

^aDepartment of Measurement and Information Systems, Budapest University of Technology and Economics, Budapest, Hungary

^bE-mail: attila.klenik@vik.bme.hu, ORCID: 0000-0003-2051-2823

^cE-mail: pataricza.andras@vik.bme.hu, ORCID: 0000-0002-6516-129X

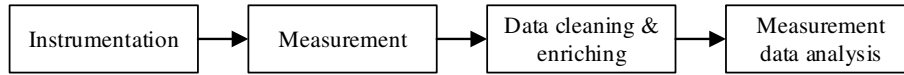


Figure 1: Typical performance evaluation steps

data acquisition; representation; analysis; visualization and reporting – with a proper model-driven engineering (MDE) support. Different performance analysis tasks – such as bottleneck identification and latency anomaly root cause analysis – can be considered domain-specific refinements [28] of the analysis step, defining further, embedded sub-workflows.

However, the *technical metrology* of performance evaluation poses specific challenges. The technical systems under test (SUT) are usually very complex, both in the terms of their architecture and potential state space. Still, performance engineering became increasingly important, as many systems have to fulfill soft real-time requirements. Moreover, poor performance dimensioning (stemming from architectural design or misconfiguration) can lead to service-level violations, or the malfunctioning of the system, even in the case of short overloads.

The paper proposes an *activity and observability-focused ontological approach* for the model-based guidance of SUT- and measurement-related, technical performance analysis tasks (Fig. 1): instrumentation; measurement; data cleaning and enriching; and measurement data analysis.

Instrumentation, the insertion of sensors into the system, plays an important role in system *observability*, i.e., the degree to which an observer can reconstruct the internal state of a system based on its outputs. However, sensor placement must balance multiple requirements: non-intrusivity whenever possible; development time/cost; and sufficient amount of resulting measurement data to work with.

On one hand, increasing the number of sensors might provide a deeper insight into the system, but application-specific sensors require a careful development to assure the integrity of the measurement results without distorting temporal metrics. On the other hand, under-instrumentation confines the granularity of root cause analysis and consequently the indication for mitigating bottlenecks. Moreover, it can leave faulty behavior undetected.

Correspondingly, instrumentation needs a careful trade-off between the relevance and redundancy of the measurements. The proposed approach aids the designer or analyst in formally arguing about the observability of the system, or in selecting a sufficient sensor placement.

During the *measurement*, data acquisition has to cope with the heterogeneity of data sources generating observation logs in very different formats. Data source models [27, 52] support the semantic fusion and representational homogenization of the sources and the following ETL (extract, transformation, load) steps.

The proposed ontology guides the ETL process towards a *relation-oriented and activity-focused representation* of measurement data, building on widely used concepts. The common format may serve as a gateway toward other temporal modeling

frameworks (e.g. the OWL Time ontology¹ from the World Wide Web Consortium), metrology-related technologies (e.g., the OpenTelemetry² project from the Cloud Native Computing Foundation), or other analysis techniques (such as process mining [56]).

Clean and detailed data is a prerequisite for many performance analysis tasks, such as bottleneck identification. Large-scale system observations constitute as *big data*, but more importantly, as *multi- or many-dimensional data*. Data is harvested from multiple layers of numerous system components, ranging from boundary-level response times to infrastructure-level resource utilization metrics. Bottleneck identification in such a context is a complex diagnostic process. It is a priori unknown how deep the analysis must drill down to uncover root causes of performance anomalies.

The proposed approach makes data validation a *systematic process* by sharding and inspecting the data set based on the modeled activities and corresponding services. Thus data omission errors, for example, can be easily pinpointed even in larger data sets. Moreover, the activity and observability models coupled with various temporal rules provide a framework for automatically *deriving further temporal information*, even if not explicitly observed.

The *analysis* of the gathered multi-dimensional data necessitates *exploratory data analysis* (EDA). EDA is, by its nature, a highly adaptive and iterative process for identifying a model of the system behavior. Usually, a domain expert is in charge of guiding the *drill-down process* if something peculiar is detected from the point of view of the application. The exhaustiveness and quality of this exploration process heavily depend on the domain knowledge of the expert, the automation of the elementary steps, and a proper navigation along the process and the data [52, 14, 45, 63].

The *hierarchical nature* of the proposed activity ontology and the corresponding service/deployment information make the drill-down process intuitive and systematic. The domain expert can dissect higher-level activities as needed, until a possible cause is found for a peculiar behavior. Then they can correlate the time range of the behaviour with various workload metadata and/or resource utilization metrics to find its root cause (may it be a bottleneck of the system, or a resource saturation issue). Furthermore, the drill-down steps are guided by concepts independent of the actual SUT, making the process reusable for different systems, a viable candidate for automation, or to be performed by a less experienced domain expert.

A complex case study demonstrates the benefits of the proposed approach:

- The *HLF blockchain* platform's consensus *activities* and their *observability* are modelled in a reusable and modular way.
- The activities of the standard TPC-C performance benchmark³ are modeled and combined with the HLF model.

¹<https://www.w3.org/TR/owl-time/>

²<https://opentelemetry.io/>

³<http://tpc.org/tpcc/default5.asp>

- The formal measurement inference capabilities of the ontology are demonstrated, coupled with a systematic data validation process.
- An ontology-guided, EDA-based, hierarchical bottleneck identification process is demonstrated on measurement data observed while executing the TPC-C workload (generated by the Hyperledger Caliper⁴ benchmarking tool) on HLF.

The paper is structured as follows. Sec. 2 introduces the proposed general approach for the performance analysis of complex systems. Sec. 3 presents related MDE approaches for activity modeling and surveys the state of the art HLF performance researches. Sec. 4 introduces the elements of the proposed ODK used for complex activity modeling and automatic observability inference, along with formal semantics. Sec. 5 presents the case study of compositional modeling of the HLF consensus process and the TPC-C benchmark execution using the ODK. Sec. 6 demonstrates the various applicability of the resulting system models in aiding complex measurement data validation and analysis tasks. Sec. 7 concludes the paper.

2 The proposed model-guided analysis approach

The cornerstone of a performance analysis process is the observability of the activities performed by a system. On a high level, the beginning, duration, and end of system tasks are the basis of common metrics, like incoming task rate, throughput, and latency. Bottleneck identification, however, requires more interconnected data to work with, including the well-defined composition semantic of complex activities. Moreover, such observations are also crucial for building precise, well-parametrized models for efficient performance prediction [8, 15].

Our contribution is an ontology-guided workflow for the *systematic, drill-down performance analysis* of multi-dimensional measurement data. Moreover, the supporting *ontology development kit* (ODK) is provided for ensuring the quality and sufficiency of measurement data, enhanced with composition semantics for facilitating bottleneck identification processes. The ODK supports the modeling of activities, their relations, and whether their execution is observable outside of the system. Furthermore, it provides a formal foundation for rigorous measurement data analysis task, e.g., bottleneck identification.

The proposed approach is outlined in Fig. 2 and detailed in subsequent sections:

1. Model the important activities of the system components, focusing on their relations and hierarchical composition.
2. Model the explicit observability of activities to assist observability inference.
3. Extend the model with additional elements (by bridging to other ontologies, for example) to support further design, DevOps, or analysis tasks, as needed.

⁴<https://www.hyperledger.org/use/caliper>

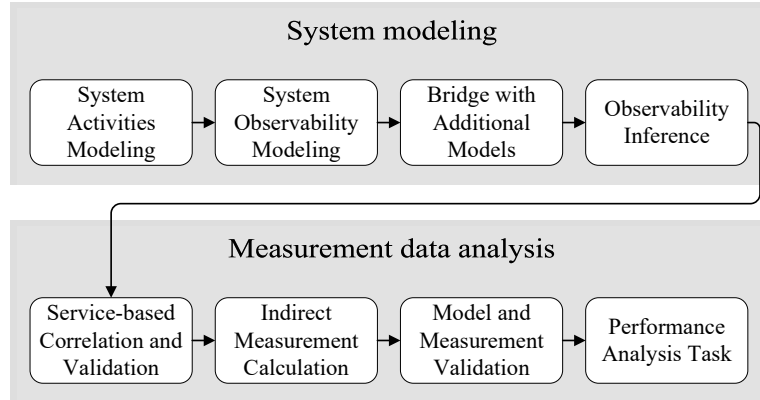


Figure 2: The proposed workflow for model-guided performance evaluation

4. Automatically enrich the current model with additional observability information by using an OWL reasoner (detailed in Sec. 4.5).
5. Correlate and validate distributed request traces from the SUT to uncover data omission or similar errors.
6. Calculate additional, indirect measurement using the ontology (or other derived) model as guide.
7. Validate the conformance of the measurement data and the model to ensure the correctness of further analysis tasks.
8. Perform the desired analysis task based on rigorously cleaned and validated data, and using the ontology model as guide.

Following the outlined steps allows for a rigorous performance analysis of the SUT. Note that the model construction steps only need to be performed once (if done properly), then the component models can be reused and recombined to fit further performance analysis scenarios for different SUT setups. Moreover, the modeling and analysis parts of the workflow can be performed by different domain experts, lowering the entry barrier for the overall performance analysis of a given SUT.

3 Related work

The section presents the related work on creating activity execution models and surveys the state of the art regarding HLF performance evaluations. The limitations of the presented literature motivated our contribution to bring MDE approaches closer to the domain of performance evaluation of complex systems.

3.1 Activity modelling

System activities are usually observed through individual events (e.g., logs, notifications) or sensors. An important requirement of activity modeling – relating to performance analysis – is to allow the systematic reconstruction of detailed timelines from the available partial observations, facilitating data analysis. Furthermore, having well-defined modeling semantics and building blocks allow the assessment of a wide range of systems.

Our experience with EDA and bottleneck identification outlined the following requirements for model-based support:

- *formal modeling* of complex *activity hierarchies* and relations;
- explicit modeling of system *observability* (i.e., sensor placement);
- *systematic derivation* of additional temporal knowledge;
- *extensibility* for incorporating further service/infrastructure models;
- *composability* and *reusability* of different activity models.

Similar approaches exist in the domain of business process analysis using on-line analytical processing (OLAP) [1, 39, 42, 48]. However, our approach has to comply with the additional requirements of technical metrology, like allowing the performance evaluation of general system activities despite limited observability of tasks, and facilitation of the adaptation of metrology principles.

Modeling the execution of activities also has a long tradition in software development, both as design phase artifacts for validation, and as inputs to automatic task orchestration systems. Business process models (building on the BPMN⁵ standard) or activity diagrams in UML⁶ or SysML⁷ are prime examples of high-level activity modeling languages.

Such visual languages facilitate the modeling of activity control flows, imposing certain temporal constraints (e.g., activity *A* must be executed *before* activity *B*). However, the enforcement of such constraints must be validated during analysis time or runtime. Such validation necessitates the detailed observation of activities to allow rigorous temporal constraint checks. Moreover, the available high-level languages lack an intuitive support of modeling observability.

Additional formal approaches can also aid the *design* and *verification* of complex systems or processes. Temporal logics – such as Temporal Logic of Actions (TLA, TLA+ [30, 31]), Propositional Temporal Logic (PTL [46]), Interval Temporal Logic (ITL [3]) – enable the specification and verification of time-dependent system behavior [9]. Furthermore, different probabilistic or stochastic process algebra approaches [22] can aid the design and verification of concurrent, distributed systems, including the communication and synchronization of their independent

⁵<https://www.omg.org/spec/BPMN/2.0/>

⁶<https://www.omg.org/spec/UML/>

⁷<https://www.omg.org/spec/SysML/>

components. Such process algebras include the Performance Evaluation Process Algebra (PEPA [16]), or extended versions of the Communicating Sequential Processes (CSP [32]) language.

The above formal approaches provide an expressive power sufficient for the detailed design, verification and subsequent implementation of concurrent systems. However, their application requires knowledgeable experts in the formal verification domain. Moreover, the intended purpose of the presented approaches is to aid the correct design and implementation process of complex systems, thus providing more facilities than needed for validating and analyzing performance measurements of the already implemented systems.

Accordingly, the goal of the current work is to provide a kernel modelling framework (with simple vocabulary and relations pertaining to system activities) that is easier to use, given a standard information technology background. The semantic mapping of the referenced formal approaches (and the corresponding, readily available formal design models during system development) to the presented modelling framework is left as future work.

Ontology-like formal approaches also gain ground in general system modeling tasks (e.g., the upcoming OMG SysML v2 Kernel Modeling Language⁸), thus our contribution relies on ontologies, preparing for future interoperability. Knowledge representation-based approaches can also aid the visual analysis of network traffic [61] or the semantic fusion of data originating from different sources [60]. Moreover, ontology-based approaches can reason about the occurrence of composite activities [43, 12, 21, 33, 49].

The referenced activity modeling works have several elements in common. They utilize Allen's interval algebra [2] for describing temporal relations, allowing bridging to other similar solutions. However, they reverse-engineer/infer the activity model based on the observation of performed activities, similarly to process mining [56]. Model mining is unavoidable in contexts where the "schedule" of executed activities is non-deterministic, such as in smart homes or in smart warehouses [12, 49].

However, when the execution of activities must conform to a predefined specification, model mining becomes unnecessary. The paper proposes a *model-first* approach to construct an ontology-based composite activity model, which will later provide a strong foundation for the systematic performance evaluation and bottleneck analysis of the target system. Accordingly, the model becomes an input to the analysis tasks, and not an output.

3.2 Hyperledger Fabric performance analysis

The complex consensus process of HLF [5] (detailed and modeled in Sec. 5) made its performance evaluation a hot research topic. Related works can be divided mainly into the following categories based on their goals:

⁸<https://github.com/Systems-Modeling/SysML-v2-Release>

1. Performance *evaluation and characterization*: [47, 7, 55, 38, 18, 34, 51, 19, 58, 29, 40, 24, 41, 4, 13, 57, 50, 10]
2. Performance *optimization*: [55, 17, 25, 37]
3. Formal consensus *modelling*: [53, 54, 26, 64, 62]

Category 1 receives most of the attention, which is identifying the performance characteristics of HLF. The evaluations employ empirical sensitivity analyses to measure the change in key performance indicators (such as throughput and end-to-end latency) when applying different network scales, configurations, and workloads.

The concern of Category 2 is the performance enhancement of HLF. Researches either transparently optimize certain consensus steps or propose changes to the architecture (and correspondingly the consensus process) itself. The researches of Category 2 also rely on empirical performance analysis to confirm bottlenecks and evaluate the effectiveness of optimizations.

Works in Category 3 build formal behavior models of the consensus process. Model parameter identifications also rely on empirical performance evaluations. Finally, the parameterized model allows for cost-efficient sensitivity analyses capable of covering a large configuration and parameter space, without actual further empirical analyses.

Table 1 further refines Categories 1–3 based on the following aspects:

- *Network scaling* (NS): the research performed multiple evaluations while varying the scale of the Fabric network, e.g., the number of orderer and/or peer nodes (horizontal scaling), or their allocated resources (vertical scaling).
- *Configuration sweep* (CS): the Fabric network was evaluated in multiple operational modes, e.g., using different transaction ordering protocols (Kafka- or Raft-based implementations), varying the configuration of a specific ordering protocol (target block size or block time), or changing the implementation of other elements (the choice for state database, or the used chaincode language).
- *Workload scaling* (WS): certain attributes of the workload were changed among multiple evaluations, e.g., the incoming rate of requests, or the ratio of read and write intensive requests.
- *Consensus step optimization* (CSO): the research proposed improvements to certain steps of the consensus process that are transparent (i.e., non-breaking, backward compatible changes) to other network components or users, e.g., the parallelization of inner component tasks (like transaction validation).
- *Consensus process optimization* (CPO): the proposed improvements significantly change the overall consensus process, i.e., the changes require the adaptation of certain APIs, breaking the existing solutions.
- *Consensus modelling* (CM): the research formally modelled and evaluated certain aspects of the Fabric consensus, e.g., calculated expected transaction latencies based on the approximated processing times of subtasks.

Table 1: Categorization of Fabric performance evaluation-related works (**NS**: Network Scaling, **CS**: Configuration Sweep, **WS**: Workload Scaling, **CSO**: Consensus Step Optimization, **CPO**: Consensus Process Optimization, **CM**: Consensus Modelling)

| Related Works | NS | CS | WS | CSO | CPO | CM |
|----------------------|----|----|----|-----|-----|----|
| [55] | ✓ | ✓ | ✓ | ✓ | | |
| [7, 57, 50] | ✓ | ✓ | ✓ | | | |
| [13] | ✓ | ✓ | | | | |
| [38, 58, 40, 10] | ✓ | | ✓ | | | |
| [34] | ✓ | | | | ✓ | |
| [51] | | ✓ | ✓ | ✓ | | |
| [17] | | ✓ | ✓ | | ✓ | |
| [19, 37] | | ✓ | ✓ | | | |
| [25] | | ✓ | | ✓ | | |
| [4] | | | ✓ | | ✓ | |
| [47, 18, 29, 24, 41] | | | ✓ | | | |
| [53, 54, 26, 64, 62] | | | | | | ✓ |

A common requirement for all categories and aspects is the rigorous empirical performance evaluation of HLF based on the analysis of measurement data. Superficial analyses may lead to incorrect hypotheses or misidentified model parameters, invalidating the results of the evaluation.

Accordingly, a *systematic, rigorous, and easy to follow* analysis process (even for complex systems) is needed to achieve relevant results. Moreover, the *correctness* and *richness* of measurement data can further increase the quality of gained insights.

4 Activity and observability modeling framework

The section introduces the formal foundations and building blocks of the proposed ODK for constructing complex activity models. Moreover, it details the observability modeling and automatic observability inference mechanisms that are the cornerstones of a rigorous performance data analysis.

4.1 Formal foundations

The ODK is constructed using the Web Ontology Language⁹ (OWL2), adhering to some constraints (resulting in an OWL-DL ontology) that make the OWL direct semantics compatible with the model-theoretic semantics of the *SRCIQ* description logic [23]. This restriction provides useful computational properties for the

⁹<https://www.w3.org/TR/owl2-syntax/>

language, backed by extensive literature and tooling support, such as OWL-DL reasoners [44].

OWL2 provides facilities such as object and data properties, literals, individuals, and classes to model relations among different concepts and resources. Classes can have associated relationship constraints that must hold for every individual belonging to the class. It is important to note that OWL employs the open-world assumption, meaning that if something is not asserted as knowledge, it is taken as unknown, rather than as untrue. The OWL2 structural specification¹⁰ further details the available language constructs and their meaning.

The paper utilizes the OWL-DL notations of Table 2 to describe the elements of the ODK and their semantics.

Table 2: OWL-DL notations

| OWL construct | Notation | OWL construct | Notation |
|-----------------------|-----------------------------|------------------------|-----------------------------|
| <i>Class</i> | $C1, C2$ | <i>SubClassOf</i> | $C1 \subseteq C2$ |
| <i>IntersectionOf</i> | $C1 \cap C2$ | <i>UnionOf</i> | $C1 \cup C2$ |
| <i>Thing</i> | T | <i>Property</i> | P |
| <i>PropertyRange</i> | $T \subseteq \forall P.C1$ | <i>EquivalentClass</i> | $C1 \equiv C2$ |
| <i>AllValuesFrom</i> | $C1 \subseteq \forall P.C2$ | <i>SomeValuesFrom</i> | $C1 \subseteq \exists P.C2$ |

The temporal constructs of the ODK build on Allen’s interval algebra. Let us consider activity instances $a = (a^b, a^e, a^d) \in A$ of an activity class with a beginning time instant $a^b \in \mathbb{N}^+$, an ending time instant $a^e \in \mathbb{N}^+$, and a non-zero duration $a^d \in \mathbb{N}^+$, where $a^b < a^e$, and $a^b + a^d = a^e$ for every $a \in A$, measured on a logical clock for the simplicity of the notation.

If a property (i.e., directed relation) P holds between activity instances a and b , we denote it by $a \in P.b$, where $a \in A$, $b \in B$ activity classes. The shorthand notation $A \subseteq P.B$ specifies the relation P as constraint between activity classes A and B , implying $\forall a \in A : \exists b \in B, a \in P.b$.

The ODK defines the following Allen interval relations as OWL properties: *after*, *before*, *meets*, *metBy*, *starts*, *startedBy*, *finishes*, and *finishedBy*. Accordingly, if an activity type A is always followed by an activity type B , the axiom $A \subseteq \textit{meets}.B$ will be part of the ontology. Note, that the ODK contains only the Allen relations that provide precise or useful activity composition semantics. Accordingly, the *during*, *overlaps*, and *equal* relations (and their inverses) are not utilized directly, but can be derived from the modeled relationships in a straightforward way.

¹⁰<https://www.w3.org/TR/owl2-syntax/>

4.2 Component overview

The ODK contains a hierarchy of smaller ontologies – each with well-defined responsibilities – to promote composability (Fig. 3).

The *Activity* ontology (Sec. 4.3) allows the modeling of system activity relations. For example, the following set of assertions partially describe an activity decomposition (Fig. 4): $Processing \subseteq SequentialActivity$, $Substep_i \subseteq AtomicActivity$, and $Processing \subseteq hasSubactivity.Substep_i$.

The *Observability* ontology (Sec. 4.4) provides classes to "annotate" the activities with further information regarding their degree of observability. For example, $Processing \subseteq EndMeasured$ denotes that the end of *Processing* activities are explicitly observed/measured through logs, or system events.

The *Structural constraints* ontology provides well-formedness axioms for activity composition. The open-world semantic of OWL2 makes it cumbersome to convey traditional (closed-world) modelling intentions to a set of ontology axioms. For example, it is not enough to just state that a subactivity is the starting activity of its parent. Correct modeling also requires the statement that the starting activity is not preceded by any other activity (otherwise it could not be the first subactivity of its parent).

The structural constraints ontology provides several axioms that can automatically detect (using an ontology reasoner) such inconsistencies or potentially missing axioms. The description of constraints, however, is outside the scope of this paper.

The *inference* ontologies (Sec. 4.5) extend the observability ontology with equivalence axioms that can automatically flag (during reasoning) activity classes

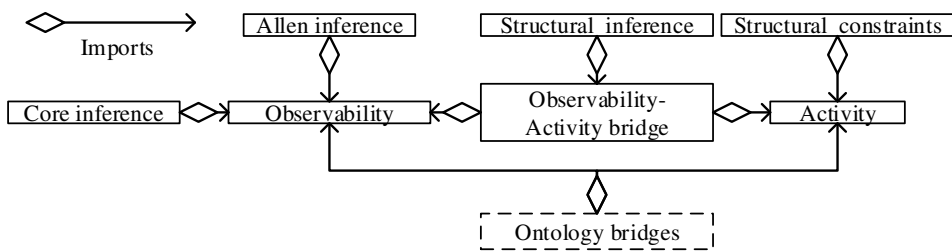


Figure 3: Ontologies in the ODK

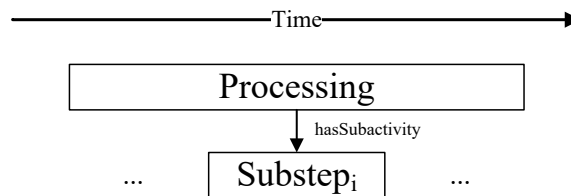


Figure 4: Example of a partial activity decomposition

based on their degree of observability. Such observability flags are propagated during reasoning along the activity relations, resulting in a complete observability description of the entire activity hierarchy.

The structural constraint and inference ontologies can be referred to as *aspect ontologies* in general: they separate orthogonal modelling concerns in a modular way and can be used to enrich a base model (similarly to aspects is aspect-oriented programming). Accordingly, a modeler can work using light-weight and simple ontology concepts (activities and observability), and only perform possibly heavyweight computations/reasoning periodically by including the aspects.

The recommended modelling approach of the ODK is to decompose the complete system model into smaller ontologies (Fig. 5) for maximum flexibility and reusability.

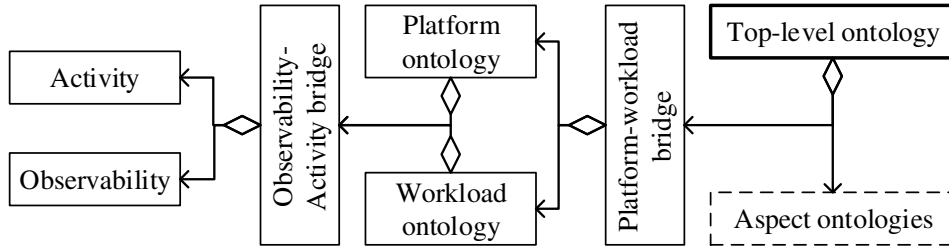


Figure 5: Ontology structure of a multi-component use case

A platform ontology models the composition of executed activity steps and their observability utilizing the ODK core vocabulary. A specific platform is usually just a means to execute a higher-level business scenario, which steps should be modeled as a platform-independent workload ontology whenever possible. This separation allows the flexible evaluation of different architectural/platform design choices by providing a platform-workload bridging ontology for specific scenarios.

The final element of the stack is a top-level (possibly automatically constructed) ontology that unites the pure system model with the chosen aspects of the ODK. Various OWL-DL reasoners can validate and enrich the top-level ontology, resulting in a detailed knowledge representation of the system that will serve as a basis for later performance analysis tasks. A concrete case study following the presented approach is detailed in Sec. 5 through modeling TPC-C benchmark execution on HLF networks.

4.3 Modeling activity hierarchies

Activity (*ACT*) hierarchies are defined with atomic (“leaf”) elements and composite elements supporting further refinement (Fig. 6). The ODK provides the following *Activity* subclasses for modeling activity composition through subsumption relations:

- *AtomicActivity* (*AA*) represents elementary steps without further refinement;

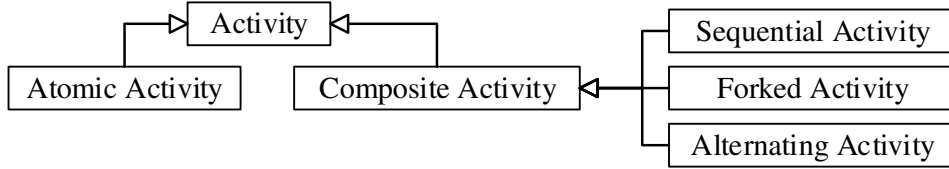


Figure 6: Hierarchy of activity types

- *CompositeActivity* (*CA*) allows further refinement of activities through the following subclasses, representing different composition semantics:
 - *SequentialActivity* (*SA*) allows refinement into a sequence of subactivities;
 - *ForkedActivity* (*FA*) allows refinement into parallel subactivities;
 - *AlternatingActivity* (*TA*) allows refinement into subactivities without additional control flow constraints.

The core classes on the same hierarchy level are disjoint ($AA \cap CA = \emptyset$, $SA \cap FA \cap TA = \emptyset$). However, uncategorized activities and additional composition semantics are allowed to promote extendability, i.e., $AA \cup CA \neq ACT$ and $SA \cup FA \cup TA \neq CA$.

The following high-level relations (OWL object properties) provide the basis for constructing complex hierarchies of activities:

- $Substep_i \subseteq hasParentActivity.Parent$, denoting that activity type $Substep_i$ has a parent (encapsulating) activity of type $Parent$.
- $Parent \subseteq hasSubactivity.Substep_i$, denoting that activity type $Parent$ has a subactivity (a refined substep) of type $Substep_i$.
- $Substep_i \subseteq hasSiblingActivity.Substep_j$, denoting that $Substep_i$ has the same parent activity type as $Substep_j$, i.e., $\exists Parent$ such that $Substep_i \subseteq hasParentActivity.Parent$, and $Substep_j \subseteq hasParentActivity.Parent$

The composite activity subclasses denote the typical (de)composition constructs for activity executions:

Sequential activities (*SA*) group together a sequence of subactivities that follow traditional sequential execution semantics. Moreover, refined relations are introduced (with Allen interval-like semantics, as mapped in Table 3) to further enrich parent-subactivity and sibling relations. The "synonyms" for the Allen relations were introduced to hint at the compositional nature of the activities (and not just their relative temporal placement), aiding modelers with traditional activity modeling backgrounds.

Note, that "gapped" relations indicate an incomplete timeline, hindering later analyses, and probably warranting additional instrumentation. However, the

Table 3: ODK and Allen relation mappings for sequential composition

| Parent Relation \supseteq | Subrelation \equiv | Allen relation |
|-----------------------------|----------------------------------------|-------------------|
| <i>hasParentActivity</i> | <i>startsParentActivity</i> | <i>starts</i> |
| | <i>finishesParentActivity</i> | <i>finishes</i> |
| <i>hasSubactivity</i> | <i>startedBySubactivity</i> | <i>startedBy</i> |
| | <i>finishedBySubactivity</i> | <i>finishedBy</i> |
| <i>hasSiblingActivity</i> | <i>hasImmediatePredecessorActivity</i> | <i>metBy</i> |
| | <i>hasImmediateSuccessorActivity</i> | <i>meets</i> |
| | <i>hasGappedPredecessorActivity</i> | <i>after</i> |
| | <i>hasGappedSuccessorActivity</i> | <i>before</i> |

ODK inference rules can be easily extended to detect "unknown", albeit observable activities whenever possible.

Forked activities (*FA*) group together parallel subactivities that are executed independently of each other. An associated synchronization/join semantic class (\subseteq *hasSyncSemantic.SyncSemantic*) can be used to model the condition when the parent activity is deemed finished.

The ODK defines two synchronization semantics: when *all* (*WaitForAll* \subseteq *SyncSemantic*), or *any* (*WaitForAny* \subseteq *SyncSemantic*) of the subactivities must finish to consider the parent activity done. Extending ontologies can define further semantics, e.g., waiting for the majority of subactivities.

Alternating activities (*TA*) are decomposed into a set of subactivities, disregarding control flow restrictions in cases when the control flow of subactivities is irrelevant. *TA* is a tool of abstraction for concentrating only on the "weight" (i.e., duration) of a subactivity, and not on its scheduling.

A typical use case is the modeling of in-process execution times and database access times of a task, disregarding execution semantics among the substeps: *Task* \subseteq *TA*, *InProcExec*, *DbAccess* \subseteq *hasParentActivity.Task*. Modeling the exact activity flow of computation and database access can be cumbersome for some use cases. Moreover, it may be sufficient during performance analysis to consider only the time/duration spent with each processing types, instead of focusing on their exact, possibly rapidly alternating order.

4.4 Modeling observability

Once the activity model is complete, the next step is modeling which activity temporal aspects (beginning, duration, and/or end) are measured directly in the system (i.e., modeling the placement of sensors and instrumentation) using the *Observability* ontology concepts.

The core concepts can be grouped into three main categories (Fig. 7):

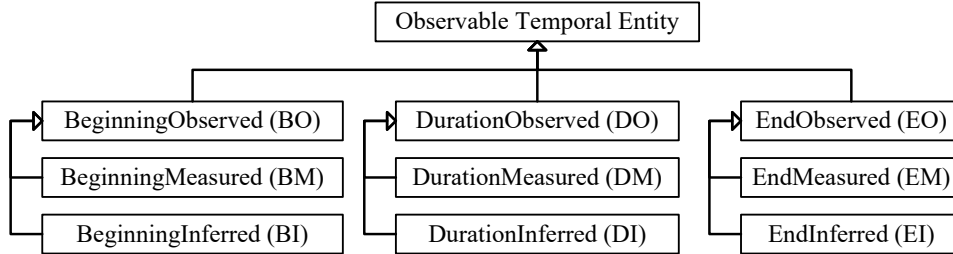


Figure 7: Observability ontology components, with abbreviations

1. observable data in the abstract sense (e.g., $Task \subseteq EndObserved$), denoting that the temporal data is available in some way (measured or inferred);
2. directly measured data (e.g., $Task \subseteq EndMeasured \subseteq EndObserved$), denoting that the data is explicitly measured;
3. and inferred data (e.g., $Task \subseteq Rule \subseteq EndInferred \subseteq EndObserved$).

The following class abbreviations are used in some places for readability (Fig. 7): *BO*, *BM*, *BI*, *DO*, *DM*, *DI*, *EO*, *EM*, and *EI*.

The modeler must "annotate" each activity class if one or more of its temporal aspects are directly measured in the system. For example, if the system logs the end time of an activity *Processing*, then the modeler can add the following axiom to the ontology: $Processing \subseteq EndMeasured$, also implicitly stating that $Processing \subseteq EndMeasured \subseteq EndObserved$. Such annotations will serve as a priori knowledge to the reasoner later. Moreover, additional instrumentation knowledge can be encoded in the ontology if modelers subsume the **Measured* classes (e.g., name of the logging component, format, reference to source code, etc.).

The general classes for observable data (*BO*, *DO*, and *EO*) provide an abstraction layer that hides the exact source of observability. Inference rules will reference this abstract level to handle and propagate explicit and inferred observability uniformly (Sec. 4.5).

The *BI*, *DI*, and *EI* classes are the extension points of the observability ontology, i.e., the superclasses for implementing observability inference rules, as detailed next.

4.5 Observability inference

Given a partially observable activity model, an OWL-DL reasoner can infer further observable temporal aspects based on rules utilizing Allen interval and structural relations. The ODK provides general inference rule ontologies that encode how explicit measurements (i.e., observability) can be propagated along activity relations.

Users of the ODK can build their custom activity ontology of the SUT (containing the hierarchy of activities and their relations), including the explicit observability of activities (i.e., which activities are directly measured through sensor

instrumentation). Then, inputting the inference rule and user ontologies to a reasoner (e.g., in the form of a top-level ontology that imports the previous ontologies) reveals how certain activity temporal aspects (time of beginning, duration, and/or end) can be calculated through related activities. For example, the reasoning process could infer the following knowledge: the *end time* of activity *Processing* can be inferred from the *end time* of its finishing subactivity *Subactivity₄* (based on *Rule_i*).

The observability inference is implemented using the class equivalence construct of OWL2. The rules are modeled as OWL classes (e.g., $RuleX \subseteq EndInferred$) with corresponding equivalence axioms as *criteria* (describing an anonymous class in OWL in the form of $RuleX \equiv criteria$). The axioms of criteria usually encode some kind of temporal data propagation rule among activities, while referencing the abstract observability of the involved activities.

When a reasoner infers that an activity type *Processing* satisfies the criteria (i.e., subsumes the corresponding anonymous class), *Processing* becomes part of the class hierarchy of the corresponding **Observed* class. E.g., if we have a rule (*RuleX*) about inferring the end time of an activity based on some *criteria* (i.e., $RuleX \equiv criteria$), then the following axiom will be added the ontology if *Processing* "matches" *criteria*:

$$Processing \subseteq criteria \equiv RuleX \subseteq EndInferred \subseteq EndObserved$$

Accordingly, *Processing* will be categorized as an *EndObserved* class, allowing the propagation of the newly inferred knowledge through other rules, continuing until no new knowledge can be inferred.

4.6 Inference rules

The inference mechanism is demonstrated through the simple constraint between the beginning time, duration, and end time of any activity instance: $a^b + a^d = a^e, \forall a \in A \subseteq ACT$. This constraint is the basis of the three core inference rules (Eqs. 1–3) provided by the ODK: if two of the temporal aspects are observable, then the third is inferrable. Rules are encoded through equivalent class axioms and subsume the proper *BI*, *DI* or *EI* inference extension points.

$$A \subseteq (DO \cap EO) \equiv Rule1 \implies A \subseteq Rule1 \subseteq BI \subseteq BO \quad (1)$$

$$A \subseteq (BO \cap EO) \equiv Rule2 \implies A \subseteq Rule2 \subseteq DI \subseteq DO \quad (2)$$

$$A \subseteq (BO \cap DO) \equiv Rule3 \implies A \subseteq Rule3 \subseteq EI \subseteq EO \quad (3)$$

The ODK contains numerous additional inference rules based on Allen interval and structural relations. The following conjunctions of criteria (in the form of

$A \subseteq (criteria_1 \cup \dots \cup criteria_n) \implies A \subseteq BO/DO/EO$ succinctly encode the additional rules for inferring beginnings, durations, and ends, respectively:

$$A \subseteq ((FA \cap \forall hasSubactivity.BO) \cup (\exists starts.BO) \cup (\exists startedBy.BO) \cup (\exists metBy.EO) \cup (\exists hasParentActivity.(BO \cap FA))) \implies A \subseteq BO \quad (4)$$

$$A \subseteq ((\exists hasParentActivity.(TA \cap DO) \cap \forall hasSiblingActivity.DO) \cup (TA \cap \forall hasSubactivity.DO)) \implies A \subseteq DO \quad (5)$$

$$A \subseteq ((\exists meets.BO) \cup (\exists finishes.EO) \cup (\exists finishedBy.EO) \cup (FA \cap \forall hasSubactivity.EO)) \implies A \subseteq EO \quad (6)$$

The ODK allows the declaration of additional rule classes by simply subsuming the appropriate *Inferred* classes.

Note that an activity A can match multiple rules. For example, if two temporal aspects are observable, A matches one of Eqs. 1-3. However, now all three aspects are observable, so A matches all three rules. In general, the matching inference rules between temporal data define a data flow network, facilitating various data analysis tasks, as detailed in Sec. 6.

4.7 ODK extendability

The ODK operates with high-level and abstract concepts in order to allow extendability with additional concepts, increasing the flexibility and usability of the model in subsequent analysis tasks. Fig. 8 contextualizes the different ODK capabilities in the typical MDE workflow.

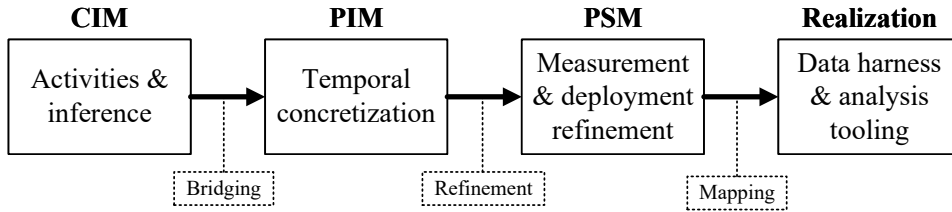


Figure 8: Envisioned MDE flow of activity modelling

The core activity concepts and inference rules comprise a computation-independent model (CIM) for describing observability in a temporal representation-agnostic way. The models at this level (e.g., in Sec. 5) only state knowledge like the beginning of a *Processing* activity is measured. Information about *how* that measurement is acquired, and in *what format*, is omitted. Moreover, inference rules define only the data dependency of calculated measurements – again, omitting the exact computational steps.

The first step towards an actual realization of the analysis process is to enrich the core model with temporal data and corresponding relations. E.g., the OWL

Time ontology¹¹ defines *Intervals* as "a temporal entity with an extent or duration." Furthermore, the ontology defines the *hasBeginning* relation (among others) between temporal entities (such as intervals) and arbitrary time instants. A simple *bridging* between the two ontologies (e.g., *Activity* \equiv *Interval*) enriches the activities with actual temporal data formats. The associated temporal data concretize the manner of measurement calculations, but still neglects the exact source and harness of measurement data, thus acting as a platform-independent model (PIM).

Two ODK aspects support the refinement of PIMs into platform-specific models (PSM). On one hand, additional ontologies can refine classes like *BeginningMeasured* to include the source of the measurement data. For example, an extending ontology could define a *BeginningLogged* subclass of *BeginningMeasured*, providing details about the log format, source software component, and the semantic structure of the message, all aiding the extraction of measurement information in a log processing pipeline.

On the other hand, the ODK provides an *executedBy* relation to associate an *Activity* type with a *Service* type (e.g., *Endorsement* activities are *executedBy* *PeerServices*). Extending ontologies can build on this relation to further model the deployment information related to a certain environment where the SUT is operated. For example, a deployment ontology could maintain information about a HLF network, where each service instance is located on a certain Kubernetes¹² node, in a cluster comprised of several virtual machines, hosted on specific hardware components.

Finally, a technology stack realizing the actual data analysis flow can utilize all levels of the final, rich model to uncover the root cause of an anomalous activity duration/latency (partially demonstrated in Sec. 6.4), even if stemming from the lowest level of the infrastructure.

5 Case study: Modeling TPC-C on Fabric

Performance benchmarks serve as platform-agnostic workload specifications representative for a given domain, facilitating the comparison of different backend platform implementations under reproducible conditions. The benchmark plays the role of a platform-independent model (PIM) in MDE terminology, while the emulated clients and database engine make it platform-specific. The section introduces a compositional model of the TPC-C workload executed on HLF, using the presented ODK concepts as case study.

5.1 Modeling the TPC-C benchmark

TPC-C is a mature online transaction processing (OLTP) benchmark inspired by the typical activities of a wholesale supplier. TPC-C uses a mix of five transaction

¹¹<https://www.w3.org/TR/owl-time/>

¹²<https://kubernetes.io/>

types – with varying complexity – to be executed against a rich database schema (HLF in the case study).

The execution of a TPC-C transaction by an emulated client has the following, strictly sequential composition of steps: (1) the client selects a transaction type (*Menu selection*); (2) then fills the required inputs for the request (*Fill inputs*); (3) then the database engine executes the transaction (*Execute TX*); (4) and finally the client takes some time to think about the next transaction (*Think time*) before starting the next cycle. The model of this client cycle plays the role of the workload ontology in Fig. 5.

Accordingly, the activity model (Fig. 9) declares a top-level/root sequential activity, having four subactivities. The *Menu selection*, *Fill inputs*, and *Think time* subactivities are atomic activities that simply emulate user behavior through artificial delays with specified distributions.

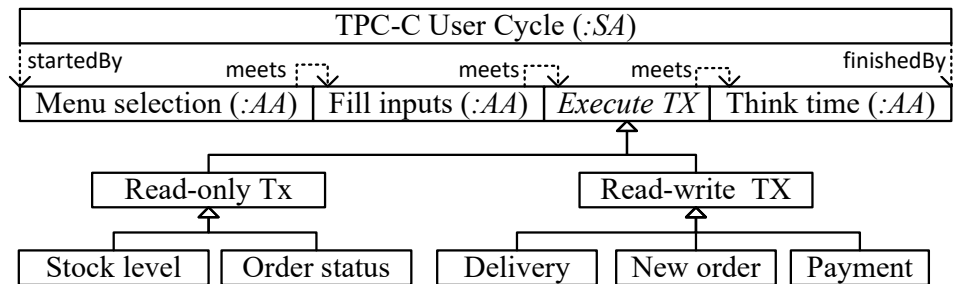


Figure 9: TPC-C transaction execution scheduling

The exact composition of the *Execute TX* activity is specific to the database engine, thus its type does not subsume any of the *Activity* ontology classes. The exact type binding is the task of a platform-workload bridging ontology that maps the platform request execution activities to the *Execute TX* activity.

The TPC-C transaction types are further categorized based on whether they are read-only, or read-write requests, making the bridging easier to platforms that differentiate between the execution of the two categories (like HLF does).

5.2 Modeling the Hyperledger Fabric consensus

During the benchmark measurement, a HLF network served as the “database engine.” The novelty of HLF is its execute-order-validate style consensus mechanism, breaking with the traditional order first approaches [5]. However, its performance characterization is still incomplete. The case study models the detailed HLF consensus mechanism, enriched with client-side observations provided by the Hyperledger Caliper workload generator.

The concepts and consensus steps of HLF are detailed in [5] or in the official documentation.¹³ The section focuses only on the composition of activities (and

¹³<https://hyperledger-fabric.readthedocs.io/en/release-1.4/tflow.html>

not on their technical descriptions) to demonstrate that deep domain knowledge is not required during the guided performance analysis tasks. Note that creating the model, however, requires familiarity with the modeled platform, but ideally it is the responsibility of the designers or platform experts to create such a model.

Fig. 10 details the high-level, sequential steps of the HLF transaction life-cycle. The model plays the role of the platform ontology in Fig. 5.

Clients first assemble and send a transaction proposal to one or multiple peers for *parallel* simulation/endorsement and wait for *all* results (*Awaiting Endorsement* activity) to arrive, modeled by an associated *WaitForAll* synchronization semantic. Once the results are available, the client then sends them to the ordering service and waits for a notification from the network that the transaction was successfully committed or not (*Awaiting Ordering and Validation* activity).

The ordering and validation phase is modelled by two consecutive subactivities: *Block inclusion* and the client *Awaiting Validation* from *any* peer (denoted by a *WaitForAny* synchronization semantic). It is important to note, that the *Awaiting Validation* activity is not a dedicated, explicitly observable activity of the client. It is artificially introduced for convenience to separate the ordering and the validation steps for detailed analysis. This choice demonstrates that the activity model is constructed in a way to facilitate detailed performance analysis, rather than be a technically faithful representation of the platform.

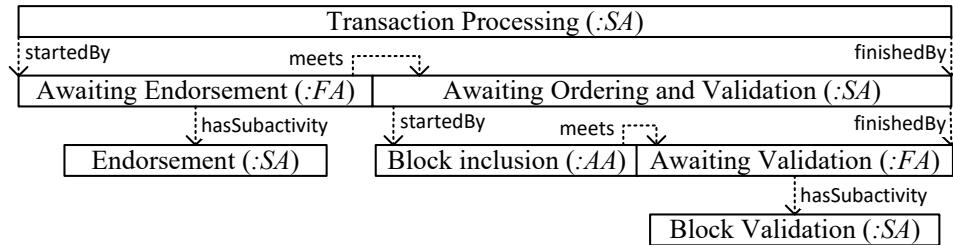


Figure 10: High-level steps of the HLF consensus

The endorsement activity (Fig. 11) consists of the peer receiving the proposal, calling the required chaincode, then returning the result to the client. On the platform level, the *Chaincode Call* activity type is not specified to enable refinement by different use cases, detailed in the next section.

The block validation and commit process of peers is modelled by a hierarchy of activity sequences (Fig. 12). The validation begins by the ordering service delivering the new block to the peer (*Getting Block*). Then the peer checks the block payload and fetches any private data (a privacy feature of HLF) required for further validation (*Check Payload* and *Fetch pvt. data* activities).

The *State validation and commit* step is refined into further subactivities. First, the state modifications of transactions are validated (*State validation*). Then the raw block content is committed to the blockchain storage (*Block Commit*). Next the state modification of valid transactions are committed to the world state database (*State Commit*). Finally, the history database is updated with the data accesses of

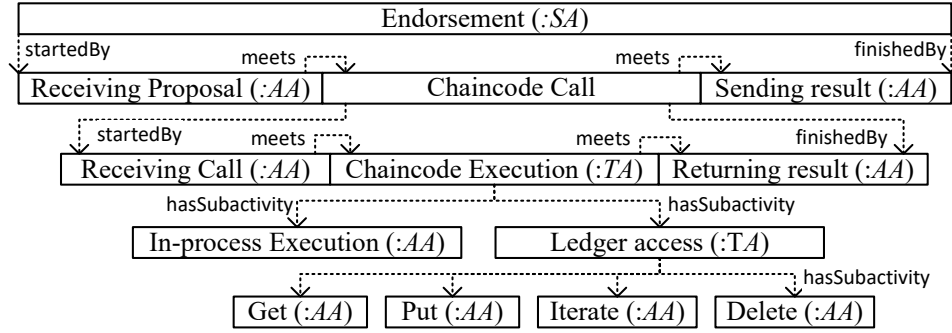


Figure 11: Steps of the endorsement activity

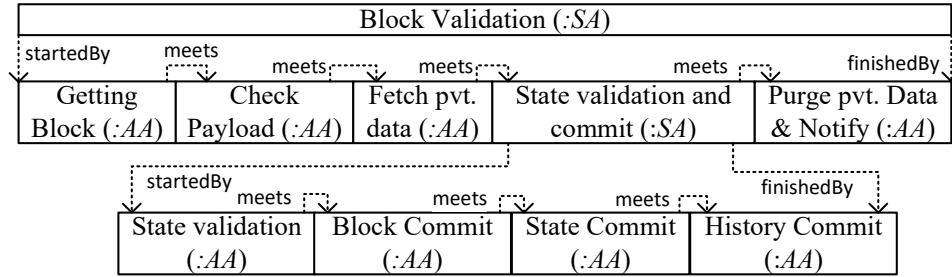


Figure 12: Steps of the validation and commit activity

committed transactions (*History Commit*).

Finally, the peer purges stale private data and sends a notification about the block commit to subscribed clients. Once a client receives a notification about a block/transaction, the transaction life-cycle is considered complete.

5.3 TPC-C and HLF bridge ontology

The case study contains a final ontology that maps/bridges the TPC-C and HLF concepts, achieving the *TPC-C on HLF* model. The mapping plays the role of the platform-workload bridge ontology in Fig. 5.

On one hand, the TPC-C case study chaincode was instrumented to measure the raw execution time of the chaincode. This allows the observation of peer-chaincode communication activities and differentiate between in-process execution and ledger access times (lower part of Fig. 11). The exact control flow of the chaincode is not modelled, alternating activities are used instead to focus only on the duration of subactivities, and not on their order.

On the other hand, the bridge also refines the *Execute TX* class of the TPC-C ontology. Due to the *Read-only TX* and *Read-write TX* class hierarchy, the following equivalence axioms are enough to specify that the workload is executed on HLF: i) *Read-write TX* \equiv *Transaction Processing* and ii) *Read-only TX* \equiv *Query*

Processing (which is a simplified version of transaction processing, containing only the endorsement activity hierarchy, without further ordering or validation).

6 Systematic measurement data analysis

At this point, the workload and platform ontologies are combined, and the measured activities are annotated with the appropriate *BM/DM/EM* observability classes. Inputting the ODK and user-defined ontologies to an OWL-DL reasoner will propagate the measured activity aspects throughout the rest of the model by flagging activities with different inference rule classes. The added classes denote how the beginning, duration and end of a flagged activity can be calculated based on other activity observations.

The case study imported the bridged TPC-C and Fabric (Sec. 5.3), and inference ontologies (Sec. 4.5) into a top-level ontology (as depicted by Fig. 5). Then, the Hermit OWL-DL reasoner [35] (present as a built-in reasoner plugin in the Protégé ontology tool [36]) performed a subsumption relation inference. The resulting (output) ontology contained additional knowledge about how to derive certain temporal aspects of TPC-C and Fabric activities (even when directly unobserved by sensors) based on the explicit measurement data.

The added rule classes define *relations* between the temporal data of different activities. The following subsections provide examples for how such relations can be exploited to:

1. correlate and validate the distributed measurement data;
2. derive further, directly not measured (i.e., indirect) temporal data;
3. validate the conformance of measurement data to the activity model;
4. and systematically guide the bottleneck analysis tasks.

6.1 Correlate and validate measurement data

Online services today exhibit a shift towards micro-service architectures to facilitate different DevOps tasks (e.g., rapid continuous delivery and deployment) and increase certain extra-functional properties of systems (e.g., availability, maintainability, fault tolerance, scalability). Accordingly, an end user request will traverse many services and corresponding components during processing. The same phenomenon is inherently present in distributed, peer-to-peer systems, such as HLF.

In most cases a unique correlation/trace identifier is associated with each request to facilitate its tracing across component boundaries. HLF, for example, associates a unique transaction identifier (TX ID) with each client request, calculated from the client's identity and the time the transaction was constructed. When network components provide logs about certain transaction steps, they also log the corresponding TX ID along with the trace data.

A prerequisite of reconstructing a detailed activity timeline of transactions is the collection and correlation of such distributed traces. Novel observability frameworks, such as OpenTelemetry,¹⁴ may provide means to collect traces across component boundaries. For example, services utilizing OpenTelemetry can also send the collected traces (as metadata) along with the requests to other system components. Even though such approaches can centralize trace collection to a certain level, it is a rather intrusive instrumentation choice, hindering adoption by existing systems (such as HLF¹⁵).

Instead, many systems opt to provide request trace data utilizing their already existing logging capabilities. In this case, distributed transaction traces must be collected and correlated using a separate monitoring stack, which presents its own challenges (but at least it is separate from the core system functionality). Having a detailed activity model for distributed transaction processing (such as the HLF consensus process) can facilitate the correlation and availability check of traces.

A HLF network setup usually contains the following trace sources (Fig. 13):

- optional end-to-end traces logged by the client (Caliper, in this case), with an associated TX ID;
- optional traces logged by chaincodes (one for each executing peer), with an associated TX ID;
- chaincode call traces logged by the peer nodes (one for each executing peer), with an associated *shortened* TX ID (first 8 characters only);
- block validation and commit traces logged by the peer nodes (one for each peer), with an associated block ID;
- and block creation traces logged by the leader orderer node, with an associated block ID.

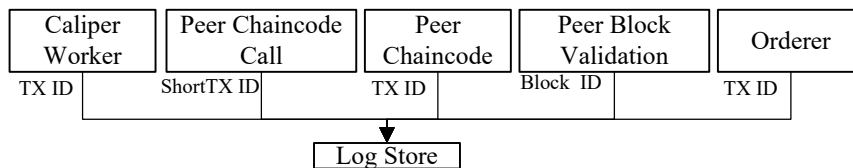


Figure 13: Different trace sources of a transaction

The activity model of the case study defined the measured temporal data of activities, associated with the service types logging them. Accordingly, the prerequisite trace correlation step simply followed the structure of the model to check whether all supposedly measured data are available from all sources.

¹⁴<https://opentelemetry.io/>

¹⁵<https://hyperledger.github.io/fabric-rfcs/text/0000-opentelemetry-tracing.html>

The check revealed two anomalies:

1. there were 4334 transactions with missing traces;
2. there were 2 transactions with more traces than required.

Case 1 had an interesting symmetry in it: there were 2167 transactions where *all* Caliper-side traces were missing; and there were 2167 transactions where *all* other (non-Caliper) traces were missing. This led to the hypothesis that one half is actually corresponding to the other half.

Since the non-Caliper traces constituted a complete data set on their own, the focus of investigation was Caliper's Fabric integration. Further transaction meta-data analysis revealed that all mismatched traces were HLF queries. Finally, the investigation revealed a *bug in Caliper's query submitting logic*.¹⁶ Caliper created a TX ID for the request, but did not pass it along to the HLF SDK, which in turn created a new (and different) TX ID, unknown to Caliper. This resulted in client-side traces having a different TX ID than HLF-side traces.

Case 2 was a similarly peculiar anomaly. Two transactions had chaincode call traces from peers that did not even execute those transactions. Closer inspection revealed that the shortened TX IDs contained a duplicate, i.e., two different TX IDs had the same shortened (8 characters) versions. Accordingly, the pairing of traces was not unique, two transactions got each others chaincode call traces.

Luckily, the correct traces could be restored without data loss through temporal correlation: the conflicting transactions were executed well apart in time. However, if all peers would have executed those transactions, then the short TX ID conflict would have gone unnoticed until later in the analysis workflow (Sec. 6.3). The anomaly showed that reducing the information carried by trace correlation IDs is highly discouraged.

6.2 Deriving indirect measurement data

The final activity model of the HLF consensus refines a transaction into 28 hierarchical steps even if only a single peer endorses and validates transactions. In general, the number of activities corresponding to a transaction is $5 + 13 * E + 10 * V$, where $E \in \mathbb{N}^+$ is the number of endorsing peers for a transaction, and $V \in \mathbb{N}^+$ is the total number of peers in the network (since every peer validates transactions).

Moreover, each activity has three associated temporal data: its beginning time, duration, and end time. Accordingly, the volume of temporal data can quickly increase with the network size and the number of analysed transactions. For the sake of readability, let us assume that only a single peer endorses and validates transactions, resulting in 84 potentially observable temporal data for the 28 activities of *each* transaction.

Figs. 14–16 depict each activity and their corresponding temporal data (beginning, duration, and end). Black-filled shapes mark the directly measured data

¹⁶<https://github.com/hyperledger/caliper/issues/1187>

points. Using a component-off-the-shelf (COTS) HLF as SUT and Caliper as workload generator, there are 18 directly measured data points:

- Caliper marks: the beginning of a transaction; the end time when all endorsements arrive; and the end time when a notification is received about a committed block/transaction.
- Orderer nodes mark the end time when a new block is created.
- Peer nodes mark: the beginning, duration and end of a chaincode call; the end time when a block is received from an orderer; the end time and duration for checking the payload of a new block; and the end time and duration (including durations of some substeps) for validating and committing a block.
- The TPC-C chaincode implementation marks the start time, duration, and end time of the actual chaincode program execution.

The arrows in Figs. 14–16 symbolize the direction of measurement data propagation, i.e., $A \rightarrow B$ means that data B is calculated from data A (and possibly from other data in cases like $A \rightarrow B \leftarrow C$). The arrows essentially represent inference rules in the model, e.g., stating that the start time of an activity can be calculated from its end time and duration (like in the case of the *State validation and commit* activity).

As shown in the figures, the directly measured temporal aspects are sufficient to completely observe the entire activity hierarchy through measurement propagation. If that were not the case, then the missing data propagation paths would identify the places where the SUT needs additional sensor instrumentation to allow for more detailed observability.

Note that Figs. 14–16 are just a single, simplified view of a more complex data flow network determined by the applicable inference rules. The rigorous formal analysis of such data flow networks (in the context of temporal data) is subject to future work.

Moreover, the example assumes a single-peer HLF network. If the network consists of more than one peer, then the single-peer assumption is achieved by reducing the replicated endorsement and validation activities to a single instance by disregarding the non-bottleneck instances:

1. Since transaction endorsements have a *WaitForAll* synchronization semantic, keep only the longest running (i.e., the slowest) *Endorsement* activity and its subactivities.
2. Since block validations have a *WaitForAny* synchronization semantic, keep only the shortest running (i.e., the fastest) *BlockValidation* activity and its subactivities.

At this point, a data analyst can use the formal data flow network to systematically derive new temporal data about the SUT's activities. However, an additional validation step is still needed to ensure not only the cleanness and richness, but also the correctness of the measurement data (or the model).

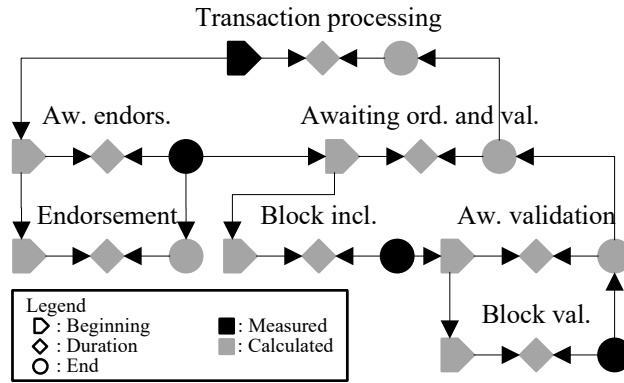


Figure 14: Measurement propagation for high-level HLF activities

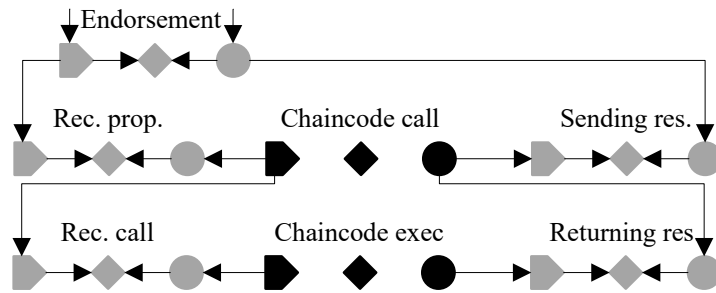


Figure 15: Measurement propagation for the endorsement-related activities

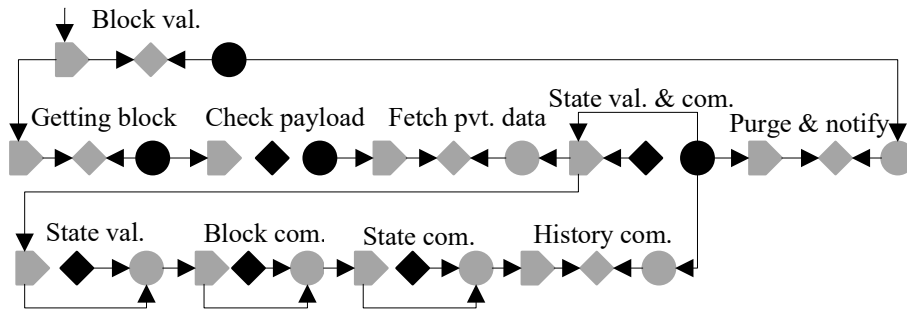


Figure 16: Measurement propagation for the correct validation-related activities

6.3 Model and measurement data validation

Validating the measurement data is an important step to ensure the correctness of data analysis findings and insights. The proposed model-guided approach necessitates the following validation steps before proceeding to the performance analysis tasks:

1. checking the conformance of measurement data to the activity model;
2. and checking the consistency of the measurement data itself.

6.3.1 Detecting modelling errors

The following scenario demonstrates how model conformance checks can reveal activity modelling errors. Such errors can be common if the model is reverse-engineered by others than the platform developers (like in this case study).

For example, HLF peers log the *State validation and commit* activity details using the following message format: [mychannel] Committed block ... in 26ms (state_validation=3ms block_and_pvtdata_commit=16ms state_commit=3ms).

Accordingly, a previous version of the consensus model refined the *State validation and commit* activity as having only three subactivities (state validation, block commit, and state commit, as indicated by the log format). Fig. 17 shows the temporal data propagation for the initial version.

Note how the (directly unobserved) beginning time of the *State validation* subactivity can be calculated in two different ways (highlighted arrows in Fig. 17): i) based on sibling activity data; ii) and/or directly from parent activity data. There should not be any difference between the two paths in the case of a correct model and instrumentation. Validating this assumption requires checking whether the beginning times of the *State validation* subactivities coincide with the beginning times of the *State validation and commit* parent activities for every transaction, as required by the *startedBySubactivity* relation among the two activity classes.

However, performing the check revealed that the equality constraint was violated for *every* transaction. The *State validation* activities always started later than their parent activities, indicating the presence of a hidden subactivity. Moreover, the magnitude of the missing time was sometimes non-negligible (Fig. 18), i.e., it could not be considered a measurement noise, thus warranting further investigation.

As it turns out, the format of the log message was misleading and not all relevant subactivities were listed in the message. The source code inspection of HLF revealed that there is an other non-negligible subactivity performed during *State validation*

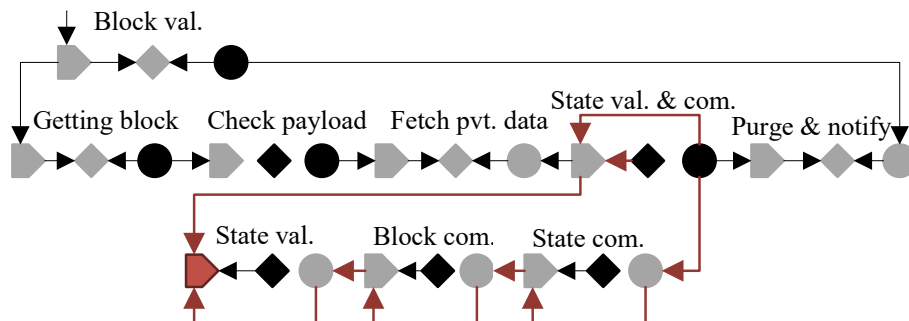


Figure 17: Measurement propagation for the initial validation-related activities

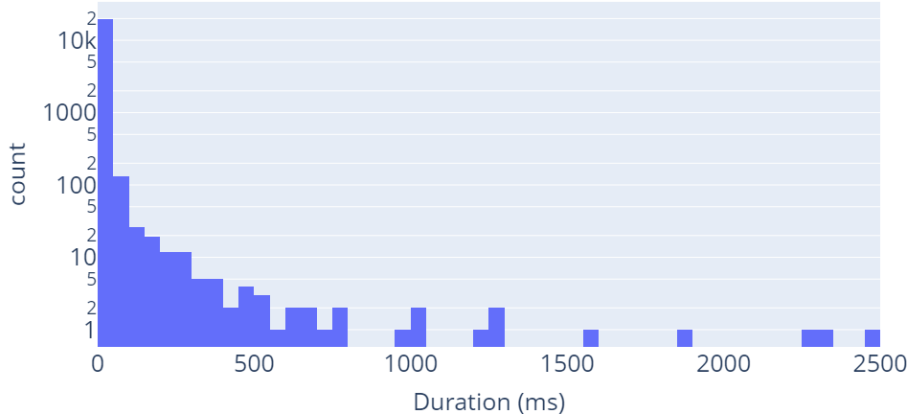


Figure 18: Frequency distribution of missing validation time durations

and *commit*, namely committing the state modifications of a transactions to a history database. Accordingly, the final model of the HLF consensus was extended with the *Commit history* subactivity (Fig. 16).

Note that measurement noises are a common occurrence in complex, especially high-throughput or overloaded systems. The measurement data conformance check also revealed some inconsistencies around the *Check payload* activity. Calculating the beginning time of the activity from its own end time and duration yielded a different result than propagating the end time of its immediate predecessor *Getting block* activity. Even though the propagation path is short and simple, it still violated the modeled activity relationship. However, the magnitude of the amount of unaccounted time for the activity (Fig. 19) is negligible.

One probable explanation could be that the missing time is a side-effect of the logging mechanism: the measured duration was calculated based on times $start_{calc}$ and end_{calc} , while the logging library marked the log message with an $end_{log} > end_{calc}$ timestamp, and end_{log} was taken as the measured end time by the log processing pipeline. An other explanation could be that negligible activities were performed between the two modeled activities that can be safely ignored during performance analysis.

6.3.2 Detecting measurement errors

The systematic data propagation can also aid the detection of measurement (or measurement setup) errors. The missing subactivity issue manifested itself as unaccounted time in the transaction timeline. The other important symptom of inconsistent measurement data is negative durations.

The analysis showed negative *Receiving proposal* activity durations upon measurement data validation. The duration in question is a derived measurement. Its value is indirectly calculated as the difference between the beginning time of calling a chaincode (*Chaincode call* activity) and the beginning time of creating

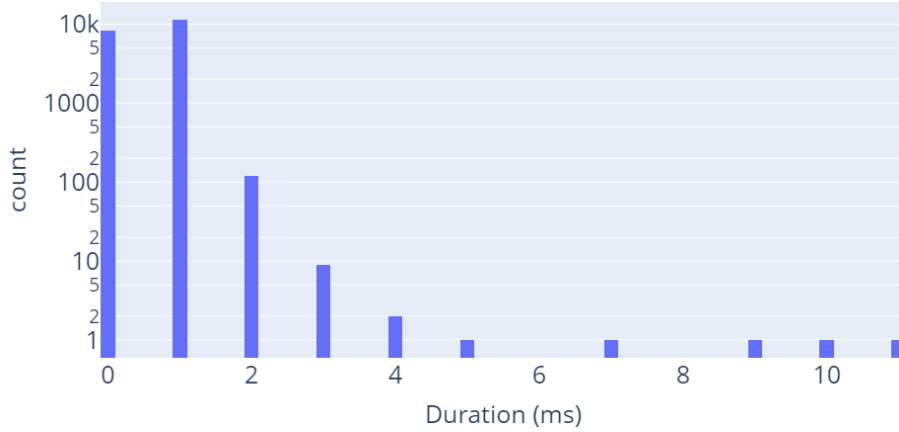


Figure 19: Frequency distribution of missing payload check time durations

a transaction (*Transaction processing* activity), both data being direct measurements. A negative duration result would mean that the chaincode is called before the transaction is even constructed, which is a serious *event causality violation*.

Note that the two direct measurements (the bases of the duration calculation) originate from two different (physical) components in the distributed network: the beginning of *Transaction processing* is captured by Hyperledger Caliper (i.e., the HLF client), while the beginning of the Chaincode call is logged by the HLF peer nodes. Fig. 20 shows the *Receiving Proposal* durations for each transaction over the time of the SUT measurement and reveals a curious trend: the anomalous durations smoothly oscillate around zero over time, i.e., negative durations are not that isolated and sporadic. Moreover, Fig. 20 depicts the activity data of each transaction after non-bottleneck endorsement activities have been eliminated, as outlined in Sec. 6.2. Correspondingly, different *Receiving Proposal* activity durations may originate from different peer nodes of the network.

Combining the observations with the outlined assumptions results in the following working hypothesis: the system clock of a peer node periodically drifts out of sync from the other components. Measurement setup investigations later revealed that network nodes used a default, light-weight time synchronization service instead of a more sophisticated one that provides higher precision.

Measurement errors of such a low magnitude was deemed negligible in the previous section (Fig. 19). However, in this case, the presence of event causality violations shadows the usually insignificant magnitude of the actual measurement error. For example, process mining approaches can produce significantly different results in the presence of such causality violations.

Considering only the atomic activities of the HLF consensus model results in the low-level sequence of steps of the transaction life-cycle. Inputting the measurement data of such activities into a process mining algorithm should result in the process of Fig.21, assuming that the measurement data reflects the correct causality of

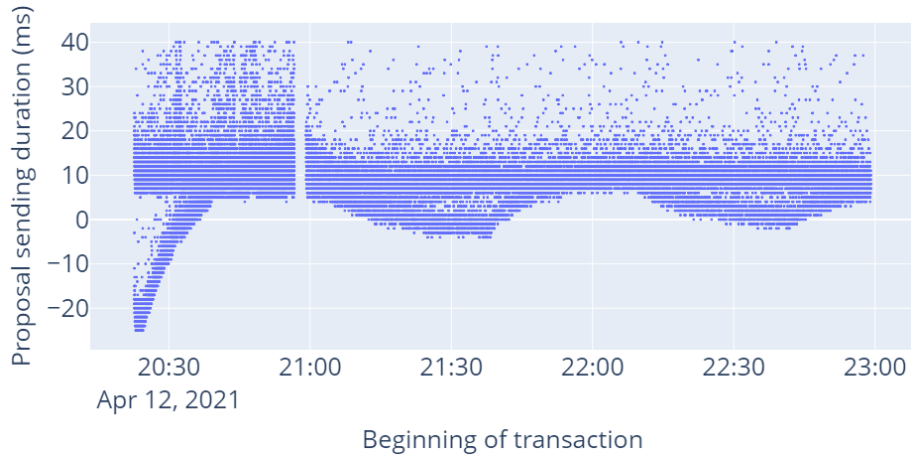


Figure 20: Effect of misaligned system clocks over time

events. However, the presence of causality violations in the input temporal data can lead to an incorrect process model (Fig. 22). Such models can hinder the correct understanding and insights of the SUT (that would be the goal of process mining) even for experienced HLF domain experts.

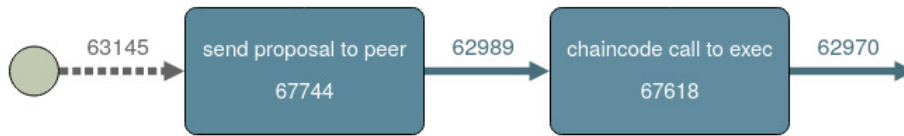


Figure 21: Process mining result without causality violations

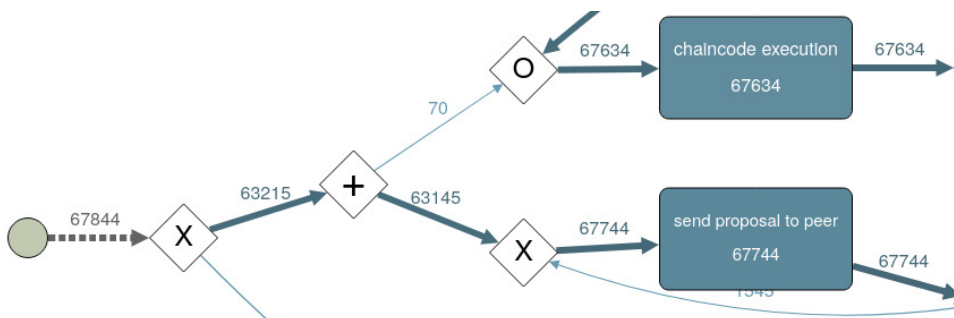


Figure 22: Process mining result with causality violations

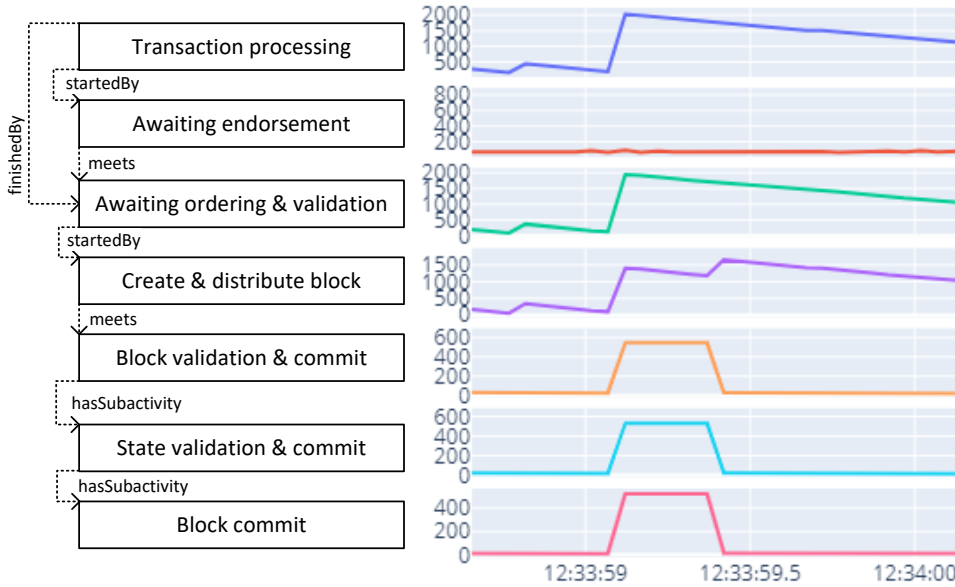


Figure 23: A partial HLF activity time series hierarchy demonstrating a latency anomaly.

Correspondingly, systematically cleaned and validated data is a must if the data analysis workflow incorporates formal approaches. The proposed approach and supporting ontology models enable rigorous (and possibly automated) measurement data validation before performing further performance analysis tasks.

6.4 Guided bottleneck identification

The primary goal and advantage of the proposed approach is that by the time the data analysts reach the actual performance analysis task, the available measurement data is validated, cleaned, and structured among semantically precise relations. The last section demonstrates how bottleneck identification and the root cause analysis of latency anomalies become intuitive and easy-to-perform tasks, given the proper input data.

Let us assume that an end-to-end latency spike is detected on the client-side, classified as an anomaly (the exact anomaly detection methods are out of the scope of this paper). Fig. 23 demonstrates how the hierarchical activity data aids the root cause analysis of the anomaly, uncovering bottlenecks contributing to the latency spike.

The analysis employs a drill-down approach using the parent-subactivity hierarchy relations to gradually pinpoint significant latency contributors. At first, the latency of the high-level transaction processing subactivities are considered. Since endorsement times seem constant during the anomaly, the endorsement activity is

dismissed as bottleneck and root cause. The ordering and validation subactivity, however, exhibits the same latency trend as the end-to-end anomaly. Correspondingly, it becomes the next activity of interest.

At this point, the subactivity latency trends show an interesting pattern. Neither the block creation, nor the block validation subactivities show the same trend as their anomalous parent activity. However, both indicate deviation from their previous baseline latency characteristics. Accordingly, the hierarchical exploratory process supports the identification of *multivariate root causes*.

Block creation is a leaf activity element in the HLF consensus model, thus further root cause analysis along this path would require additional instrumentation or the detailed inspection of corresponding computing resource utilizations. The other prominent root cause path is the block validation activity. Further drill-down steps reveal that the atomic block commit activity caused the latency spike in this path. However, it must be noted that while the block commit anomaly is a short transient spike, the block creation latency needs more time to settle, hinting at some system statefulness and *memory in the performance domain* (probably resulting from a queuing mechanism).

Nevertheless, the hierarchical and systematic approach allows the intuitive and quick identification of bottleneck activities of the SUT. Given the activities of interest, the next analysis steps include the correlation of bottleneck activity latencies with the corresponding component *resource utilizations*, or with the *characteristics of the workload*. Such correlations can answer the question whether the anomaly is caused by resource limits, or it is not really an anomaly, but a change in the presumed workload affected the expected performance characteristics of the SUT. Such analysis, however, is outside the scope of this paper.

7 Conclusion

The increasing volume and dimensionality of performance measurement data necessitate the rigorous model-based support of data analysis tasks, such as bottleneck identification. While traditional DevOps approaches already benefit from MDE, performance data analysis lacks such support.

The paper proposed an ontology-guided workflow (and presented the corresponding ODK) for modeling the composition of complex platform activities and their explicit observability. The ODK also supplies numerous inference rules to reason about the implicit observability of activities, creating a rich model serving as a strong formal basis for later performance analysis tasks.

A representative case study demonstrated the advantages of the approach: a model-guided drill-down bottleneck identification process for a TPC-C benchmark workload executed on a HLF network. The current work aims at the integration of domain-specific knowledge in performance analysis into a core ontology, providing a strong formal foundation for measurement data analysis and performance monitoring of complex systems.

References

- [1] Abelló, Alberto, Romero, Oscar, Pedersen, Torben Bach, Berlanga, Rafael, Nebot, Victoria, Aramburu, María José, and Simitsis, Alkis. Using semantic web technologies for exploratory OLAP: A survey. *IEEE Trans. on Knowledge and Data Engineering*, 27(2):571–588, 2015. DOI: [10.1109/TKDE.2014.2330822](https://doi.org/10.1109/TKDE.2014.2330822).
- [2] Allen, James F. Maintaining Knowledge about Temporal Intervals. *Communications of the ACM*, 26(11):832–843, 1983. DOI: [10.1145/182.358434](https://doi.org/10.1145/182.358434).
- [3] Allen, James F. and Ferguson, George. Actions and events in interval temporal logic. *Journal of logic and computation*, 4(5):531–579, 1994. DOI: [10.1007/978-0-585-28322-7_7](https://doi.org/10.1007/978-0-585-28322-7_7).
- [4] Androulaki, Elli, De Caro, Angelo, Neugschwandtner, Matthias, and Sorniotti, Alessandro. Endorsement in Hyperledger Fabric. In *Proceedings — 2nd IEEE Int. Conf. on Blockchain*, pages 510–519. Institute of Electrical and Electronics Engineers Inc., 2019. DOI: [10.1109/Blockchain.2019.00077](https://doi.org/10.1109/Blockchain.2019.00077).
- [5] Androulaki, Elli et al. Hyperledger Fabric: A distributed operating system for permissioned blockchains. In *Proc. of the Thirteenth EuroSys Conf.*, 2018. DOI: [10.1145/3190508.3190538](https://doi.org/10.1145/3190508.3190538).
- [6] Ardagna, Claudio A., Bellandi, Valerio, Ceravolo, Paolo, Damiani, Ernesto, Bezzi, Michele, and Hebert, Cedric. A Model-Driven Methodology for Big Data Analytics-as-a-Service. In *Proc. of the IEEE 6th Int. Congress on Big Data*, pages 105–112, 2017. DOI: [10.1109/BigDataCongress.2017.23](https://doi.org/10.1109/BigDataCongress.2017.23).
- [7] Baliga, Arati, Solanki, Nitesh, Verekar, Shubham, Pednekar, Amol, Kamat, Pandurang, and Chatterjee, Siddhartha. Performance characterization of Hyperledger Fabric. In *Crypto Valley Conf. on Blockchain Technology*, pages 65–74, 2018. DOI: [10.1109/CVCBT.2018.00013](https://doi.org/10.1109/CVCBT.2018.00013).
- [8] Becker, Steffen, Koziol, Heiko, and Reussner, Ralf. The Palladio component model for model-driven performance prediction. *J. of Systems and Software*, 82(1):3–22, 2009. DOI: [10.1016/j.jss.2008.03.066](https://doi.org/10.1016/j.jss.2008.03.066).
- [9] Bellini, Pierfrancesco, Mattolini, Riccardo, and Nesi, Paolo. Temporal logics for real-time system specification. *ACM Computing Surveys (CSUR)*, 32(1):12–42, 2000. DOI: [10.1145/349194.349197](https://doi.org/10.1145/349194.349197).
- [10] Bergman, Sara, Asplund, Mikael, and Nadjm-Tehrani, Simin. Permissioned blockchains and distributed databases: A performance study. *Concurrency and Computation: Practice and Experience*, 32(12):e5227, 2020. DOI: [10.1002/cpe.5227](https://doi.org/10.1002/cpe.5227).

- [11] Cerveira, Frederico, Kocsis, Imre, Barbosa, Raul, Madeira, Henrique, and Pataricza, András. Exploratory data analysis of fault injection campaigns. In *2018 IEEE International Conference on Software Quality, Reliability and Security (QRS)*, pages 191–202, 2018. DOI: [10.1109/QRS.2018.00033](https://doi.org/10.1109/QRS.2018.00033).
- [12] Chen, Liming, Nugent, Chris D., and Wang, Hui. A knowledge-driven approach to activity recognition in smart homes. *IEEE Trans. on Knowledge and Data Engineering*, 24(6):961–974, 2012. DOI: [10.1109/TKDE.2011.51](https://doi.org/10.1109/TKDE.2011.51).
- [13] Foschini, Luca, Gavagna, Andrea, Martuscelli, Giuseppe, and Montanari, Rebecca. Hyperledger Fabric blockchain: Chaincode performance analysis. In *IEEE Int. Conf. on Communications*, 2020. DOI: [10.1109/ICC40277.2020.9149080](https://doi.org/10.1109/ICC40277.2020.9149080).
- [14] Garg, Supriya, Nam, Julia Eunju, Ramakrishnan, I.V., and Mueller, Klaus. Model-driven visual analytics. In *VAST'08 - IEEE Symp. on Visual Analytics Science and Technology, Proc.*, pages 19–26, 2008. DOI: [10.1109/VAST.2008.4677352](https://doi.org/10.1109/VAST.2008.4677352).
- [15] Garlan, D., Cheng, S.W., Huang, A.C., Schmerl, B., and Steenkiste, P. Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer*, 37(10):46–54, 2004. DOI: [10.1109/MC.2004.175](https://doi.org/10.1109/MC.2004.175).
- [16] Gilmore, Stephen and Hillston, Jane. The PEPA workbench: A tool to support a process algebra-based approach to performance modelling. In *International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, pages 353–368. Springer, 1994. DOI: [10.1007/3-540-58021-2_20](https://doi.org/10.1007/3-540-58021-2_20).
- [17] Gorenflo, Christian, Lee, Stephen, Golab, Lukasz, and Keshav, Srinivasan. FastFabric: Scaling Hyperledger Fabric to 20,000 transactions per second. In *IEEE Int. Conf. on Blockchain and Cryptocurrency*, pages 455–463, 2019. DOI: [10.1002/nem.2099](https://doi.org/10.1002/nem.2099).
- [18] Gupta, Himanshu, Hans, Sandeep, Aggarwal, Kushagra, Mehta, Sameep, Chatterjee, Bapi, and Jayachandran, Praveen. Efficiently processing temporal queries on Hyperledger Fabric. In *Proc. IEEE 34th Int. Conf. on Data Engineering, ICDE 2018*, pages 1435–1440. Institute of Electrical and Electronics Engineers Inc., 2018. DOI: [10.1109/ICDE.2018.00167](https://doi.org/10.1109/ICDE.2018.00167).
- [19] Hao, Yue, Li, Yi, Dong, Xinghua, Fang, Li, and Chen, Ping. Performance analysis of consensus algorithm in private blockchain. In *IEEE Intelligent Vehicles Symposium, Proceedings*, pages 280–285. Institute of Electrical and Electronics Engineers Inc., 2018. DOI: [10.1109/IVS.2018.8500557](https://doi.org/10.1109/IVS.2018.8500557).
- [20] Hashem, Ibrahim Abaker Targio, Yaqoob, Ibrar, Anuara, Nor Badrul, Mokhtar, Salimah, Gani, Abdullah, and Khanb, Samee Ullah. The rise of "big data" on cloud computing: Review and open research issues. *Information Systems*, 47:98–115, 2015. DOI: [10.1016/j.is.2014.07.006](https://doi.org/10.1016/j.is.2014.07.006).

- [21] Helaoui, Rim, Niepert, Mathias, and Stuckenschmidt, Heiner. Recognizing interleaved and concurrent activities using qualitative and quantitative temporal relationships. In *Pervasive and Mobile Computing*, volume 7, pages 660–670. Elsevier B.V., 2011. DOI: [10.1016/j.pmcj.2011.08.004](https://doi.org/10.1016/j.pmcj.2011.08.004).
- [22] Hermanns, Holger, Herzog, Ulrich, and Katoen, Joost-Pieter. Process algebra for performance evaluation. *Theoretical Computer Science*, 274(1-2):43–87, 2002. DOI: [10.1016/S0304-3975\(00\)00305-4](https://doi.org/10.1016/S0304-3975(00)00305-4).
- [23] Horrocks, Ian, Kutz, Oliver, and Sattler, Ulrike. The even more irresistible SROIQ. In *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning, KR'06*, page 57–67. AAAI Press, 2006.
- [24] Inagaki, Tatsushi, Ueda, Yohei, Nakaïke, Takuya, and Ohara, Moriyoshi. Profile-based detection of layered bottlenecks. In *Proc. of the 2019 ACM/SPEC Int. Conf. on Performance Engineering*, pages 197–208. ACM, 2019. DOI: [10.1145/3297663.3310296](https://doi.org/10.1145/3297663.3310296).
- [25] Javaid, Haris, Hu, Chengchen, and Brebner, Gordon. Optimizing validation phase of Hyperledger Fabric. In *IEEE Computer Society's Annual Int. Symp. on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems*, pages 269–275, 2019. DOI: [10.1109/MASCOTS.2019.00038](https://doi.org/10.1109/MASCOTS.2019.00038).
- [26] Jiang, Lili, Chang, Xiaolin, Liu, Yuhang, Mišić, Jelena, and Mišić, Vojislav B. Performance analysis of Hyperledger Fabric platform: A hierarchical model approach. *Peer-to-Peer Networking and Applications*, 13(3):1014–1025, 2020. DOI: [10.1007/s12083-019-00850-z](https://doi.org/10.1007/s12083-019-00850-z).
- [27] Klein, John, Gorton, Ian, Alhmod, Laila, Gao, Joel, Gemici, Caglayan, Kapoor, Rajat, Nair, Prasanth, and Saravagi, Varun. Model-driven observability for big data storage. In *13th Working IEEE/IFIP Conf. on Software Architecture, WICSA 2016*, pages 134–139, 2016. DOI: [10.1109/WICSA.2016.27](https://doi.org/10.1109/WICSA.2016.27).
- [28] Kocsis, Imre, Klenik, Attila, Pataricza, András, Telek, Miklós, Deé, Flórián, and Cseh, Dávid. Systematic performance evaluation using component-in-the-loop approach. *Int. Journal of Cloud Computing*, 7(3-4):336–357, 2018. DOI: [10.1504/IJCC.2018.095401](https://doi.org/10.1504/IJCC.2018.095401).
- [29] Kuzlu, Murat, Pipattanasomporn, Manisa, Gurses, Levent, and Rahman, Saifur. Performance analysis of a Hyperledger Fabric blockchain framework: Throughput, latency and scalability. In *2nd IEEE Int. Conf. on Blockchain*, pages 536–540, 2019. DOI: [10.1109/Blockchain.2019.00003](https://doi.org/10.1109/Blockchain.2019.00003).
- [30] Lamport, Leslie. The temporal logic of actions. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 16(3):872–923, 1994. DOI: [10.1145/177492.177726](https://doi.org/10.1145/177492.177726).

- [31] Lamport, Leslie. *Specifying concurrent systems with TLA+*. In *Calculational System Design*, pages 183–247. IOS Press, Amsterdam, 1999.
- [32] Lowe, Gavin. Probabilistic and prioritized models of timed CSP. *Theoretical Computer Science*, 138(2):315–352, 1995. DOI: [10.1016/0304-3975\(94\)00171-E](https://doi.org/10.1016/0304-3975(94)00171-E).
- [33] Meditskos, Georgios, Dasiopoulou, Stamatia, Efstathiou, Vasiliki, and Kompatsiaris, Ioannis. Ontology patterns for complex activity modelling. In *Proc. International Workshop on Rules and Rule Markup Languages for the Semantic Web*, volume 8035 of *Lecture Notes in Computer Science*, pages 144–157. Springer, 2013. DOI: [10.1007/978-3-642-39617-5_15](https://doi.org/10.1007/978-3-642-39617-5_15).
- [34] Miyamae, Takeshi, Honda, Takeo, Tamura, Masahisa, and Kawaba, Motoyuki. Performance improvement of the consortium blockchain for financial business applications. *Journal of Digital Banking*, 2(4):369–378, 2018.
- [35] Motik, Boris, Shearer, Rob, and Horrocks, Ian. Hypertableau reasoning for description logics. *Journal of Artificial Intelligence Research*, 36(1):165–228, 2009. DOI: <https://doi.org/10.1613/jair.2811>.
- [36] Musen, Mark A. The protégé project: A look back and a look forward. *AI matters*, 1(4):4–12, 2015. DOI: [10.1145/2757001.2757003](https://doi.org/10.1145/2757001.2757003).
- [37] Nakaike, Takuya, Zhang, Qi, Ueda, Yohei, Inagaki, Tatsushi, and Ohara, Moriyoshi. Hyperledger Fabric performance characterization and optimization using GoLevelDB benchmark. In *IEEE Int. Conf. on Blockchain and Cryptocurrency*, pages 1–9, 2020. DOI: [10.1109/ICBC48266.2020.9169454](https://doi.org/10.1109/ICBC48266.2020.9169454).
- [38] Nasir, Qassim, Qasse, Ilham A., Abu Talib, Manar, and Nassif, Ali Bou. Performance analysis of Hyperledger Fabric platforms. *Security and Communication Networks*, pages 1–14, 2018. DOI: [10.1155/2018/3976093](https://doi.org/10.1155/2018/3976093).
- [39] Neumayr, Bernd, Anderlik, Stefan, and Schrefl, Michael. Towards ontology-based OLAP: Datalog-based reasoning over multidimensional ontologies. In *Int. Conf. on Information and Knowledge Management*, pages 41–48, 2012. DOI: [10.1145/2390045.2390053](https://doi.org/10.1145/2390045.2390053).
- [40] Nguyen, Minh Quang, Loghin, Dumitrel, Tuan, Tien, and Dinh, Anh. Understanding the scalability of Hyperledger Fabric. In *45th International Conference on Very Large Data Bases*, 2019.
- [41] Nguyen, Thanh Son Lam, Jourjon, Guillaume, Potop-Butucaru, Maria, and Thai, Kim Loan. Impact of network delays on Hyperledger Fabric. In *INFOCOM 2019 — IEEE Conf. on Computer Communications Workshops 2019*, pages 222–227, 2019. DOI: [10.1109/INFCOMW.2019.8845168](https://doi.org/10.1109/INFCOMW.2019.8845168).
- [42] Niemi, Tapio and Niinimäki, Marko. Ontologies and summarizability in OLAP. In *Proc. of the ACM Symp. on Applied Computing*, pages 1349–1353, 2010. DOI: [10.1145/1774088.1774378](https://doi.org/10.1145/1774088.1774378).

- [43] Okeyo, George, Chen, Liming, Wang, Hui, and Sterritt, Roy. A hybrid ontological and temporal approach for composite activity modelling. In *Proc. of the 11th IEEE Int. Conf. on Trust, Security and Privacy in Computing and Communications*, pages 1763–1770, 2012. DOI: [10.1109/TrustCom.2012.34](https://doi.org/10.1109/TrustCom.2012.34).
- [44] Parsia, Bijan, Matentzoglou, Nicolas, Gonçalves, Rafael S., Glimm, Birte, and Steigmiller, Andreas. The OWL Reasoner Evaluation (ORE) 2015 competition report. *Journal of Automated Reasoning*, 59(4):455–482, 2017. DOI: [10.1007/s10817-017-9406-8](https://doi.org/10.1007/s10817-017-9406-8).
- [45] Perer, Adam and Shneiderman, Ben. Systematic yet flexible discovery: Guiding domain experts through exploratory data analysis. In *Proc. Int. Conf. on Intelligent User Interfaces*, pages 109–118. ACM Press, 2008. DOI: [10.1145/1378773.1378788](https://doi.org/10.1145/1378773.1378788).
- [46] Pnueli, Amir. The temporal semantics of concurrent programs. *Theoretical Computer Science*, 13(1):45–60, 1981. DOI: [10.1016/0304-3975\(81\)90110-9](https://doi.org/10.1016/0304-3975(81)90110-9).
- [47] Pongnumkul, Suporn, Siripanpornchana, Chaiyaphum, and Thajchayapong, Suttipong. Performance analysis of private blockchain platforms in varying workloads. In *26th Int. Conf. on Computer Communications and Networks*, pages 1–6, 2017. DOI: [10.1109/ICCCN.2017.8038517](https://doi.org/10.1109/ICCCN.2017.8038517).
- [48] Prat, Nicolas, Megdiche, Imen, and Akoka, Jacky. Multidimensional models meet the semantic web: Defining and reasoning on OWL-DL ontologies for OLAP. In *Int. Conf. on Information and Knowledge Management, Proc.*, pages 17–24, 2012. DOI: [10.1145/2390045.2390049](https://doi.org/10.1145/2390045.2390049).
- [49] Riboni, Daniele and Bettini, Claudio. OWL 2 modeling and reasoning with complex human activities. *Pervasive and Mobile Computing*, 7(3):379–395, 2011. DOI: [10.1016/j.pmcj.2011.02.001](https://doi.org/10.1016/j.pmcj.2011.02.001).
- [50] Shalaby, Salma, Abdellatif, Alaa Awad, Al-Ali, Abdulla, Mohamed, Amr, Erbad, Aiman, and Guizani, Mohsen. Performance evaluation of Hyperledger Fabric. In *IEEE Int. Conf. on Informatics, IoT, and Enabling Technologies*, pages 608–613, 2020. DOI: [10.1109/ICIoT48696.2020.9089614](https://doi.org/10.1109/ICIoT48696.2020.9089614).
- [51] Sharma, Ankur, Schuhknecht, Felix Martin, Agrawal, Divya, and Dittrich, Jens. How to databasify a blockchain: the case of Hyperledger Fabric. *arxiv.org*, 2018. DOI: [10.48550/arXiv.1810.13177](https://doi.org/10.48550/arXiv.1810.13177), Preprint.
- [52] Streit, Marc, Schulz, Hans-Jörg, Lex, Alexander, Schmalstieg, Dieter, and Schumann, Heidrun. Model-driven design for the visual analysis of heterogeneous data. *IEEE Trans. on Visualization and Computer Graphics*, 18(6):998–1010, 2012. DOI: [10.1109/TVCG.2011.108](https://doi.org/10.1109/TVCG.2011.108).

- [53] Sukhwani, Harish, Martínez, José M., Chang, Xiaolin, Trivedi, Kishor S., and Rindos, Andy. Performance modeling of PBFT consensus process for permissioned blockchain network (Hyperledger Fabric). In *IEEE Symp. on Reliable Distributed Systems*, pages 253–255, 2017. DOI: [10.1109/SRDS.2017.36](https://doi.org/10.1109/SRDS.2017.36).
- [54] Sukhwani, Harish, Wang, Nan, Trivedi, Kishor S., and Rindos, Andy. Performance modeling of Hyperledger Fabric (permissioned blockchain network). In *17th IEEE Int. Symp. on Network Computing and Applications*, 2018. DOI: [10.1109/NCA.2018.8548070](https://doi.org/10.1109/NCA.2018.8548070).
- [55] Thakkar, Parth, Nathan, Senthil, and Viswanathan, Balaji. Performance benchmarking and optimizing Hyperledger Fabric blockchain platform. In *26th IEEE Int. Symp. on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, pages 264–276, 2018. DOI: [10.1109/MASCOTS.2018.00034](https://doi.org/10.1109/MASCOTS.2018.00034).
- [56] van der Aalst, Wil et al. Process mining manifesto. In *Lecture Notes in Business Information Processing*, volume 99, pages 169–194, 2012. DOI: [10.1007/978-3-642-28108-2_19](https://doi.org/10.1007/978-3-642-28108-2_19).
- [57] Wang, Canhui and Chu, Xiaowen. Performance characterization and bottleneck analysis of Hyperledger Fabric. In *Int. Conf. on Distributed Computing Systems*, pages 1281–1286, 2020. DOI: [10.1109/ICDCS47774.2020.00165](https://doi.org/10.1109/ICDCS47774.2020.00165).
- [58] Wang, Shuo. Performance evaluation of Hyperledger Fabric with malicious behavior. In *Lecture Notes in Computer Science*, volume 11521, pages 211–219, 2019. DOI: [10.1007/978-3-030-23404-1_15](https://doi.org/10.1007/978-3-030-23404-1_15).
- [59] Wirth, Rüdiger. CRISP-DM: Towards a standard process model for data mining. In *Proceedings of the Fourth International Conference on the Practical Application of Knowledge Discovery and Data Mining*, volume 1, pages 29–39, 2000. <http://www.cs.unibo.it/~danilo.montesi/CBD/Beatriz/10.1.1.198.5133.pdf>.
- [60] Wun, Alex, Petrovi, Milenko, and Jacobsen, Hans Arno. A system for semantic data fusion in sensor networks. In *Proc. of the 2007 Inaugural Int. Conf. on Distributed Event-Based Systems*, volume 233, pages 75–79, 2007. DOI: [10.1145/1266894.1266907](https://doi.org/10.1145/1266894.1266907).
- [61] Xiao, Ling, Gerth, John, and Hanrahan, Pat. Enhancing visual analysis of network traffic using a knowledge representation. In *IEEE Symp. on Visual Analytics Science and Technology, VAST 2006*, pages 107–114, 2006. DOI: [10.1109/VAST.2006.261436](https://doi.org/10.1109/VAST.2006.261436).
- [62] Xu, Xiaoqiong, Sun, Gang, Luo, Long, Cao, Huilong, Yu, Hongfang, and V.Vasilakos, Athanasios. Latency performance modeling and analysis for Hyperledger Fabric blockchain network. *Information Processing & Management*, 58(1):102436, 2021. DOI: [10.1016/j.ipm.2020.102436](https://doi.org/10.1016/j.ipm.2020.102436).

- [63] Yang, Di, Rundensteiner, Elke A., and Ward, Matthew O. Analysis guided visual exploration of multivariate data. In *VAST IEEE Symp. on Visual Analytics Science and Technology 2007*, pages 83–90, 2007. DOI: [10.1109/VAST.2007.4389000](https://doi.org/10.1109/VAST.2007.4389000).
- [64] Yuan, Pu, Zheng, Kan, Xiong, Xiong, Zhang, Kuan, and Lei, Lei. Performance modeling and analysis of a Hyperledger-based system using GSPN. *Computer Communications*, 153:117–124, 2020. DOI: [10.1016/j.comcom.2020.01.073](https://doi.org/10.1016/j.comcom.2020.01.073).

Received 19th December 2021