

DOI: 10.15514/ISPRAS-2021-33(3)-9



Method of Performance Analysis of Time-Critical Applications Using DB-Nets

A.M. Rigin, ORCID: 0000-0003-4081-9144 <amrigin@edu.hse.ru>
S.A. Shershakov, ORCID: 0000-0001-8173-5970 <sshershakov@hse.ru>

HSE University,
20, Myasnitskaya st., Moscow, 101000, Russia

Abstract. These days, most of time-critical business processes are performed using computer technologies. As an example, one can consider financial processes including trading on stock exchanges powered by electronic communication protocols such as the Financial Information eXchange (FIX) Protocol. One of the main challenges emerging with such processes concerns maintaining the best possible performance since any unspecified delay may cause a large financial loss or other damage. Therefore, performance analysis of time-critical systems and applications is required. In the current work, we develop a novel method for a performance analysis of time-critical applications based on the db-net formalism, which combines the ability of colored Petri nets to model a system control flow with the ability to model relational database states. This method allows to conduct a performance analysis for time-critical applications that work as transactional systems and have log messages which can be represented in the form of table records in a relational database. One of such applications is a FIX protocol-based trading communication system. This system is used in the work to demonstrate applicability of the proposed method for time-critical systems performance analysis. However, there are plenty of similar systems existing for different domains, and the method can also be applied for a performance analysis of these systems. The software prototype is developed for testing and demonstrating abilities of the method. This software prototype is based on an extension of Renew software tool, which is a reference net simulator. The testing input for the software prototype includes a test log with FIX messages, provided by a software developer of testing solutions for one of the global stock exchanges. An application of the method for quantitative analysis of maximum acceptable delay violations is presented. The developed method allows to conduct a performance analysis as a part of conformance checking of a considered system. The method can be used in further research in this domain as well as in testing the performance of real time-critical software systems.

Keywords: performance analysis; time-critical applications; db-nets; FIX protocol; software modeling; software testing

For citation: Rigin A.M., Shershakov S.A. Method of Performance Analysis of Time-Critical Applications Using DB-Nets. Trudy ISP RAN/Proc. ISP RAS, vol. 33, issue 3, 2021, pp. 109-122. DOI: 10.15514/ISPRAS-2021-33(3)-9

Acknowledgements. This work is supported by the Basic Research Program at the HSE University.

Метод анализа производительности критичных по времени приложений с помощью DB-Nets

A.M. Ригин, ORCID: 0000-0003-4081-9144 <amrigin@edu.hse.ru>
С.А. Шершаков, ORCID: 0000-0001-8173-5970 <sshershakov@hse.ru>

Национальный исследовательский университет «Высшая школа экономики»,
101000, Россия, г. Москва, ул. Мясницкая, д. 20.

Аннотация. В настоящее время большинство критичных по времени бизнес-процессов выполняются с использованием компьютерных технологий. В качестве примера можно рассмотреть финансовые процессы, включая торговлю на фондовых биржах, использующие такие протоколы передачи информации, как Financial Information eXchange (FIX) Protocol. Один из основных вызовов, возникающих относительно таких процессов, – это поддержка наилучшей производительности, так как любая задержка, не установленная спецификацией, может привести к большим финансовым потерям и иному ущербу. Следовательно, необходимо проводить анализ производительности критичных по времени систем и приложений. В данной работе предложен новый метод для анализа производительности критичных по времени приложений, основанный на формализме db-net. Этот формализм позволяет моделировать поток управления системой с использованием цветных сетей Петри, а также моделировать состояния реляционной базы данных. Метод позволяет проводить анализ производительности критичных по времени приложений, которые работают как транзакционные системы и создают логи с сообщениями, представимыми в форме записей в таблице реляционной базы данных. Примером таких приложений является коммуникационная система для торговли на фондовой бирже, работающая на основе протокола FIX. Эта система рассматривается в данной работе для демонстрации применимости предложенного метода. В то же время существует множество подобных систем в различных предметных областях, и предложенный метод может быть также применён и для анализа производительности таких систем. Для тестирования и апробирования метода разработан программный прототип. Он основан на расширении программного инструмента Renew – симулятора сеточных сетей Петри. Прототип протестирован на лог, содержащем сообщения протокола FIX, предоставленном разработчиком решений для тестирования программного обеспечения одной из мировых фондовых бирж. Показано применение метода для количественного анализа превышений максимально допустимых задержек между сообщениями. Разработанный метод позволяет выполнять анализ производительности как часть проверки соответствия свойств системы заданной модели (conformance checking). Метод может быть использован как для научно-исследовательских целей, так и для анализа производительности реальных информационных систем.

Ключевые слова: анализ производительности; критичные по времени приложения; db-nets; протокол FIX; моделирование ПО; тестирование ПО.

Для цитирования: Ригин А.М., Шершаков С.А. Метод анализа производительности критичных по времени приложений с помощью DB-Nets. Труды ИСП РАН, том 33, вып. 3, 2021 г., стр. 109-122 (на английском языке). DOI: 10.15514/ISPRAS-2021-33(3)-9.

Благодарности. Работа выполнена в рамках Программы фундаментальных исследований НИУ ВШЭ.

1. Introduction

Nowadays, most of time-critical business processes are performed using computer technologies. Nuclear reactor control, medical equipment control, spaceship control are some obvious examples of such processes. However, different financial processes including trading on stock exchanges also can demand strict performance requirements.

In the previous century, trading on stock exchanges was primarily performed through phone calls and with use of paper-based order books [1]. Working this way did not allow traders to compete for the best price that is generally offered during very short period. This was the reason of beginning of automatization of trading on stock exchanges. In order to guarantee compatibility of software systems of different traders, brokers, and exchanges, there were financial protocols for electronic communication between trading participants created. Financial Information eXchange (FIX)

Protocol maintained by the FIX Trading Community [2] is one of the most known and widely used protocols of such type. There exist different approaches to encode messages transferring with the FIX protocol. In this paper we focus on the *FIX TagValue Encoding*, which is the main standard of encoding FIX messages [2].

The FIX protocol allows traders, brokers, and exchanges to create and fill (execute) orders for buying or selling securities in several milliseconds using electronic communication channels such as Internet [2]. It is a great driver for competence in the global stock markets, however it creates new challenges for financial software vendors. One of such challenges is maintaining the best possible performance. Any unspecified delay may cause a large financial loss for a trader due to the best price is missed. Such delays may create unequal and unfair conditions for different participants, lead to local or global economic problems as well as public scandals and reputational problems for the exchange or some traders or brokers.

Financial protocol-based communication systems are considered in this work to demonstrate applicability of the proposed method for time-critical systems performance analysis. However, there are plenty of similar systems existing for different domains, and the method can also be applied for a performance analysis of these systems.

Any FIX message consists of a set of tag-value pairs [3]. In fact, it means that we can represent these messages in the form of records of a table in some relational database. Therefore, some methods of system modeling, which rely on relational database states, can be considered here. The same is valid not only for messages of the FIX protocol, but for any messages of transactional systems that are represented as sets of tag-value pairs.

In 2020, we developed a software simulator for the db-net formalism [4] introduced by Montali and Rivkin in 2017. This formalism is represented by the layer with modified colored Petri net modeling a control flow of a process system, and two inner layers for working with an attached relational database modeling a persistent storage [5] as shown in Fig. 1. This simulator is developed as a plugin for Renew (Reference Net Workshop) software tool which is a Java-based reference net simulator [6].

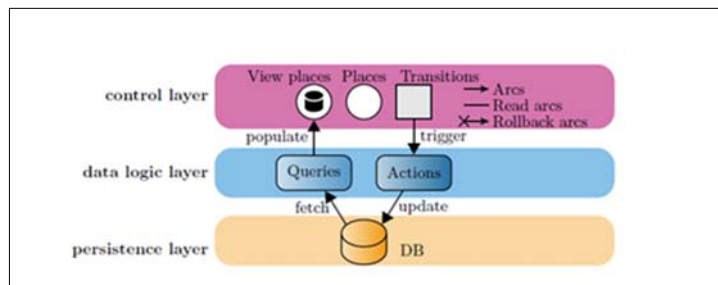


Fig. 1. The db-net structure [5].

Generally, the lowest layer of the db-net (the persistence layer) is represented by an ordinary relational database [5]. However, it can be replaced with any other information storage, which is accessible through a custom relational DML interface that is to be implemented.

One can model a tag-value message sending by using the “insert” database operation, where tags are represented as attributes of a relational table and values are represented as attributes of a record in the table. A tag-value message receiving can be modeled similarly using the “select” database operation.

In the current time, there are some performance analysis research works focused on distributed software systems such as [7, 8], however performance analysis using db-nets has its advantages for transactional systems which send and receive messages that are representable in the form of records in relational tables.

Firstly, it allows to integrate performance analysis into conformance checking of a system. A performance property can be considered together with other checked properties of the system to check all of them simultaneously. Therefore, it allows to abstract away from the performance and to combine performance analysis of transactional systems with other methods for their verification and validation, based on Petri nets and their modifications, especially db-nets (e.g., checking safety, liveness, fairness, and similar properties). Moreover, colored Petri net models, that are automatically generated from event logs using process discovery algorithms, may be extended with db-net elements and time constraints, and used for performance analysis.

Secondly, this method allows to apply well-known approaches used in the relational database domain to the wide set of transactional systems supporting time-critical applications.

All the above provides the motivation for the research.

The purpose of the research is development of a method of performance analysis of time-critical applications using db-nets.

The objectives of the research are as follows.

- 1) Developing a method for performance analysis of time-critical applications using db-nets.
- 2) Developing a software prototype for performance analysis of time-critical application logs using db-net models.
- 3) Checking the method by testing the developed software prototype on a test log of FIX messages provided by a software developer of testing solutions for one of the global stock exchanges.

The rest of the paper is organized as follows. The Section 2 presents the theoretical foundations and concepts of the work and the developed method. In the Section 3, the developed software prototype and its testing are described. After this, the main points of the paper are summarized in the conclusion.

2. Performance Analysis Using DB-Nets

2.1 DB-Nets

The db-net formalism is a modification of the colored Petri net, which allows to model a system control flow together with relational database states. The db-net consists of three layers: (1) the control layer, (2) the data logic layer which connects the control layer and the persistence layer together, and (3) the persistence layer [5]. The scheme of db-net structure is shown in Fig. 1.

The persistence layer allows to store the persistent data and is formally defined by a relational database schema and constraints that declare the data consistency rules [5].

The data logic layer is defined by two sets: (1) set of *queries* for retrieving records from a database in the persistence layer and (2) set of *actions* for insertion and deletion of records in the persistence layer database. Each action includes sets of added and deleted *facts* (records in relational tables) [5]. The control layer allows to model a system control flow and is defined by a colored Petri net with the following modifications [5].

- 1) Queries defined in the data logic layer are assigned to places of a colored Petri net in the control layer. Such places are called view places. View places cannot contain tokens (resources modeled in a Petri net) such as other places, but they produce new tokens by retrieving data from the persistence layer through assigned queries.
- 2) Actions defined in the data logic layer are assigned to transitions of the net in the control layer. When a transition with the assigned action is fired (executed), the action is performed on a database in the persistence layer.
- 3) In addition to traditional Petri net arcs, there exist read arcs and rollback arcs in the db-net control layer. The former is used for connecting view places with transitions and the latter is used for defining a flow for a case of rollback of an action due to violation of the data consistency rules in a database of the persistence layer after performing the action.

The db-net control layer's net and persistence layer's database schema example for the taxi booking software system is shown in Fig. 2.

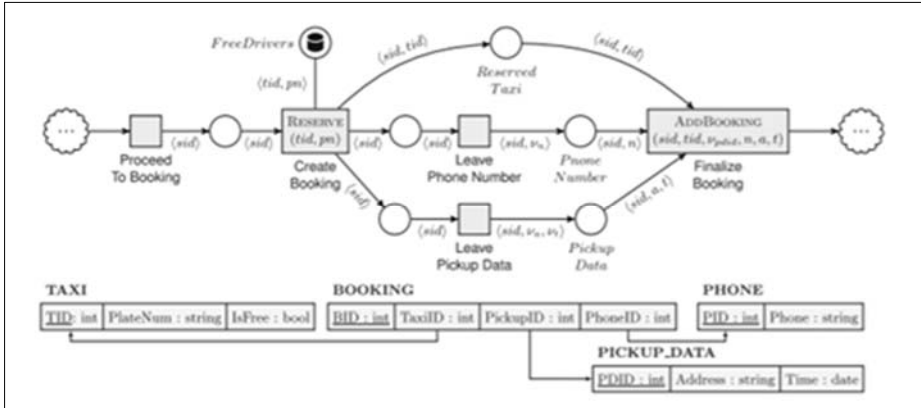


Fig. 2. The control layer's net and persistence layer's database schema example for the taxi booking software system [5].

2.2 Conformance Checking

Conformance checking allows to verify that a considered system satisfies desirable properties through ensuring that an event log produced by the system fits a designed model [9]. These properties include safety, liveness, fairness, and similar ones. For example, safety properties guarantee that the system does not achieve certain undesirable states.

The method proposed in this paper allows to check performance simultaneously with checking other properties of a system. A performance property is considered in this work as a safety property for satisfying that a system should not achieve the state where a delay between two messages exceeds a maximum acceptable one. Therefore, performance can be checked in a process model designed for checking other properties by extending the model with information about time constraints [9].

Since the db-net formalism extends colored Petri net, it is possible to check all properties using a db-net if these properties can be checked using a traditional colored Petri net. The performance in the proposed method is checked using db-net elements.

The set of properties $S = \{s_1, s_2, \dots, s_n\}$, where $s_i, i = \overline{1, n}$ is the i -th checked property, is considered as an example. Some of these properties may be performance properties. The set of performance properties is $P = \{p_1, p_2, \dots, p_m\}$, where $P \subseteq S$. Each performance property $p_j, j = \overline{1, m}$ contains time constraints in the form of a maximum acceptable delay for pairs of messages of particular types as specified in the proposed method (the subsection 2.3). Each property $p_j, j = \overline{1, m}$ such that $p_j \in P$ is checked by the proposed method using db-nets. Other properties $s_i, i = \overline{1, n}$ such that $s_i \in S \setminus P$ are checked by other methods utilizing colored Petri nets and db-nets.

As a result, performance analysis can be conducted as a part of conformance checking of a system, where performance properties are among of all checked properties. This allows to abstract away from the performance properties and check all properties simultaneously.

2.3 Method of Performance Analysis Using DB-Nets

The developed method implies analyzing messages sent or received by the application or its modeled part (*request* and *response messages*, respectively) and stored in a log of the application. We analyze

those messages for which a maximum delay between sent message and received response is restricted. The method utilizes the db-net formalism.

The method consists of two parts: (1) set of requirements for implementing the method in a software tool and (2) sequence of stages and steps for using the method after being implemented.

2.3.1 Implementing the Method in a Software Tool

The following set of requirements specifies how the method of performance analysis using db-nets is to be implemented as a software tool. These requirements extend general principles of the db-net behavior, which are described in [5].

- 1) When a *request message* is inserted in an action assigned to a db-net transition, it should be stored in the memory (RAM or persistent storage) for further retrieving when the corresponding response message is retrieved.
- 2) When a *response message* is retrieved by a "select" query assigned to a db-net view place and the connected by a read arc db-net transition contains parameters for performance analysis as specified in the step 6 of the stage 1 of the method (the Section 2.3.2), the following sequence of steps is to be executed:
 - a) The corresponding request message (with the same *id* attribute value) is to be retrieved through the specified query from the memory/storage (as specified in the item 1 of the current set of requirements).
 - b) If there is no stored corresponding request message, then this sequence is to be stopped and the token with the response message is to be moved to the places connected by output arcs.
 - c) The *sending timestamps* of the request and response messages are to be parsed using a specified pattern or a regular expression.
 - d) A *delay* that is a difference (in milliseconds) between these two *sending timestamps* is to be calculated. If it exceeds the specified *maximum acceptable value of a delay*, then the validation is to be considered as failed – information about the *id* and *message type* of the problematic messages is to be displayed or stored in the report (depending on the requirements and implementation), for the first violation or for each violation (also depending on the requirements and implementation).
- 3) If there are several response messages for one request message, only the first response message is considered.
- 4) If the simulation is finished (no transitions can be *fired* – executed) and the validation did not fail, then such validation is considered as succeeded.

2.3.2 Use of the Method

After implementing the software tool, the method is to be used by following the sequence of steps divided into three stages, as follows.

Stage 1. Modeling a DB-Net. A db-net that matches a system/a modeled part of a system is to be modeled using the following steps.

- 1) A scope of the modeled system is to be defined. It should include considered components of the system which send *request messages* (messages sent by the system or its considered component) and get responses to them (*response messages*), and considered types of *request messages* and corresponding types of *response messages*. From now on, we will call a modeled system/part of the system a *time-critical application* (or just an *application*).
- 2) It is necessary to make sure that the application works as a transactional system and satisfies the ACID (atomicity, consistency, isolation, durability) properties [10], and a log with its request and response messages can be represented in the form of tables in a relational database. It means that each message includes a set of *tags* (attributes) together with their *values*. *Tags* are represented as attributes of a relational table, messages are represented as records of the

table, and *values* are represented as attributes of a record in the table. Types of messages and parts of the application which do not satisfy these properties, if any, are to be removed from the scope.

- 3) A persistence layer of the modeled db-net is to be defined. To do this, a relational database schema is to be created and populated with necessary tables. The table attributes reflect the tags of considered request and response messages.
- 4) A data logic layer of the modeled db-net is to be defined. The «insert» queries, which model insertion of the request messages into the modeled relational database, are to be specified. The «select» queries, which model retrieving the request and response messages from the modeled relational database, should similarly be specified.
- 5) A model of a system control flow (a control layer of the modeled db-net) is to be defined. After that, «insert» and «select» queries from the modeled data logic layer are assigned to transitions and view places, respectively.
- 6) For each db-net transition connected by a read arc with a view place that is assigned with a «select» query for retrieving the response messages, the following parameters for conducting a performance analysis are to be specified:
 - a) The name of a variable in the control layer that stores a value of the *id* attribute of a response message, which allows to find a corresponding request message by the same value of the same *id* attribute.
 - b) The name of a variable in the control layer that stores a value of the *sending timestamp* attribute of a response message.
 - c) An ordering number of the *sending timestamp* attribute of a message in results of a "select" query for retrieving the corresponding request message, that is mentioned in the item "f" of the current list.
 - d) A pattern or a regular expression for parsing the *sending timestamp* string in a message.
 - e) An ordering number of the *message type* attribute of a message in results of a "select" query for retrieving the corresponding request message, that is mentioned in the item "f" of the current list.
 - f) The name of a declared "select" query for retrieving the corresponding request message.
 - g) The *maximum acceptable value of a delay* between *sending timestamps* of corresponding request and response messages (in milliseconds).

Stage 2. Preprocessing the Log. Preparing a log of the application includes the following steps.

- 1) It is necessary to make sure that the messages in a log are represented in a form satisfying properties described in the step 2 of the stage 1. Any messages that are not represented in a valid form as well as broken messages are to be removed.
- 2) The log should be prepared in a format compatible with a software tool implementing the method.

Stage 3. Conducting a Performance Analysis Using DB-Nets. A simulation of the modeled db-net is to be run in the software tool implementing the method.

2.4 Example of Performance Analysis Using DB-Nets for the FIX Protocol

The developed method is illustrated by an example modeling a trading order creation with use of the FIX protocol. The example includes the analysis of two types of FIX messages: (1) *create_order_single* (*msg_type* = "D") which is used for request messages sent from a trader or a broker to the exchange, to create an order for buying or selling securities, and (2) *execution_report* (*msg_type* = "8") which is used for response messages sent from the exchange to the trader or the broker as a confirmation of the order creation (or information about the order rejection with clarification of a reason). For each message, the attributes *msg_type*, *cl_ord_id* and *sending_time* are considered in the model. The *msg_type* attribute defines a type of the message. The

corresponding request/response messages are connected by a key (*id*), whose role is played by the *cl_ord_id* attribute. The *sending_time* attribute is a sending timestamp of the message.

The db-net modeling this example is shown in Fig. 3. A schema of a relational database in the db-net persistence layer includes a *msg* relational table for storing FIX messages. The table contains *msg_type*, *cl_ord_id* and *sending_time* attributes. The *create_order_single* action models the "insert" DML query for insertion of the *msg_type*, *cl_ord_id* and *sending_time* attributes of the *create_order_single* FIX message. The *create_order_single* and *execution_report* queries model the "select" SQL query for retrieving the same attributes of the *create_order_single* and *execution_report* FIX messages, respectively. The *create_order_single_corr_req* query models the "select" SQL query for retrieving the same attributes of the *create_order_single* FIX message by the given *cl_ord_id*. It is used for retrieving the corresponding request message for a previously retrieved response message.

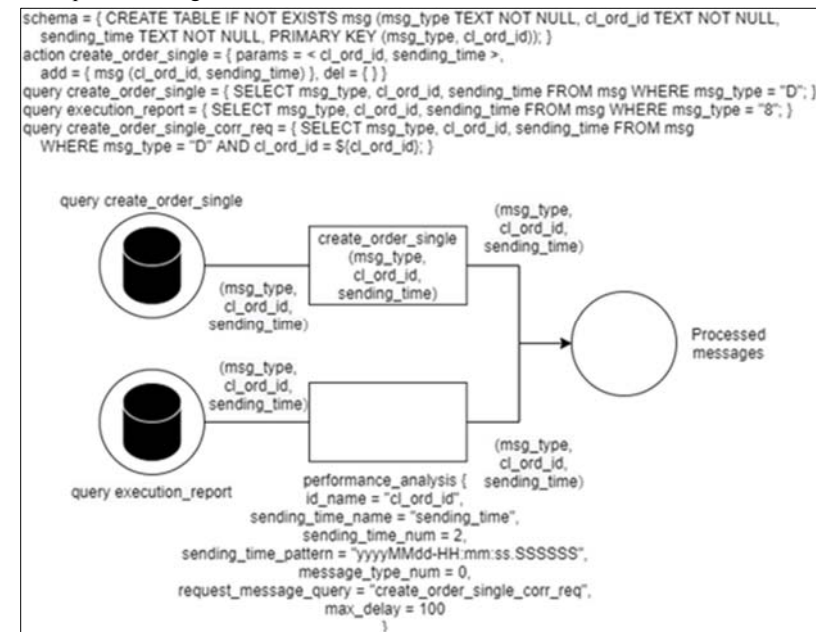


Fig. 3. Example of a db-net model for a performance analysis of a FIX protocol-based system.

The view place assigned with *create_order_single* query (fig. 3) is responsible for retrieving messages of *create_order_single* type. The following transition executes the *create_order_single* action, modeling insertion of the messages into the *msg* table. Then the transition transfers the messages to the *Processed messages* place.

The view place assigned with *execution_report* query (fig. 3) is responsible for retrieving messages of *execution_report* type. After retrieving an *execution_report* message, the following transition retrieves the corresponding *create_order_single* message (with *msg_type* = "D" and the same *cl_ord_id*) using the *create_order_single_corr_req* query and calculates a delay between these two messages as a difference between their sending timestamps (the *sending_time* attribute). If the calculated delay exceeds *max_delay* (it is 100 ms in the example), then the validation fails. Otherwise, the *execution_report* message is transferred to the *Processed messages* place. After all messages are retrieved from the log, and validation did not fail, the validation of the log is considered succeeded.

We consider two following FIX messages (these messages are presented below in the human-readable form, not in the original FIX tag-value form): (1) *create_order_single* (*msg_type* = “D”, *cl_ord_id* = “12345”, *sending_time* = “20190218-02:14:45.490000”) and (2) *execution_report* (*msg_type* = “8”, *cl_ord_id* = “12345”, *sending_time* = “20190218-02:14:45.492787”). Firstly, the *create_order_single* message is retrieved by the view place assigned with the *create_order_single* query. The following transition performs the *create_order_single* action with the “insert” DML query for this message and transfers the message to the *Processed messages* place. Secondly, the *execution_report* message is retrieved by the view place assigned with the *execution_report* query. By the *cl_ord_id* = “12345” attribute value of the message, the following transition retrieves the corresponding *create_order_single* message (with *msg_type* = “D” and the same *cl_ord_id* = “12345”) using the *create_order_single_corr_req* query and calculates a delay between these two messages as a difference between their sending timestamps (the *sending_time* attribute). This delay equals 3 ms (rounding up). A maximum acceptable delay linked with the transition is defined to 100 ms. The delay does not exceed the maximum acceptable delay, so the validation does not fail, and the *execution_report* message is transferred to the *Processed messages* place. However, if the *sending_time* attribute value of the *execution_report* message was, for example, “20190218-02:14:45.592787”, then the delay would be equal to 103 ms (rounding up) and the maximum acceptable delay would be exceeded which would lead the validation to fail.

3. Software Prototype

3.1 Software Prototype Features and Implementation

For testing and illustrating abilities of the method, the latter is implemented in the form of a software prototype. For doing this, we developed the db-net software simulator (Renew DB-Nets Plugin) in 2020 [4] and then extended it with features for conducting a performance analysis of time-critical applications using the proposed method. The simulator has a form of a plugin for Renew software tool which is a Java-based reference net simulator [6]. The simulator has a graphical user interface as shown in the screenshot in fig. 4.

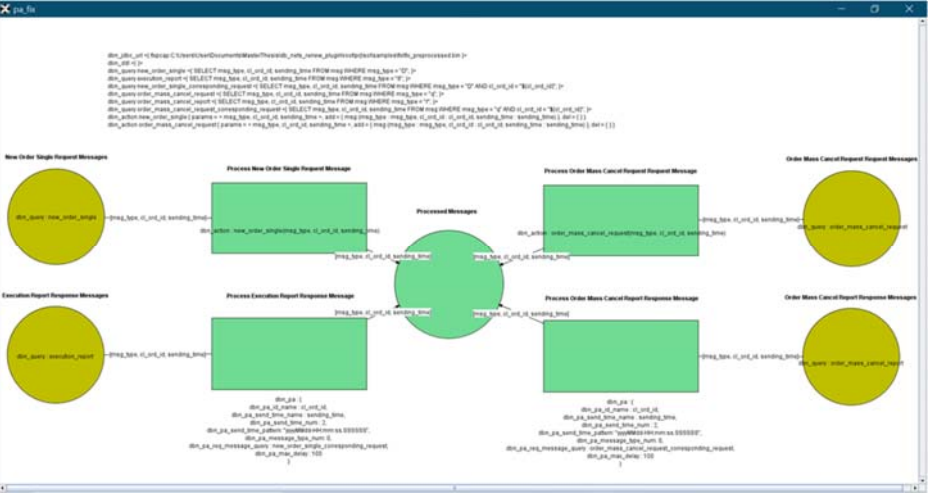


Fig. 4. Screenshot of a graphical user interface of the developed software prototype.

The prototype allows to (1) model a db-net for a considered system, (2) specify parameters for conducting a performance analysis of time-critical applications, as described in the step 6 of the stage 1 of the described method (the Section 2.3.2), (3) conduct a performance analysis of an

application in parallel with a db-net model simulation using the proposed method and (4) work with a FIX log (raw binary data of the FIX protocol packages captured as a Wireshark PCAP file [11] with further filtering) through a relational DML interface.

An implementation of the developed db-net simulator is described in [4]. This implementation is based on an implementation of Renew software tool, a reference Petri net simulator. The Renew code which was suitable for the db-net behavior is reused. Other code is overridden by a custom db-net implementation. Classes representing elements of the db-net control layer are inherited from Renew classes representing similar elements of traditional colored Petri nets and necessary methods are overridden. The prototype is implemented as a pure plugin for Renew tool, without modifying existing Renew source code [4]. The plugin code, UML class diagram and documentation are available in the project GitHub repository¹.

For working with a FIX log through a relational DML interface, the alternative implementation of the database connection interface is created. It is used if the JDBC URL in a db-net model starts from the “fixpcap:” prefix. All messages that are read from file through this connection are stored in RAM (in the *java.util.HashMap* container, where keys, which are pairs of *message type* and *id*, are stored in a hashtable). When the message is being retrieved through this connection, it is firstly searched in RAM. If it is found in RAM, it is returned and removed from RAM. If it is not found in RAM, then the file is scanned until finding this message (and all scanned messages are stored in RAM). This approach allows to scan each line of the file only once and to minimize the RAM usage.

For goals of a performance analysis, the prototype follows the set of requirements described in the subsection 2.3.1. When the first *maximum acceptable delay* violation is detected while simulating a db-net model, the dialog window with an information message describing this violation is shown and the corresponding CSV report is created. All *maximum acceptable delay* violations that are detected during the current simulation are written into the created CSV report. The format of a CSV report is presented in Table. 1.

Table. 1. Columns of the CSV Report

Column Name	Description	Type	Example
#	Order number of the row in the CSV report (starting from 1)	Integer	1
Request Message Type	Type of the request message	String	D ²
Message ID	ID of the request and response message pair	String	15504
Delay	Difference (in milliseconds) between request and response message sending timestamps	Integer	493
Max Delay	Maximum acceptable delay	Integer	100
Diff	Difference between detected delay and maximum acceptable delay	Integer	393

¹ Link: https://github.com/Glost/db_nets_renew_plugin

² In the FIX Protocol, the D message type is used for the *New Order Single* messages.

3.2 Testing the Prototype on the FIX Log and Quantitative Analysis of Maximum Acceptable Delay Violations

The developed software prototype is tested on a log with FIX protocol messages, which is represented by the raw binary data extracted from a Wireshark PCAP file with some FIX protocol messages captured in the testing environment. The file is provided by a software developer of testing solutions for one of the global stock exchanges.

The screenshot in fig. 4 shows the db-net model for performance analysis applied to the FIX protocol messages for the *New Order Single* scenario (request message: *New Order Single*, message type: "D"; response message: *Execution Report*, message type: "8") and the *Order Mass Cancel Request* scenario (request message: *Order Mass Cancel Request*, message type: "q"; response message: *Order Mass Cancel Report*, message type: "r"). The total number of processed messages in this model equals 321671.




Fig. 5. Quantitative analysis of maximum acceptable delay violations for maximum acceptable delay values from 1000 ms to 9000 ms.




Fig. 6. Quantitative analysis of maximum acceptable delay violations for maximum acceptable delay values from 3100 ms to 3900 ms.

Using this model, the quantitative analysis of *maximum acceptable delay* violations is conducted based on the CSV reports with information about violations. The plots in Fig. 5 show (1) counts of *maximum acceptable delay* violations and (2) percentages (ratios) of message pairs with *maximum acceptable delay* violations (where 100 % is all processed message pairs), in the db-net model described above, with a breakdown to the request message types ("D" is used for the *New Order Single* messages and "q" is used for the *Order Mass Cancel Request* messages) for maximum acceptable delay values from 1000 ms to 9000 ms.

The significant decrease in count of violations between *maximum acceptable delay* values 3000 ms and 4000 ms is notable. The plots in Fig. 6 show the same metrics for *maximum acceptable delay* values from 3100 ms to 3900 ms. We can conclude that the most of delays larger than 1 second are between 3 and 4 seconds.

Such quantitative analysis is an example of possible applications of the developed method. For instance, requirements and service level agreements (SLAs) can be specified and adjusted basing on some statistics on ratio of message pairs violating each *maximum acceptable delay*. This information with a breakdown to the request message types allows to focus on improving the speed of the most critical scenarios.

4. Conclusion

In the current work, a novel method of performance analysis of time-critical applications based on the db-net formalism is developed. This method allows to integrate performance analysis into conformance checking of a system. Therefore, it allows to abstract away from performance and to combine performance analysis of transactional systems with other methods for their verification and validation, based on Petri nets and their modifications, especially db-nets (e.g., checking safety, liveness, fairness, and similar properties). Colored Petri net models, that are automatically generated from event logs using process discovery algorithms, may be extended with db-net elements and time constraints, and used for performance analysis. Moreover, the method allows to apply well-known approaches used in the relational database domain to a wide set of transactional systems supporting time-critical applications.

A software prototype implementing the method is developed. The prototype is checked on a test log with FIX messages provided by a software developer of testing solutions for one of the global stock exchanges. A quantitative analysis of maximum acceptable delay violations is conducted based on this log. This demonstrates how the method can be applied for similar analysis.

The developed method can be used in research in this domain as well as in testing performance of real time-critical software systems. Further steps include extending the method for use with hierarchical Petri nets and more complex variants of performance analysis of transactional systems. Approbation of the method for integrating performance analysis into conformance checking of a real software system is planned.

The developed software prototype is to be improved for being more usable. This will make the prototype a new software tool in the pool of open-source solutions for conformance checking and performance analysis.

References

- [1]. Harris L. Back Office Operations. Trading and Exchanges: Market Microstructure for Practitioners Oxford Univiversity Press, 2003, chapter 7, section 7.2.2, pp. 148-149.
- [2]. Introduction, FIX Trading Community, Available at: <https://www.fixtrading.org/online-specification/introduction/>, accessed 28.03.2021.
- [3]. FIX TagValue Encoding, FIX Trading Community, Available at: <https://www.fixtrading.org/standards/tagvalue-online/>, accessed 28.03.2021.
- [4]. Rigin A., Shershakov S. Data and Reference Semantic-Based Simulator of DB-Nets with the Use of Renew Tool. Lecture Notes in Computer Science, vol. 12602, 2021, pp. 453-465, DOI: 10.1007/978-3-030-72610-2_34.

- [5]. Montali M., Rivkin A. DB-Nets: On the Marriage of Colored Petri Nets and Relational Databases. Lecture Notes in Computer Science, vol. 10470, 2017, pp. 91-118.
- [6]. Renew – The Reference Net Workshop. Renew.de, Available at: <http://www.renew.de/>, accessed 28.03.2021.
- [7]. Vetter J. Performance analysis of distributed applications using automatic classification of communication inefficiencies. In Proc. of the 14th international conference on Supercomputing (ICS '00), 2000, pp. 245-254.
- [8]. Marsan M. A., Bianco A. et al. A LOTOS extension for the performance analysis of distributed systems. IEEE/ACM Transactions on Networking, vol. 2, no. 2, 1994, pp. 151-165.
- [9]. van der Aalst W., Adriansyah A., van Dongen B. Replaying history on process models for conformance checking and performance analysis. WIREs Data Mining and Knowledge Discovery, vol. 2, no. 2, 2012, pp. 182-192.
- [10]. Haerder T., Reuter A. Principles of transaction-oriented database recovery. ACM Computing Surveys, vol. 15, no. 4, 1983, pp. 287-317.
- [11]. 5.2. Open Capture Files, Wireshark.org, Available at: https://www.wireshark.org/docs/wsug_html_chunked/ChIOOpenSection.html, accessed 28.03.2021.

Информация об авторах / Information about authors

Антон Михайлович РИГИН получил степень магистра в области системной и программной инженерии в 2021 г. в Национальном исследовательском университете «Высшая школа экономики» (Москва, Россия). Его исследовательские интересы включают программную инженерию, извлечение и анализ процессов (process mining), верификацию программного обеспечения, алгоритмы и структуры данных и их применение в задачах индексирования и хранения данных в системах управления базами данных.

Anton Mikhailovich RIGIN received his master's degree in System and Software Engineering from the National Research University Higher School of Economics (Moscow, Russia) in 2021. His research interests include software engineering, process mining, software verification, algorithms and data structures and their usage in problems of data indexing and storage in database management systems.

Сергей Андреевич ШЕРШАКОВ получил степень кандидата компьютерных наук Национального исследовательского университета «Высшая школа экономики» (Москва, Россия) в 2020 году. В настоящий момент он является доцентом департамента больших данных и информационного поиска и научным сотрудником научно-учебной лаборатории процессно-ориентированных информационных систем (Лаборатории ПОИС) факультета компьютерных наук Высшей школы экономики. В число научных интересов входят извлечение и анализ процессов (process mining), верификация программного обеспечения, архитектуры информационных систем и преподавание программной инженерии.

Sergey Andreevich SHERSHAKOV received his PhD degree in Computer Science from the National Research University Higher School of Economics (Moscow, Russia) in 2020. He is currently an Associate Professor at the Big Data and Information Retrieval School and a research fellow at the Laboratory of Process-Aware Information Systems (PAIS Lab) of the Faculty of Computer Science at the HSE University. His research interests include process mining, software verification, information system architectures and teaching software engineering.