

UNIVERSITY OF GRONINGEN

FACULTY OF SCIENCE AND ENGINEERING
DEPARTMENT OF COMPUTING SCIENCE

**Design and Implementation of a
Process-Aware System
for Flexible Batch Processing**

Author

Dimitrios LASKARATOS

Supervisor

pr. dr. Dimka KARASTOYANOVA



university of
 groningen

faculty of science
and engineering

Contents

Abbreviations	4
Abstract	5
1 Introduction	6
I Models, Process Engines and Batch Processing	6
II Related Work	8
III Present Work	11
2 Architecture and Implementation	12
I Platform Selection	12
I.1 Camunda BPM System Architecture	12
I.2 Camunda Basic Workflow	13
II Extension Architecture	14
II.1 Overview	14
II.2 Engine	15
II.3 Batch Controller	17
II.4 Web Applications	18
III Implementation	19
III.1 Process Engine Modification	19
III.2 Batch Controller	31
III.3 Batch Management Tool	44
3 Results	46
4 Discussion	54
5 Conclusion	57
I Future Work	57
Acknowledgements	58
References	59

List of Figures

1	BPMN process model	6
2	BPEL format	7
3	Batch region in Camunda Modeller modeller batch extension by Pufahl et al. [1]	8
4	Illustration of the modelling strategy [2]	9
5	Illustration of the deployment strategy [2]	10
6	Illustration of the execution strategy [2]	10
7	High level architecture of batch processing extension [2]	11
8	Camunda components [3]	13
12	(1): Process model - (2.1): Start event - (2.2): Start variables - (3.1): Activity - (3.2): Script binding - (4.1): Generated .war file - (4.2): Process deployed - (5): Process in Cockpit - (6.1): Deployed process in Tasklist - (6.2): Process instantiation - (7): Process instance	14
13	Extension Architecture	15
14	Engine Adapter function within the architecture	16
15	Batch Controller within the architecture	17
16	Batch Management Tool	18
17	Blocked Instances tab	19
18	Routes file	33
19	Sample process #1 for demonstration	46
20	Sample process #2 for demonstration	46
21	Batch Management Tool	48
22	Starting an instance	48
23	Result of test instance in the Batch Controller	48
24	Result of test instance in Apache Tomcat log	49
25	Creating model for process "Stock Evaluation"	49
26	Two instances blocked	49
27	Blocked instances at the Instances tab	49
28	Batch Controller log showing both instances are resumed	50
29	Apache Tomcat log showing both instances are resumed	50
30	All three instances are blocked	50
31	Two instances now remain in the batch	51
32	Resuming one instance the batch is canceled after countdown	51
33	Result of execution on Tomcat	51
34	The batch is executed	52
35	Result of executed batch on Tomcat	52
36	User task input	52
37	Batch model for UserTaskBatch process	53
38	Instance is blocked	53
39	Batch model for case 5	53
40	Batch executed immediately after after reducing the maximum number of instances	53

List of Tables

1	Process Engine modification functionality	20
2	Batch Controller functionality	32
3	Requirements	55

Listings

1	Method declaration with REST path	20
2	Method definition	21
3	Initialization script	22
4	Batch Activity Retrieval Step Implementation	23
5	Setting batch flag to activity	24
6	Batch flag set during runtime	25
7	Blocking mechanism	26
8	BatchEntity definition	27
9	MessageProducer definition	29
10	REST method handling the resume command	30
11	Engine Adapter method to resume instance	30
12	Startup script	32
13	Create model	33
14	Update model	34
15	Delete model	35
16	Block Message handling	36
17	Assigning an instance to a cluster - BatchClusterManager class	37
18	Assigning an instance to a cluster - BatchCluster class	38
19	Timer - BatchClusterManager class	39
20	TimerTask - BatchCluster class	39
21	Loading activity to execute and thread spawning - BatchExecutionMan- ager class	40
22	Thread Execution - InstanceExecutor class	40
23	Batch finalization - BatchCluster class	41
24	REST method	42
25	REST method	42
26	Resume single instance method	43
27	Cannot remove instance from running batch	43
28	Delete model request	44
29	Resume process instance request	45
30	Caclulate Value implementation	46
31	Change Owner implementation	47

Abbreviations

Abbreviation	Meaning
BPM	Business Process Management
BPMN	Business Process Model and Notation
BPMS	Business Process Management System

Abstract

Business Process Management (BPM) is important to the day-to-day operations of large organizations that constantly seek to optimize their workflow. BPM systems (BPMS) are numerous and offer many diverse services to these organizations, based on their needs. However, these systems all have one thing in common: they do not perform batch processing. Even though there is an abundance of literature approaching this problem theoretically along with several implementations, they are hindered by the fact that they are static and limited in their flexibility. This project attempts an implementation of a flexible batch processing that is independent of both the modelling and the execution phase of BPM. This implementation is a system on its own, involving multiple components that interact with each other, as well as the engine. The results after modifying the target BPMS are promising, showing that the flexibility aspect discussed is certainly possible. Concurrently, it offers a view on extensive future work on several other aspects, like further extending the batch components in order to generalize, i.e., to apply to all types of activities in a process.

1 Introduction

I Models, Process Engines and Batch Processing

Business Process Management is an integral part of large organizations, whose operations often require multiple actions from several different departments to carry out a complex task. This complexity is not always visible from a high-level view of the task, which hides the intertwined activities needed for its completion. In time, along with the growth of the organization, the task may evolve and become even more complex. This may lead to several sub-activities to need to change or even become obsolete, but the scale of the task itself might make it hard to detect minor details. Eventually, the time and cost of the task will increase for the organization. The goal of BPM is to not only document the processes that are crucial for the completion of the task and analyze their behaviour, but also discover processes that are running in the background, and require additional resources, potentially preventing the task from completing faster [4]. The main artifact of BPM is the conceptual model, depicting the life cycle of the target process from the moment it is initiated, to the moment it is completed. This model is constructed using the Business Process Model and Notation (BPMN) format. BPMN is based on graphical flowcharts and consists of five core elements, namely the "start" and "end" symbols, the "gateway" symbol, the "activity" symbol and the "sequence flow" symbol. A brief BPMN example, greatly resembling UML, is visible in Figure 1, depicting a simple ordering and shipping of a product. The "start" and "end" events are the moment the order is placed and the moments the order is either rejected or accepted, respectively. The rectangles represent the activities that take place during the process, which might as well be processes themselves, and the arrows dictate the sequence of events in the process. Lastly, the diamond shape is the "gateway" that can be either a decision making point or a point where two separate activities must take place in parallel. In both cases, there are multiple execution paths possible. As mentioned, the complexity of the model depends on the complexity of the process itself, but can also depend on the level of detail that the modeler requires. As a result, a process model may scale up to hundreds of internal activities, sub-processes and possible execution paths. Additionally, other entities (processes, external services) that communicate with the process might also have to be included, adding to the total complexity of the model. This can be compared to a program with multiple execution paths, that can spawn procedures, each with their own execution paths, eventually resulting in several different outputs.




Figure 1: BPMN process model

```

<receive name="receiveItinerary"
  partnerLink="Customer"
  portType="tag:TravelService"
  operation="orderTrip"
  variable="Order">
  <source linkName="Itinerary-to-Hotel"
    transitionCondition="overnight='true'"/>
  <source linkName="Itinerary-to-Flight"/>
</receive>

```

Figure 2: BPEL format

The process model can only offer static analysis of the processes and resources involved, which is why the next step is to submit this model to a business process engine for execution. The modelling tool, which in most cases is engine specific, generates an XML-like file that can be parsed by the engine to create executable code out of its semantics. It is worth mentioning that in some cases, engines generate a similar representation of the elements called Business Process Execution Language (BPEL) which is visible in Figure 2. BPEL is used to dictate the interaction of specific elements (eg. an activity) of the process workflow with external web services.

Batch Processing

An important part of BPM is the notion of batch processing [5] which is very common in day-to-day operations of organizations. Batch processing stems from the idea that several instances of a process may exist at the same time, which, by definition, require the same resources for their operation. As a result, these instances will eventually have to be executed at the same activity, provided a normal execution (i.e. no faults during execution that can result in different activities) takes place. In such a case, the multiplicity of the process will result in the commitment of a large number of these resources, as well as increase the time needed to conclude the execution of all the processes [6]. That is, if the execution is conducted sequentially. By using batch processing, the process instances are aggregated in groups called "batches" right before a certain activity starts, in order to execute them as a group. This reduces the cost and processing time compared to when these processes are handled individually [7]. Consider the example of blood testing in a medical center. The blood samples taken from the patients must be carried to the laboratory for testing. It doesn't make sense for the employee to take each new sample to the lab individually because that would be a waste of time. Instead, the employee waits until a number of samples have been accumulated before submitting them to the lab. The same idea applies to BPM where the main process execution halts until the required conditions have been met to begin group processing [2].

From a modelling perspective it is sufficient to only denote the batch activity in the process model in BPMN as shown in Figure 3 [8]. However, the BPM engine that parses the process models in preparation for execution, have to also be extended accordingly to parse the batch region notation. Additionally, the activation of a batch cluster, meaning, the actual execution of the batch needs to meet the conditions mentioned above. [7] and [2] have proposed as a threshold either a time constraint, eg. one hour before the batch is executed anyway, or the number of instances needed to be aggregated before execution, eg. 10 or 20. Further information on that work can be found in the Related Work section.

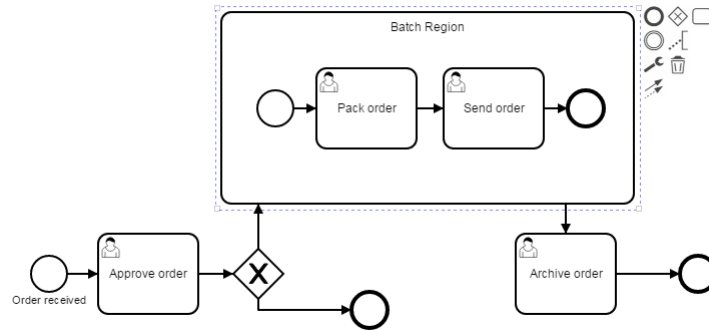


Figure 3: Batch region in Camunda Modeller modeller batch extension by Pufahl et al. [1]

II Related Work

Existing BPM engines do not actually implement batch processing and are limited to a static and sequential execution and resource handling. While the idea of batch processing has been around for quite long [6], only in recent years several attempts have been made to introduce batch processing in BPM engines, either on a theoretical level or on a practical one.

Weske et al. [2018] [6] describes in depth the prototypical implementation of a batch extension by [1], from both a theoretical and practical point of view. Several requirements that need to be satisfied in order for the extension to be successful are identified:

1. **Batch processing over several activities:** Batching must be able to be applied to a set of connected activities instead of just one.
2. **Grouping:** Cases must be able to be grouped in batches based on a common characteristic.
3. **Resource capacity:** A restriction must be put on the number of process instances that can be aggregated before batched execution.
4. **Batch assignment:** Process instances whose number has exceeded the maximum allowed for batch execution must be assigned to different batches.
5. **Activation mechanism:** A mechanism is needed to define the optimal activation time for a batch job to start.
6. **Execution strategy:** A decision has to be made about whether the manner of execution of the batch is parallel or sequential.

Karastoyanova et al. [2018] [2] serves as the basis for this work, as it provides the theoretical background of the batch processing extension. This work discusses three strategies that can support batch processing on an engine and then proposes three flexibility needs that the extension should implement. These three strategies require certain artifacts in order to work. These artifacts are the process model, the batch model and the connector, the latter in some cases. The batch model is a file that specifies the conditions on which a batch

cluster is activated and the connector defines the relations between a batch model and a process activity. Lastly, a system architecture to further support the previous is provided.

The three strategies are the *Modelling strategy*, the *Deployment strategy* and the *Execution strategy*. The Modelling strategy states that the batch region of the process is defined during the modelling phase on the process model, as shown on Figure 4. The model is then deployed on the engine for execution and once control reaches the batch activity, it is executed based on the predefined conditions of the batch model. This strategy is the simplest of the three since the requirements it imposes on the system are minimal compared to the other two. On the other hand, any change in the batch configuration requires the process to be modelled again and redeployed.

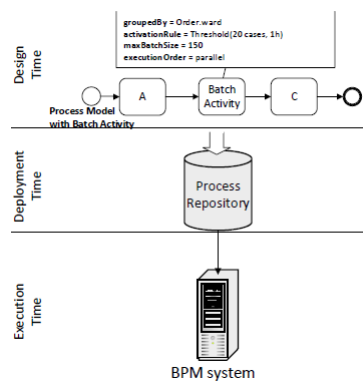


Figure 4: Illustration of the modelling strategy [2]

In the Deployment strategy the batch region is no longer defined in the process model file, but is instead specified in the connector artifact, completely separate from the model file. The process model, the connector and the batch model all have to be stored in the BPM repository as a group as shown in Figure 5. The requirements on the system are more extensive compared to the Modelling strategy, but using the Deployment strategy, the modularity, reusability and flexibility of the artifacts are improved. For example, the modeller needs not be concerned with specifying batch regions or any of the batch configuration since these are stored in separate files. Per contra, changes to the batch model require it to be modified and redeployed which eventually decreases the flexibility of the system and has to be avoided.

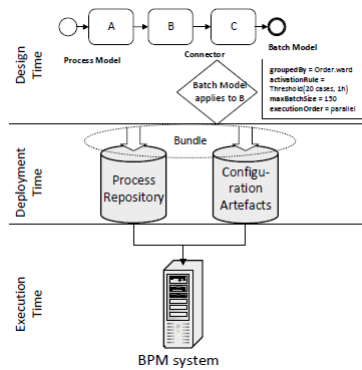


Figure 5: Illustration of the deployment strategy [2]

In the Execution strategy, as with the Deployment strategy, the process model, the batch model and the connector have to be generated by the modeller, but are not deployed on the BPM system as a bundle during deployment. Instead, each artifact can be submitted separate from the others and/or at different times. For example, a batch model can be assigned to a process after the process has been deployed on the BPM, or even during its execution. An illustration can be found in Figure 6. Since a batch model can be assigned and connected to multiple activities across several process instances in real time, it negates the previous disadvantage of the Deployment strategy. Additionally, it is possible to remove said connections during execution. The system requirements for this strategy are again more extensive, but the three flexibility needs are supported.

Note that the connector artifact presented above might seem as an entity that needs to be handled as the batch model. In reality, it is an abstract concept that is merely as a link between a process, an activity and a batch model. It's role will become more clear in Section 2.

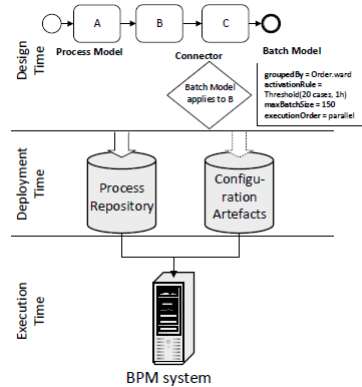


Figure 6: Illustration of the execution strategy [2]

The flexibility needs for batch processing as identified in the work of [2] are as follows:

- **Evolution:** Either the process model and/or the batch configuration must be able to change (eg. due to business changes) without the need of redeployment of all artifacts.

- **Looseness:** Batch configuration must be decoupled from the modelling phase to allow certain aspects of batch work to remain undefined if need be.
- **Adaptation of Batch Activity:** Batch configuration elements must be adaptive to react to exceptions during process execution or other reasons.

In the context of [2], a system architecture to support the batch extension is provided in Figure 7. This architecture will be used as a basis of this work. Note that this architecture contains a high level view of the components involved and the role of each will be thoroughly discussed in Section 2

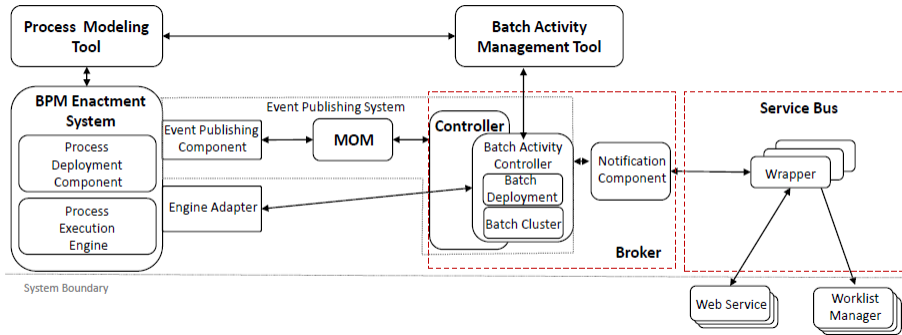


Figure 7: High level architecture of batch processing extension [2]

In **Pufahl [2018] [1]** there exists a prototypical concrete implementation of a batch extension that largely follows the Modelling strategy presented in [2] and briefly discussed previously. In this work the batch properties have to be specified in the process model and deployed together on the Camunda Engine. This implementation applies only to activities of type User Task, where the process instances are aggregated based on the user assigned to the task.

III Present Work

The purpose of this thesis is to extend a BPM engine so that it can support flexible batch processing based on the work of [1] and [2]. More specifically, the goal is to combine the theoretical background of [2] and, possibly, elements of the Java implementation of batch processing of [1] to enable the BPM engine to support real-time manipulation of the batch artifacts, by taking into account the flexibility needs, strategy and architecture that have been discussed previously. Because of time constraints this implementation will only take into account activities that execute an automated script. In the context of Camunda this is equivalent to invoking a web service.

In Section 2 we will deal with the platform selection criteria, followed by a brief explanation of the workflow components, and lastly, the architecture and the implementation details of the extension. In Section 3 we will present the results based on case studies conducted with the new implementation and subsequently discuss them in Section 4. Finally, in Section 5 we will discuss possible future research opportunities based on the results

obtained and the current landscape of BPM.

2 Architecture and Implementation

I Platform Selection

There are several Workflow Management Systems (WfMS) that are currently supporting the latest version of BPMN (version 2.0) [9]. Of the ones whose source code is publicly available, two are selected with which the writer has had substantial experience. These are Camunda BPM Community version and Apache ODE. Apache ODE proved difficult to use because of the limited support and documentation that was available online and, additionally, needed Eclipse as a platform which is widely considered outdated and unreliable. Considering the fact that Camunda was the basis with which Dr. Pufahl [1] developed the batch extension and that a large and active community of users is maintained online that can provide support, it was decided that Camunda will also be used for the flexible batch processing implementation.

I.1 Camunda BPM System Architecture

Since Camunda will serve as a host platform for the extension, we will describe the high-level architecture of the system, component-wise, before moving on to explain the extension components.

The first part of the process life cycle is the modelling tool, which, in most cases produces an XML file with the .bpmn extension. Other possible extensions include CMMN (Case Management Model and Notation) and DMN (Decision Model and Notation). Camunda comes with its own modelling tool called Camunda Modeller that offers additional features, such as specifying user or automated tasks, that are specific to the engine and the applications of the system. In Figure 8 this applies to the "Modeller".

A distinct characteristic of the Camunda Modeller is that it enables the user to specify the behaviour of the activities during the modelling phase as well as any variables that are manipulated during the process workflow. A notable example, which will be described in detail later, involves the creation of an activity that requires the execution of Java code. This activity is called a "Script Task". In order for Camunda to execute the Java code specified by this activity, both the process model and the Java class have to be built together and packaged in a .war file. This procedure is called Process Definition Deployment.

As stated in Section 1, this .war file is then uploaded to a File Repository which is accessed by the Engine itself, which parses the .bpmn file, translating the BPMN artifacts and other semantics into Java objects, internal to the engine. The engine is a Java application that uses the REST API, on top of an Apache Tomcat server to communicate with external applications. The engine handles all aspects of the process life-cycle, including starting and stopping a process, and provides an API for such manipulations by the users. The engine uses the H2 database to store process entities by default, but can also be configured to use

additional ones like DB2, MySQL, and others.

Finally, user, technical and administrative tasks are accessible from a range of Angular.js web pages, provided through the Tomcat server. These applications offer a variety of tools like a list of tasks for the end user, details regarding the current state of process instances, the ability to start multiple process instances, adding or deleting users and many others. These applications, along with any other external Java application (eg. web service) can invoke engine routines through the REST API.

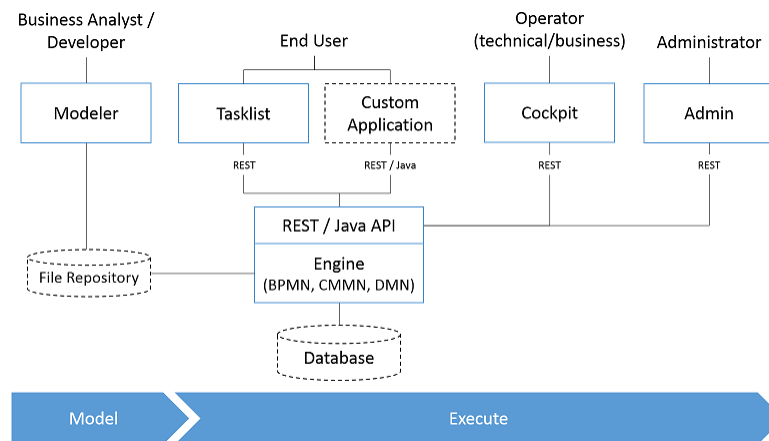


Figure 8: Camunda components [3]

I.2 Camunda Basic Workflow

In Figure 12 presented is the basic workflow for the creation, deployment and execution of a process in Camunda Engine. The first step is to create the process in the modelling tool as shown in (1), using the BPM Notation discussed previously. Optionally, there is the possibility to specify the starting variables by selecting the start event (2.1) and subsequently add variables by specifying their name and type (2.2). Lastly, when the process contains script activities as in the Figure, it is mandatory to bind a script to it. In this case a Java class is bound to activity "Calculate Value", by selecting the activity (3.1) and then providing the full path of the Java class (3.2).

By building the process model classes bound to the activities in an IDE using the Camunda maven archetypes, a .war file is generated as shown in (4.1). When this file is copied to the "webapps" folder of the running Apache Tomcat server, the engine deploys the application as shown in (4.2).

By opening the Cockpit web application of Camunda, the deployed process definition is now visible (5). In order to start an instance of this process, the user must then go to the Tasklist application and select the process from the list as shown in (6.1). If the process specified starting variables, the values have to be provided in the form in (6.2). By pressing the "Start" button, the engine starts an instance of the process.

Going back to Cockpit and selecting the process from the list, the started instance is now visible, along with its current state and variables (7).

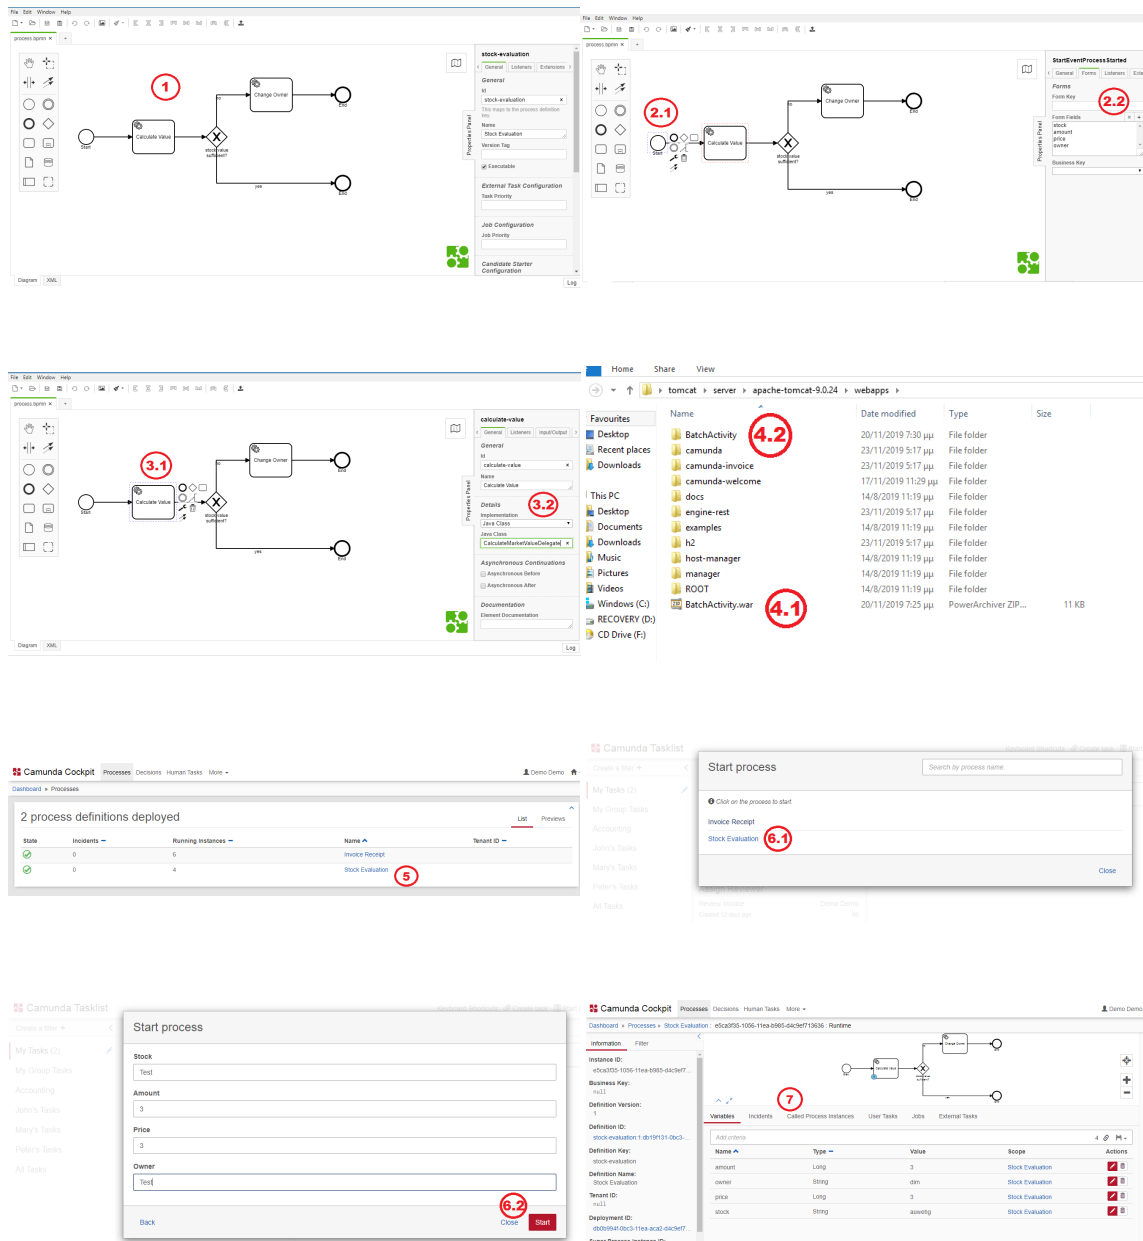


Figure 12: (1): Process model - (2.1): Start event - (2.2): Start variables - (3.1): Activity - (3.2): Script binding - (4.1): Generated .war file - (4.2): Process deployed - (5): Process in Cockpit - (6.1): Deployed process in Tasklist - (6.2): Process instantiation - (7): Process instance

II Extension Architecture

In this section the architectural components will be discussed in greater detail.

II.1 Overview

As stated in Section II, the high-level architecture presented in Figure 7 will be the backbone of the component architecture described in detail in this section. Since the architecture presented in Figure 7 does not refer to a specific BPMN engine, certain adjustments

are necessary in order for this work’s architecture to apply to Camunda. The proposed and implemented architecture is visible in Figure 13. It consists of three major components, namely the Camunda engine itself, the Batch Controller and the User Interface. It is complimented by a message queue that is implemented by a RabbitMQ node, and a MySQL instance database for storing the batch properties and conditions. The Web Service at the right of the figure is not part of the architecture, but is merely there to infer that the Batch Controller may potentially interact with third-party services.

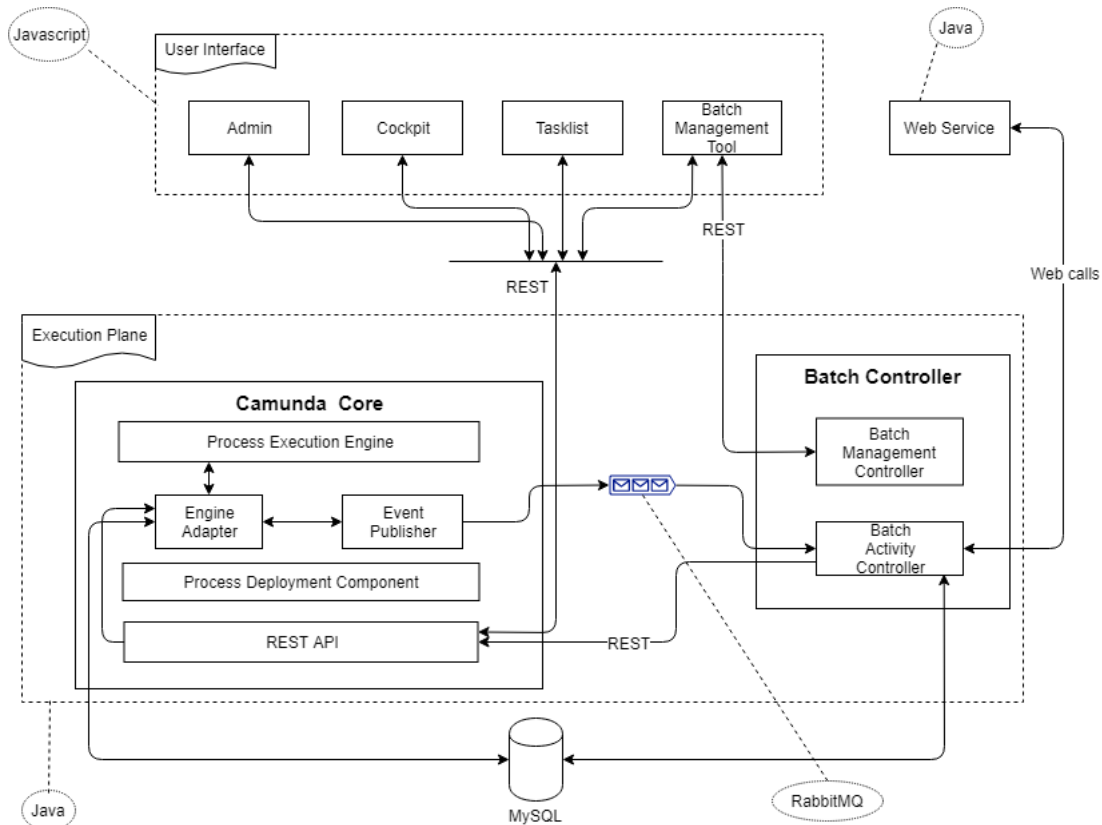


Figure 13: Extension Architecture

II.2 Engine

The software that Camunda is made up of, consists of multiple packages, the most important of which are the Process Engine and REST API. In Figure 13 it may seem that Camunda Core contains many different modules that interact with each other. In reality, only the Process Execution Engine and REST API are separate modules, while the remaining are included in the Process Execution Engine part. This separation in the figure is only done for illustration purposes, as there is no physical distinction between these entities.

The **Process Execution Engine** is responsible for virtually every operation that is needed in Camunda. It parses newly deployed processes, starts, runs or suspends processes, tracks the state of every running process and manages their entire life-cycle. It is a complex and large piece of Java code that conducts many internal calls and interacts with databases and the underlying server (Apache Tomcat, JBoss, Wildfly). Since Camunda does not implement batch processing or anything closely related, and it is also not possible to use its

REST API to halt process instances, modification of the engine code is necessary.

The **Engine Adapter** is a sub-module of the Process Engine that contains the functionality to handle process instances before and after batch execution. The Engine Adapter directly interacts with Process Engine elements, the REST API and the Event Publisher. A brief explanation of the Engine Adapter's function is presented in Figure 14. A more detailed explanation is available in Section III.

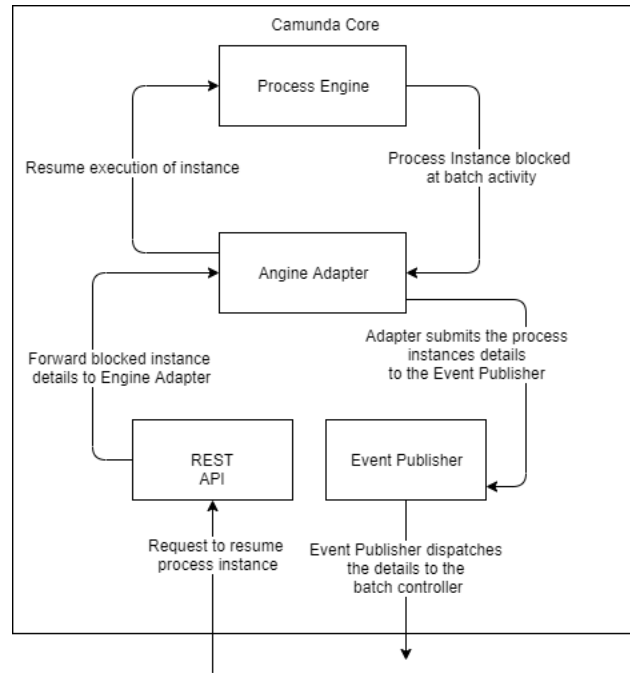


Figure 14: Engine Adapter function within the architecture

Figure 14 shows how the Engine Adapter interacts within the scope of Camunda Core. Once a process instance reaches an activity that has been flagged for batching, its execution is halted. Its current state and the activity are stored in an Engine Adapter object internally, and the Adapter then proceeds to send the details to the Event Publisher. In turn, the Event Publisher sends these details to the message queue. The instance remains in this halted state until an external signal is sent to the REST API which then calls the Adapter to resume its execution. Once the instance starts executing again, the Process Engine is solely responsible for its life-cycle.

The **Event Publisher** is a very simple message queue client, whose sole purpose is to publish messages to the message queue server, a third party software running locally.

The **REST API** is, as mentioned, a separate module that contains all the classes that handle HTTP communication with external components. It routes requests from Camunda's web page applications, as well as the Batch Controller, as can be seen in Figure 13. After the requests are routed to the right method, Process Engine functions are then called to process them. As a result, external applications do not have explicit access to Process Engine functions.

II.3 Batch Controller

The Batch Controller is similar to Camunda Core, in that, there is no actual distinction between the components depicted in Figure 13, but is rather a piece of code with multiple functionalities. Additionally, Batch Management Controller and Batch Activity Controller correspond to REST API and Process Execution Engine respectively, since their role in the component is quite similar to their counterparts. The Batch Controller is a play! application written in Java. It is essentially a web server that uses an underlying REST system for HTTP communication.

A more accurate view of how the various modules of the Batch Controller interact with their environment is available in Figure 15. In Section III, each module's role and function will be discussed in more detail.

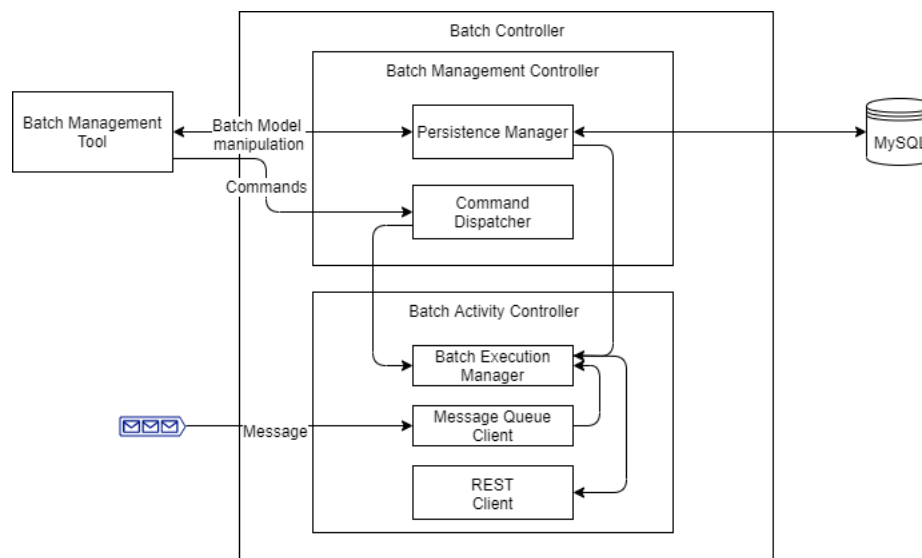


Figure 15: Batch Controller within the architecture

As with Camunda Core, the parts of Batch Controller in Figure 15 are not separate modules, but elements of the same piece of code. In that regard, the **Batch Management Controller** corresponds to the REST API of Camunda Core, as it contains the functionality to handle HTTP requests from the Batch Management Tool. The **Persistence Manager** manages the storage of the request data to the MySQL database, as well as their retrieval, when needed. The **Command Dispatcher** is simply a loose term to describe methods that unpackage the data of the HTTP request from the web page and call functions from the Batch Activity Controller to handle them.

The **Batch Activity Controller** is more complex than Batch Management Controller, operation-wise. It manages the life-cycle of a batch, from the moment the first process instance is halted, to the moment the batch concludes its execution and the instances resume their execution again. It uses the **Message Queue Client** to receive updates, halted instances or exceptions related to these from Camunda Core. Lastly, the **REST Client** makes HTTP requests to the REST API of Camunda in order to resume process instance executions.

II.4 Web Applications

The web applications of the system, namely Cockpit, Tasklist, Admin and Batch Management are the user interface of the system, written in HTML CSS and JavaScript. They provide the means to manipulate process instances, process definitions, users, batch models, monitoring tools etc. Every communication from and to the applications is done via HTTP requests to Camunda (through the web server) or the Batch Controller.

The **Batch Management Tool** is a very simple interface that supports basic operations on a batch model: creation, update and deletion. In Figure 16, at the top of the page, there is a list presenting all available batch models (1). Additionally, the fields below (2) correspond to the batch properties as suggested by [2]. From first to last, these are: the name of the model, the maximum number of instances in a batch, the minimum number of instances before a batch is executed, the time (in hours, minutes, or seconds) that has to pass before the batch is executed, the manner of execution (parallel or sequential), the process name, the activity name and the grouping attribute. Additionally, there is a second tab which lists all currently blocked process instances with information on the process they belongs and which activity they are blocked on. The instances tab in Figure 17.

The screenshot shows the 'Batch Management' web application. At the top, there's a header with the title 'Batch Management' and links for 'About' and 'Examples'. A dark red sidebar on the left contains two buttons: 'Models' and 'Instances'. The main content area is titled 'Models' and features a large, empty light green box labeled with a circled '1'. Below this box is a form for creating or modifying a batch model. The form includes several input fields: 'Model Name' (with a circled '2'), 'Max Batch Size', 'Activation Rules', 'Execution Order', 'Process', 'Activity', and 'Grouping Attributes'. Each field has a red 'x' icon next to it. There are 'Clear' and 'Delete' buttons on the right side of the form, and a 'Create or Modify' button at the bottom right.

Figure 16: Batch Management Tool

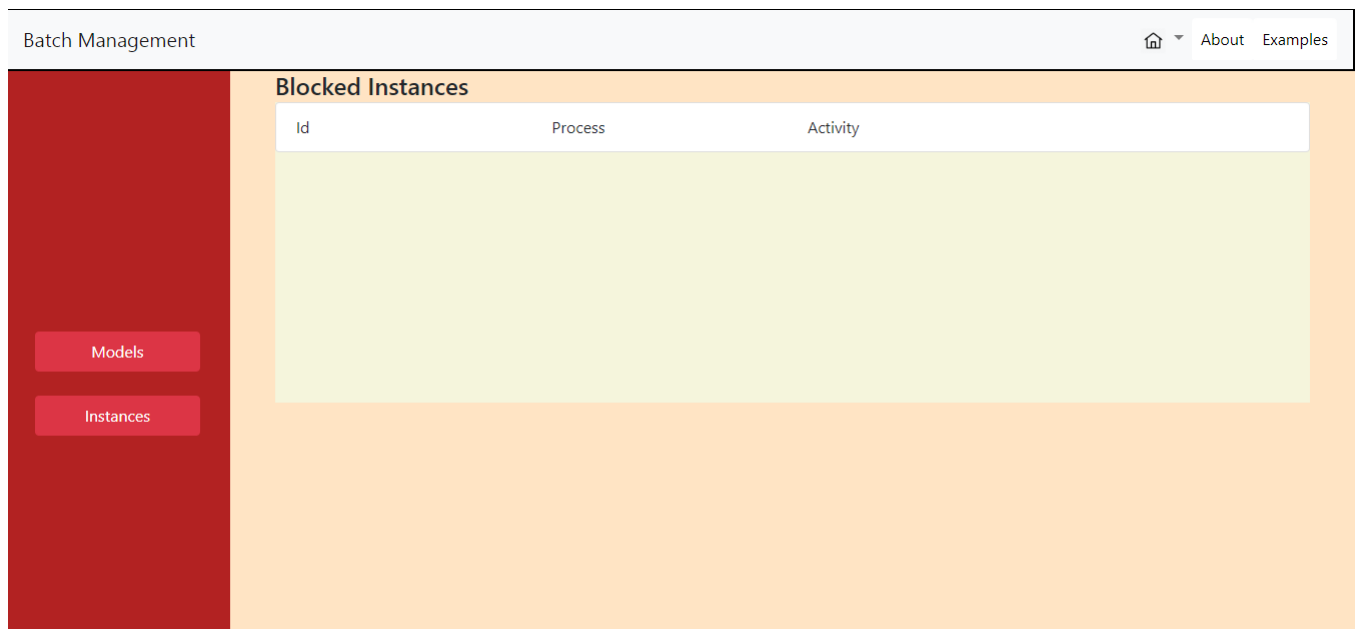


Figure 17: Blocked Instances tab

III Implementation

This section will discuss how the batch extension of the system is implemented, with regard to the three major components presented in Section II. As already mentioned in the Introduction, this system will attempt an implementation of the most flexible batch strategy described in [2], that is, the Execution Strategy. According to this strategy and the architecture discussed in the preceding sections, the implementation will involve the following steps: the process engine modification, the batch controller development and the Batch Management interface development. Additionally, the strategy and architecture impose certain requirements to the system, that the implementation will have to address. These requirements can be grouped into two categories: requirements stemming from the strategy, and requirements stemming from the batch processing theory itself. As such the identified requirements are the following:

1. Fulfillment of the three flexibility needs as presented in Section II (strategy).
2. Interruption of the process may not alter its normal execution. This effectively means that after resuming, the process must continue its execution as it would normally do without the interruption (theory).
3. The batching mechanism must be able to aggregate processes, execute them, and submit them to the engine with the new data (theory).
4. Processes must be able to continue their execution from the interrupted activity, if the batch is canceled (theory).

III.1 Process Engine Modification

Arguably, the most challenging part of this work, was the Process Engine Modification, which required extensive source code investigation, in order to identify places that code

RELATED REQUIREMENTS	1, 2, 4
----------------------	---------

can be injected or existing code must be modified. According to the theory and requirements presented earlier, the modification of the engine must implement the following functionality of Table 1.

#	Functionality	Requirement	Notes
1	The engine must load information regarding existing batch models on startup	1	Flexibility Need 1 & 2 (II)
2	The engine must flag batch activities during runtime	1	Flexibility Need 1
3	A blocking mechanism is necessary to halt the execution of a process when a batch activity is reached.	2, 4	
4	The state of a blocked process must be stored internally.	2, 4	
5	The engine must be able to send blocked processes and updates to the Batch Controller.	2, 4	
6	Process instances that were blocked, must be able to continue their execution from the batch activity, or the next activity in the process.	2, 4	
7	The engine must provide the necessary information for the specification of a batch model.	1	Flexibility Need 2 & 3

Table 1: Process Engine modification functionality

Functionality #7 was the first to be implemented, as it was decided that batch models have to be present before implementing the blocking mechanism. Since batch models are specified in the Batch Management Tool, certain properties of the model, specifically the Process and the Activity, have to be provided by the engine through an HTTP request from the page in Figure 16.

As mentioned already, the REST API of Camunda is a separate module, which made the decision of where to put the code to handle the aforementioned HTTP request easy. The idea of how to implement said functionality is as follows: 1) Retrieve all process definitions currently deployed on Camunda, 2) Retrieve all activities that each process contains, 3) Package each process with its activities in a new object, 4) Serialize all new objects into a JSON object and send as response to browser. In Listing 2 is the implementation of steps 1 and 2. The list of *ProcessDefinitionExtract* objects is then serialized into a JSON object and sent back to the requester (browser).

```

1 @GET
2 @Path("/activities")

```

```

3 @Produces(MediaType.APPLICATION_JSON)
4 List<ProcessDefinitionExtract> getProcessActivityNames(@Context
  UriInfo uriInfo);

```

Listing 1: Method declaration with REST path

```

1 @Override
2 public List<ProcessDefinitionExtract> getProcessActivityNames(
  UriInfo uriInfo){
3
4     List<ProcessDefinitionDto> list = getProcessDefinitions(
  uriInfo, null, null);
5     List<ProcessDefinitionExtract> extractList = new ArrayList<>();
6
7     for(ProcessDefinitionDto def: list){
8         String name = def.getName();
9         String id = def.getId();
10        List<ProcessActivityExtract> activityList =
  getProcessActivitiesByProcessId(id);
11        extractList.add(new ProcessDefinitionExtract(id, name,
  activityList));
12    }
13
14    return extractList;
15 }
16
17 private List<ProcessActivityExtract>
  getProcessActivitiesByProcessId(String processId){
18
19     List<ProcessActivityExtract> list = new ArrayList<>();
20     Collection<ModelElementInstance> taskInstances;
21     InputStream processModelIn = null;
22
23     try {
24         processModelIn = getProcessEngine().getRepositoryService().
  getProcessModel(processId);
25         byte[] processModel = IoUtil.readInputStream(processModelIn
  , "processModelBpmn20Xml");
26         InputStream stream = new ByteArrayInputStream(processModel)
  ;
27         BpmnModelInstance modelInstance = org.camunda.bpm.model.
  bpmn.Bpmn.readModelFromStream(stream);
28         ModelElementType taskType = modelInstance.getModel().
  getType(Task.class);
29         taskInstances = modelInstance.getModelElementsByType(
  taskType);
30     } catch (AuthorizationException e) {
31         throw e;
32     } catch (ProcessEngineException e) {
33         throw new InvalidRequestException(Status.BAD_REQUEST, e, "
  No matching definition with id " + processId);
34     } finally {
35         IoUtil.closeSilently(processModelIn);
36     }
37
38     for(ModelElementInstance task: taskInstances){
39         list.add(new ProcessActivityExtract(task.getAttributeValue(

```

```

40     "id"), task.getAttributeValue("name"));
41     }
42     return list;
43 }

```

Listing 2: Method definition

Something of note is that the model of each process definition must be retrieved from storage and streamed as a byte array before converted to usable Java objects, as can be seen in Listing 2. Other methods were tried but did not permit filtering process elements by type and were also considerably slower.

The next step was to update the engine about existing batch models or newly created/updated/deleted ones (#1 and #2). For this purpose, a global `Set<String>` variables was created in the Engine Adapter to hold all activity Ids that have a batch model assigned. For the sake of simplicity and reduced latency in identifying batch activities during runtime, two assumptions were made:

1. Each activity in a process has a unique name and Id.
2. No two activities with the same Id exist across all processes currently deployed to the engine.

The premise of the second assumption may, of course, be permitted, if an activity Id is correlated with a process Id, with the latter being definitely unique in the engine. In this case a different container would be required, eg. a `Map<String, Set<String>` to hold a list of process Ids, with each having a set of unique activities Ids. While this would indeed work in a possible production environment, it was not deemed necessary for the needs of this project (extra memory etc.).

The set of activity Ids is modified in two occasions: on engine startup and during runtime, when a batch model is create, deleted, or updated. In the first case, code has be included in a startup script in order to initialize the set as soon as possible. In the code of Camunda there is an initialization script that executes operations steps and is suitable for such an occasion. In Listing 3, the step `MySQLBatchModelRetrievalStep` retrieves all activity Ids from a MySQL database and inserts them in the set. The implementation of the step in Listing 4.

```

1     protected void deployBpmPlatform(LifecycleEvent event) {
2
3         final StandardServer server = (StandardServer) event.getSource();
4
5         containerDelegate.getServiceContainer().createDeploymentOperation("
        deploy BPM platform")
6             .addAttachment(TomcatAttachments.SERVER, server)
7             .addStep(new TomcatParseBpmPlatformXmlStep())
8             .addStep(new DiscoverBpmPlatformPluginsStep())
9             .addStep(new StartManagedThreadPoolStep())
10            .addStep(new StartJobExecutorStep())
11            .addStep(new PlatformXmlStartProcessEnginesStep())
12            .addStep(new MessageProducerStep())

```

```

13     .addStep(new MysqlBatchModelRetrievalStep())
14     .execute();
15
16     LOG.camundaBpmPlatformSuccessfullyStarted(server.getServerInfo());
17
18 }

```

Listing 3: Initialization script

```

1     public class MysqlBatchModelRetrievalStep extends
DeploymentOperationStep {
2
3     private static Connection connection;
4
5     public String getName() {
6         return "Retrieving batch models";
7     }
8
9     @Override
10    public void performOperationStep(DeploymentOperation
operationContext) {
11
12        connect();
13        EngineAdapter.INSTANCE.setList(getData());
14    }
15
16    public void connect() {
17
18        final String url = "jdbc:mysql://localhost:3306/test";
19
20        try {
21            connection = DriverManager.getConnection(url, "root", "root
");
22        }
23        catch (Exception e) {
24            MessageProducer.publish("Connection error: "+e.toString(),
Queue.exception);
25        }
26    }
27
28    public Set<String> getData() {
29
30        Set<String> list = new HashSet<>();
31
32        try {
33            Statement stmt = connection.createStatement();
34            ResultSet rs = stmt.executeQuery("select activityId from
models");
35
36            while (rs.next()) {
37                list.add(rs.getString("activityId"));
38            }
39        }
40        catch (Exception e) {
41            MessageProducer.publish("Data retrieval error: "+e.toString
(), Queue.exception);
42        }

```



```

43
44     return list;
45 }
46 }
47
48 }

```

Listing 4: Batch Activity Retrieval Step Implementation

Before moving to the implementation of the second case, it is important to explain the two ways with which Camunda starts a process instance. In a given engine session, an internal storage service called DeploymentCache containing process definitions that have already been parsed and stored with their internal Java object representation. This means that a process that has been instantiated once during the session, does not need to be parsed again every time it needs to be instantiated. When a process definition model is parsed for the first time during the session, an ActivityImpl object is created representing an activity in the model. The code was modified to check whether there is a batch model for this activity, by searching for the activity Id in the Set discussed previously. If the Id exists, then a flag "block", which is added to the ActivityImpl class as a field, is set to true. The modification in Listing 5 (lines 4-6).

```

1     public ActivityImpl createActivity(String activityId) {
2         ActivityImpl activity = new ActivityImpl(activityId,
3             processDefinition);
4         if (activityId!=null) {
5             if (EngineAdapter.INSTANCE.findBatchActivityById(activityId)) {
6                 activity.setBlock(true);
7             }
8             if (processDefinition.findActivity(activityId) != null) {
9                 throw new PvmException("duplicate activity id '" + activityId +
10                    "'");
11             }
12             if (BACKLOG.containsKey(activityId)) {
13                 BACKLOG.remove(activityId);
14             }
15             namedFlowActivities.put(activityId, activity);
16         }
17         activity.flowScope = this;
18         flowActivities.add(activity);
19     }

```

Listing 5: Setting batch flag to activity

In case the process definition exists in the DeploymentCache, meaning ActivityImpl objects already exist for each activity, then it is sufficient to retrieve the definition from the DeploymentCache, then retrieve the specific activity from the definition and set the flag to true. Moreover, as a flexibility aspect, the administrator may decide that a specific process instance that has been halted should be removed from the batch and resumed independently from the rest. This is why the object ExecutionEntity, representing an active process instance, should also have a "block" flag. The code in Listing 6 sets the flags in both the ActivityImpl and ExecutionEntity of already running instances. If the instance has already reached and executed the batch activity, then it is no longer affected by the flags and continues its execution. If another instance has not yet reached the batch activity, then it is

halted when it does. This way the flexibility need Evolution (II) is satisfied. it should be noted that the function in Listing 6 is a REST function that receives data from the batch controller such as the value of the flag (true/false) and activity Id.

```
1      public void setBatchActivity(BatchActivityDto dto){
2
3          String activityId = dto.getActivityId();
4          String definitionId = dto.getProcessDefinitionId();
5          boolean block = dto.isBlock();
6
7          ProcessEngineConfigurationImpl config = EngineAdapter.
getConfiguration();
8          Context.setProcessEngineConfiguration(config);
9          Context.setCommandContext(new CommandContext(config));
10         DeploymentCache cache = config.getDeploymentCache();
11         ProcessDefinitionEntity process = cache.
findDeployedProcessDefinitionById(definitionId);
12
13         if(process != null){
14             try{
15                 process.getActivities()
16                     .stream()
17                     .filter(item -> item.getActivityId().equals(
activityId))
18                     .findFirst().ifPresent(batchActivity ->
batchActivity.setBlock(block));
19
20                 List<ExecutionEntity> executions = EngineAdapter.
getExecutionList()
21                     .values()
22                     .stream()
23                     .filter(item-> item.getProcessDefinitionId().
equals(definitionId))
24                     .collect(Collectors.toList());
25
26                 for(ExecutionEntity item : executions){
27                     item.getProcessInstance().setBlock(block);
28                 }
29             }
30             catch (Exception e){
31                 MessageProducer.publish(this.getClass().getName()+"->"+
e.toString(), Queue.exception);
32             }
33             finally{
34                 Context.removeProcessEngineConfiguration();
35                 Context.removeCommandContext();
36             }
37         }
38         else{
39             if(block){
40                 EngineAdapter.addBatchActivity(activityId);
41             }
42             else{
43                 EngineAdapter.removeBatchActivity(activityId);
44             }
45         }
}
```

Listing 6: Batch flag set during runtime

The blocking mechanism at #3 also required extensive investigation which revealed that when a process reaches an activity is executes and operation called "PvmAtomicOperationActivityExecute" that implements the execute() method of the interface "CoreAtomicOperation". The code of this method, visible in Listing ??, contains the blocking mechanism at lines 5-15. The idea is that when this activity is reached, the flags of the activity and the instance are checked, and if both are true, the general execution of the instance is interrupted by a "return" statement. The engine behaves as if the instance has finished and then removes entities needed for further execution, namely the CommandInvocationContext, the ProcessEngineConfiguration and the ContextClassLoader. As a result, these entities have to be stored separately along with the general information of the process instance. This information is stored in an object called BatchEntity which is defined in Listing 8. This satisfies functionality #4 of Table 1.

As an additional note, the Id of the instance is used to retrieve the BatchEntity in order to resume the instance, but sometimes the Id of the ExecutionEntity may refer to a child execution of the parent, meaning the process instance. Retrieving this ExecutionEntity and attempting to resume it would fail, since the child does not contain the necessary information that is contained in the father. The solution is in line 11 of the listing, which is to check whether the parentId of the ExecutionEntity is null. If it is null, then the ExecutionEntity is the parent and we can use the Id for the BatchEntity. If parentId is not null, then the ExecutionEntity is a child execution, and the parentId needs to be used in the BatchEntity. In any case, the correlation later, always uses the Id of the parent.

```

1  public void execute(PvmExecutionImpl execution) {
2
3  final ActivityImpl activity = execution.getActivity();
4
5  if(activity.isBlock() && execution.getProcessInstance().isBlock()){
6      ProcessEngineConfigurationImpl config = Context.
getProcessEngineConfiguration();
7      CommandInvocationContext commandInvocationContext = Context.
getCommandInvocationContext();
8      ClassLoader executionContextClassLoader = Thread.currentThread().
getContextClassLoader();
9      execution.setContextClassLoader(executionContextClassLoader);
10     EngineAdapter.getBatchList().add(new BatchEntity(execution,
config, commandInvocationContext));
11     String id = execution.getParentId() != null ? execution.
getParentId() : execution.getId();
12     EngineAdapter.INSTANCE.prepareProcessInstanceForBatch((
ExecutionEntity)execution);
13     MessageProducer.publish("Activity '"+activity.getName()+"' is
blocked from executing... Process Instance id is: "+id, Queue.info);
14     return;
15 }
16
17 //If this process instance was blocked before, then set the context
class loader
18 if(activity.isBlock()){

```

```

19     ClassLoaderUtil.setContextClassLoader(execution.
getContextClassLoader());
20     }
21
22     execution.activityInstanceStarted();
23
24     execution.continueIfExecutionDoesNotAffectNextOperation(new
Callback<PvmExecutionImpl, Void>() {
25         @Override
26         public Void callback(PvmExecutionImpl execution) {
27             if (execution.getActivity().isScope()) {
28                 execution.dispatchEvent(null);
29             }
30             return null;
31         }
32     }, new Callback<PvmExecutionImpl, Void>() {
33
34         @Override
35         public Void callback(PvmExecutionImpl execution) {
36
37             ActivityBehavior activityBehavior = getActivityBehavior(
execution);
38
39             ActivityImpl activity = execution.getActivity();
40             LOG.debugExecutesActivity(execution, activity, activityBehavior
.getClass().getName());
41
42             try {
43                 MessageProducer.publish("Now executing activity: "+activity.
getName(), Queue.info);
44                 activityBehavior.execute(execution);
45             } catch (RuntimeException e) {
46                 throw e;
47             } catch (Exception e) {
48                 throw new PvmException("couldn't execute activity <" +
activity.getProperty("type") + " id=\"" + activity.getId() + "\"
...>: " + e.getMessage(), e);
49             }
50             return null;
51         }
52     }, execution);
53 }

```

Listing 7: Blocking mechanism

```

1     public class BatchEntity {
2
3         private PvmExecutionImpl execution;
4         private ProcessEngineConfigurationImpl config;
5         private CommandInvocationContext commandInvocationContext;
6
7         public BatchEntity(PvmExecutionImpl execution,
ProcessEngineConfigurationImpl config,
8             CommandInvocationContext
commandInvocationContext) {
9
10            this.execution = execution;

```

```

11     this.config = config;
12     this.commandInvocationContext = commandInvocationContext;
13 }
14
15 public String getId(){
16     return this.execution.getId();
17 }
18
19 public String getParentId(){
20     return this.execution.getParentId();
21 }
22
23 public PvmExecutionImpl getExecution(){
24     return this.execution;
25 }
26
27 public ProcessEngineConfigurationImpl getProcessEngineConfiguration
28 () {
29     return this.config;
30 }
31
32 public CommandInvocationContext getCommandInvocationContext(){
33     return this.commandInvocationContext;
34 }
35
36 protected void compareVariables(Map<String, Object> newVariables){
37     Map<String, Object> currentVariables = execution.getVariables()
38 ;
39     Object o = null;
40     for(String str : currentVariables.keySet()){
41         Class<?> clazz = currentVariables.get(str).getClass();
42         Class<?> newClazz = newVariables.get(str).getClass();
43         if(clazz == Long.class && newClazz == Integer.class){
44             o = newVariables.get(str);
45         }
46         try{
47             o = (long) (int) o;
48             newVariables.put(str, o);
49         }
50         catch(Exception e){
51             MessageProducer.publish(this.getClass().getName()+" "+
52 e.toString(), Queue.exception);
53         }
54     }
55 }
56 }

```

Listing 8: BatchEntity definition

Functionality #5 is enabled with the use of the MessageProducer, a RabbitMQ client that pushes messages to three different queues based on the message type. These queue are: the information queue, which stores messages regarding general information about the engine, process instance states, activities that are being executed etc, the exception queue

which stores exceptions and error during engine runtime and the block queue which stores JSON object of blocked processes. The definition of the MessageProducer is available in Listing 9. Serialization is accomplished using the Jackson library.

```
1 public class MessageProducer {
2
3 private ConnectionFactory factory;
4 private Connection connection;
5 private static Channel channel;
6 private static final Logger LOGGER = Logger.getLogger(
MessageProducer.class.getName());
7 private static MessageProducer instance = new MessageProducer();
8
9 private MessageProducer() {
10
11     factory = new ConnectionFactory();
12     factory.setHost("localhost");
13
14     try{
15         connection = factory.newConnection();
16         channel = connection.createChannel();
17         channel.exchangeDeclare("batching", "direct", false, false,
null);
18         channel.queueDeclare("info", false, false, false, null);
19         channel.queueDeclare("block", false, false, false, null);
20         channel.queueDeclare("exception", false, false, false, null
);
21         channel.queueBind("info", "batching", "info");
22         channel.queueBind("block", "batching", "block");
23         channel.queueBind("exception", "batching", "exception");
24     }
25     catch (Exception e) {
26         LOGGER.severe(e.getMessage());
27     }
28 }
29
30 public static MessageProducer getInstance() {
31     return MessageProducer.instance;
32 }
33
34 public Channel getChannel() {
35     return this.channel;
36 }
37
38 public static void publish(String message, Queue queue) {
39
40     try {
41         channel.basicPublish("batching", queue.getName(), null,
message.getBytes("UTF-8"));
42     }
43     catch (Exception e) {
44         LOGGER.severe(e.getMessage());
45     }
46 }
47 }
```

Listing 9: MessageProducer definition

Finally, functionality #6 is realized in two steps: The Batch Controller has to make an HTTP request containing the process instance(s), the variables and a boolean value indicating how the instance(s) may proceed with execution. The function handling the HTTP method is visible in Listing 10. After deserializing the JSON object, the operation that has to be executed is determined by the boolean value. If true, the instance resumes execution from the batch activity with the operation `ACTIVITY_EXECUTE` which executes the activity the instance is at. If false, the operation `ACTIVITY_LEAVE` is executed, which leaves the current (batch) activity and goes to the next.

The last step requires the Engine Adapter to retrieve the `BatchEntity` based on the Id from the HTTP request and proceed to execute the operation. However, deserialization of the JSON yielded differences between the variables types sent from the Batch Controller and the variable types in Camunda, Integer and Long respectively. The solution was to repair this disparity resulting in exceptions, by converting Integer to Long. The code for resuming the method is visible in Listing 11.

```

1  public void resumeProcessInstanceId(String json){
2
3  Map<String, Object> variables;
4  PvmAtomicOperation operation;
5
6  try{
7      ResumeProcessInstanceDto dto = mapper.readValue(json,
8  ResumeProcessInstanceDto.class);
9      String processInstanceId = dto.getProcessInstanceId();
10     variables = dto.getVariables();
11     if(dto.getResumeAtCurrentActivity()){
12         operation = PvmAtomicOperation.ACTIVITY_EXECUTE;
13     }
14     else{
15         operation = PvmAtomicOperation.ACTIVITY_LEAVE;
16     }
17     EngineAdapter.INSTANCE.resumeProcessInstanceId(
18 processInstanceId, variables, operation);
19 }
20 catch(Exception e){
21     MessageProducer.publish(this.getClass().getName()+" "+e.toString
22 (), Queue.exception);
23 }

```

Listing 10: REST method handling the resume command

```

1  public void resumeProcessInstanceId(String processInstanceId, Map
2  <String, Object> variables, PvmAtomicOperation operation){
3
4      BatchEntity entity = getBatchEntity(processInstanceId);
5      PvmExecutionImpl execution = getExecution(entity);
6
7      ProcessEngineConfigurationImpl config =
8  getProcessEngineConfiguration(entity);
9      Context.setProcessEngineConfiguration(config);
10     Context.setCommandContext(new CommandContext(config));

```

```

9     Context.setCommandInvocationContext(getCommandInvocationContext
    (entity));
10
11     if(variables!= null && !variables.isEmpty()){
12         entity.compareVariables(variables);
13         if(execution.getProcessInstance() != null){
14             execution.getProcessInstance().setVariables(variables);
15         }
16         else{
17             execution.setVariables(variables);
18         }
19     }
20
21     execution.getProcessInstance().setBlock(false);
22
23     try{
24         MessageProducer.publish("Resuming process instance with id:
    "+ execution.getProcessInstance().getId(), Queue.info);
25         operation.execute(execution);
26     }
27     catch(Exception e){
28         MessageProducer.publish(this.getClass().getName()+" "+e.
    getMessage(), Queue.exception);
29     }
30     finally{
31         Context.removeProcessEngineConfiguration();
32         Context.removeCommandContext();
33         Context.removeCommandInvocationContext();
34         ProcessEngineContextImpl.set(false);
35     }
36 }

```

Listing 11: Engine Adapter method to resume instance

It should be noted that the same principle applies for resuming multiple instances (eg. when a batch has finished), the only difference being that the HTTP request now containing a list of process instance Ids with their variables, instead of a single one.

III.2 Batch Controller

RELATED REQUIREMENTS	1, 3, 4
----------------------	---------

The Batch Controller component roughly corresponds to the Controller in Figure 7, presented in [2]. The Batch Controller is responsible for facilitating the management of batch models, and to a lesser extend halted processes, following operations from the Batch Management Tool. These include: creating, updating and deleting models, assigning halted processes to batch clusters and managing the clusters' life-cycle, managing the execution of a batch and acting as a broker between the Batch Management Tool and Camunda in limited cases.

Similarly to the Camunda Process Engine, the following functionality can be identified, that has to be implemented by the Controller.

The Batch Controller uses a script to initialize components early on start-up. These include the BatchClusterManager, the MessageReceiver, the DBManager and the list of

#	Functionality	Requirement	Notes
1	Create/Update/Delete batch models	1	Flexibility need 2 - Looseness (II)
2	Assignment of processes to batches	3	
3	Execution of a batch	3	
4	Send batch processes with variables back to Camunda	3, 4	
5	Cancel batch	1, 4	Flexibility need 1 - Evolution
6	Remove individual instance from batch	4	Flexibility need 1 - Evolution

Table 2: Batch Controller functionality

models and blocked processes. The script is available in Listing 12.

```

1  public class Startup {
2
3  public static DBManager db;
4  public static List<Model> models;
5  public static final Logger log = LoggerFactory.getLogger("BATCH
6  CONTROLLER");
7  public static MessageReceiver receiver;
8  public static List<ProcessInstance> globalList;
9
10
11 public Startup() {
12
13     //init components
14     new BatchClusterManager();
15     models = new ArrayList<>();
16     db = new DBManager();
17     receiver = new MessageReceiver();
18     globalList = new ArrayList<>();
19     db.fillLists();
20 }

```

Listing 12: Startup script

As stated in previous sections, the Batch Controller uses an underlying REST component to route HTTP requests to the right handler. Play! framework provides a simple way to specify the request handler, by writing a line in the "routes" file of the project with format: [Http method] [handler url path] [path of the method]. In Figure 18, is the routes file of the Batch Controller. All HTTP request from the Controller are handled by functions which return a Result object indicating whether the request has succeeded. Next, the method makes an internal call to the corresponding method of the DBManager class which makes direct contact with the database to execute the appropriate queries. For each operation of functionality #1 in Table 2, the two methods are listed in Listings [], [] and []. First is the method of the request handler, followed by the DBManager method.

```

1 # Routes
2 # This file defines all application routes (Higher priority routes first)
3 # ~~~~
4
5 # An example controller showing a sample home page
6 GET / controllers.ModelController.getModel
7
8 GET /get-blocked-processes controllers.ProcessController.sendList
9
10 POST /create/:json controllers.ModelController.createModel(json: String)
11
12 POST /delete/:id controllers.ModelController.deleteModel(id: String)
13
14 POST /update/:json controllers.ModelController.updateModel(json: String)
15
16 POST /resume/:id controllers.ProcessController.resume(id: String)

```

Figure 18: Routes file

```

1 public Result createModel(String json){
2
3     Result result = null;
4     String success;
5
6     try{
7         Model model = (Model)JsonBuilder.fromJson(json, Model.class
8     );
9         success = Boolean.toString(Startup.db.create(model.
10     createValueString()));
11         if(success.equals("true")){
12             Startup.models.add(model);
13             result = ok("Model successfully created");
14             String jsonRest = "{\"processDefinitionId\" : \""+model
15     .getProcessId()+"\", \"activityId\" : \""+model.getActivityId()+"\",
16     \"block\" : \"true\"}";
17             String method = "process-definition/set-batch-activity"
18     ;
19             RESTClient.newRequest(jsonRest, method);
20         }
21     }
22     catch(Exception e){
23         result = internalServerError("Server error: Could not
24     create model\n Reason: "+e.toString());
25     }
26
27     return result;
28 }
29
30 public boolean create(String modelValues){
31
32     boolean success = false;
33
34     try{
35         Statement stmt=connection.createStatement();
36         int rs=stmt.executeUpdate("insert into models (name,
37     maxBatchSize, minBatchSize, timeLimit, time," +
38     " executionOrder, process, processId, activity,
39     activityId, attributes, Id)" +
40     " values"+modelValues);
41         success = (rs==1);
42     }
43     catch(Exception e){
44         Startup.log.error("Create: " +e.getMessage());

```

```

37     }
38
39     return success;
40 }

```

Listing 13: Create model

```

1  public Result updateModel(String json){
2
3      Result result = null;
4      String success;
5      try{
6          Model model = (Model)JsonBuilder.fromJson(json, Model.class
7      );
8          success = Boolean.toString(Startup.db.update(model));
9          if(success.equals("true")){
10             listIt = Startup.models.listIterator();
11             while(listIt.hasNext()){
12                 Model m = listIt.next();
13                 if(m.getId().equals(model.getId())) listIt.set(
14 model);
15             }
16             result = ok("Model successfully updated");
17         }
18         String jsonRest = "{\"processDefinitionId\" : \""+model.
19 getProcessId()+"\", \"activityId\" : \""+model.getActivityId()+"\",
20 \"block\" : \"true\"}";
21         String method = "process-definition/set-batch-activity";
22         RestClient.newRequest(jsonRest, method);
23     }
24     catch(Exception e){
25         result = internalServerError("Server error: Could not
26 update model\n Reason: "+e.toString());
27     }
28
29     return result;
30 }
31
32 public boolean update(Model model){
33
34     boolean success = false;
35
36     try{
37         PreparedStatement ps = connection.prepareStatement("update
38 models set " +
39
40             "name=' "+model.getName()+"', " +
41             "maxBatchSize="+model.getMaxBatchSize()+" , "+
42             "minBatchSize="+model.getMinBatchSize()+" , "+
43             "timeLimit="+model.getTimeLimit()+" , "+
44             "time=' "+model.getTime().toString()+"', "+
45             "executionOrder=' "+model.getOrder().toString()+"', "
46
47 +
48             "process=' "+model.getProcess()+"', "+
49             "processId=' "+model.getProcessId()+"', "+
50             "activity=' "+model.getActivity()+"', "+

```

```

44         "attributes='"+model.getAttributes()+"' "+
45         "where Id='"+model.getId()+"';");
46         success = (ps.executeUpdate()==1);
47     }
48     catch(Exception e){
49         Startup.log.error(e.getMessage());
50     }
51
52     return success;
53 }

```

Listing 14: Update model

```

1     public Result deleteModel(String id){
2
3         Result result = null;
4
5         try{
6             Model m = null;
7             String response = Boolean.toString(Startup.db.delete(id));
8             if (response.equals("true")){
9                 modelsIt = Startup.models.iterator();
10                while(modelsIt.hasNext()){
11                    m = modelsIt.next();
12                    if(m.getId().equals(id)){
13                        modelsIt.remove();
14                        break;
15                    }
16                }
17
18                if(m != null){
19                    result = ok("Model successfully deleted");
20                    String jsonRest = "{\"processDefinitionId\" : \""+m
21                    .getProcessId()+"\", \"activityId\" : \""+m.getActivityId()+"\", \"
22                    block\" : \"false\"}";
23                    String method = "process-definition/set-batch-
24                    activity";
25
26                    RestClient.newRequest(jsonRest, method);
27                    BatchClusterManager.removeAndFinalize(m);
28                }
29
30                else{
31                    result = internalServerError("Server error: Could not
32                    delete model");
33                }
34            }
35            catch(Exception e){
36                result = internalServerError("Server error: Could not
37                delete model\n Reason: "+e.toString());
38            }
39            return result;
40        }
41    }
42
43    public boolean delete(String id){
44
45    }

```

```

40     boolean success;
41     try{
42         Statement stmt=connection.createStatement();
43         int rs=stmt.executeUpdate("delete from models where id='"+
id+"'");
44         success = (rs==1);
45     }
46     catch(Exception e){
47         success = false;
48
49         e.printStackTrace();
50     }
51     return success;
52 }

```

Listing 15: Delete model

Create: The createModel() method receives a JSON object containing the model information specified from the Batch Management Tool. This String is deserialized into a Model object using the Jackson library and the create() method of the DBManager inserts it into the MySQL database. If the operation is successful, the new model is also inserted to the global model list defined in Listing 12. Subsequently, the Batch Controller sends an HTTP request to Camunda informing it of the new batch model as explained in Listing 6.

Update: The updateModel() method operates similarly to the createModel(), with a difference being that after the successful database query, the batch model has to be retrieved from model list in order to be updated locally.

Delete: the deleteModel() method only receives a single String with the Id of the model to be deleted. After the successful database query, the models list has to be iterated again in order to remove the model.

As mentioned in Section II, the system uses three different message queues for information, exceptions and blocked processes from Camunda to the Batch Controller. The MessageReceiver of the Controller listens to three queues and uses a different handler to handle the messages from each of them. While for the "information" and "exception" queues the handler only print the message in the Log, the handler of the "block" queue is functioning differently. In Listing [] is the receiving end of the "block" queue and the handling of he message by the BlockMessageHandler.

```

1     channel.basicConsume("block", false, "BlockConsumer",
2         new DefaultConsumer(channel){
3             @Override
4             public void handleDelivery(String consumerTag,
5                                     Envelope envelope,
6                                     AMQP.BasicProperties
7                                     properties,
8                                     byte[] body)
9                                     throws IOException
10            {
11                String routingKey = envelope.getRoutingKey
() ;
                String contentType = properties.

```

```

12     getContentType();
13         long deliveryTag = envelope.getDeliveryTag
14         ();
15         String message = new String(body, "UTF-8");
16         BlockMessageHandler.handle(message);
17         channel.basicAck(deliveryTag, false);
18     }
19     });
20
21     public static void handle(String message) {
22         ProcessInstance instance;
23         try{
24             instance = (ProcessInstance) JsonBuilder.fromJson(message,
25             ProcessInstance.class);
26             instance.repairVariables();
27             ProcessController.addToList(instance);
28             Startup.globalList.add(instance);
29             BatchClusterManager.assignToCluster(instance);
30         }
31         catch (Exception e) {
32             e.printStackTrace();
33         }
34     }

```

Listing 16: Block Message handling

For the implementation of functionality #2, the String message from the queue has to be deserialized into a ProcessInstance object and its variables repaired (see mod. of Process Engine). In order to assign the new instance to a batch cluster, the BatchClusterManager implements the following procedure. The String identifier has to be constructed from the instance and the corresponding batch model. This identifier called a "cluster key identifier", is synthesized by three parts: the process Id, the activity name and the value of the grouping attribute. They key then has the form '[process_Id]/[activity_name]:[attribute_value]'. This string is then compared against a list of type Map<String, Integer> containing the keys of existing batch clusters. If the key is not present in the list, a new batch cluster is created. The Integer value of each entry in the Map signifies the cluster order, meaning the number of cluster of the same type that exist. This is because several cluster may exist that have the same key. For example if 5 process instances have been started and blocked, with value "test" in the grouping attribute, but the batch model only allows 3 in every batch, then these 5 instances will be split in two batches with 3 and 2 instances each respectively. When the first cluster is full, the second must be created to accommodate the remaining two instances. Additionally, every cluster has a unique id which is comprised of the cluster identifier, plus the cluster number. Lastly, the received instance is inserted to the newly created cluster or the existing one if possible. The implementation of this procedure in Listings 17 and 18.

```

1     public static void assignToCluster(ProcessInstance instance) {
2         BatchCluster cluster = findOrCreate(instance);
3         cluster.tryAdd(instance);
4         instance.setClusterId(cluster.getId());
5     }
6
7     private static BatchCluster findOrCreate(ProcessInstance instance) {
8         BatchCluster cluster = null;

```

```

9      Model model = instance.getBatchModel();
10     String key = constructClusterKeyIdentifier(instance, model);
11     int clusterOrder = clusterKeys.getOrDefault(key, 0);
12
13     if(clusterOrder == 0){
14         cluster = new BatchCluster(model);
15         cluster.setKey(key);
16         addCluster(cluster);
17         clusterKeys.put(key, ++clusterOrder);
18         startBatchTimer(cluster);
19         Startup.log.info("New batch with id: "+cluster.getId()+"
started");
20     }
21     else{
22         String clusterId = key+"-"+clusterOrder;
23         cluster = batchClusters.get(clusterId);
24     }
25
26     return cluster;
27 }
28
29 public static String constructClusterKeyIdentifier(ProcessInstance
instance, Model model){
30     String key;
31     String processDefinitionComponent;
32     String processDefinitionIdParts[] = model.getProcessId().split(
":");
33
34     //If the process definition id contains the process version,
include it in the processDefinitionComponent string
35     if(processDefinitionIdParts.length == 3){
36         processDefinitionComponent = processDefinitionIdParts[0]+":
"+processDefinitionIdParts[1];
37     }
38     else{
39         processDefinitionComponent = processDefinitionIdParts[0];
40     }
41     String activityComponent = model.getActivity();
42     String attribute = model.getAttributes();
43     String attributeValue = (String)instance.variables.get(
attribute.toLowerCase());
44     key = processDefinitionComponent+"/"+activityComponent+": "+
attributeValue;
45
46     return key;
47 }

```

Listing 17: Assigning an instance to a cluster - BatchClusterManager class

```

1     public void tryAdd(ProcessInstance instance){
2
3         if(this.currentCapacity < this.max){
4             this.instances.add(instance);
5             increaseCurrentCapacity();
6         }
7         else{
8             BatchCluster newCluster;

```

```

9         try{
10             newCluster = new BatchCluster(this.batchModel);
11             newCluster.setKey(this.key);
12             newCluster.instances.add(instance);
13             BatchClusterManager.addCluster(newCluster);
14             int nextBatchNumber = BatchClusterManager.
getClusterKeys().get(this.key);
15             BatchClusterManager.getClusterKeys().put(this.key,
nextBatchNumber);
16             BatchClusterManager.startBatchTimer(newCluster);
17         }
18         catch(Exception e){
19             System.out.println(e.toString());
20         }
21     }
22 }

```

Listing 18: Assigning an instance to a cluster - BatchCluster class

When a batch cluster is created, a countdown begins using a Timer object. The duration of this countdown is specified in the batch model and is common for all batches stemming from the model. Once the countdown ends, the batch is executed if the minimum number of instances has been reached. Otherwise, the batch is canceled. The Timer and TimerTask are visible in Listings [] and [].

```

1     public static void startBatchTimer(BatchCluster cluster){
2         Timer timer = new Timer();
3         timers.put(cluster.getId(), timer);
4
5         try{
6             long delay = cluster.transformToLong();
7             timer.schedule(cluster.startTimer(), delay);
8         }
9         catch(Exception e){
10            Startup.log.error(e.toString()+" : "+e.getMessage());
11        }
12    }

```

Listing 19: Timer - BatchClusterManager class

```

1     public TimerTask startTimer() {
2
3         BatchCluster cluster = this;
4         TimerTask task;
5         task = new TimerTask() {
6             @Override
7             public void run() {
8                 if(currentCapacity < min){
9                     Startup.log.info("Minimum number of instances has
not been reached. Resuming individual execution...");
10                    BatchClusterManager.cancelBatch(cluster);
11                    cancel();
12                }
13                else{
14                    Startup.log.info("Execution for Batch "+cluster.id+
" has started");
15                    cluster.setRunning(true);

```



```

16         try{
17             if(cluster.getModel().getOrder() == Model.
executionOrder.PARALLEL) {
18                 CountdownLatch latch = new
CountdownLatch(cluster.instances.size());
19                 BatchExecutionManager.executeParallel(
cluster, latch);
20                 latch.await();
21             }
22             else{
23                 BatchExecutionManager.executeSequential
(cluster);
24             }
25             Startup.log.info("Execution for Batch "+
cluster.id+" has finished");
26             cluster.finalizeBatch(false);
27         }
28         catch(Exception e){
29             e.printStackTrace();
30         }
31     }
32 }
33 };
34
35 return task;
36 }

```

Listing 20: TimerTask - BatchCluster class

In case the execution order in the model is specified as parallel, each instance is executed on a different thread. The batch is concluded once all thread have finished their execution, which is determined by a CountdownLatch object. Right before the execution, the activity to execute is determined on runtime. The code in Listings 21 and 22.

```

1     public static void executeParallel(BatchCluster cluster,
CountdownLatch latch) {
2         String activity = cluster.getModel().getActivity().replace(" ",
"");
3
4         try{
5             Class<?> clazz = Class.forName(delegatePackage+activity);
6             Constructor<?> delegateConstructor = clazz.getConstructor(
Map.class);
7             for(ProcessInstance instance : cluster.getList()){
8                 new InstanceExecutor(instance.getVariables(),
delegateConstructor, latch);
9             }
10        }
11        catch(Exception e){
12            Startup.log.error(e.toString());
13        }
14    }

```

Listing 21: Loading activity to execute and thread spawning - BatchExecutionManager class

```

1     public InstanceExecutor(Map<String, Object> variables,
2     Constructor<?> constructor, CountdownLatch latch){
3         this.constructor = constructor;
4         this.variables = variables;
5         this.latch = latch;
6         thread = new Thread(this);
7         thread.start();
8     }
9
10    @Override
11    public void run() {
12
13        try{
14            this.constructor.newInstance(variables);
15            latch.countDown();
16        }
17        catch (Exception e){
18            e.printStackTrace();
19        }
20    }

```

Listing 22: Thread Execution - InstanceExecutor class

The activity that is determined in real-time is of course identical to the activity behaviour that is deployed to Camunda along with the process model. In the Batch Controller, the activity script changes the process variables, which are then ready to be serialized again and sent back to the Camunda Process Engine.

When the batch has finished execution, or is canceled, then it is up for finalization, meaning, to send the process instances back to the engine to continue their normal flow of execution. In finalization, the process instances of the batch are then packaged into a list which is then serialized and sent to Camunda through HTTP, as explained in the previous section. The finalization method takes a boolean parameter indicating how the instances are to be resumed. The finalization and REST method are available in Listing 23 and 24.

```

1     public void finalizeBatch(boolean resumeAtCurrentActivity) {
2         Map<String, Map<String, Object>> map = new HashMap<>();
3         this.setRunning(false);
4
5         for (ProcessInstance instance : this.instances) {
6             Startup.globalList.remove(instance);
7             map.put(instance.processInstanceId, instance.variables);
8         }
9
10        BatchClusterManager.removeCluster(this);
11
12        int oldOrder = BatchClusterManager.getClusterKeys().get(this.
13        key);
14        int newOrder = --oldOrder;
15        if(newOrder != 0){
16            BatchClusterManager.getClusterKeys().put(this.key, newOrder
17        );
18        }
19        else{
20            BatchClusterManager.getClusterKeys().remove(this.key);

```

```

19     }
20
21
22     String json = JsonBuilder.toJson(new
ProcessInstanceListRestExtract(map, resumeAtCurrentActivity));
23     String method = "process-instance/resumeMany";
24     RestClient.newRequest(json, method);
25 }

```

Listing 23: Batch finalization - BatchCluster class

```

1     public static int newRequest(String payload, String method){
2
3         int code = 0;
4
5         try{
6             StringEntity jsonEntity = new StringEntity(payload,
ContentType.APPLICATION_JSON);
7             HttpPost request = new HttpPost(url+method);
8             request.addHeader("content-type", "application/json");
9             request.setEntity(jsonEntity);
10            CloseableHttpResponse response = httpClient.execute(request
);
11            code = response.getStatusLine().getStatusCode();
12
13        }
14        catch(Exception e){
15
16        }
17
18        return code;
19    }

```

Listing 24: REST method

Canceling a batch means preventing the batch from executing because, eg. the minimum number of instances has not been reached. Considering the implementation of the Timer-Task in Listing 20, the only operation needed is to cancel the timer itself and finalize the batch, setting the boolean value to true (the batch is not executed, so the instances have to continue from the same blocked activity). The code is visible in Listing ??.

```

1     public static void cancelBatch(BatchCluster cluster){
2         Timer timer = timers.get(cluster.getId());
3         try{
4             timer.cancel();
5         }
6         catch(Exception e){
7             Startup.log.error(e.toString()+" : "+e.getMessage());
8         }
9         finally{
10            timers.remove(cluster.getId());
11            cluster.finalizeBatch(true);
12        }
13    }

```

Listing 25: REST method

Lastly, for reasons of increased flexibility in the system, suppose that the administrator would require a way to remove a process instance from the batch, in which case the instance would have to be resumed at the blocked activity. Assuming that a process instance cannot be removed from a running batch, the method can be written as in Listing 26. An HTTP request is sent from the Batch Management Tool with the Id of the process that needs to be removed. The method then searches for it in the global list of startup (see Listing 12) and then sends its own HTTP request to Camunda.

```

1     public Result resume(String processInstanceId) {
2
3         Result result;
4         int code;
5         ProcessInstance instance = null;
6
7         try{
8             instance = Startup.globalList
9                 .stream()
10                .filter(item -> item.getProcessInstanceId().equals(
11                    processInstanceId) || (item.getParentId() != null && item.
12                    getParentId().equals(processInstanceId))
13                .findAny()
14                .orElse(null);
15            if(instance != null){
16                instance.removeFromCluster();
17                Startup.globalList.remove(instance);
18                instance.setResumeAtCurrentActivity(true);
19                ProcessInstanceRestExtract extract = new
20                ProcessInstanceRestExtract(instance);
21                String json = JsonBuilder.toJson(extract);
22                String method = "process-instance/resumeOne";
23                code = RestClient.newRequest(json, method);
24
25                if(code > 299){
26                    throw new NoHttpResponseException("Server error");
27                }
28            }
29            else{
30                throw new ProcessInstanceNotFoundException("Process
31                instance not found");
32            }
33
34            result = ok("Process successfully resumed");
35        }
36        catch(Exception e){
37            result = internalServerError(e.toString());
38        }
39
40        return result;
41    }

```

Listing 26: Resume single instance method

```

1     public static void removeCluster(BatchCluster cluster) {
2
3         try{

```

```

4         if(cluster.isRunning()){
5             throw new RemoveRunningBatchException();
6         }
7         cluster.decrease();
8         batchClusters.remove(cluster.getId());
9     }
10    catch(Exception e){
11        Startup.log.error(e.toString());
12    }
13 }

```

Listing 27: Cannot remove instance from running batch

III.3 Batch Management Tool

RELATED REQUIREMENTS	1, 4
----------------------	------

The implementation of the Batch Management Tool is not directly involved with the four requirements defined earlier in the section, as it mainly deals with the user interface functionality. Its principal functions are the listing, creation, update and deletion of models, as well as the listing of blocked processes in a separate tab as shown in Section II.4. For the implementation of the above functions, the tool makes several HTTP requests to the Batch Controller, with the exception of the list of process and list of activities which are requested from Camunda. In most cases, if the request succeeds, the result will be visualized in the browser along a pop-up message stating as such. Otherwise, the message will display an error message (eg. "Internal Server Error"), with no change in the browser. Two examples of such requests can be seen in Listings 28 and 29.

```

1     function deleteModel(){
2         if(selectedId !== null){
3             if (confirm("Delete this model?")) {
4                 var xhr = new XMLHttpRequest();
5                 xhr.open("POST", "http://localhost:9000/delete/"+selectedId
6                 , true);
7                 xhr.onload = function (e) {
8                     if (xhr.readyState === 4 && xhr.status === 200) {
9                         var item = document.getElementById(selectedId);
10                        item.parentNode.removeChild(item);
11                        var idx = modelList.indexOf(modelList.find(x => x.
12                        id === selectedId));
13                        delete modelList[idx];
14                        clearPropertyFields();
15                        alert(xhr.responseText);
16                    }
17                }
18                else {
19                    alert("Server error");
20                }
21            }
22            xhr.onerror = function (e) {
23                alert("Could not reach server");
24            };
25            xhr.send(null);

```

```

24     }
25     }
26     else{
27         alert("No model is selected!");
28     }
29 }

```

Listing 28: Delete model request

```

1     function resumeProcess (processId) {
2
3     var request = new XMLHttpRequest();
4     request.open("POST", "http://localhost:9000/resume/"+processId,
5     true);
6     request.onload = function() {
7         if (request.status === 200 && request.status <= 299) {
8             var process = document.getElementById(processId);
9             process.parentNode.removeChild(process);
10            alert (request.responseText);
11        }
12        else{
13            var process = document.getElementById(processId);
14            process.parentNode.removeChild(process);
15            alert (request.responseText);
16        }
17    }
18    request.onerror = function (e) {
19        alert ("error: " +request.status);
20    };
21    request.send(null);
22 }

```

Listing 29: Resume process instance request

3 Results

For the needs of demonstrating the system's operations, two sample processes have been created in figures 19 and 20.

The first process has two activities, Calculate Value and Change Owner and also an OR gateway. Both activities are of type "Script" and their implementation can be found in Listings 30 and 31 respectively. The following start variables were also set: stock (String), amount (Long), price (Long), owner (String). The Calculate Value activity calculates the value of the stock based on the amount, the individual price and a random multiplier and stores the result in newValue. newValue is then inserted into the process variables. The OR gateway checks whether newValue is sufficiently high, and if so, the process ends. Otherwise, a new owner is randomly selected from a predefined pool of names. Lastly, both newValue and owner variables are then printed to Tomcat log and the process ends.

The second process in Figure 20 was created to demonstrate that the flexibility need Loose-ness has been satisfied. As described in previous sections, this flexibility aspect dictates that the batch model can be specified during process enactment. In Camunda, one way to show this is to create a user task before the activity that we want to batch over. When engine control reaches the user task, the rest of the process instance operation is put on hold until the a human agent completes the task. During this time, the process instance is considered active and running, providing the opportunity to define a batch model.

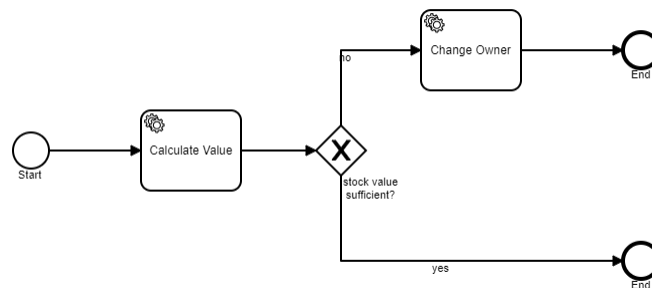


Figure 19: Sample process #1 for demonstration




Figure 20: Sample process #2 for demonstration

```
1 public class CalculateMarketValueDelegate implements JavaDelegate  
2 {
```

```

3  @Override
4  public void execute(DelegateExecution execution) throws Exception {
5
6      long amount = (long)execution.getVariable("amount");
7      long price = (long)execution.getVariable("price");
8      double random = Math.random();
9      long newPrice = (long)random*price + (price+40);
10     long newValue = amount*newPrice;
11
12     execution.setVariable("newValue", newValue);
13
14 }
15
16 }

```

Listing 30: Caclulate Value implementation

```

1     public class ChangeOwnerDelegate implements JavaDelegate {
2
3     private final Logger LOGGER = Logger.getLogger(ChangeOwnerDelegate.
4         class.getName());
5
6     @Override
7     public void execute(DelegateExecution execution) throws Exception {
8
9         String namePool[] = {"Olof Palme", "James McRandom", "Elia S.
10        Whatever", "Natalie Doe", "Albert Einstein"};
11        Random random = new Random();
12        int index = random.nextInt(5);
13        String newOwner = namePool[index];
14        execution.setVariable("owner", newOwner);
15
16        LOGGER.info("New value of stock is: "+execution.getVariable("
17        newValue"));
18        LOGGER.info("New owner is: "+execution.getVariable("owner"));
19    }
20 }

```

Listing 31: Change Owner implementation

After starting Camunda BPM platform and the Batch Controller, the Batch Management Tool looks like Figure 21. Without specifying a batch model for the above process, starting an instance as shown in Figure 22, the result is visible in Figures 23 and 24. As is evident, the process starts and finishes immediately as there is not batch activity specified. The final result in Figure 24 show that the owner changed from "test" to "Natalie Doe".

By specifying a sample batch model for this process as shown in Figure 25, and setting the minimum number of instances at 3, the countdown duration at 30 seconds and the grouping attribute to "owner", it's time to start multiple instances.

Test Case 1

As a first case, two instances are started with the same owner, let that be "random". Both are now blocked at activity "Calculate Value" and also assigned to the same cluster as they

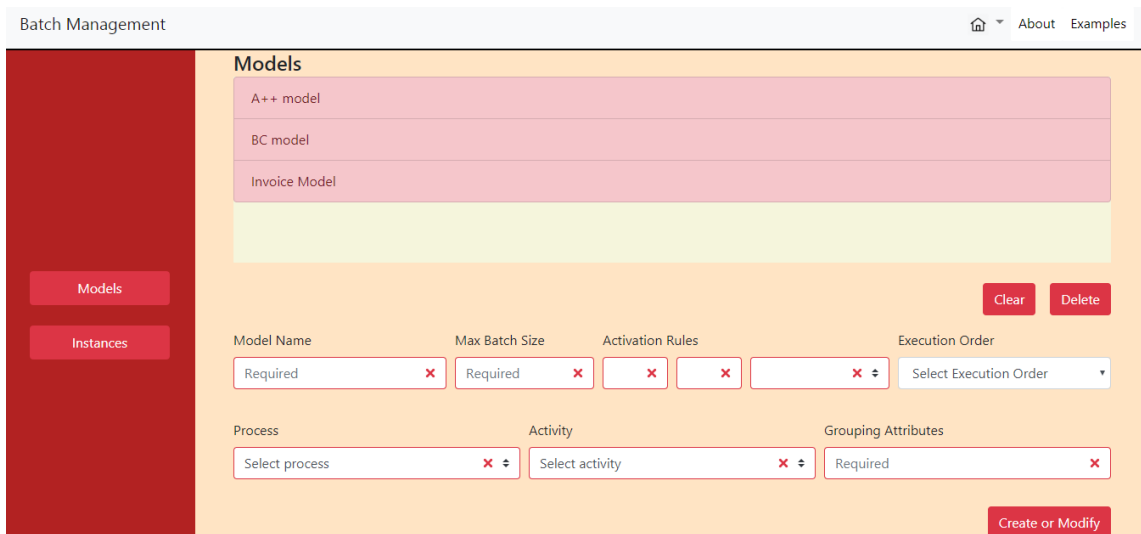


Figure 21: Batch Management Tool

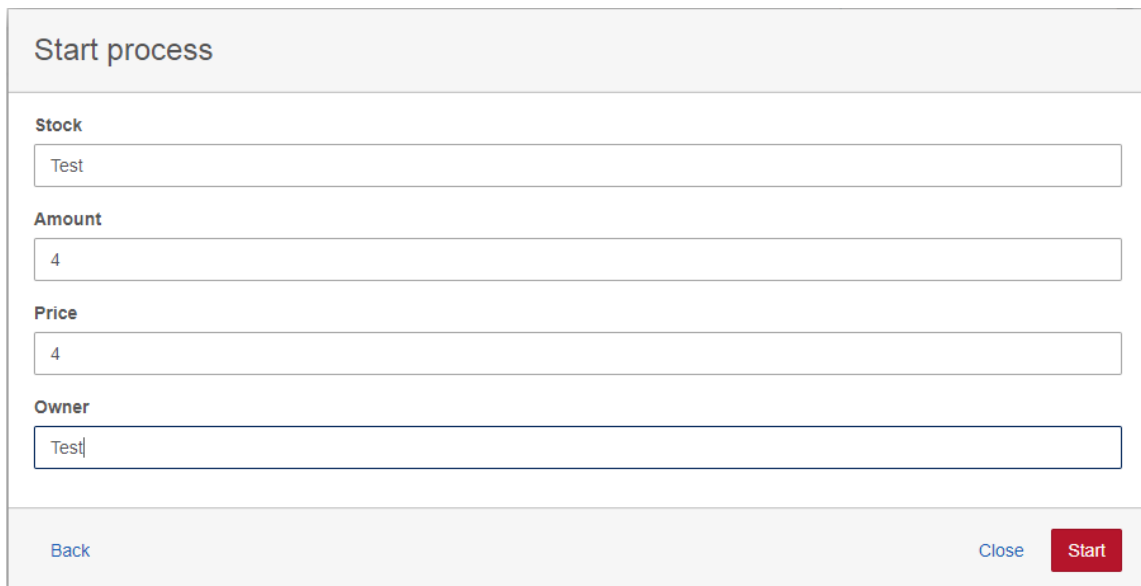


Figure 22: Starting an instance

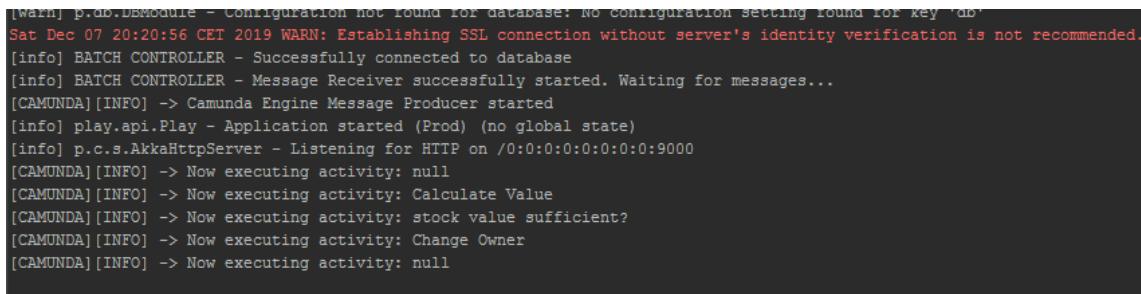


Figure 23: Result of test instance in the Batch Controller

should, as can be seen in Figure 26. Additionally, the Instances tab at the web page now looks like Figure 27. After a wait of 30 seconds without starting a third process, the batch

```

nda.bpm.engine.rest.mapper.JacksonConfigurator from Application class org.camunda
a.bpm.engine.rest.impl.application.DefaultApplication
07-Dec-2019 20:29:28.274 INFO [http-nio-8080-exec-2] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New value of stock is: 176
07-Dec-2019 20:29:28.274 INFO [http-nio-8080-exec-2] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New owner is: Natalie Doe

```

Figure 24: Result of test instance in Apache Tomcat log

is canceled and the processes sent back to Camunda, where their execution resumes at "Calculate Value" activity, as can be seen in Figures 28 and 29.

Figure 25: Creating model for process "Stock Evaluation"

```

[info] BATCH CONTROLLER - New batch with id: stock-evaluation:1/Calculate Value:random-1 started
[CAMUNDA] [INFO] -> Activity 'Calculate Value' is blocked from executing... Process Instance id is: c63312b8-192d-11ea-be87-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: null
[error] BATCH CONTROLLER - java.lang.NullPointerException
[error] BATCH CONTROLLER - java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer
[CAMUNDA] [INFO] -> Activity 'Calculate Value' is blocked from executing... Process Instance id is: cb5bdfc6-192d-11ea-be87-d4c9ef713636

```

Figure 26: Two instances blocked

Blocked Instances			
Id	Process	Activity	
c63312b8-192d-11ea-be87-d4c9ef713636	Stock Evaluation	Calculate Value	<button>Resume</button>
cb5bdfc6-192d-11ea-be87-d4c9ef713636	Stock Evaluation	Calculate Value	<button>Resume</button>

Figure 27: Blocked instances at the Instances tab

Test Case 2

By starting three instances, which is the minimum in order to execute the batch, and resuming one of them from the Instances tab of the Management Tool before the countdown ends, the instance resumes normally with rest following shortly after since the batch is

```
[info] BATCH CONTROLLER - Minimum number of instances has not been reached. Resuming individual execution...
[CAMUNDA] [EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA] [EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: cb5bdfc6-192d-11ea-be87-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: Calculate Value
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: Change Owner
[CAMUNDA] [INFO] -> Now executing activity: null
[CAMUNDA] [EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA] [EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: c63312b8-192d-11ea-be87-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: Calculate Value
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: Change Owner
[CAMUNDA] [INFO] -> Now executing activity: null
```

Figure 28: Batch Controller log showing both instances are resumed

```
07-Dec-2019 21:12:08.365 INFO [http-nio-8080-exec-1] com.camunda.batch.BatchActivity.ChangeOwnerDelegate.execute New value of stock is: 176
07-Dec-2019 21:12:08.365 INFO [http-nio-8080-exec-1] com.camunda.batch.BatchActivity.ChangeOwnerDelegate.execute New owner is: Olof Palme
07-Dec-2019 21:12:08.380 INFO [http-nio-8080-exec-1] com.camunda.batch.BatchActivity.ChangeOwnerDelegate.execute New value of stock is: 215
07-Dec-2019 21:12:08.380 INFO [http-nio-8080-exec-1] com.camunda.batch.BatchActivity.ChangeOwnerDelegate.execute New owner is: Elia S. Whatever
```

Figure 29: Apache Tomcat log showing both instances are resumed

canceled. Figure 30 shows all three instances are blocked. After resuming one of them, two are left in the batch (Figure 31) which is not enough to execute it. After 30 seconds have passed, the batch is canceled, and the remaining two are also individually executed by Camunda (Figures 32 and 33).

Test Case 3

Lastly, starting again three instances and waiting 30 seconds, the batch is executed and the instances are then returned to Camunda where the execution continues at the OR gateway as expected. The result in Figures 34 and 35.

Id	Process	Activity	
0ff03e68-19b1-11ea-af12-d4c9ef713636	Stock Evaluation	Calculate Value	Resume
13d17046-19b1-11ea-af12-d4c9ef713636	Stock Evaluation	Calculate Value	Resume
183391f4-19b1-11ea-af12-d4c9ef713636	Stock Evaluation	Calculate Value	Resume

Figure 30: All three instances are blocked

Test Case 4

For the next test case, let us start an instance of the process in Figure 20. Initially, there is no batch model for this process. Once started, the instance waits for input by a human agent as shown in Figure 36. While the instance is in this waiting state, we can create

Blocked Instances			
Id	Process	Activity	
13d17046-19b1-11ea-af12-d4c9ef713636	Stock Evaluation	Calculate Value	<button>Resume</button>
183391f4-19b1-11ea-af12-d4c9ef713636	Stock Evaluation	Calculate Value	<button>Resume</button>

Figure 31: Two instances now remain in the batch

```
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: 0ff03e68-19b1-11ea-af12-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: Calculate Value
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: Change Owner
[CAMUNDA] [INFO] -> Now executing activity: null
[info] BATCH CONTROLLER - Minimum number of instances has not been reached. Resuming individual execution...
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: 183391f4-19b1-11ea-af12-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: Calculate Value
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: null
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: 13d17046-19b1-11ea-af12-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: Calculate Value
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: Change Owner
[CAMUNDA] [INFO] -> Now executing activity: null
```

Figure 32: Resuming one instance the batch is canceled after countdown

```
a.bpm.engine.rest.impl.application.DefaultApplication
08-Dec-2019 12:51:53.394 INFO [http-nio-8080-exec-3] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New value of stock is: 220
08-Dec-2019 12:51:53.394 INFO [http-nio-8080-exec-3] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New owner is: Albert Einstein
08-Dec-2019 12:51:56.416 INFO [http-nio-8080-exec-4] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New value of stock is: 129
08-Dec-2019 12:51:56.416 INFO [http-nio-8080-exec-4] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New owner is: Albert Einstein
08-Dec-2019 12:52:30.718 INFO [http-nio-8080-exec-9] org.jboss.resteasy.logging.
impl.Slf4jLogger.info Deploying javax.ws.rs.core.Application: class org.camunda.
bpm.cockpit.impl.web.CockpitApplication
```

Figure 33: Result of execution on Tomcat

a batch model as shown in Figure 37, similarly to the previous cases. Once the model is created, the instance may continue the operation by pushing the Complete button in Figure 36. The instance now blocks, as expected, at the Web Service activity that was specified in the batch model, shown in Figure 38.

Test Case 5

For the last case, we shall create a batch model in Figure 39. Then we will start 2 instances of the Stock Evaluation process. The cluster that has been created will wait for 10 hours before it is executed. In the meantime, the model will be modified from a maximum number of instances 4 to 3. Then after starting a third instance, we can see in Figure ?? that the batch has started execution immediately, as expected.

```

[error] BATCH CONTROLLER - java.lang.NullPointerException
[error] BATCH CONTROLLER - java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer
[info] BATCH CONTROLLER - New batch with id: stock-evaluation:1/Calculate Value:dim-1 started
[CAMUNDA] [INFO] -> Activity 'Calculate Value' is blocked from executing... Process Instance id is: 13456035-19b5-11ea-a008-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: null
[error] BATCH CONTROLLER - java.lang.NullPointerException
[error] BATCH CONTROLLER - java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer
[CAMUNDA] [INFO] -> Activity 'Calculate Value' is blocked from executing... Process Instance id is: 16fc74d3-19b5-11ea-a008-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: null
[CAMUNDA] [INFO] -> Activity 'Calculate Value' is blocked from executing... Process Instance id is: 1a6dcd1-19b5-11ea-a008-d4c9ef713636
[error] BATCH CONTROLLER - java.lang.NullPointerException
[error] BATCH CONTROLLER - java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer
[info] BATCH CONTROLLER - Execution for Batch stock-evaluation:1/Calculate Value:dim-1 has started
[info] BATCH CONTROLLER - Execution for Batch stock-evaluation:1/Calculate Value:dim-1 has finished
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: 1a6dcd1-19b5-11ea-a008-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: Change Owner
[CAMUNDA] [INFO] -> Now executing activity: null
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: 13456035-19b5-11ea-a008-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: Change Owner
[CAMUNDA] [INFO] -> Now executing activity: null
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA] [EXCEPTION] -> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA] [INFO] -> Resuming process instance with id: 16fc74d3-19b5-11ea-a008-d4c9ef713636
[CAMUNDA] [INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA] [INFO] -> Now executing activity: Change Owner
[CAMUNDA] [INFO] -> Now executing activity: null

```

Figure 34: The batch is executed

```

ss org.camunda.bpm.engine.rest.impl.application.DefaultApplication
08-Dec-2019 13:20:41.754 INFO [http-nio-8080-exec-11] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New value of stock is: 165
08-Dec-2019 13:20:41.754 INFO [http-nio-8080-exec-11] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New owner is: Natalie Doe
08-Dec-2019 13:20:41.785 INFO [http-nio-8080-exec-11] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New value of stock is: 120
08-Dec-2019 13:20:41.785 INFO [http-nio-8080-exec-11] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New owner is: Albert Einstein
08-Dec-2019 13:20:41.801 INFO [http-nio-8080-exec-11] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New value of stock is: 124
08-Dec-2019 13:20:41.801 INFO [http-nio-8080-exec-11] com.camunda.batch.BatchActi
vity.ChangeOwnerDelegate.execute New owner is: Elia S. Whatever
08-Dec-2019 13:21:26.501 INFO [http-nio-8080-exec-7] org.jboss.resteasy.logging.
impl.Slf4jLogger.info Deploying javax.ws.rs.core.Application: class org.camunda.
bpm.cockpit.impl.web.CockpitApplication

```

Figure 35: Result of executed batch on Tomcat

< > ↗
Add Comment +

Approve

UserTaskBatch 🔗

📅 Set follow-up date
🕒 Set due date
👤 Add groups
👤 Demo Demo ×

Form
History
Diagram
Description

Approve?

Save
Complete

Figure 36: User task input

Clear Delete

Model Name: User model ✓ Max Batch Size: 2 ✓ Activation Rules: 2 ✓ 10 ✓ SECONDS ✓ Execution Order: Parallel

Process: UserTaskBatch v1 ✓ Activity: Web Service ✓ Grouping Attributes: test ✓

Create or Modify

Figure 37: Batch model for UserTaskBatch process

```
[info] play.api.Play - Application started (Prod) (no global state)
[info] p.c.s.AkkaHttpServer - Listening for HTTP on /0:0:0:0:0:0:0:0:9000
[CAMUNDA][INFO] -> Now executing activity: null
[CAMUNDA][INFO] -> Now executing activity: Approve
[CAMUNDA][INFO] -> Activity 'Web Service' is blocked from executing... Process Instance id is: 3507b53e-354e-11ea-baa4-d4c9ef713636
[info] BATCH CONTROLLER - New batch with id: UserTaskBatch:1/Web Service:null-1 started
```

Figure 38: Instance is blocked

Clear Delete

Model Name: Stock model ✓ Max Batch Size: 4 ✓ Activation Rules: 2 ✓ 10 ✓ HOURS ✓ Execution Order: Parallel

Process: Stock Evaluation v1 ✓ Activity: Calculate Value ✓ Grouping Attributes: owner ✓

Create or Modify

Figure 39: Batch model for case 5

```
[CAMUNDA][INFO] -> Activity 'Calculate value' is blocked from executing... Process Instance id is: 6ebbb5ae-3551-11ea-baa4-d4c9ef713636
[CAMUNDA][INFO] -> Now executing activity: null
[CAMUNDA][INFO] -> Activity 'Calculate Value' is blocked from executing... Process Instance id is: 72c7ef2c-3551-11ea-baa4-d4c9ef713636
[error] BATCH CONTROLLER - java.lang.NullPointerException
[error] BATCH CONTROLLER - java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer
[CAMUNDA][INFO] -> Now executing activity: null
[error] BATCH CONTROLLER - java.lang.NullPointerException
[error] BATCH CONTROLLER - java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer
[info] BATCH CONTROLLER - Execution for Batch stock-evaluation:1/Calculate Value:dim-1 has started
[info] BATCH CONTROLLER - Execution for Batch stock-evaluation:1/Calculate Value:dim-1 has finished
[CAMUNDA][INFO] -> Activity 'Calculate Value' is blocked from executing... Process Instance id is: af43071a-3551-11ea-baa4-d4c9ef713636
[CAMUNDA][EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA][EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA][INFO] -> Resuming process instance with id: af43071a-3551-11ea-baa4-d4c9ef713636
[CAMUNDA][INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA][INFO] -> Now executing activity: Change Owner
[CAMUNDA][INFO] -> Now executing activity: null
[CAMUNDA][EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA][EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA][INFO] -> Resuming process instance with id: 72c7ef2c-3551-11ea-baa4-d4c9ef713636
[CAMUNDA][INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA][INFO] -> Now executing activity: Change Owner
[CAMUNDA][INFO] -> Now executing activity: null
[CAMUNDA][EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.NullPointerException'
[CAMUNDA][EXCEPTION]-> 'org.camunda.bpm.engine.adapter.BatchEntity: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.Integer'
[CAMUNDA][INFO] -> Resuming process instance with id: 6ebbb5ae-3551-11ea-baa4-d4c9ef713636
[CAMUNDA][INFO] -> Now executing activity: stock value sufficient?
[CAMUNDA][INFO] -> Now executing activity: Change Owner
[CAMUNDA][INFO] -> Now executing activity: null
```

Figure 40: Batch executed immediately after after reducing the maximum number of instances

4 Discussion

In the beginning of Section III, the four requirements that were defined, essentially dictate the measure of success of this project. It was also stated early in this work, that a flexible approach for batch processing has to adhere to two sources: the work of Karastoyanova et al. [2] describing the most flexible strategy, and the nature of batch processing itself. The architecture and implementation presented in sections II and III were done having these requirements in mind. Table 3 shows whether the requirements were fulfilled and also which functionality was directly responsible.

Due to the architecture itself, the batch model and the process are completely independent from one another, which satisfies the Evolution requirement. Updating either of them does not affect the other in any way, since the updating the model only requires the Batch Management Tool, the Batch Controller and MySQL database. Updating a process only requires the Camunda Modeller and the Process Engine for deployment.

The Looseness property is also satisfied by the fact that because of the architecture, the batch model and the process are independent. It was shown early in the Results section that a batch model can be specified after a process has been instantiated. In the subsequent demonstration, the next instance that was started was blocked in the activity specified in the batch model. Although it was not explicitly exhibited, the same instance would be blocked even if the model was created (or update) after the instance had been started. This is because, the ActivityImpl object described in Section III is common for all instances and as a result, setting its flag to true would affect all future instances of that process. Coincidentally, this also satisfies the Adaptation of Batch Activity requirement since the batch model can be updated during runtime by specifying an activity to batch over. The flag of the activity is then set to true which now makes it a batch activity.

The process interruption was showcased in the Implementation section with the introduction of the blocking mechanism. While the mechanism succeeds in its basic function and also does not negatively affect the overall execution of the process, it may introduce bugs when taking into consideration the general workflow of Camunda. One such example is the fact that when a process resumes after being interrupted in activity A, the Cockpit application is not updated and continues to show the instance at activity A. This potentially means that traces of the instance have been left behind after it has finished and could interfere with other operations of Camunda. Furthermore, the blocking mechanism has only been implemented on simple process, meaning, a process that does not have concurrent activities (activities that are executed concurrently in different threads). This would require a different approach, in order to ensure that no zombie threads are left behind and that overall process execution is not distorted because of parent-child thread mismanagement.

The process aggregation in batches and execution seems to be working correctly, with process of the same type rounded up in batches based on the grouping attribute specified in the batch model. Some properties of the model can be updated in real-time and affect the batches referring to that model. This exempts the process and activity properties, and also the timer, since once the countdown has started it will not be possible to change the delay

of the Timer object. Execution is parallel and sequential is also working, with minimal thread management that waits for all instances to finish execution before aggregating them again in a list for Camunda. As sufficiently shown in the Result section, the processes can continue their execution from the next activity in line with new variable values specified from the Batch Controller.

For increased flexibility reasons, if the batch is canceled the instances it contains must be sent back to Camunda and resume execution at the same activity they were supposed to batch over. This is indeed the case as it has been shown in the Results in Figures 32 and 33. As an additional flexible aspect, the Batch Management Tool provides a list of the blocked process with the possibility of resuming one more of them. This may obviously affect the operation of a batch as shown in the above images.

Requirement	Fulfilled	Implementation
Evolution	✓	Update model methods in Batch Controller
Looseness	✓	Process Engine modification
Process Adaptation	✓	Architecture
Adaptation of Batch Activity	✓	Update model methods in Batch Controller
Process interruption	✓	Process Engine modification
Batch Execution	✓	Batch Controller implementation
Canceled batch and process instances execution	✓	Batch Controller implementation

Table 3: Requirements

In general, the implementation of batch processing has been applied to a specific type of activity, the Service Task type which in the Camunda system is supposed to simulate a web service call. In real-life operations, batching is applied to few tasks the most important of which is the user task and the web service interaction. Considering the first is implemented in the work of Dr. Luise Pufahl [1], this work undertook the second option, with very promising results. Due to time constraints this flexible approach could not be applied to the user task, which can be the subject of a future work. Moreover, batch processing cannot be applied to a group of activities, as of yet. This means that a batch model can only be assigned to a single activity. This may count as a reduction in flexibility, as the work of Karastoyanova et. al clearly states that a batch should also apply to a group of activities, effectively forming a batch region. Unfortunately the concept of a batch region is not applicable to this project.

Additionally, given the current implementation of the Batch Controller, the Java script that is supposed to be executed and is evaluated during runtime, must be hardcoded in classpath of the project. Again due to time constraints, a Class loader could not be implemented in order to load classes in real-time as happens in Camunda. However this is only relevant when developing the software for use in actual production and does not affect elements of batch processing.

In a brief mention earlier in the section, it was reported that the blocking mechanism in Camunda leaves traces of finished instances. Given the current implementation of certain operations in the system (Process Engine), certain aspects of the mechanism may not yield the most efficient solution, in terms of memory management.

5 Conclusion

The aim of this project was to provide a concrete implementation of the most flexible approach in batch processing as described in Karastoyanova et. al [2]. An architecture was provided to support the implementation, with three major components, the BPM Engine, the Batch Controller and the User Interface, each having been implemented from scratch or modified. For determining the success of the project, six requirements were identified, stemming from the most flexible approach and the batch processing theory. The results after running several test cases have showed that all requirements have been satisfied. In conclusion, flexible batch processing on Camunda BPM is definitely possible and with a high degree of efficiency.

I Future Work

The promising results of the project also provide an opportunity for further applications in batch processing, but also additional modifications that could enhance the efficiency and usability in the context of the current implementation. The following list contains suggestions on how to proceed with future work on the project.

- **Enable batch processing in a region.** This project has achieved batching on a single activity at a time but existing research suggests that it can also be applied to a group of activities.
- **Enable batch processing for various activity types.** This work applies batching explicitly on activities of type "Script" which is executes Java code. However, it should also apply to activities that are undertaken by human agents (User Tasks). This has been successfully implemented by dr. Luise Pufahl in [1]. A combination of both projects would be very interesting and challenging.
- **Update Cockpit application to propagate changes occurring after batching.** The Camunda Cockpit app provides information on the state of every process instance. Unfortunately, after the modification, that information is no longer displayed correctly.
- **Implementation of multiple attribute grouping.** Grouping the instances in batches is currently realized using only one attribute.
- **Improvement of the blocking/restart mechanism.** The blocking/restart mechanism of the engine modification probably can be implemented more efficiently to avoid memory overhead.
- **Enable batch processing on concurrent activities.** This implementation of batch processing does not support activities running concurrently on different threads on the engine. This area might prove to be challenging given the different handling it requires.

Acknowledgements

I would like to give thanks to pr. dr. Dimka Karastoyanova from the University of Groningen for giving me the opportunity to work on such a challenging and thought provoking project, as well as for providing the necessary theoretical background.

I would also like to thank dr. Luise Pufahl from the University of Potsdam for her important inputs and suggestions during this period.

Lastly, I would like to thank my friend Angeliki, for giving me support.

References

- [1] Luise Pufahl. *Modeling and enacting batch activities in business processes*. PhD thesis, University of Potsdam, 2018.
- [2] Luise Pufahl and Dimka Karastoyanova. Enhancing business process flexibility by flexible batch processing. In *OTM Conferences*, 2018.
- [3] Camunda docs user guide. <https://docs.camunda.org/stable/guides/user-guide/>. [Online; accessed 23-August-2019].
- [4] M. Dumas, M. La Rosa, J. Mendling, and H. A. Reijers. *Fundamentals of Business Process Management, Volume 1*. Springer, 2013.
- [5] N. Slack, S. Chambers, and R. Johnston. *Operations and Process Management: principles and practice for strategic impact*. Pearson Education, 2009.
- [6] Luise Pufahl and Mathias Weske. Batch activity: enhancing business process modeling and enactment with batch processing. *Computing*, 04 2019.
- [7] J. Liu and J. Hu. *Dynamic batch processing in workflows: Model and Implementation*. Pearson Education Future Generation Computer Systems, 2007.
- [8] Enabling batch processing in bpmn processes. <https://bpt.hpi.uni-potsdam.de/Public/BatchProcessing>. [Online; accessed 20-July-2019].
- [9] Business process model and notation. https://en.wikipedia.org/wiki/Business_Process_Model_and_Notation. [Online; accessed 20-July-2019].