UNIVERSITY OF TARTU
Institute of Computer Science
Software Engineering Curriculum

Luis Gerson Noboa Martillo

Holograph: A Tool for Assessing the Impact of Resource Assignment on Business Process Performance Based on Event Logs

Master's Thesis (30 ECTS)

Supervisor:
Marlon Dumas

Tartu 2018

Holograph: A Tool for Assessing the Impact of Resource Assignment on Business Process Performance Based on Event Logs

Abstract:
This thesis aims to identify if the variations in the performance of a business process can be explained by the resource allocation observed in an event log. This aim is pursued by closely analyzing the logs produced by a process-aware information system. The approach addresses the common problems in this area, such as overcomplicated, hard to understand output, or tools that are not specialized for the task, by building a method that considers factors such as individual performance versus group performance, the moments in which the resources are involved, and the variants in which they take part of. Through this method, the goal is to obtain a result that is meaningful from different business points of view and helps answer questions such as: which resources are better suited for a given task? Which groups of resources work together in the most efficient way? In order to evaluate the benefits and usefulness of the approach, a web application called Holograph was implemented using the proposed guidelines. The approach was validated via an experiment involving a group of IT management students with prior knowledge of process mining.

Holograph: Sündmuste logisid kasutav rakendus hindamaks ressursi jaotamise mõju äriprotsesside efektiivsusele

Lühikokkuvõte:
Antud magistritöö eesmärgiks on välja selgitada, kas sündmuste logidest saab välja lugeda ressursside kasutamise mõju äriprotsessidele. Uuringu läbiviimiseks analüüsiti protsessiteadlikke infosüsteemide genereeritud logisid. Töös tuuakse välja enamlevinud probleeid antud valdkonnas: liiga keerulise lähenemise, keerulised logid või ülesande lahendamiseks ebasobilikud rakendused. Antud töös arvestatakse järgnevaid faktoreid: individuaalne efektiivsus võrreldes grupiefektiivsusega ning millal milliseid ressursse kasutada on vaja. Kasutatava meetodi eesmärk on luua tulemus, mis aitab vastata küsimustele erinevates äriaspektides: Milliseid ressursse on parim antud ülesande lahendamiseks kasutada? Millised ressursigrupid töötavad koos kõige efektiivsemalt? Eelmainitud meetodi kasulikkuse mõõtmiseks loodi rakendus nimega Holograph, mis rakendas töös välja toodud põhimõtteid. Meetodi valideerisid eelnevalt protsessikaevandamisega kokku puutunud infotehnoloogia juhtimise üliõpilased.

1

TABLE OF CONTENTS

3

1 INTRODUCTION

A business process is a "collection of inter-related events, activities, and decision points that involve a number of actors and objects, which collectively lead to an outcome that is of value to a customer" (Dumas, Rosa, Mendling, & Reijers, 2018). Typical examples of business processes include order-to-cash processes (a process that starts when a purchase order has been received and ends when the product or service has been delivered and the payment has been received) or fault-to-resolution processes (a process that starts when a fault is reported by a customer and ends with its resolution).

An organization can outperform another organization offering similar services by optimizing their business processes (Dumas, Rosa, Mendling, & Reijers, 2018). The way in which a business process is designed and executed affects its quality, efficiency, and its compliance with regulations. Business Process Management (BPM) is "a body of methods and tools to identify, discover, analyze, redesign, execute, and monitor business processes in order to improve their performance" (Dumas, Rosa, Mendling, & Reijers, 2018).

BPM encompasses a wide range of methods and techniques. In this thesis, we focus on one emerging family of methods in the field of BPM, namely process mining.

1.1 PROCESS MINING

With the extensive amount of data that is collected by modern software solutions, the field of process mining has been gaining traction at an accelerated phase. It is a set of methods that deal with "the analysis of event records" that are generated by the execution of business processes (Dumas & García-Bañuelos, 2015). Some of the most common results include a process map, a summary of the paths that have been traversed more frequently, a graphical comparison between two related processes, or simple text that explains the differences between processes that have been compared.

This output is then used to identify patterns or behaviors that the process displays and is used accordingly to improve the business area being analyzed (Wynn, et al., 2017). As companies are interested in identifying issues that affect their operations, such as bottlenecks and other sources of delays and inefficiencies, process mining techniques can be hugely beneficial. Stakeholders can use the results obtained from this analysis and redesign processes in order to use resources more efficiently, reduce execution times, and save money.

According to Dumas et al., the different process mining techniques used at the moment by researchers can be grouped into four different categories. An overview of them can be seen in the picture below.

4
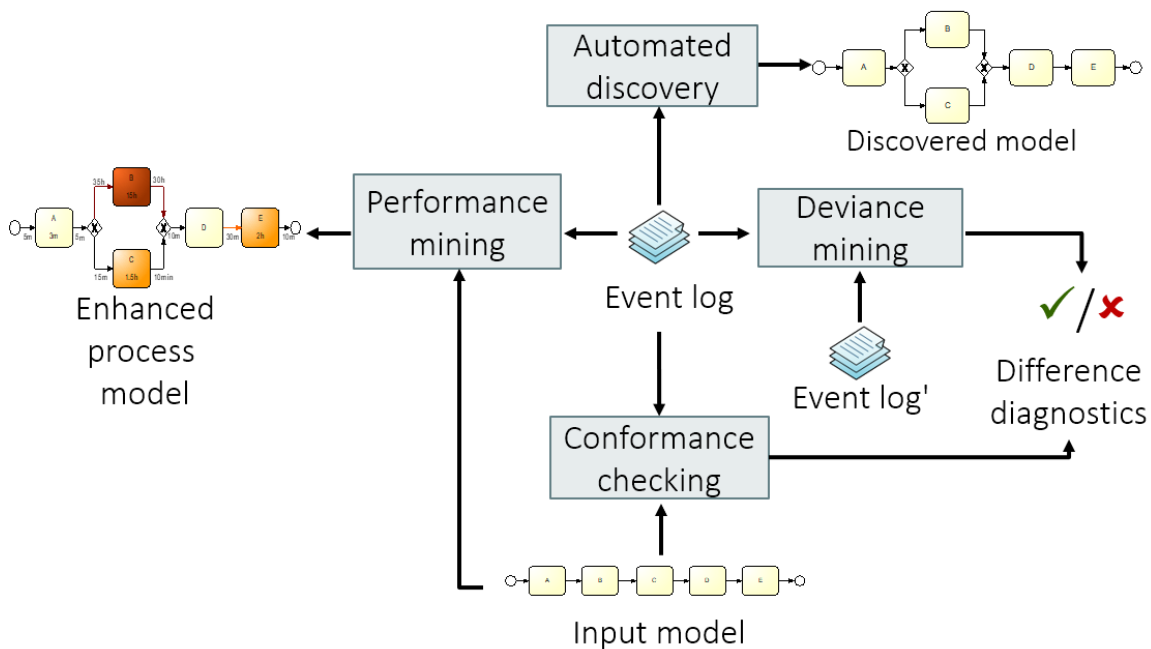
Figure 1. The four categories of process mining techniques (Dumas, Rosa, Mendling, & Reijers, 2018)

The first one is automated process discovery. This analysis accepts an event log as an input and "they automatically produce a business process model" that tries to approximate what is happening in said event log (Dumas, Rosa, Mendling, & Reijers, 2018). This is what Disco and Minit (explained later) do on their process map screen, where a graph which shows the relation between different activities inside an event log is drawn.

The second category is conformance checking. Even though this technique also accepts an event log as an input, conformance checking tools require a process model as well. The tool should check how similar the process model is (the theoretical way of executing a process) to the event log (the actual way of executing said process). This can lead to errors being discovered in the process model, errors in the way the process is executed, or other inconsistencies.

The third category is performance mining. This analysis also takes an event log and a process model and outputs a process model that has annotations that indicate performance numbers. Examples of these annotations include time, bottlenecks, and efficiency. Disco and Minit also do this by selecting the performance visualization on their interfaces, and they color-code their process maps to identify possible bottlenecks.

The final category is variant analysis. Tools that follow this technique normally take two event logs as inputs, "and produce as output a list of differences" (Dumas, Rosa, Mendling, & Reijers, 2018). The most common case for this analysis is having two event logs, one with a positive outcome and one with a negative outcome, and the business wants to know which factors lead to a case to be positive or negative.

5

The goal of this thesis is to propose a new methodology for analyzing, through event logs, the impact of resources in the execution of a business process. It is based on the performance mining category, although the proposal works by adding an event log only, as opposed as requiring an event log and a process model. This new method will be able to accurately describe the impact that a resource (or a group of resources) have on the overall performance of a process, and how the combination of resources or activities are able to impact in a positive way the processing time of a business operation.

1.2 MOTIVATION AND PROBLEM STATEMENT

Businesses are always interested in the performance of their employees. Whether it is for analyzing pay rises, promotions, restructuring of operations, or bonuses, companies need to see how individuals work. Also, analyzing the performance of employees is key for human resource management, since "it allows making objective and timely management decisions" (Stojadinović, Marinković, & Ivković, 2014).

However, as employees rarely work by themselves, but rather are part of a bigger group or process, identifying the performance of a sole worker is just one side of the story. How employees affect and are affected by other people with whom they collaborate is vital in order to understand the intricacies of a business process and to find optimizations.

Normally, the way that performance is analyzed is through the use of key performance indicators, commonly abbreviated to KPI. Key performance indicators "are used to measure the performance achievement of systems" (Springer London, 2010). It is "a way of measuring a company's progress" towards a goal that it has set for itself (Cambridge Dictionary). Through the use of KPIs, one can quickly glance at how projects or processes are going versus how they are supposed to go. They are a great way of assessing how close an endeavor is to the objectives that have been planned for it.

KPIs are normally used as a global measurement. That is, they are used to analyze the results of a project or process as a whole, rather than the contribution each individual has made to the KPI. While there are methods of implementing KPIs in a way that analyzes individual contribution, they are less used and very parametrized. While this is good to evaluate certain criteria that might be needed from the business, it misses key aspects, such as the impact that a group has on an individual user, what tasks does the resource excel at, and how the involvement in the process affects the output of the employee.

Also, their usage is sometimes mishandled and doesn't lead to better outcomes. As Parmenter writes, there are a lot of cases of business that have operated with KPIs in mind in order to increase their output, but "have found the KPIs made little or no difference to performance" (Parmenter, 2015). This is due to the organization not fully understanding the underlying problems with the business, and have failed to identify what are the factors that are truly impacting the execution of their processes.

With the rise of data mining and all of its applications, there has been an increase in the information that businesses can extract from their operations. Tools like Disco and Minit offer its customers ways of mining the information in event logs. This way, businesses can get visual representations of how their operations are conducted, who gets involved, and what is the performance of everyone involved.

These tools are primarily targeted at analyzing the execution of activities. Both offer a quick way of generating a directly-follows graph with the activities of each process, how they are interconnected, and what are the bottlenecks in the process. The user can switch between mining frequency or performance information, providing a very quick glance of the activities that are executed the most, paths that are barely executed, and where the biggest gains in time can be made.

Also, both Disco and Minit offer ways of analyzing resource information in different ways. If the user selects a resource field when importing the data log, resource information is computed for further analysis. In Disco, you can get information about frequency, median duration, mean duration, and duration range, among others.

Figure 2. Disco's process discovery screen

Minit goes further than Disco and adds a social network graph that shows how the resources give and receive information from other resources involved in the process. Its statistics are also more in-depth, since they deal with active times, waiting times, standard deviations, and influence.
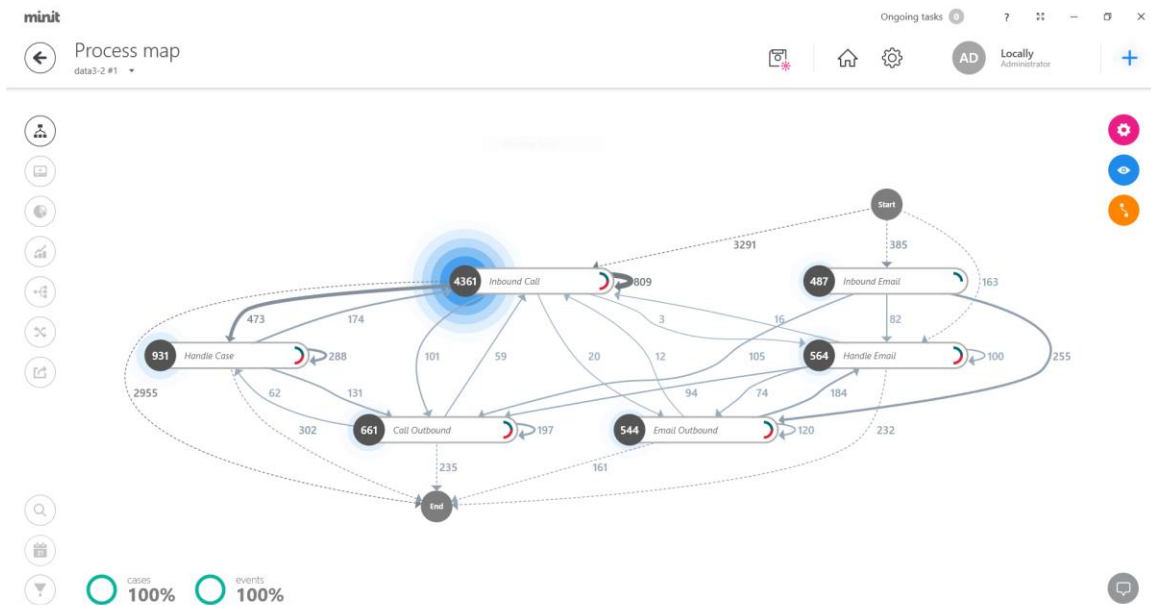
Figure 3. Minit's process discovery screen

However, these tools are not enough to explain how a resource impacts, either positively or negatively, the execution of a process. If a company wants to determine where resources are better placed, with which other employees should they work, or in what activities does a resource excel, then these tools are lacking in information.

Therefore, the main research problem that needs to be solved is identifying if the variations of the performance in a business process can be explained by the resource allocation observed in a given event log. The solution to this problem should be able to classify and show users information about how a resource impacts the completion of a process and should offer a better set of tools for this purpose than the commercial products available right now. Also, the following three research questions have been formulated based on the previously stated problem:

- RQ1. What are the current methods for analyzing resource impact on business processes and what are their flaws?
- RQ2. Given an event log, what kind of analysis needs to be done to it in order to discover the impact of a resource in the execution of an event log?
- RQ3. What improvements do the proposed method offer compared to the current methods?

1.3 CONTRIBUTION

The first contribution of this thesis comes in the form of the state of the art about different approaches to use process mining to analyze resource performance. Through this literature review, it was determined that researchers have come up with different solutions in order to evaluate employee performance based on the logs that companies collect on the execution of their business processes. However, these tools rarely have been built specifically for the analysis of resource impact, which makes them less than ideal for solving the previously stated problem.

8

The second contribution is the method of analyzing resource information on an event log. Several different components of the method are proposed in order to analyze resources from different perspectives. There are descriptions for the analysis of individual resources based on active time, variant, and flows. Also, an individual analysis is proposed that bases its principles in the different pitchers used for baseball games. Finally, there is a group performance method that describes how to analyze the behavior of resources within a group and how their involvement affects and is affected by the output of others.

This new method goes beyond existing ones by considering the resource perspective of a process (i.e. who executes which task) instead of the activities only. Through this approach, we are able to identify tendencies within the resources that directly affect the performance of the process, such as the best resources for specific tasks, the groups that work the best together, and the most suitable stage in which to involve specific resources.

The third contribution is Holograph, a web tool that implements almost all aspects of the proposed method. Holograph has a front and back-end that work together in order to visualize what kind of data can be extracted when implementing the method. It lets users upload a CSV file that follow certain criteria, and then generates graphs and numbers that indicate different parameters that can be used when taking decisions about a process.

1.4 STRUCTURE OF THE THESIS

The thesis is structured as follows. Section 2 talks about the previous work done by researchers to address the subject of analyzing resource performance and impact on a business process. Section 3 introduces the resource performance analysis approach, while Section 4 talks about Holograph, a partial implementation of the approach used to validate the usefulness of the proposed solution. Section 5 delves into the evaluation methods for the results of the thesis. Section 6 discusses the results obtained after the experiments. Section 7 concludes the thesis based on the results and outlines possible future work. References used for this work can be found in Section 8.

2   STATE OF THE ART

As this topic is of great interest for businesses, the analysis of resource performance is a subject that has been researched extensively. Previously, these tasks were traditionally done using different sets of tools that were not built with this kind of analysis in mind, which were described as "time-consuming, tedious and error-prone" (Wynn, et al., 2017). The main reasons have been the fact that they required a lot of knowledge of the tools used, and have to be done by experts in the area, since the output is normally too complicated to understand for business analysts or stakeholders.

2.1   EVENT LOGS

As technology aids or takes over the execution of business processes, the systems used are able to generate more information regarding the implementation of it. From the activity that is being executed, to the exact dates of completion, to the resource carrying out the task, these logs are full of data that can be used later to obtain valuable information. These logs are called event logs. According to Dumas, an event log is a "collection of timestamp event records" (Dumas, Process Monitoring, 2016).

Even though they normally carry a lot of information, a log can have at least three parameters in order to be considered useful:

- Trace ID
- Activity
- Timestamp

With this information, a log can be used for process mining techniques. Each of the entries in the log represents a step within an instance of a process. By getting all the entries with the same trace ID, one can figure out how the trace develops. With this, a process mining application can draw a map in order to visualize how the process goes.

According to Bolt, van der Aalst & de Leoni, a trace within a business process "records the execution of an instance of a process" (Bolt, van der Aalst, & de Leoni, 2017). A process recorded in an event log has a determined amount of possible activities. Each row signals the "start, end, abortion, or other relevant state change" within the process being executed (Dumas & García-Bañuelos, 2015). Also, the timestamp lets us get information about the time taken for each activity. It can also have other additional information that might be useful for the evaluation of a log. This information includes initial and ending timestamps, resources, and cost.

When a set of traces are identified as being equal to one another within an event log, the name "process variant" is used. In order for a trace to be classified as belonging to a variant, it should be very similar to the variant being analyzed, but also significantly different from other variants present in the process log.

There are several formats for storing logs, such as CSV, XES, and MXML. While each has its own advantages and disadvantages, the current implementation of Holograph (discussed later) works by uploading a CSV exclusively.

2.2   COMPARISON BASED ON PROCESS METRICS

In his work "A Visual Approach to Spot Statistically-Significant Differences in Event Logs Based on Process Metrics," Bolt, Leoni and van der Aalst expose that one of the most glaring mistakes that previous researchers have done when presenting their process is assume that the execution of a given process within an organization is always done in a similar way (Bolt, Leoni, & van der Aalst, 2016).

The truth is that not all processes are done in a same way. People might do the same tasks in different order or in a slightly different way but reach the same goal at the end. This is exacerbated in the cases of big organizations that have several branches in different geographical locations. Factors such as personnel, training, and even culture can affect the way the same process is executed by different branches and different people. This means that comparing processes through normal means gets more complicated, as more factors can skew results in ways that might not be desirable and would render the analysis less useful or even useless.

That is why Bolt presented a method that could compare different variants of a process "by identifying statistically significant differences" (Bolt, Leoni, & van der Aalst, 2016). The method highlights the nodes and edges that are more relevant in the process map according to its frequency. If the border of an element is thicker, it means that it is more relevant to the process than other elements. The states that are not significant enough for the process are not represented or their borders are very thin.

This analysis made based on the significance of an event means that two different processes recorded by different information systems can be compared in a consistent manner. Although this method has been engineered to compare only two processes at a time, a workaround to analyze more than two processes simultaneously is to merge the logs of two or more different business processes into one file and the compare them to another.

The method works by executing three steps. First, an annotated transition system is created based on the logs. According to the author, an annotated transition system is a structure composed of states, activities, and transitions between states, to which annotations are added (Bolt, Leoni, & van der Aalst, 2016). Then, the annotations are compared between each other in order to determine which states are relevant for the process. Finally, with all of this information, a visualization is created so that a user can effectively analyze the results generated by the tool.

The researchers developed this tool as a plugin called Process Comparator for the ProM framework. To make it easier, when analyzing multiple processes, the tool suggests similarity scores so that a user that is not knowledgeable enough in the area can easily obtain results. The plugin also lets users change different parameters to filter states and get the most relevant results based on the analysis that needs to be conducted.

2.3   PROCESS VARIANTS ANALYSIS TO UNDERSTAND KPI DIFFERENCES

In the article "Analyzing Process Variants to Understand Differences in Key Performance Indices," Pai Ballambettu arguments that with bigger organizations having to adhere to strict Service Level Agreements, there is an increased interest in monitoring the performance of the organization's workflows by evaluation of performance indexes, cost,

11

and compliance of requirements (Pai Ballambettu, Agumbe Suresh, & Chandra Bose, 2017).

Since some clients might have different needs, and some industries might have different regulations that have to be fulfilled, a department can have different workflows for the same underlying process. This leads to process variants that slow down the process or reduce the score in KPIs. Thus, the authors propose a solution to compare two or more business variants "to identify key differences among them" (Pai Ballambettu, Agumbe Suresh, & Chandra Bose, 2017).

In their approach, the developers consider the differences between process not only by the tasks present in them, but also taking in consideration factors such as frequency and flow times. Frequency counts the number of times that a process goes through a defined state, while time can measure the time that a process spends in a task, the duration of the whole process, or the time that the process spends going from one activity to the other.

The process for comparing different processes is done in five different steps. First, the tool generates a unified process map that contains all of the different variations of the process. Then, the solution detects the differences between each process map by calculating the absolute difference, the relative difference or the statistical difference. Afterwards, a difference process map is generated. This is just a process map that has annotations in it so that the differences calculated before are easily accessible. The last two steps are to compute cascaded components, and to visualize the results obtained by the process. This implementation has been made available as a plugin for ProM.

2.4   CAUSAL FACTORS OF PROCESS PERFORMANCE VARIATION

Holmes et al. have proposed a method for discovering the factors that cause different factors that affect business performance. This method is comprised of three steps and its goal is to generate "a graph of causal factors of process performance" (Hompes, et al., 2017). The approach tries to relate specific characteristics of the business to the performance of the process it executes.

This approach first decomposes the whole event log into what it calls a decomposition graph. This graph is a series of events that are related to each other, such as all the activities with the same name, all the activities done by a specific resource, etc. It can also be made as a combination of several factors.

Afterwards, the decomposition graph generated in the first step is again processed in order to create a so-called inclusion graph. This graph already holds all of the causes that could affect the performance of the process. In order to improve performance, some cleaning and streamlining is done on the graph before doing any further processing on it. However, while some of them have a logical structure that is easy to automate, a part of this pruning process is very domain-knowledge specific, making it complicated to automate and turning it in a manual endeavor.

Finally, a causality graph is generated from the inclusion graph. The nodes in the inclusion graph, called causes, are checked for their effect in the overall process. This is done through a Granger causality test. This approach was implemented in ProM, and a case study was done for it using a data set of an insurance provider in Australia.

2.5   RESOURCE BEHAVIOR ANALYSIS

One of the works closest to this thesis is summarized in the work "An Extensible Framework for Analysing Resource Behavior Using Event Logs." Written by Pika et al., this proposal aims to aid organizations "to extract knowledge about the behavior of their employees" (Pika, et al., 2014). This approach also analyzes how the behavior of specific resources changes over time. As the previous analysis, this one also includes three different steps in order to fulfill its purpose.

The first step is to define what the paper calls resource behavior indicators (RBI). These indicators are not set in stone and can vary according to the analysis that the company wants to do. The researchers have proposed some general categories, such as skills, resource utilization, preferences, productivity, and collaboration, which inside have specific calculations that can be done in order to get information from that category. In order to continue with the following steps, it is important to define the aim of the project beforehand.

After defining the indicators that will be used for the remainder of the analysis, these RBIs are extracted on a time series. The time is also not set in stone, so it depends on the researcher if they want something granular such as days or weeks, or rather see a process through long periods of time, such as months or years.

Finally, an analysis about these results is done and a visualization is presented to the user. This is done through time series charts and complemented by trend lines. The chart let compare between different aspects, such as two different resources, or against values that have been defined as benchmarks.

As will be discussed later, this process has several similarities with the approach presented in this thesis. One of these similarities is pointed out by the authors and deals with the interpretation of the results. While a specific result might seem at first to indicate a particular behavior, this is purely interpretation and another researcher could get to a different conclusion using the same set of data. Thus, a lot of care is needed when analyzing the information that this approach yields as a result.

This approach has also been implemented in ProM, and was validated by testing it with the event log of a German bank. A subset of resources was chosen and then the approach was tested with them using particular indicators regarding involvement of these resources in the completion of cases.

2.6   RESOURCE PROFILES

Another work also presented by Pika et al. tries to overcome some of the limitations that the resource behavior analysis framework had, including the effects that the behavior had on the overall process. The researchers present an extension to the framework that adds a method for analyzing if "any relationships exist between given resource behaviors and outcomes" (Pika, et al., 2017). It elaborates on the previous offering by performing regression analysis and focusing on three specific dimensions: time, case, and task. Also, resource productivity is analyzed in order to get a more complete picture of how a resource behavior is affecting the outcome of a process.

The paper defines two main concepts. The first one is resource behavior indicators, which are the measurements that can be done when analyzing resources, such as the skill or utilization of a resource within a process. They were also defined as part of their previous work, but their meaning was expanded and more thoroughly explained on this paper. Parting from that, resource profiles try to classify the behavior of resources based on those indicators. According to Pika et al., a resource profile is a "collection of RBIs that are relevant in a particular context" (Pika, et al., 2017). They have defined several different dimensions in which resources can be evaluated:

- Skills: This deals with the capabilities of the resource, such as knowledge that it has based on education, talents, among others.
- Utilization: How the resource is used throughout a process. It tries to explain factors such as if the resource being used in specific parts of the process, the amount of time the resource is involved, the overall usage of the resource compared to the total amount of cases that the company handles, among others.
- Preferences: Refers to the behavior the resource demonstrate in working terms.
- Productivity: The total output that the resource has. It can be time related, number of cases done in a specific time frame, among others.
- Collaboration: Refers to how the resource works with other resources.

These indicators vary in complexity, in the way they are calculated based on the information of the event log, and their relativeness to other indicators. Thus, they mostly depend on the kind of information that needs to be extracted for the benefit of the company. This also means that this is not an exhaustive list, but rather the indicators suggested by the researchers.

The authors used these concepts to create a framework composed of three modules. The first module is called Analyzing Resource Behavior. This section is in charge of discovering different factors such as what is the utilization of the resource, productivity, preferences, among others.

The second part of the framework is called Quantifying the Outcome of Resource Behavior. The goal of this section is to identify how certain behaviors of resources affect the outcome of the process. The idea is to analyze what is the relationship between these factors, rather than if one behavior causes an outcome.

Finally, the last module is called Evaluating Resource Productivity. This focuses on the individual performance of a resource and compares them with other resources across a specific frame of time. This also takes into consideration the inputs and outputs of a process, and tries to summarize it in a way that it is easy for managers and others interested in the results to analyze them in an easy, automated way.

The validation for this work was done through the analysis of a data set provided by an Australian company. Also, there was an evaluation done in the way of an online survey, targeted at managers. In this questionnaire, the opinions about the usefulness of the newly proposed framework was evaluated.

3   RESOURCE IMPACT ANALYSIS

This section explains the proposed method for analyzing the impact of resources that are part of a business process through mining the corresponding event log. This proposal has two main parts that will be discussed separately: individual impact and group impact. Also, examples from Holograph, the tool that was implemented based on this analysis, are shown when possible in order to exemplify the use cases of this approach.

3.1   OVERVIEW

Normally, resource performance is thought as an individual metric. However, resources are part of a bigger business process, getting input from other resources' work and generating output that has direct consequences on the performance of their peers. This makes the analysis of group impact as significant as analyzing individual impact. Because of this, the analysis is divided into two parts: individual and group impact. An overview of the different sub-analyses done is also showed.
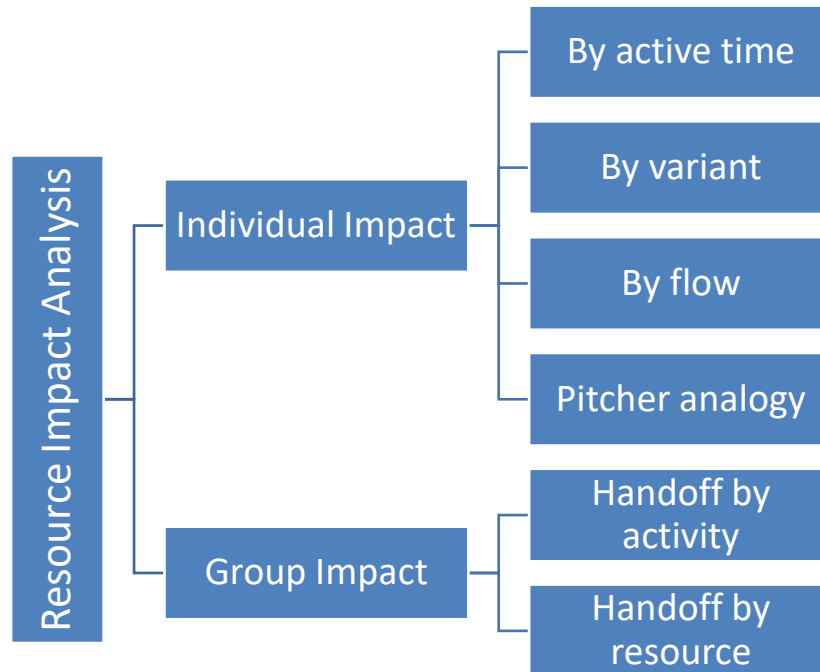
Figure 4. Resource impact analysis in detail

Impact by active time refers to the total and average processing a resource spends on a specific task. This analysis is there in order to start the individual impact analysis, but is not on its own a reliable way of determine individual impact. Analysis by variant digs into how a resource being involved in a case within a variant affects its processing time. This is useful to know if there are specific resources that make the overall variant go faster or slower when they are involved, and is key towards not only comprehending the individual impact but also gives us some clues about the impact the resource might have in the group as a whole.

15

Impact by flow is aimed at expanding on the knowledge gathered on impact by analysis, by closely analyzing specific variants and filtering them by activity. The goal is to gather information about what resources take part on which activities and how the active time of the variant is affected from that moment onwards, while comparing it to the active time before the resource took part of it. The pitcher analogy, also called individual analysis in Holograph (the tool built to validate the method), analyzes what resources are better suited to some parts of a variant. Instead of analyzing as a whole, we can see which resources excel at being in a certain point of the execution of a process.

Finally, group impact has been divided into two parts. Handoff by activity analyzes how resources collaborate between each other by selecting two specific activities. This displays the best and worst combinations of resources and would help in the analysis of determining the best handoffs occur more frequently for those tasks. Handoff by resource focuses on one resource itself, and shows the results of it with the interaction of other resources. This helps determine groups of resource that collaborate in a good way between them regardless of the activity they are working on.

These set of techniques put together can be a very good indicator regarding the impact that a resource has on the execution of a business process. Other analyses were also considered, such as analyzing the impact by departments or associating the impact to a cost metric. However, this would have narrowed down the amount of logs that an analyst can use to take advantage of the tool. With these analyses, we can be sure that an analyst would get the best results possible with the minimal amount of information required.

3.2   INDIVIDUAL IMPACT

The easiest way of analyzing the impact of a resource within an organization is through taking a look into the individual aspect of it. As discussed before, this has been done throughout the years through the implementation of KPIs and similar measurements. However, with the recent push of data mining tools, resource impact analysis is easier than ever, and gives us a quick glance on the resources that perform fastest and slowest for a determined activity.

Unfortunately, current tools don't give us the whole picture regarding individual impact. While resources could appear to be slow in a determined activity, the root cause of it might not be the resource itself, but the tools that they have to work with, how the handoff is affecting the performance of the resource, or that their aptitudes are better suited for other parts of the process.

Therefore, the proposed method gives focus to aspects that have been overlooked in current solutions, in order to understand better how a resource is affected based on different factors.

3.2.1   By active time

In order to analyze resource individual impact better, several methods are proposed. The first one is analyzing how each resource performs when doing a specific activity. For this, the active time of each activity is collected and associated to a specific resource, and then compared with other resources. In order to conduct this analysis, the log should have fields

16

for both start time and end time, and to calculate the active time, it is a simple matter of applying the formula:

$$\text{Activity active time}^1 = \text{Activity end time} - \text{activity start time}$$

There are two ways the active time can be aggregated. The easiest one is by adding all times. However, this approach yields results that might be skewed towards a resource that is highly involved in a specific activity. However, it can still be useful to determine how much total time a resource spends on a specific activity on a specific period of time.

The other path is to calculate the average time, by dividing the total time by the number of times the resource is involved in the activity. This makes the active time of each resource easier to compare with others, and gives a quick glance about which resources are taking more time performing an activity.
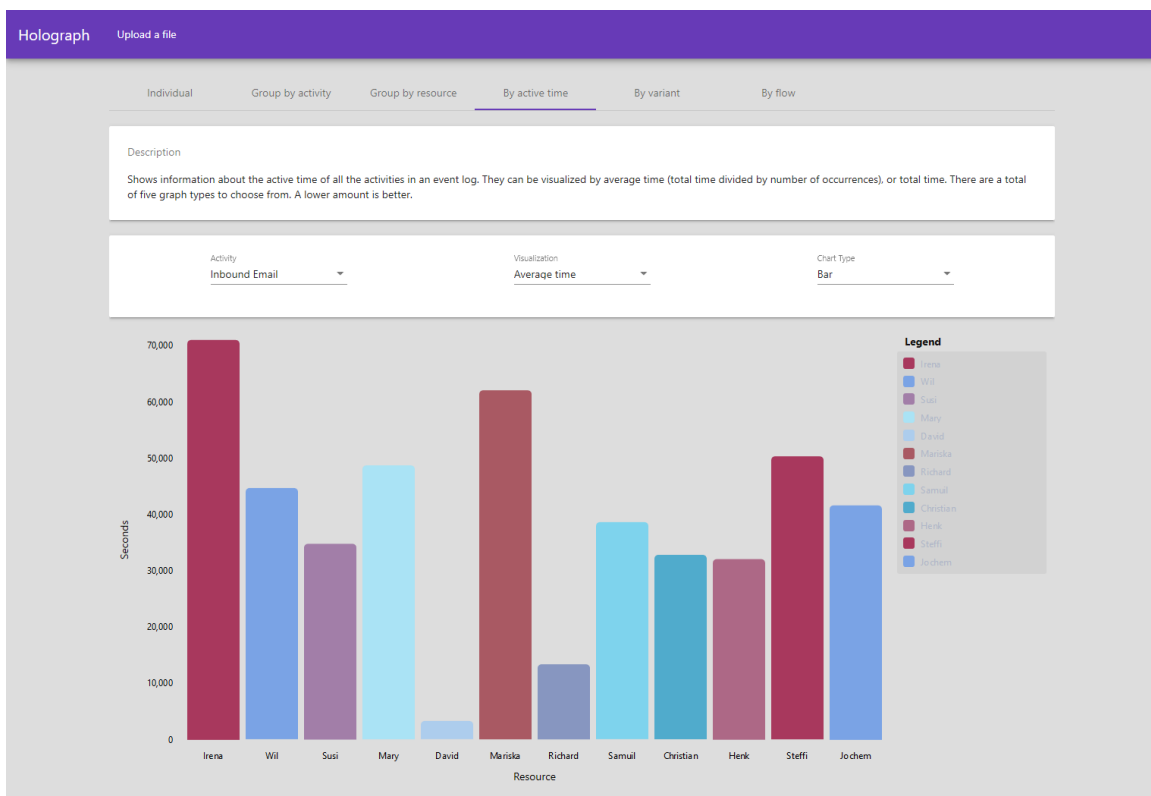
Figure 5. Impact by active time

Because the results obtained from this method will be numeric, then they are hard to analyze in their raw values. A better method of showing these values would be by displaying them in charts. A bar chart is greatly suited for this, since it makes the analysis a matter of comparing different bars. Higher bars would mean higher active times. Also, a

---

[1] This is also called activity processing time according to Fundamentals of Business Process Management (Dumas, Rosa, Mendling, & Reijers, 2018). These terms will be used interchangeably throughout the document.

visualization that exposes the percentage of active time would be useful to easily compare among different resources.

3.2.2   By variant

Even though analyzing the active time of different activities can point us in the right direction regarding adjustments that need to be done in processes, it is a very barebones way of assessing problems with resources. The analysis by active time takes into account the time a resource takes while performing an activity, but does not take into account the context in which it is produced.

Because variants contain a lot of specific cases that execute the exact set of activities in the same order, then a more robust way to analyze individual resource impact is to calculate statistics for variants. In this context, variants refer to a sequence of activities that characterizes a group of cases in the business process. This will start to give us a clue about specific situations in which resources do not perform in the way it is expected, while yielding hints about what is going wrong with the execution.

Under this approach, the variants of the event log are calculated first. This calculation is explained in detail in subsequent sections. Afterwards, the average case time duration should be calculated for all cases within a variant. This can be done by using:

Case time duration = All cases time duration / Number of cases

With this information, and knowing all of the resources involved in the variant, we can then calculate three different statistics for each resource, spread across two different categories:

- Minimum variant time (with and without resource)
- Maximum variant time (with and without resource)
- Average variant time (with and without resource)

Although minimum and maximum time are important to identify rare cases and outliers, the main focus of this analysis is to work with the average time. From this information, there are several different calculations that can be done:

- Quickest: The resources that are part of the quickest cases within a variant. Calculated by getting the minimum value among the average variant times.
- Slowest: The resources that are part of the slowest cases within a variant. Calculated by getting the maximum value among the average variant times.
- Above average: The resources that are part of the cases whose times are above the average for the variant. Calculated by getting all of the average variant times that are greater than the average.
- Below average: The resources that are part of the cases whose times are below the average for the variant. Calculated by getting all of the average variant times that are lower or equal than the average.
- Most involved: The resources who appear the most in the cases within a variant.
- Least involved: The resources who appear the least in the cases within a variant.

Through the analysis of this information, we can already start to determine some tendencies in the execution of a process. It is important to note how valuable the facts

"Above average" and "Below average" are towards the analysis of performance. Even though they are still dependent on a lot of factors, these two might start to shed a light about who is an ideal candidate to stay working on the task they are currently performing, and whose position might have to be evaluated in order to better suit to the resource's competencies.
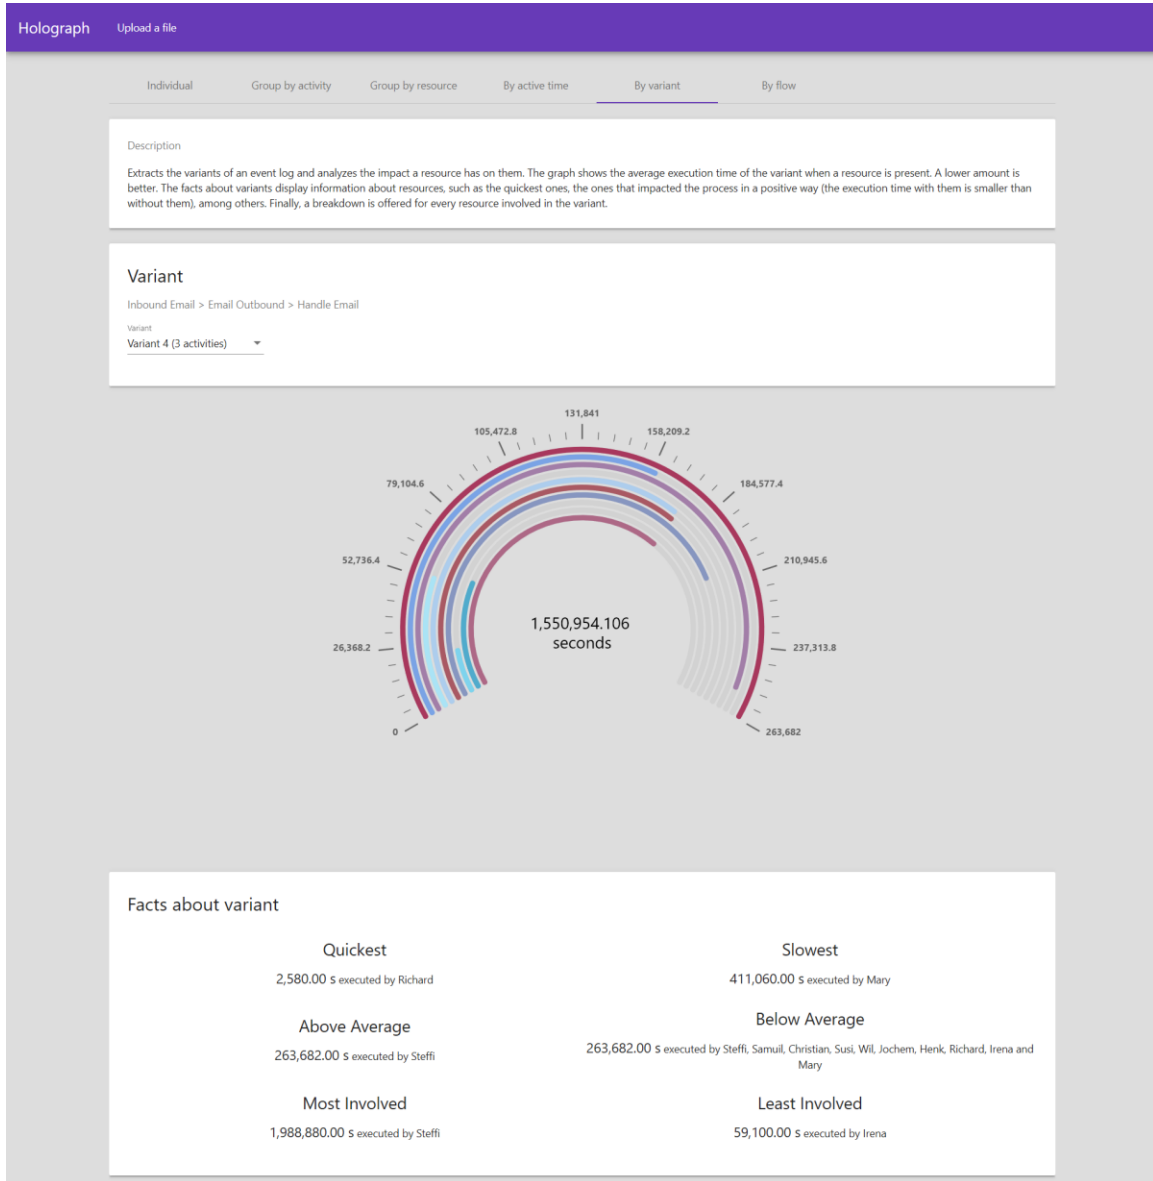

Figure 6. Analysis by variant

In the above screenshot, we can see that Mary has the slowest average processing time out of all the cases in the variant, while Richard has the fastest. We also can see all of the resources whose active times fall below the average and above the average. They are very high-level indicators, but meaningful results are starting to appear.
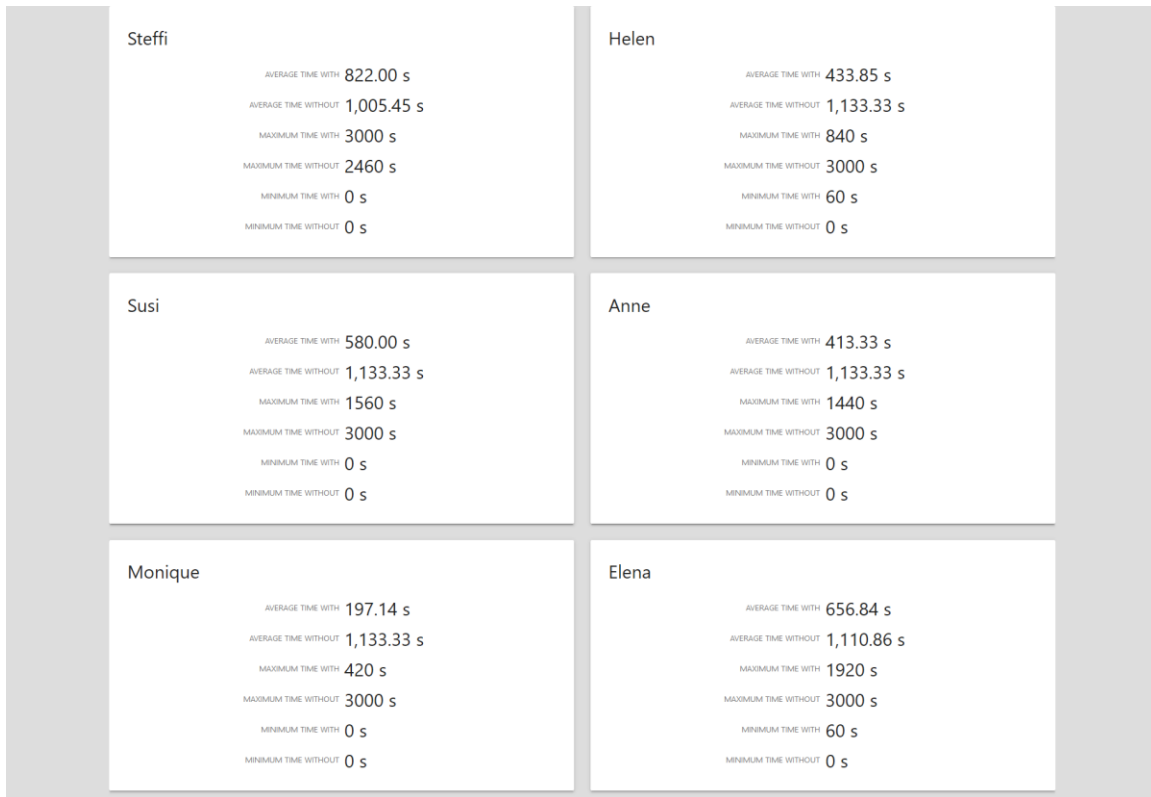
Steffi		Helen	
AVERAGE TIME WITH	822.00 s	AVERAGE TIME WITH	433.85 s
AVERAGE TIME WITHOUT	1,005.45 s	AVERAGE TIME WITHOUT	1,133.33 s
MAXIMUM TIME WITH	3000 s	MAXIMUM TIME WITH	840 s
MAXIMUM TIME WITHOUT	2460 s	MAXIMUM TIME WITHOUT	3000 s
MINIMUM TIME WITH	0 s	MINIMUM TIME WITH	60 s
MINIMUM TIME WITHOUT	0 s	MINIMUM TIME WITHOUT	0 s

Susi		Anne	
AVERAGE TIME WITH	580.00 s	AVERAGE TIME WITH	413.33 s
AVERAGE TIME WITHOUT	1,133.33 s	AVERAGE TIME WITHOUT	1,133.33 s
MAXIMUM TIME WITH	1560 s	MAXIMUM TIME WITH	1440 s
MAXIMUM TIME WITHOUT	3000 s	MAXIMUM TIME WITHOUT	3000 s
MINIMUM TIME WITH	0 s	MINIMUM TIME WITH	0 s
MINIMUM TIME WITHOUT	0 s	MINIMUM TIME WITHOUT	0 s

Monique		Elena	
AVERAGE TIME WITH	197.14 s	AVERAGE TIME WITH	656.84 s
AVERAGE TIME WITHOUT	1,133.33 s	AVERAGE TIME WITHOUT	1,110.86 s
MAXIMUM TIME WITH	420 s	MAXIMUM TIME WITH	1920 s
MAXIMUM TIME WITHOUT	3000 s	MAXIMUM TIME WITHOUT	3000 s
MINIMUM TIME WITH	0 s	MINIMUM TIME WITH	60 s
MINIMUM TIME WITHOUT	0 s	MINIMUM TIME WITHOUT	0 s

Figure 7. Details on the analysis by variant

If the aggregation is not enough, the information specific to each resource is also available to be seen, with numbers for average, maximum and minimum time.

3.2.3 By flow

Another of the proposed ways of analyzing individual impact stems from the closer examination of the impact of a resource in a specific flow. This has already been examined partially previously, but here it is done in a more specific way.

A resource's impact does not only deal with the competencies it has or the training that it has received, but also with the input that it receives from previous phases of the process. If a resource receives an input that better suits a process, or that better suits the resource's strengths, then the following steps will be made easier. However, if resources have to deal with an input that makes their work more complicated, then not only the resource will take more time completing the task, but also the output will likely be of inferior quality. As a result, subsequent tasks will suffer in the times that are needed for their completion.

Through deeper analysis of these factors, companies can not only identify resources that are really well suited for a job, but also might discover that formalities like documentation, manuals, or rules put in place by the company might not the best for the execution of a process. This could lead to changes in the company's processes, restructuring of some departments to better suit the competencies of a resource, new departments being created, or a combination of the three.

20

In order to start the calculations for this analysis, the same variants that were calculated for the previous chapter should be used. However, only the variants with two or more activities should be considered, since a variant with only one activity will not have a time before or time after.

After the variants have been filtered, all of the distinct activities should be identified. Finally, three values should be calculated for each activity: time before, time taken, and time after.

At the moment of visualizing, the implementation should let the user choose between all of the variants in the case, one of the activities within the variant, and one of the time options. The calculation is done in this way:

- Time before: The average active time registered by all resources present in the variant for all cases from the beginning of the process until the activity preceding the selected activity.

- Time taken: The average active time for the selected activity registered by all resources present in the variant.

- Time after: The average active time registered by all resources present in the variant for all cases from the activity following the selected activity until the end of the process.

By getting these calculations, one can quickly see several aspects of a process. For example, by analyzing the time before, we can identify resources that receive more complicated cases (signaled by the longer execution time) and analyze their output. If we find resources that are especially good at handling problematic instances, then we can determine if they are doing something that could be replicated throughout the company, or if it is because of their innate skills.

Otherwise, if we find resources that struggle a lot with complicated tasks, then perhaps a restructuring can be applied to move the resource to a part of the process where tasks match more closely the skillset of the employee.

Also, if we identify that, when the task takes longer in previous phases, it results in very short active times for the tasks superseding it, and overall shorter execution time, then an analysis can be done to determine the workflow of the resources to better adjust the process.
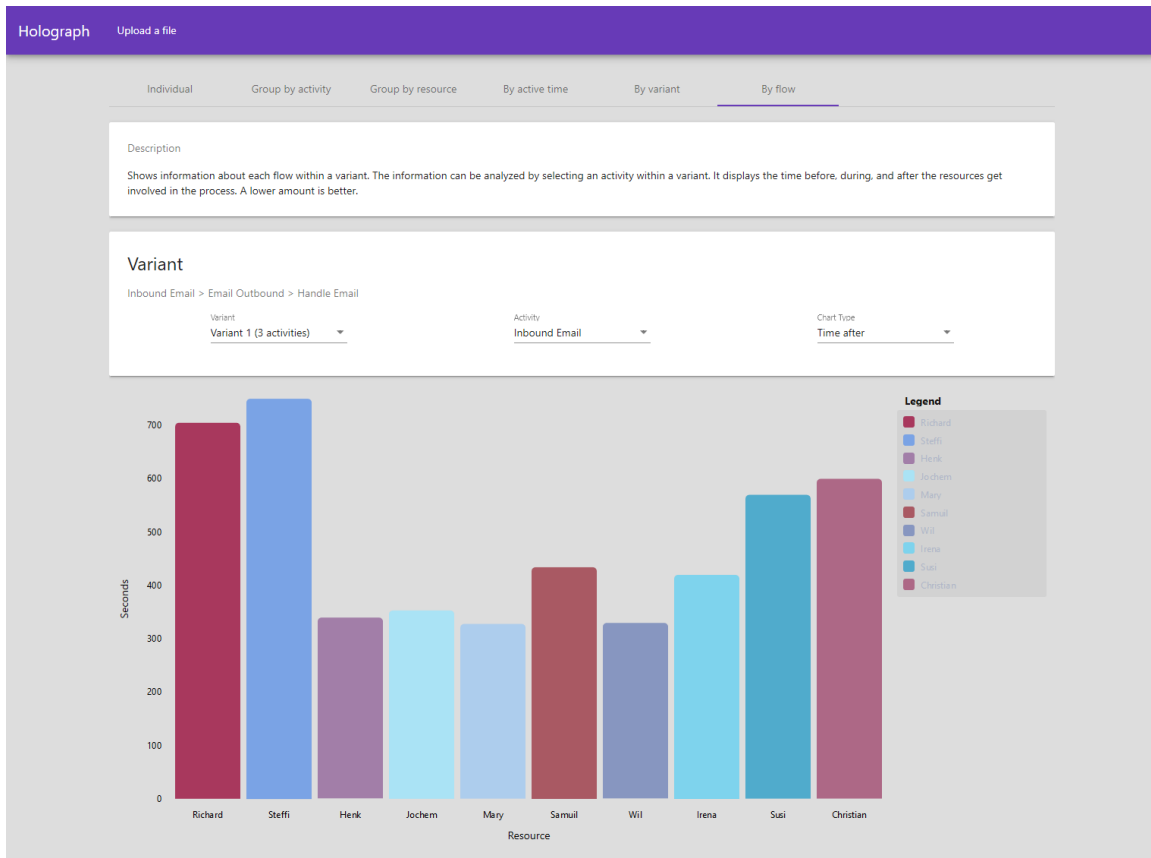
Figure 8. Impact by flow

In the above example, we can see that there are several resources, such as Richard and Steffi, whose involvement led to longer active times on average, while Mary and Wil's processes perform better after they are involved. However, if we see other activities within the same variant, we see that Wil's involvement in later tasks caused a huge amount of delay compared to others. We can compare this to Wil's active time analysis and determine what is going on with Wil's tasks and start to explore causes for this difference in behavior between different activities.

As there will be cases where resources get cases with long execution times, there will also be situations in which employees get cases with short execution times and they turn into cases with longer times. It is important to note that this should not mean that the resource is not up to the task. There is a plethora of situations that can be the root cause of this, and tasks can get complicated thanks to external factors that could be outside of the resource's control. Even if the problem lies with the resource itself, it could be because the resource is a new employee, the training has not been sufficient, or that their skillset is better suited to other tasks within the same process.

Also, as previously said, the input a resource receives is essential to the success or failure of the task. A resource struggling with an activity could mean that there is not enough work done in the previous parts of the process, resulting in an output that is very hard to work with for the resources involved in the following activities. It is imperative to analyze this information closely with other cases in order to take decisions.

If we identify that there is a task that has abnormally high values for the "Time after" parameter, then we can probably conclude that it is very problematic. Not only it can take a significant amount of additional time from the company, but also there is a risk that the whole case might end up in a negative outcome.

3.2.4 The pitcher analogy

In order to explain the next analysis, it would be better to introduce an analogy with how pitchers work in baseball. According to the Major League Baseball, there are at least three different types of pitchers: starters, relievers, and closers. Other sources also identify a set-up man position, but for this analysis, the official terms are appropriate.

The starting pitcher, as the name indicates, is the pitcher that is in charge of starting the game. Their skills have been identified as being better tailored to be useful at the beginning of the game. Also, teams normally have several starting pitchers in order to avoid fatigue from pitching very often.

The relief pitcher is the one that comes after the starting pitcher. When the starting element starts to feel tired, has done a predetermined amount of pitches, or is performing poorly, the relief pitcher is introduced into the field. Since they are less used than starting pitchers, who can take part of several innings, they have a different skillset than starting pitchers and can afford to play stronger balls. Thus, together with factors like fatigue and the status of the game, they plan their approach accordingly.

Finally, the closer is a kind of pitcher whose function is to close the game. They are introduced normally on the final inning, when the circumstances of the game are different than when the starting and relief pitchers were part of the game (Major League Baseball, 2018).

Similarly, we can classify the activities that are executed in a process in different groups. Normally, a process has several activities that could be considered as initial, others as middle, and others as the closing tasks of the process. There might be some other tasks that are considered special and need specialized resources or a different approach. It is paramount to understand how different resources can be used in each kind of situation in order to maximize the output of a process. If we identify that a resource is particularly good when working at the beginning of a process, or that thrives under unknown conditions compared to others, then we can use these resources at the appropriate times and make significant gains in performance.

To adapt this concept to the proposed approach, we can divide the activities in three different parts: initial, middle, and end. Depending on the length of the processes at specific companies, the classification can be extended into more phases to take them into account.

If we want to analyze the process more deeply, then other additional measurements can be generated, such as problematic cases, instances where very specific knowledge is required, or other industry-specific factors. However, for explaining the process more clearly, the default implementation of selecting initial, middle, and end groups will be used. This is also the selection used for Holograph.

In order to conduct this analysis, only variants with three or more activities should be used for obvious reasons. After that, the average active times for each activity should

be compiled. This should be classified by resource in order to offer a comparison between them. The user would need to have a way to choose between initial, middle, and end visualizations, and the specific factors in case there are any. The implementation should be able to let the users recognize which resources are the most impactful in each stage.

In summary, this section should have some resemblance with the above and below average representations exposed previously. However, by analyzing the information in sections of the process, it will be easier to identify overall tendencies by section and not by specific activities. This way, the significance of the magnitude of the impact is more accurate, since it analyzes closely-related tasks.
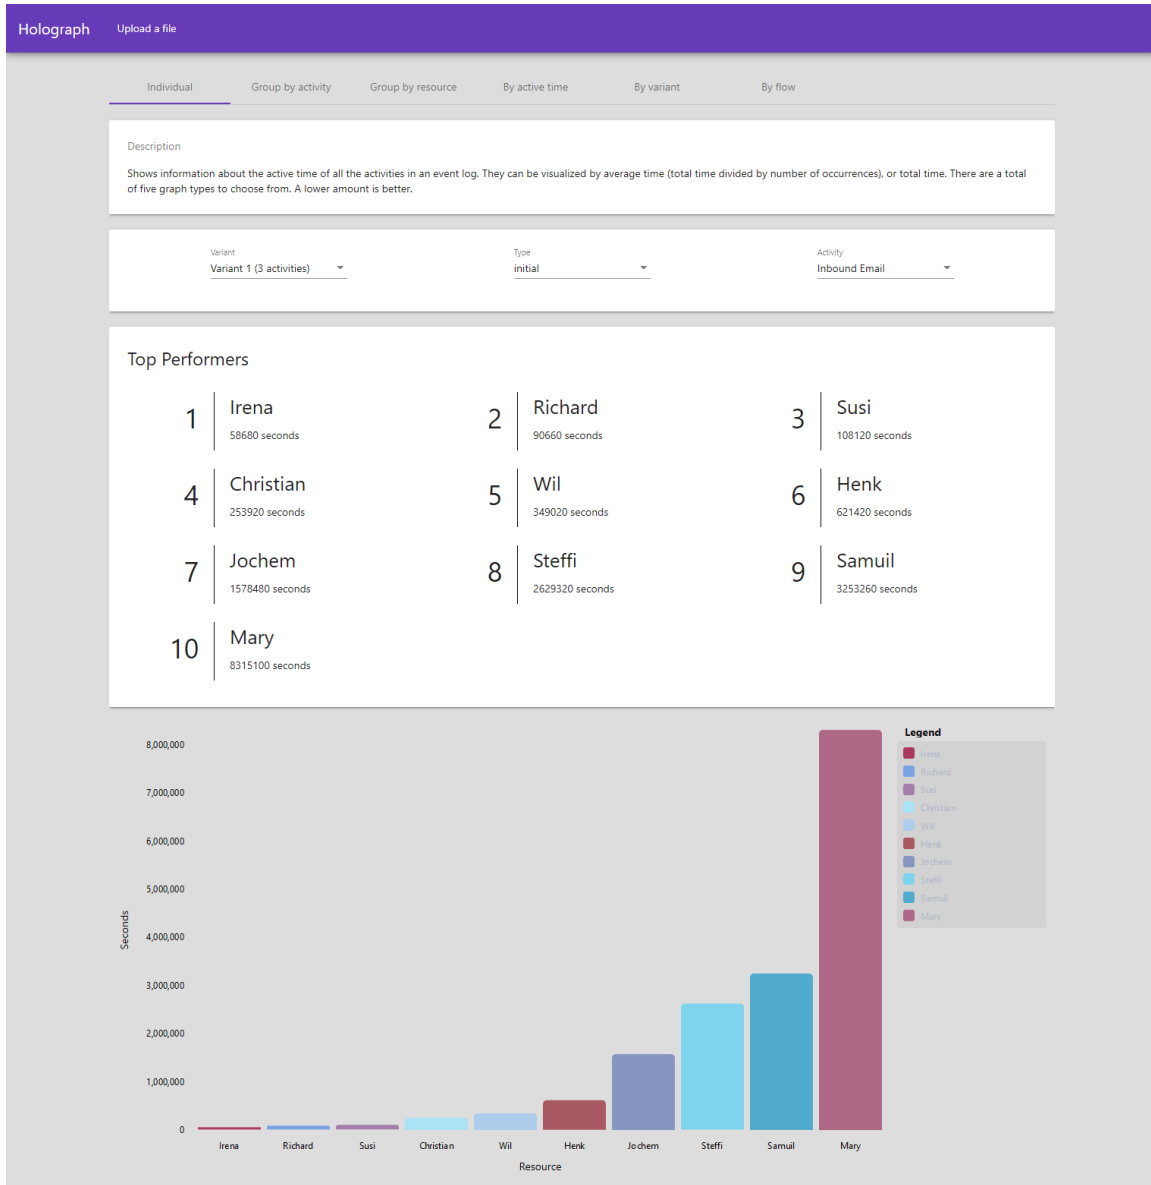

Figure 9. The pitcher analogy analysis

For example, in the previous screenshot of Holograph, we can see that Irena, Richard and Susi are the top performers at the beginning of the process. However, on later stages, the

same sandbox example shows that Richard is outperformed by several other resources, but Irena is still an excellent performer then. This might give us a small insight on tracking how Richard performs at different stages on other variants, and lead us to restructure the process a bit and have Richard take initial tasks.

Also, as previously said, for an extended implementation, an analysis of problematic cases could be added. For this, it is important to determine what constitutes a problematic case for a specific process. Some processes consider a problematic case a situation where the process takes longer than expected. For others, it is having to execute specific activities, such as account terminations, failure of money collection, or other instances where customers are affected.

Thus, this makes it non-trivial to implement, since it depends on the intricacies of the process rather than a specific set of rules that can be generally applied. However, the idea of it is that after gathering all of the problematic cases, a classification could be done, by separating the problematic cases with good outcome from the ones with bad outcome. Subsequently, the implementation could calculate if there are resources that are more likely to be present in a good outcome. This could indicate that either the training or the natural aptitudes of the employee are good for dealing with high-risk cases, and the company can act accordingly.

Because of the nature of these cases, the failure of a process might not be entirely the fault of the employee handling it. For example, in the case where a customer calls to terminate a service and does not accept a counter-proposal, the causes that lead to the customer's decision are not in any way related to the work of the call operator. Therefore, this information should be analyzed very extensively.

3.3   GROUP IMPACT
Individual performance is vital to the overall execution of a process. However, it is imperative that the group performance is also considered if one wants to understand how the work is handed from one department to the other, and how the whole process can be better than the sum of its parts.

For the analysis of the impact of a group in a process, the active time of cases should also be used, just as it was done for the individual tasks. There is one proposed analysis for group performance: the handoff. It is divided into two different sub-analyses: by activity and by resource.

3.3.1   Handoff
After the work has been done by a resource, it is paramount to understand how the output generated by it affects the whole resource. One interaction that greatly depends on that is the handoff. A term originally coined in the telecommunications sector, a handoff in our approach refers to the act of transferring the final result of an activity to the resource that executes the following activity.

In order to analyze the impact of the handoff in the active time of resources, first the data has to be classified. It can be done by getting a list of the distinct activities in the process, and then get all of the activities that follow it in all of the variants.

For getting more strict results, this can be done within variants, in order to calculate handoff in specific circumstances instead of taking a general approach.

3.3.2 Handoff by activity

The handoff by activity lets the user see what are the best resources performing two different activities one right after the other. With this, we can closely analyze how the output of one resource directly affects another activity, and quickly identify which resources are both better suited for some activities (because their output is better suited for next activities), and which resources work together the best.

The implementation should calculate what are the total values for all the interactions between resources for both selected activities, while also showing the average from them. This could be done on a log basis (if there is the need to analyze activities as a whole irrespective of their context) or on a variant basis (if the context is valuable for the analysis, and there is enough information for that).

If we see that there is an interaction that is being repeated multiple times with low active time, then we can delve deeper into it and try to analyze what exactly is being done there. Also, if we see that there is a resource that executes an activity quickly with some resources but slowly with others, then probably the output needs to be analyzed in order to determine why it is working for some cases only.
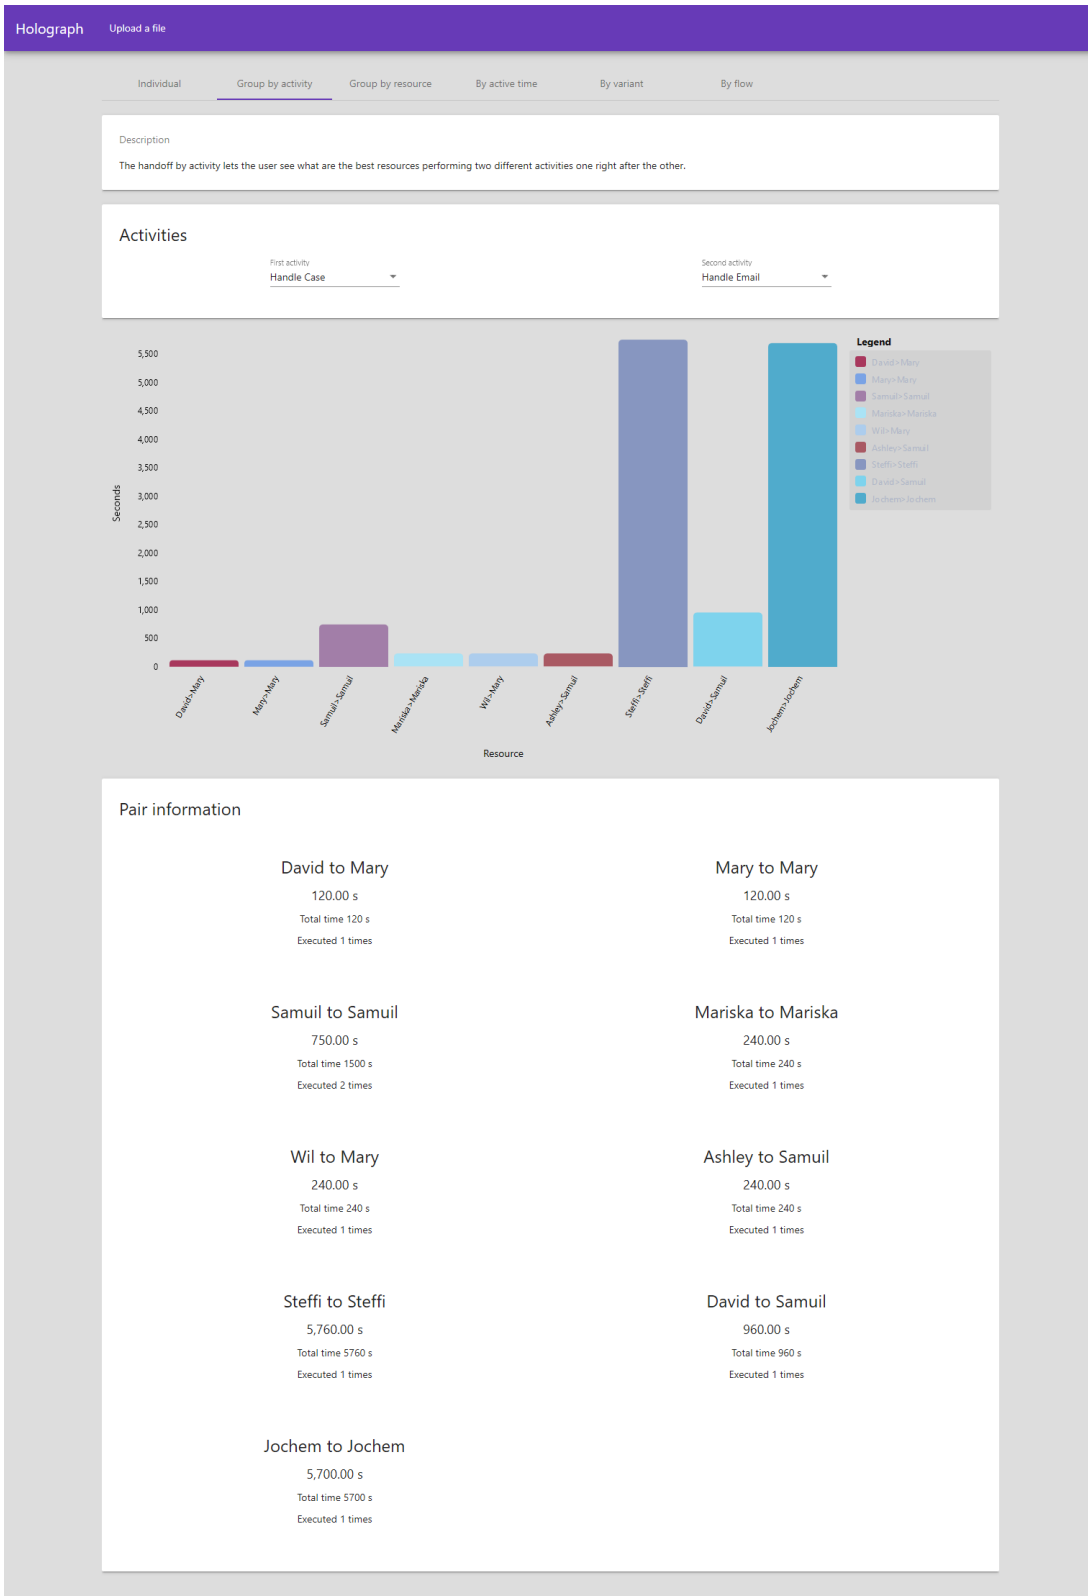
Figure 10. Analysis of handoff by activity

In Holograph, handoff by activity displays the different activities that resources execute on a bar graph. In this event log in particular, several activities were performed by the same resource, but also, we see that there are handoffs between multiple resources, and the comparison between them. The times are all laid out on a graph, so we can see the difference between different handoffs.

3.3.3   Handoff by resource

Although the handoff by activity is very valuable to analyze the interaction between resources for two specific activities, another analysis can be done here. The focus could shift from the activities to the resource itself. As the handoff by activity, this can also be done on a log or variant.

In this analysis, a resource is selected, and the resources with which it interacts with are shown. The average active time of the two activities performed by both resources should be shown. Through this information, we can have an overview of the resource itself instead of the activities it performs. It can show us information about who the resource works best with, regardless of the activity. By using this information, we can closely match resources that collaborate the best way, and easily increase efficiency in the process without changing the process itself.

3.4   LIMITATIONS OF THE APPROACH

Analyzing how a process can be done better through the mining of resource information can be a valuable source of data. It could surface patterns that could potentially improve business operation, lead to restructuring that would make the employees happier and more productive in their positions, and aid in the making of training material for new employees.

However, it is also a double-edged sword, since it could be used in negative ways. Some of the drawbacks of the proposed method are analyzed in more detail here, and should be the starting point for further research made in this area.

3.4.1   Blame on resources exclusively

For example, if a resource is found to be taking more time to complete their tasks, one of the easiest paths to follow is to wrongly blame it on incompetency and take rash, undeserved decisions. Also, it can be wrongly assumed that the active time of an activity is solely the responsibility of a resource.

If an activity is delayed or is executed slower than usual, normally there are a lot of factors that contribute to that together. This includes imprecise input, dependencies on third-parties, or conditions that cannot be controlled. That is why it is imperative that the results obtained here take into account all of the different possibilities and factors that affect a process.

3.4.2   Focus on active time only

Since the bulk of activities that take a long time are executed by humans, then it makes sense to analyze the active time of activities. However, one factor that is not analyzed by this approach is the waiting time between activities. One can naively dismiss this period since no work is done there, but that would be an error.

Waiting times for a process can be fundamental in discovering ways of implementing improvements for it. By analyzing the intricacies of the handoff between activities (and resources), big optimizations can be applied to a process. For example, if a process needs the signature of an individual that is rarely available for doing the task, then the whole process takes a very long time, although the active time for this will be very short. By restructuring the procedures to stop depending on resources that are rarely available, the process can gain a lot in total time, which could lead to better outcomes, and all without improvements to the active time.

Also, the way that the active time is calculated means that logs with timestamps only, instead of start and end times, cannot be calculated easily. One option would be to calculate the active time as the difference between the timestamp of the next event in the case minus the current event. However, this would account for the waiting time as well, which completely distorts the results, since, after the resource finishes executing the task, he or she has no power to influence the following waiting time.

Another option would be to get it as the difference between the current task minus the previous task. This would eliminate the waiting time cause by another resource. However, it would still include waiting time, which may wrongly yield longer times and would distort the results. Also, in the case where an employee has a queue of tasks, the waiting time would increase for all of them, wrongly penalizing the resource on multiple occasions. However, if the analysts take into account this limitation, then this approach would be the least harmful for the analysis because it takes into account the times recorded by the same resource.

3.4.3   Low focus on outcome of the process or team dynamics

When executing a process, naturally one of the most important aspects of it is to optimize its speed. However, this sole dimension neglects other intangible information, such as the outcome of it or the harmony within a team.

If a process is done more slowly but it is almost guaranteed to reach a positive outcome, while other team has a fast approach that is more prone to failures, probably it would be a better idea to slow done the fast team in order to increase its success rate. However, an analysis that bases its conclusions only on active time will fail to take this into account, and will display the slower team as the less efficient among them.

In addition, if a team performs more slowly but creates a good environment by doing so, it could be more effective in the long run. Teams that work fast but, by doing so, fail to create a good working environment between them might reach a point in which they might suffer burnout or intra-team conflicts.

29

4 IMPLEMENTATION OF HOLOGRAPH, A RESOURCE IMPACT ANALYSIS TOOL

To validate the results of the proposed method, an implementation of the method is the best way. For this, a web implementation of the methods previously exposed was done. This tool was done with the specific purpose of testing if the proposal does actually yield the results that are expected from it.

For building this tool, some considerations and limitations were put in place. Although this tool was done with general purpose in mind, and attempts to use it would probably yield good results, it was optimized for a specific set of files. Also, this implementation was not done with performance in mind, so results take time to be generated. There is no cache of these results, so each time a log needs to be analyzed, the results are generated from scratch. Finally, the tool accepts only CSV files and assumes they have a header row.

4.1 TECHNICAL SPECIFICATIONS

The app consists of two parts: a front-end and a back-end. The front-end collects a CSV and sends it to the back-end for its processing. The front-end then calls to specific endpoints providing the CSV identifier, and receives information in JSON format to be displayed.

These are more specifications regarding the front-end implementation:

- Framework: Angular 5.0.0
- Node 8.9.4
- Angular CLI 1.6.8
- Typescript 2.5.3
- GitHub Link: https://github.com/gersonnoboa/Holograph_FrontEnd
- Heroku Link: http://holograph.herokuapp.com/home

A sandbox is provided to the user by accessing the Holograph Front-End's home page in Heroku, and then clicking on the "Enter Sandbox" button.

These are more specifications regarding the back-end implementation:

- Framework: Flask 0.12.2
- Python 3.6.4
- GitHub Link: https://github.com/gersonnoboa/Holograph_BackEnd
- Heroku Link: http://holograph-backend.herokuapp.com

4.2 OVERVIEW OF HOLOGRAPH FRONT-END

Holograph Front-End is a front-facing application developed in Angular to evaluate the veracity and usefulness of the proposed method for evaluating resource impact in an event log. The site works closely with Holograph Back-End. The app can be divided into three parts: file upload, parameter generation, and results. The hierarchy of classes can be seen in the following graph:
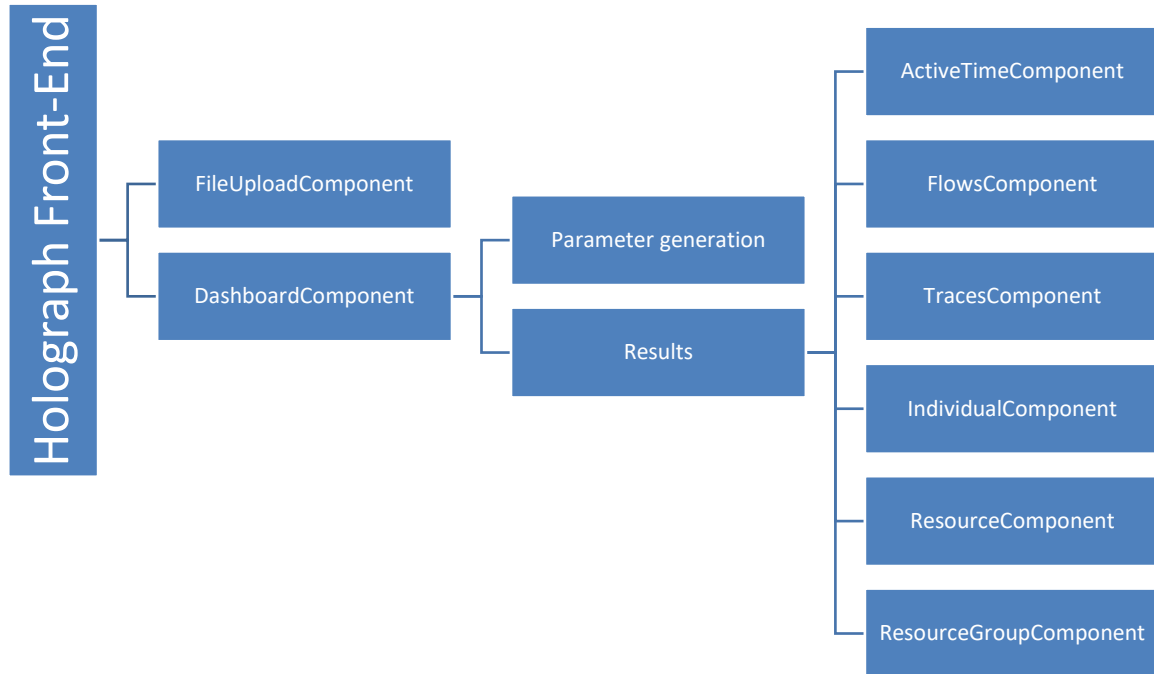
Figure 11. Holograph Front-End's classes

4.2.1 File upload

In order for the application to work correctly, it is necessary to provide a CSV file. The requirements for it are the following:

- The file needs to have a header row.
- The file needs to be separated by a comma or a semicolon
- The log must have a case ID, activity, resource, start time, and end time

The application uploads the supplied CSV file and gets a unique identifier generated by the back-end, which must be used for the remainder of the process.

The uploading of files to the server is controlled with the FileUploadComponent class. As every class in Angular, it has an HTML and CSS file associated to it. Also, a FileUploadService has been added in order to handle the POST for the file itself. The HTML declares a DIV tag to which a file can be dropped.

When the user releases a file to this area, the Angular component performs the necessary validations. Similar to the back-end, the application verifies if the file is not empty and if the format of the file is CSV, XES, MXML, TXT or XML. It is important to note that currently the software supports the mining of CSV only for now because the Python back-end has not been developed yet to support the parsing of other file formats.
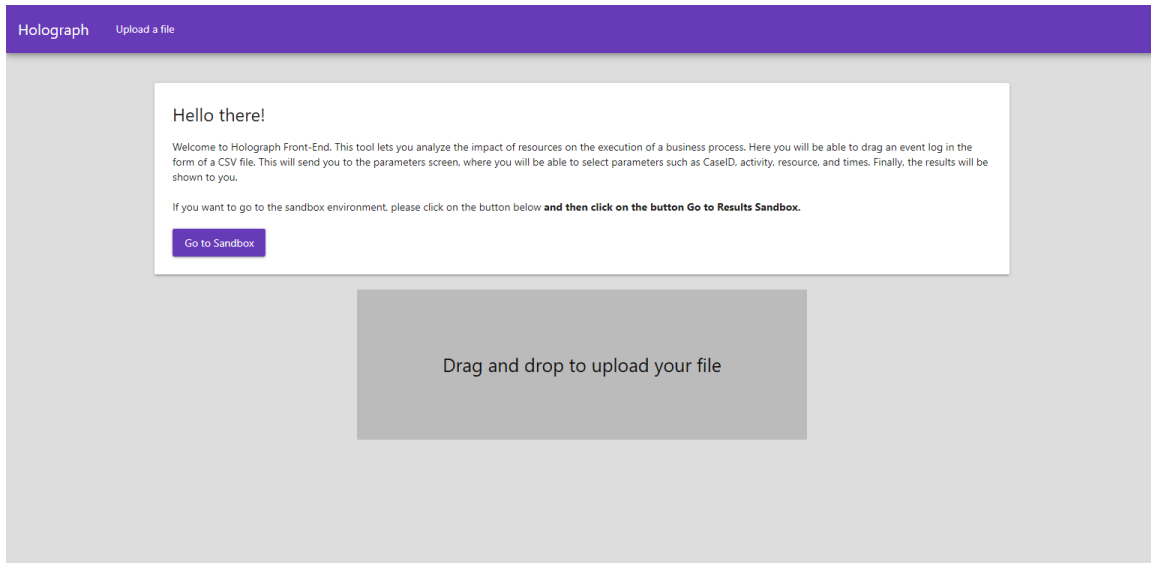
Figure 12. Holograph Front-End's file upload screen

4.2.2 Parameter generation

After the file has been uploaded, Holograph needs to know which of the rows correspond to case ID, activity, resources, start time and end time. Thus, it asks from the user this information. The application provides the rows extracted from the first row of the file so that it is easier for the user to choose one. Also, the application displays the first row of the file, so that the user can double-check that the information is correct. The DateTime format is also required in order to parse the dates of the file.

All of this is done through the class DashboardComponent. This is the main component of the whole application, since it is the one that is in charge of displaying the next two sections. As with the FileUploadComponent, a service has been added to the component called DashboardService, which handles the request of file headers. This request returns a String array with the file headers, and the first row of information of the file. Holograph Front-End displays this row to the user to make it easier for him or her to decide which row are to be used. This is done through the DashboardComponent class.

The application displays a Horizontal Stepper provided by Material in order to ask for the parameters for Case ID, Activity, Resource, and DateTime Format. Afterwards, it also asks what the type of date is present in the log, in order to do the calculations accordingly. Even though the three options are "Has active time only," "Has start and end date," and "Has timestamp only," the application currently supports "Has start and end date" only. However, the necessary steps have been done to support other types of logs too.

All of the information given by the user is saved into a FormGroup object, which is then sent to each of the components that form the result part of the application.
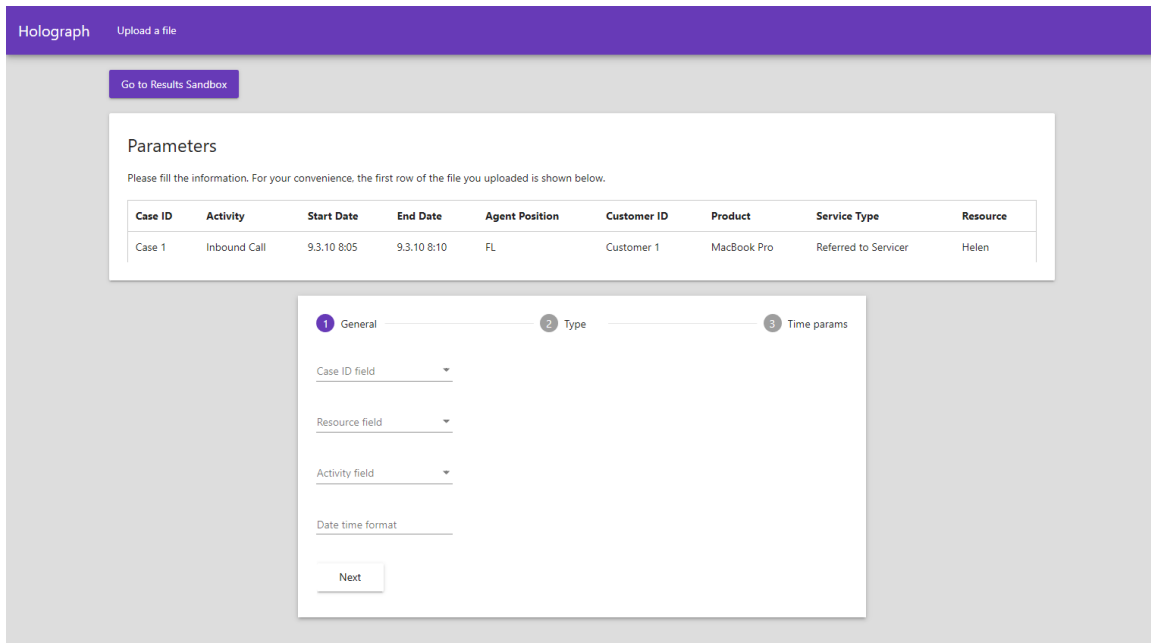
Figure 13. Holograph Front-End's parameter generation screen

This information is then sent to the back-end in JSON format. This will start the back-end processing and generation of the results. Holograph shows one tab for each of the different analysis that were proposed: active time, variant, flow, pitcher, and handoffs. The handoff analysis has been divided into group by activity and resource.

4.2.3   Results

The results are formed by six different components, all grouped into a Tab Group provided by Material. They are the following:

- ActiveTimeComponent
- FlowsComponent
- TracesComponent
- IndividualComponent
- GroupComponent
- GroupResourceComponent

All of them work in a very similar way. They all have a service associated with them (if the component is called ActiveTimeComponent, then the service is called ActiveTimeService) that is in charge of communicating with the back-end. It is important to note that the front-end requests all the information at once, instead of requesting it on-demand.

Even though they all display data, the components themselves expect to receive vastly different data sets, since they are tailored to the specific analysis that they are performing. The only similarity between them is that they are always received as JSON objects. All of the components show a loading progress indicator until the data transfer is completed.

After getting all the data for each component, the information is displayed according to what has been outlined in the approach in the previous section. This is done with the aid of NGX-Charts, a charting package that provides several different chart elements which are ideal for the purpose of this application. Since the charting library expects the data in a specific format, the obtained data is converted to the required format first, and then passed to the charting component. When the user requires a visualization change through drop-downs, the data for the charting component is generated again. However, the whole data set obtained from Python is never reloaded.
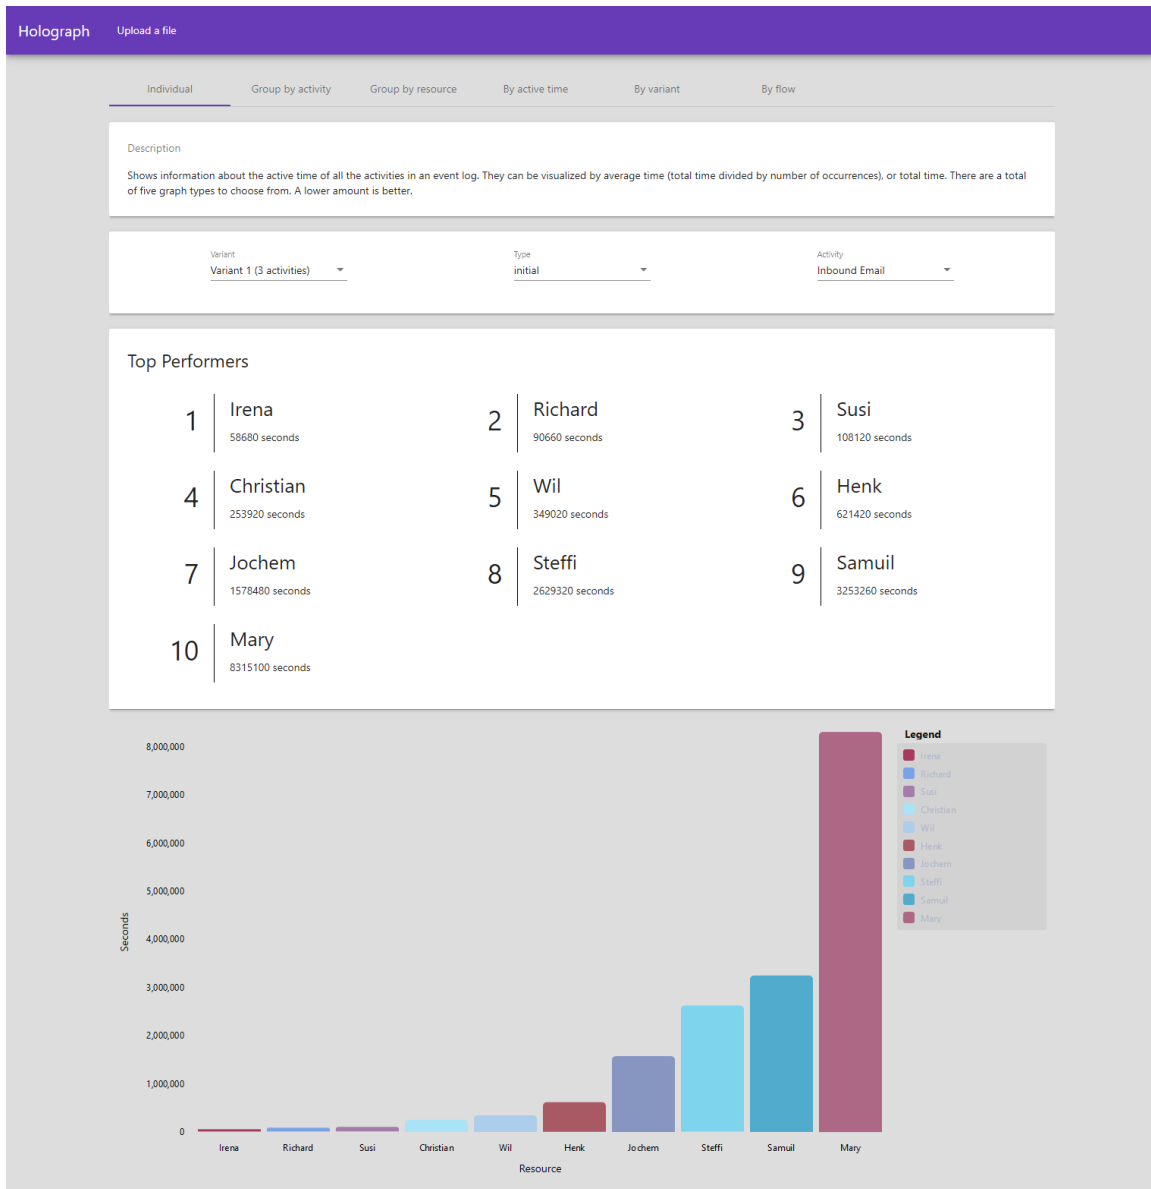

Figure 14. Holograph Front-End's individual analysis screen

In the case of active time, the application shows three different dropdowns: one for activity, one for visualization type (average or total), and one for the chart type (bar, pie,

advanced pie, pie grid, and number cards). Because of the nature of this analysis, the bar chart is the one that works the best to compare results between different activities.

For the variant analysis, only the variant has to be chosen, and then Holograph displays all the relevant results. First, it shows a gauge chart that displays the average time for all cases for each employee. Then, all of the facts generated are displayed with their value and resources. Finally, the breakdown of each resource is shown with the average, maximum and minimum values, both for the cases with the resource and without the resource.

The flow analysis has three dropdowns: variant, activity, and type (chart before, taken and after). The page shows a bar chat with the information of the resource for that variant, activity and type. Also, the application shows the activities that the variant encompasses.

The pitcher analysis displays the variants, the type of the analysis (initial, middle or final) and the activity to research. Below, the resources that take part in this combination are ranked from best to worst, showing their active time in seconds. Finally, at the bottom, there is a bar chart that displays this same information in a visual way.

The group statistics are divided into two parts: handoff by activity and by resource. In the handoff by activity, the information about the handoff from one activity to the other is displayed. Here, the application lets the user choose between two activities that are guaranteed to be consecutive. Then, a bar chart with the information is displayed. This represents the active time that takes for two resources to perform their activities in succession. Finally, the information is shown in numbers. Information that is displayed includes starting and ending resource, the average combined times, the total time, and the times the handoff occurs for the specific combination.

Handoff by resource has been implemented in the back-end. However, this information is still not present in the front-end.

4.3   OVERVIEW OF HOLOGRAPH BACK-END

Holograph Back-End is the service that provides the information needed for displaying the different visualizations that Holograph Front-End supports. It was developed in Python with the Flask framework. Although it was made with Holograph Front-End in mind, since it is a service that receives and provides data, any application could be built on top of it. Just as the front-end implementation, Holograph Back-End can be divided into three parts: file upload, general mining, and results.

Because Holograph Back-End is the one providing the data for the front-end, its architecture is more complex. The hierarchy of classes can be seen here.

35
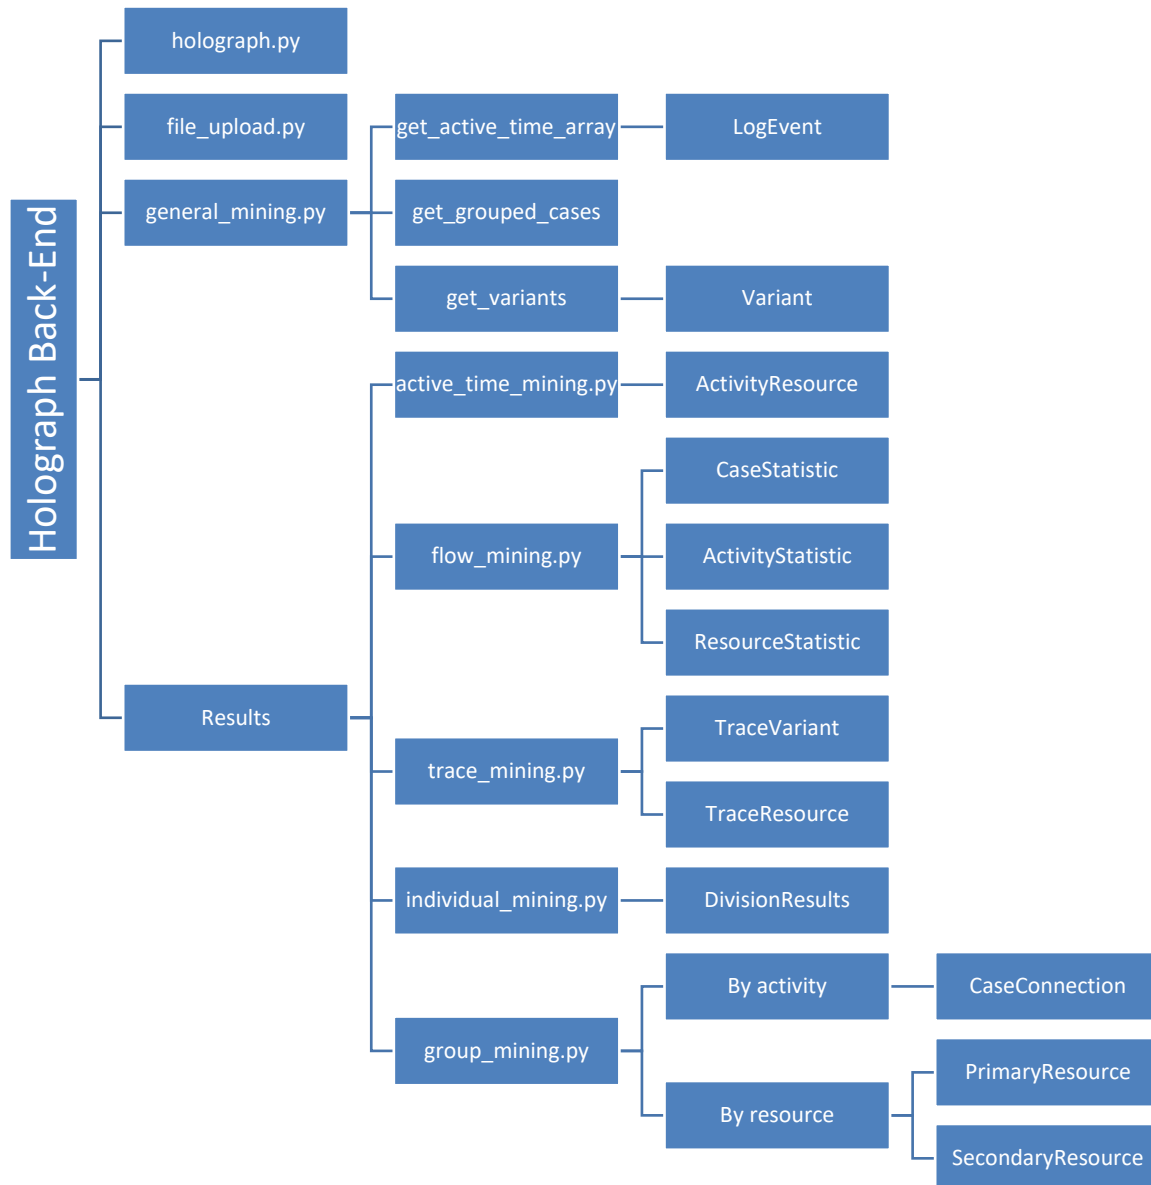
Figure 15. Holograph Back-End's files and classes

4.3.1   Main file

A file called holograph.py is used as an entry point for the application. Here, the app itself is created and this is the file that runs on the server and redirects to other parts of the application upon request of specific endpoints.

Also, configuration properties have been stored on this file. The upload path of the CSV files has been defined under the UPLOAD_PATH property. The allowed extensions for files have been defined in the allowed_extensions variable under the upload_file() method. Other than the uploading of the file, all of the requests are get methods, with an expected querystring with the necessary parameters for the correct functioning of the app.

CORS has also been implemented here in order to facilitate the communication between the back-end and the front-end.

4.3.2   File upload

The application receives the file information and, after performing the necessary checkups, it saves the file on the server with a unique identifier. This, along with an array of the file header, and the first row of data is returned to the client in order for it to display the parameter generation page. It is important to note that Holograph was created with CSV file support in mind only. However, because of how the application has been built, additional file extensions, like XES and MXML, could be supported in the future.

File upload is the only POST method of the application. It handles the file validation and upload. By default, it saves the file to the uploads directory inside the root folder of the application. The file is saved with the same file extension as the uploaded file. The native saving mechanism provided by Python is used.

Upon completion of the process, the application either returns an error string or the name of the file. It is important that the name is returned, since the front-end sends this string to the back-end when requesting mining operations.

4.3.3   General mining

This section is comprised of one file: general_mining.py. Inside this file, three main functions are the ones that execute the general mining activities: get_active_time_array, get_grouped_cases, and get_variants.

When getting the active time array through the corresponding function, it expects to receive two parameters. The first is the filename, which was generated by Python in the file upload process. The second is a Javascript object that contains the needed parameters in order to execute the mining. These are the following:

- CaseID
- Activity
- Resource
- Type (ActiveTime, StartAndEndDate, or Timestamp)
- ParameterOne
- ParameterTwo (only needed if the type of the log equals StartAndEndDate)
- DateTimeFormat

After getting the required fields converted to Python objects, the application goes row by row extracting the information for mining. Each loop pass creates a LogEvent object, which holds information for case identifier, activity, resource, start time, and end time.

Because Python (specifically, the datetime package) uses a very particular format for its DateTime objects, the parsing of date objects is a bit complex. The function tries_date_time_retrieval is used for this, and expects the string provided directly by the data set, a priority format string, and a secondary format string. The priority string is the format string as is provided in the parameters object sent by the front-end. The secondary string is a format string generated by Python using the supplied DateTimeFormat that is

37

better suited for parsing using Arrow and strptime, since it changes the day and year elements to uppercase.

Regardless of the string that is used, Arrow is always tried first, since it provides better support for obscure date formats. If that fails, Holograph reverts to the datetime package and tries the strptime function. For this function to work correctly, the application first converts the format string to something that is usable by strptime. Specifically, it converts regular elements like days (dd) and months (MM), etc. to their equivalent in Python (in this example, (%d and %m).

Finally, before returning the array with the information, it sorts the objects using its startTime property in an ascending order. This is done to aid with the variant extraction.

For getting the variants out of the event log, all of the events are first classified according to the caseID property. With this information, Holograph determines if the case is part of an existing variant, or if the creation of a new one is needed. Every time a variant is created, the associated activities to it are stored in a String array, while the case, which is an array of events, is stored in an array with all of the other cases.

In order to get the variants, the application first get the cases by ID through the get_grouped_cases method. The events that form a case are all grouped in a list, and they are all saved in another list. Finally, the cases are sorted by the start date of their events in order to provide them in chronological order. This method is not only used for inner procedures of the general_mining class, but also as a part of the calculation of results for group mining.

The result of both of these sections are Python objects, either in the form of arrays or custom classes. Although they could be extended to be accessed on their own and generate a JSON representation of the results, this was not part of the scope of the application. For now, then, these methods can only be used and accessed from within the rest of the Python code and they cannot be accessed through any API endpoint.

4.3.4 Results

As previously mentioned, the results section is divided into six different analyses: active time, flow, trace, individual, group by activity, and group by resource. The endpoints for accessing each of the analysis are as follows:

Analysis	Endpoints
Active time	/active-time
Flows	/flows
Traces	/traces
Individual	/individual
Group by activity	/group-activity
Group by resource	/group-resource

Each analysis has been developed under its own file, and they all require as input parameters the arguments sent from the front-end (discussed in the previous subsection), and the filename. In order to avoid mentioning it for every result, it is important to note

that all results are converted to a JSON through the JSONPickle package at the end of the processing.

Out of the different analysis supported by Holograph, active time is the most elementary one. It gets the active time array generated in general_mining.py and gets all the distinct activities in the event log. It loops through all activities and it gets all of the events that involve that specific activity. The active time is added on a per resource basis thanks to the creation of an ActivityResource object for each pair of activities and resources that execute them.

In the case of the flow analysis, the variants generated by general_mining.py are first retrieved. The application analyzes each event part of all the cases inside all variants. The resource is obtained for each event, and three parameters are extracted: time before, time taken, and time after. For time before, all of the times from previous events are added. Time taken takes the current activity and calculates its time. Finally, to get the time after, all the events from subsequent events are extracted and their times added.

This classification is done through the ResourceStatistic class. An array of them is stored in an ActivityStatistic object, which stores all the activity information for all resources. Finally, CaseStatistic stores an array of all ActivityStatistc instances within a variant.

In order to facilitate the calculation, one of the limitations of this calculation is that, for time before and time after, the application does not consider the active time only, but rather the whole time from the start of the process until the end of the previous activity (in the case of time before), and from the beginning of the subsequent activity until the end of the last activity (in the case of time after). If more precise results need to be obtained, one area of improvement would be to calculate the active time for each activity, and then add their active times for the time before and after.

Trace mining is fairly straightforward as well. Just like the previous analysis, the variants are obtained from general_mining.py. All of the resources present in the variant are obtained by analyzing the cases, and then for each resource, the following facts are calculated:

- Quickest
- Slowest
- Below average
- Above average
- Most involved
- Least involved

The exact ways of calculating this have been discussed previously. All of this information is saved in the TraceResource object. The information for all resources inside a variant are stored in a TraceVariant object. An array of these objects is the one that is converted to JSON and returned as a result of the execution of this endpoint.

Individual mining is also a calculation that is not complex in nature. Before starting, it also gets the variants from general_mining.py. Afterwards, it performs the same

operations as flow mining, and gets the same results minus the JSON conversion. All of the variants within the results are looped, and the Division object is introduced. This object calculates the optimal way of dividing the activities in a variant in three. It is optimized to get the highest amount of activities in the middle section. For example, if the variant has five activities, the object will divide the log in activities of 1, 3, and 1. If the log has seven activities, it will divide the activities in parts of 2, 3, and 2.

With this information, the statistics are sectioned into their own divisions and the results are stored into a DivisionResults object. An array of these objects is return when consulting this endpoint.

Finally, group mining starts by getting the grouped cases from general_mining. With this information, there are two different results the application can obtain: by activity or by resource.

Since the group by activity gets all of the pair combinations of a specific activity, this part of the application uses the CaseConnection class to store all two-sided connections. It loops through all the cases and gets the current event and the next event in order to form a connection between its resources. If there is no existing connection, then the application creates one. If there is, then the total time and the number of occurrences is added to the total.

The resource analysis also loops through all the cases provided by general_mining.py, but instead of focusing on the relation between resources that perform specific activities, it focuses on the interaction between a resource and other resources regardless of the activities that they are developing together. For this, the PrimaryResource and SecondaryResource classes are used. A PrimaryResource object contains one or more SecondaryResource objects, which is the one that stores the information about the interaction with the primary resource. Just as the CaseConnection class, information about total time and number of occurrences is also recorded.

Apart from the information sent when consulting the endpoint (such as the name of the fields and the name of the file to be analyzed), currently there is no way of customizing the results that Holograph Back-End provides.

4.4 DEPLOYMENT

Since both tools are developed using fairly common tools (Angular and Python with Flask), deployment wasn't a problematic endeavor. Both approaches have great support for being deployed into Heroku with ease. At first, a dedicated virtual machine on the University of Tartu's network was considered for the deployment of both tools. However, because of the experimental nature of Holograph, Heroku was considered enough for both parts of the application.

Because the tool is not being used by other parties at the moment, both parts are deployed on Heroku's free tier, which includes the following features:

- Multiple deployment methods
- Automated patching
- Sleeps after 30 minutes of inactivity

- 1000 free dyno hours per month (amount of hours the dyno can be active on a calendar month)
- 512 MB of RAM
- One worker

The only limitation that impacts Holograph in a significant way is the sleeping feature. When a Heroku dyno is woken up from inactivity, it takes longer than normal to respond. Combined with Holograph's lenient approach on performance optimizations, the first request to an inactive worker makes the back-end respond slower than usual.

Both dynos in Heroku have been linked to the respective GitHub repositories, and automatic deployments have been enabled. This means that both applications are deployed automatically into Heroku after their master branch in GitHub has been updated.

5 EVALUATION

In this section, a description of the evaluation method is presented. This applies to both the proposed method and Holograph, the implementation done to validate the results.

5.1 EXPERIMENTAL CONTEXT

As stated in the first section, in order to assess the effectiveness of the study, the following research questions needed to be answered:

- RQ1. What are the current methods for analyzing resource impact on business processes and what are their flaws?
- RQ2. Given an event log, what kind of analysis needs to be done to it in order to discover the impact of a resource in the execution of an event log?
- RQ3. What improvements do the proposed method offer compared to the current methods?

Since the result of this thesis is a new way of analyzing the impact of resources in the execution of business processes, the first question deals with the research of the current alternatives in the market. By analyzing what has been proposed by researchers before, we can evaluate new opportunities for improvement and address problems of the previous researchers.

After proposing a solution that try to avoid the shortcomings of other tools, while also providing an innovative way of comparing logs, the approach has to be measured in its usefulness to businesses. Is the result good enough to be used by companies that want to address problems in their current processes? Can the results be useful enough so that the companies execute redesign plans based on the results of the proposals? What is the confidence the tool gives so that decision making is influenced by this approach?

Finally, the approach has to be evaluated in its effectiveness with users. This means that they not only should be able to find better results through it, but also the implementations should be built in a way that a user has an easier time using it and be quicker while operating it. The ideal tool would bring improvements in both the results that are output by it and in user experience.

5.2 EXPERIMENTAL DESIGN

In order to conduct experiments and iterate on the features that the approach and the tool should have, there are two different dimensions that were considered. The first one is datasets, since they are needed in order to continuously test Holograph. By choosing good datasets, not only we can obtain more information about what kind of situations Holograph must be prepared for from a development standpoint, but also do an internal validation of the tool and assess if the results that are being obtained are useful.

The second one is the actual evaluation done by people unrelated to the development of the approach and the tool. This has been done through a questionnaire with 12 questions designed to extract information about the current state of Holograph, its strengths and weaknesses compared to Disco, and possible paths of future development.

5.2.1 Datasets

The datasets used needed to have traces with a significant process that has resources heavily involved. Ideal datasets for the experiments would be the ones in which the resources heavily affect the outcome of the process. The minimum requirements for the dataset are:

- Case ID
- Activity
- Resource
- Start time
- End time

Especially important are the last two parameters. It is mandatory that all datasets used should have a start and end time. This is essential in order to calculate the active time. Although it could also be supplied in other ways, such as having its own field in the dataset, the implementation of Holograph does not support that arrangement.

For future implementations, the possibility of using a timestamp only instead of a start and end time was also discussed previously. It was determined that the best way of doing it would be to treat the timestamp of the previous task and the one of the current tasks as the time taken by the resource. However, it was also mentioned that results wouldn't be as exact as having the explicit start and end times, so it is discouraged.

In this case, the most prominent log used was one provided with the Minit tool in their example project. This log has information about an invoice approval process, and boasts 66074 entries of information across 297 different variants.

Other logs that were also used through the development of Holograph were found on public websites and online repositories that deal with Business Process Management.

Widely-recognized logs like the different annual versions of the BPI Challenge were not able to be used because they lack a start and end time column. Instead, they rely on start and end time for their logging.
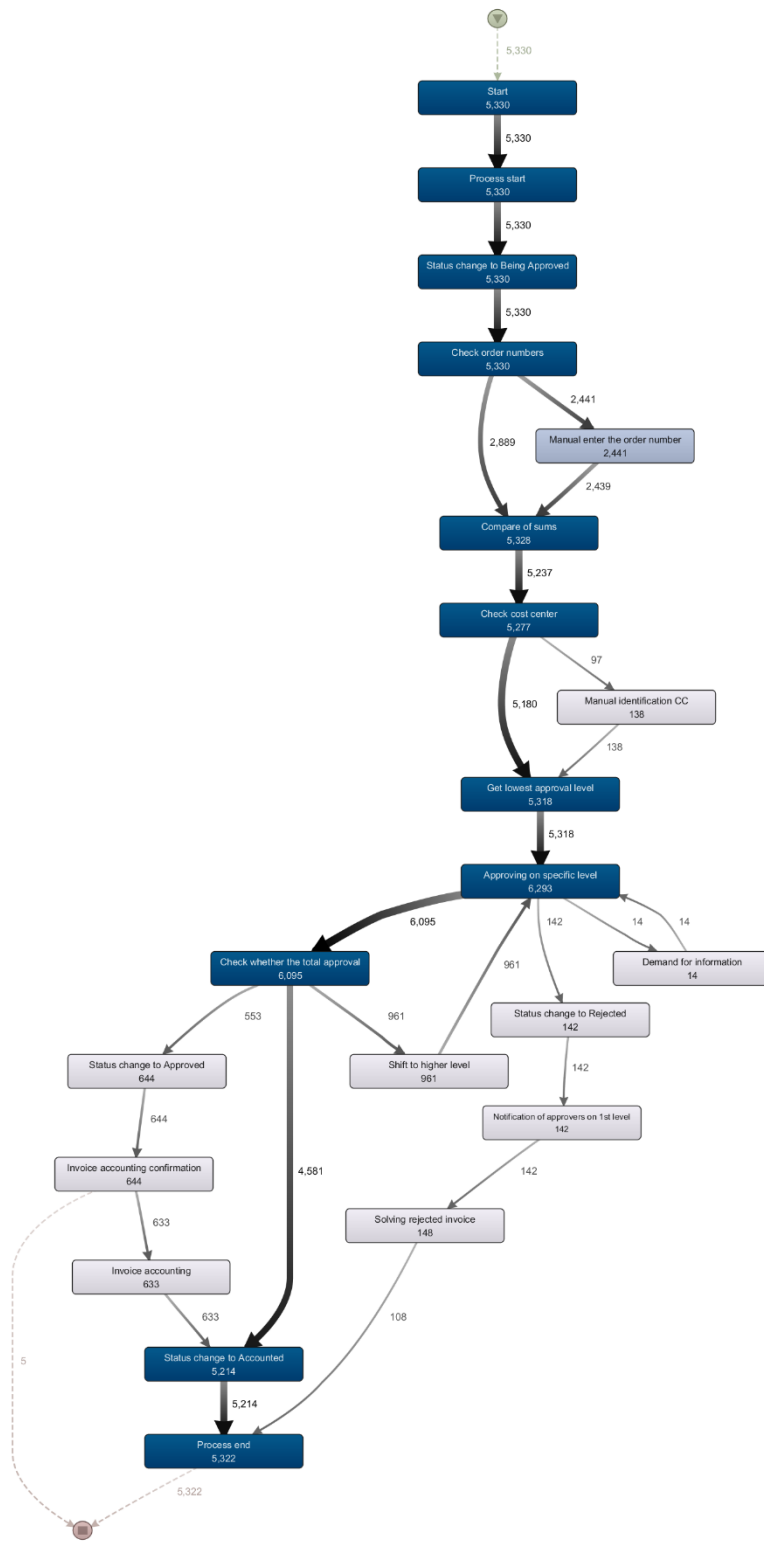
43

Figure 16. Minit's event log

5.2.2   Implementation Evaluation Measures

The main measures to be extracted from the experimentation with previously researched tools and the proposed implementation are the following:

- Correctness: How accurate are the results provided the tool. The proposed software should provide results that are in line to what the event log displays.
- Confidence: How confident is the user regarding the results that are being displayed. The tool must present results in a way that aids the user in the decision-making process of the company and said user must be sure that the discovered patterns are the correct ones.
- Usability: How easy or complicated does the researched individuals find the proposed tool. This includes how the information is displayed, if the displayed information is appropriate for the kind of research that is being developed, among others.

Since this is a new, experimental tool, factors that would normally be evaluated for commercial tools, such as user interface and performance, are not measured. The aim of the implementation is to provide a new approach to work with, so, although some optimizations were done, several parts of the code could be optimized to make response times shorter, especially for the Python back-end.

5.2.3   Implementation Evaluation Procedure

In order to evaluate the effectiveness of the tool, an experimental approach will be used. Two groups of people familiar with data mining tools were selected in order to execute the experiment in a way that knowledge would not get in the way of the times and performance. The 38 subjects were part of a business process management course taught in Colombia by professor Marlon Dumas. All of the participants are pursuing a master's degree in Management of Information Technology, and are white-collar workers with low managerial positions.

Both groups were given a physical questionnaire with a total of 12 questions. A copy of the questionnaire can be found in the Appendix section. The aim of this questionnaire was to get feedback about how Holograph Front-End presents its information, the clearness, and the usability of it, and compare it to what Disco proposes as a solution for similar situations.

In order to prepare the questionnaire, the previously discussed log obtained from Minit was chosen and sent to both Holograph and Disco. Since Disco does not have a function to mine resource information, when uploading the file, the Resource field was set as an activity parameter. This way, the in-depth analysis normally provided for activities was available for resources. Holograph was used as it was intended and no additional configuration was needed.

5.2.4   Implementation Evaluation Questionnaire

Because the location where the questionnaire was done, the questions were translated to Spanish. These are the 12 questions asked in English:

1. In general, the average cycle time of a process is determined primarily by the execution time of the activities in the process. (Scale)

2. In general, the average cycle time of a process is determined primarily by the time the handoffs between the involved resources take. (Scale)
3. Analyzing the impact of a group (like a department or team) is as important as analyzing individual impact on a business process. (Scale)
4. Which analysis do you find better to analyze the performance of resources across different activities within a process? (Binary choice)
5. Which analysis gives more targeted information about how the performance of a resource compares to other resources? (Binary choice)
6. Which analysis would help you more in taking a decision about the business process, either by changing the resources that execute tasks or by changing the task itself to make it easier? (Binary choice)
7. Which analysis do you find better to analyze the impact of a resource across a business process? (Binary choice)
8. Which tool do you think is better for comparing between two or more resources executing a variant? (Binary choice)
9. In order to take business decisions about a process, do you prefer to have aggregated information that shows a trend, detailed information that doesn't show trends, or both? (Multiple choice)
10. Which analysis do you find better to analyze the handoff between different resources across a business process? (Binary choice)
11. Which analysis gives you more information in order to take a decision about handoffs within a business process, either by changing the resources that take part in a handoff or by restructuring the process to avoid handoffs between specific activities? (Binary choice)
12. If you want to analyze the impact that a specific handoff has on a business process, do you find more useful to know the amount of times that a handoff occurs, the average execution time that the handoff, or both? (Multiple choice)

Questions that are marked as Scale had five choices, in which the options went from strongly agree to strongly disagree. The ones that are marked as binary choice had two choices: Visualization 1 and Visualization 2. Which visualization referred to which tool was defined by the version that the subject got, but if, for example, the individual had the version with Holograph screenshots first, then all binary questions of the questionnaire referred to Holograph as Visualization 1 and to Disco as Visualization 2. This was done to prevent the results from being skewed towards one of the tools because of its order. The selection of which subject got one version or the other was done randomly. Finally, questions marked as multiple choice had the options exposed in the question itself. They let the subject answer between option A, option B, or both.

The first three questions deal with general opinion about the execution of processes and the impact that a resource (or a group of resources) can have in the execution of a business process. Questions 4 through 6 display information about resource performance. A screenshot of resource statistics was displayed for Disco, and the active time analysis was displayed for Holograph. Questions 7 to 9 deal with the impact of resources in flows. A screenshot of the flows within a variant is shown in Disco, while the flow analysis is shown for Holograph. Questions 10 to 12 deal with the analysis of handoff impact. A screenshot of a directly-follows graph with resources is shown for Disco, and the group

performance by activity is shown for Holograph. Finally, the subjects were asked to write the time at which they finished the questionnaire.

After the questionnaire was done, the results of both was analyzed in order to evaluate how close they are to the expected result. The experiment should be considered a success if it provides insights on what do other professionals identify as useful in Holograph and if their opinions are aligned with what the approach is aimed to do. Ideally, the subjects would prefer the implementation of Holograph over the tools that Disco provides. However, an inclination towards Disco would also provide valuable information about refinements to the approach and the tool, which would be very valuable towards both the development of Holograph and future projects that use these ideas as a base for new implementations.

6 RESULTS

After gathering the filled questionnaires explained in the experimental design, the results of them were computed and analyzed for extracting information useful for the proposed approach and for Holograph, the implementation of said approach.

6.1 GENERAL OPINION

General opinion questions dealt with specific queries regarding the factors that influence the execution of business processes. Specifically, they dealt with factors such as active time, handoffs, and resources.

Question	In general, the average cycle time of a process is determined primarily by the execution time of the activities in the process.
Answers	(A) Completely agree (B) Partially agree (C) Neither agree nor disagree (D) Partially disagree (E) Totally Disagree (F) I don't know

Question 1

	A	B	C	D	E	F
Subjects	8	13	1	3	13	0
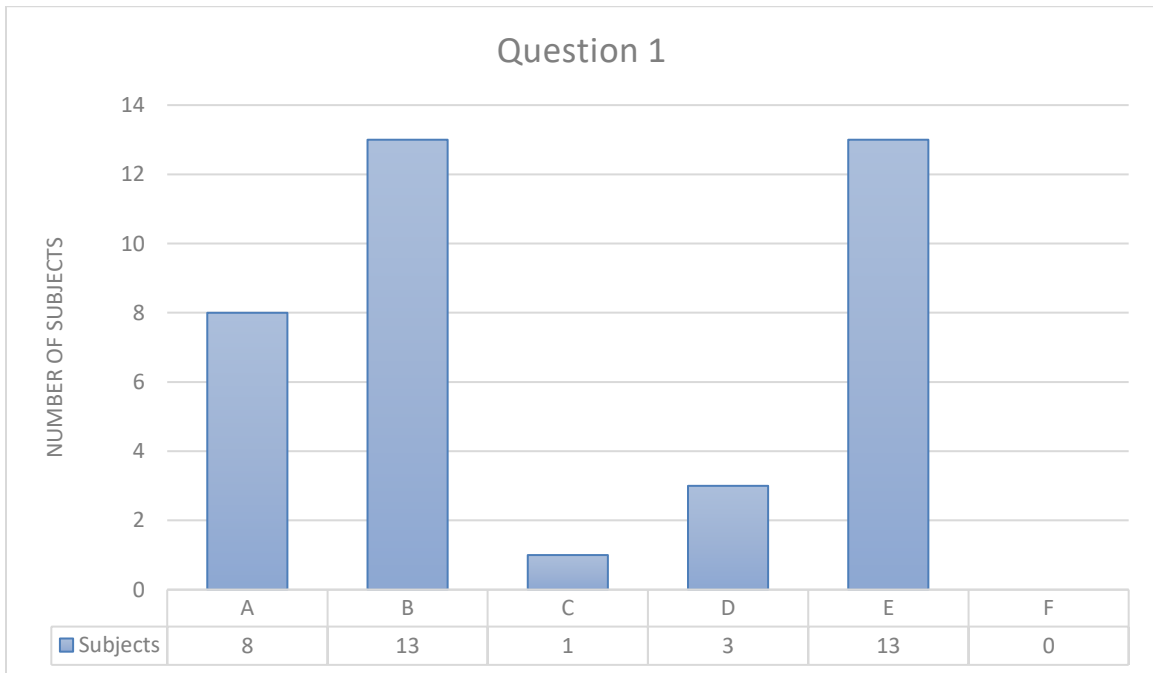
Figure 17. Answers for Question 1

In this graph, we can see a strong preference towards the agreement of the statement of the question. A total of 21 people was skewed towards the positive answer, while 16 preferred to answer negatively. It is also apparent that people that answered positively wasn't fully convinced, since the majority of answers come from the "Partially agree" option.

Question	In general, the average cycle time of a process is determined primarily by the time the handoffs between the involved resources take.
Answers	(A) Completely agree (B) Partially agree (C) Neither agree nor disagree (D) Partially disagree (E) Totally Disagree (F) I don't know
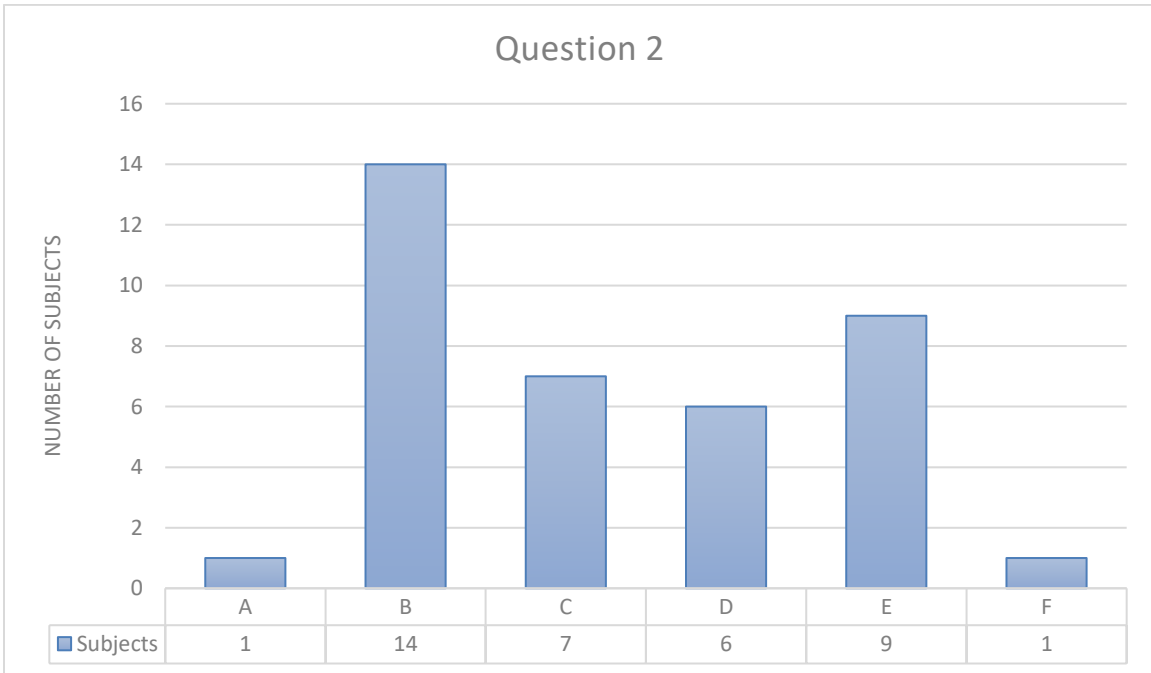
Figure 18. Answers for Question 2

This question presents even more indecision than the first one, with answers all across the options. A curious fact is that if the positive and negative options (A + B and C + D) are combined, we get a result of 15 for both. It is safe to say that this question is highly controversial, and this starts to spark some questions regarding the perceived usefulness of the handoff analysis in Holograph.

| Question	Analyzing the impact of a group (like a department or team) is as important as analyzing individual impact on a business process.
Answers	(A) Completely agree (B) Partially agree (C) Neither agree nor disagree (D) Partially disagree (E) Totally Disagree (F) I don't know

49
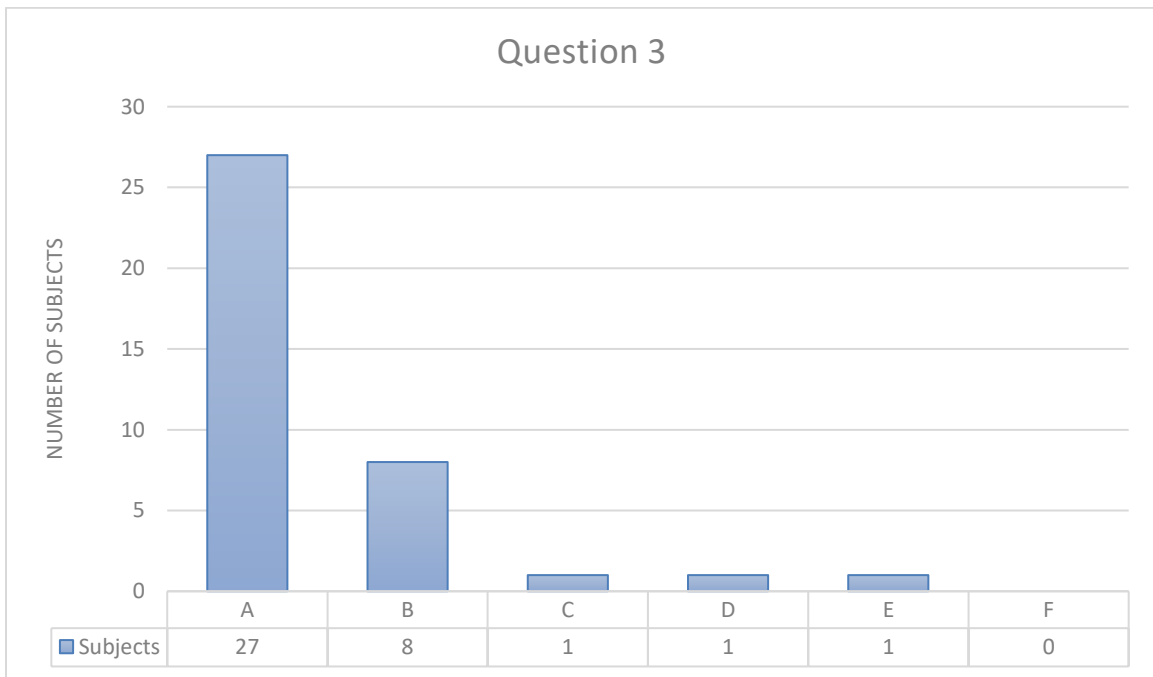
Figure 19. Answers for Question 3

Unlike the previous two questions, the third question presents highly positive results, with 35 individuals selecting a positive answer. This might seem to conflict with the overall indecision on question 2. This observation will be discussed in the Discussion section of this chapter.

6.2 RESOURCE PERFORMANCE

Resource performance presented the subject with screenshots of performance analysis made in Holograph and in Disco, and asked them which one did they prefer. The questions dealt with subjects like visualization preference and confidence in the tool for decision making.

Question	Which analysis do you find better to analyze the performance of resources across different activities within a process?
Answers	(A) Visualization 1 (B) Visualization 2
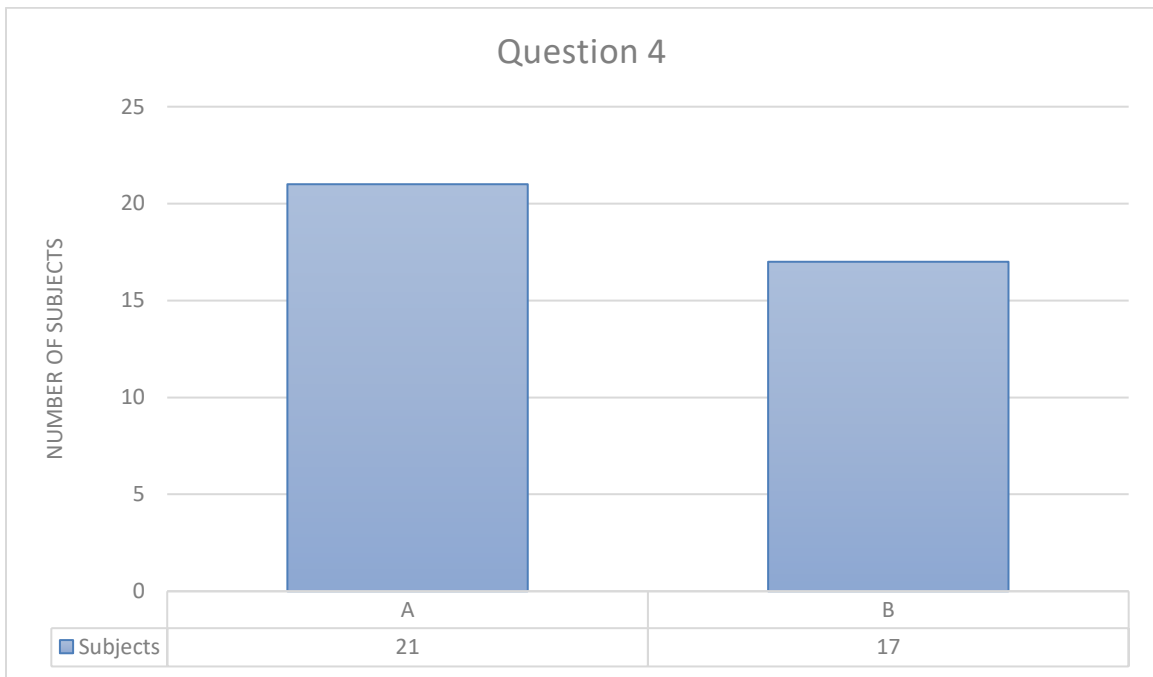
Figure 20. Answers for Question 4

With 21 votes, the surveyed individuals preferred Holograph's way of analyzing the performance of a resource. However, with 17 votes, Disco and its tools for analyzing performance were also highly favored.

Question	Which analysis gives more targeted information about how the performance of a resource compares to other resources?
Answers	(A) Visualization 1 (B) Visualization 2
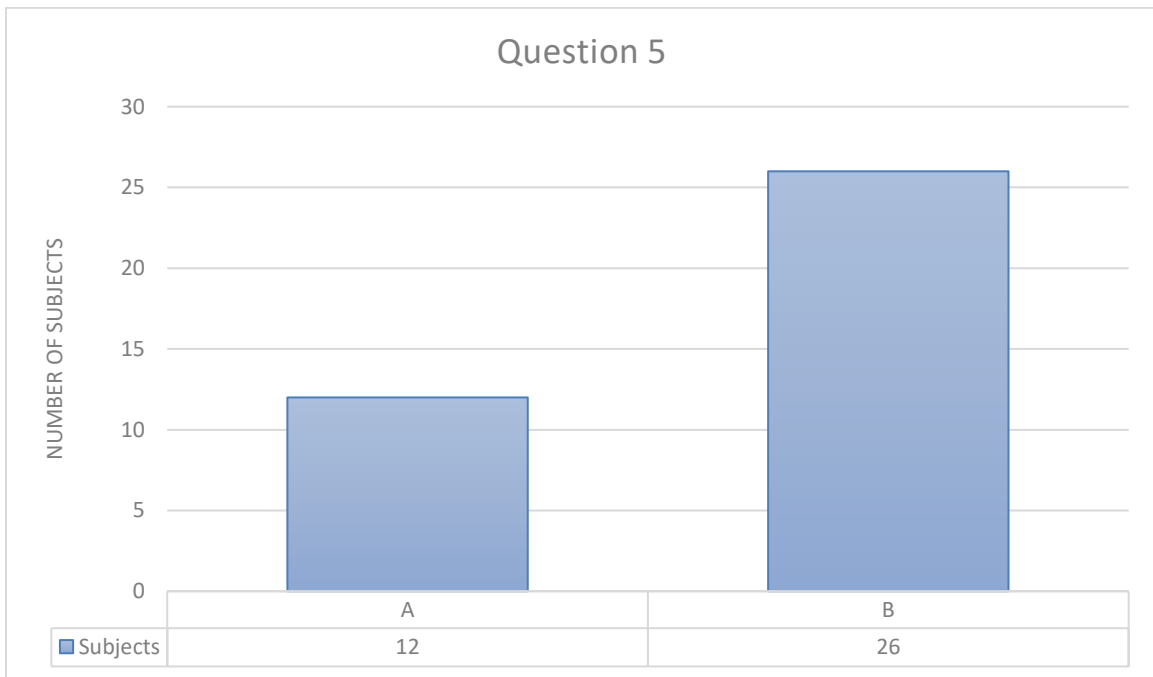
Figure 21. Answers for Question 5

More than 68% of the subjects preferred Disco's way of displaying detailed information compared to Holograph's tendency of presenting a summary of it. Even though Holograph's way is preferred when dealing with aggregate information for activities, Disco's way of detailing information seems to be more useful when comparing among resources.

Question	Which analysis would help you more in taking a decision about the business process, either by changing the resources that execute tasks or by changing the task itself to make it easier?
Answers	(A) Visualization 1 (B) Visualization 2
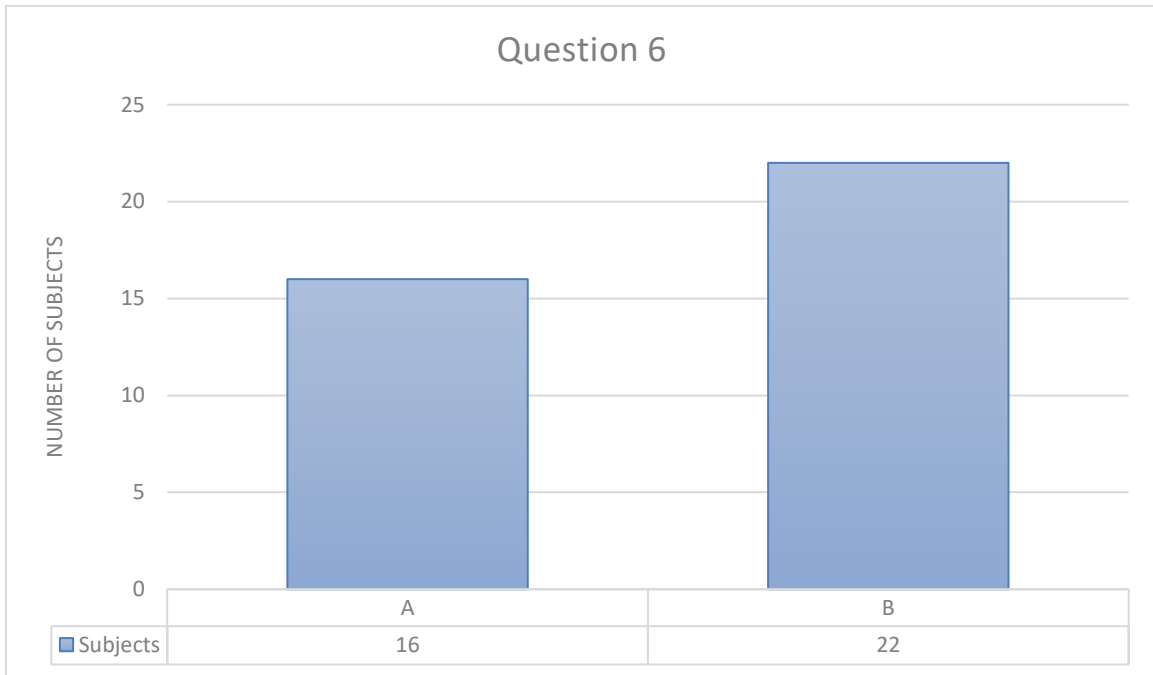
Figure 22. Answers for Question 6

This is the first out of two questions that were asked regarding the confidence that a subject would have in the results that Holograph displays and if they could be used to reach a decision regarding a business process. Holograph failed to capture the majority in either of them. This phenomenon will be discussed in more detail in the Discussion section.

6.3 IMPACT OF RESOURCE IN FLOWS

This section dealt with the impact that a resource can have in a business process. The subject was presented with screenshots of Holograph and Disco in which this analysis was performed, and questions about preferences were asked. Also, there was a question about preference of information being displayed, in order to make decisions for future developments of the tool.

Question	Which analysis do you find better to analyze the impact of a resource across a business process?
Answers	(A) Visualization 1 (B) Visualization 2
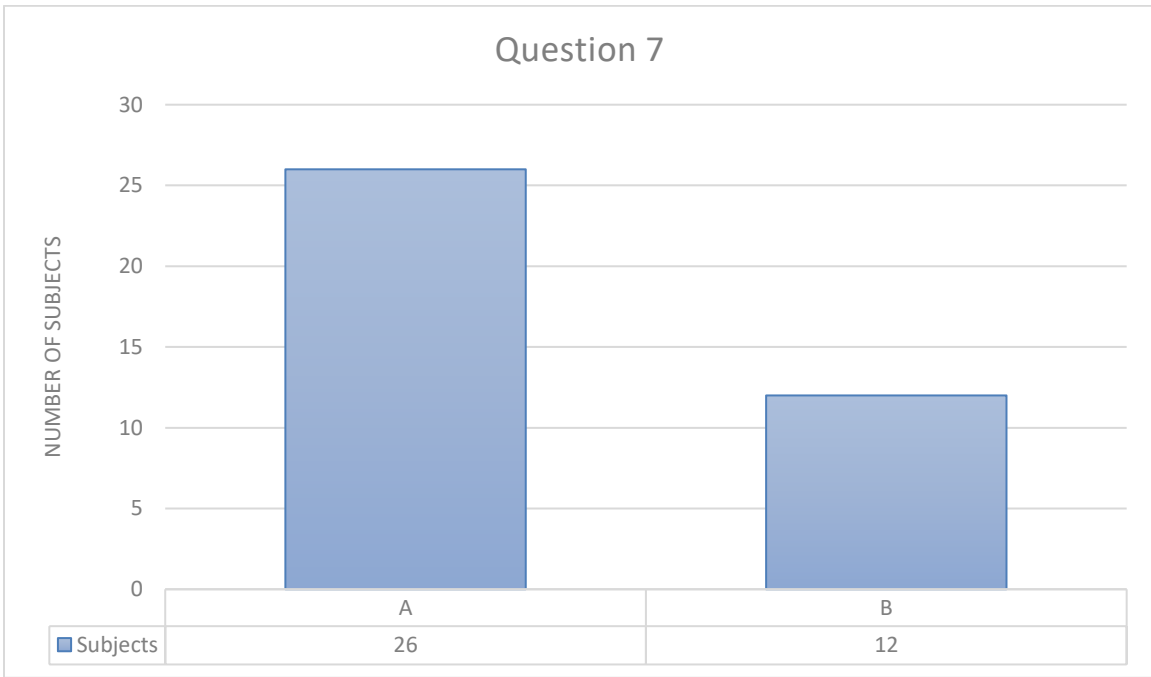
Figure 23. Answers for Question 7

This section is rather positive for Holograph, since this is where it has shined in comparison to Disco, according to the subjects. With 26 votes, Holograph's detail in the flows screen was more appreciated than Disco's approach.

Question	Which tool do you think is better for comparing between two or more resources executing a variant?
Answers	(A) Visualization 1 (B) Visualization 2
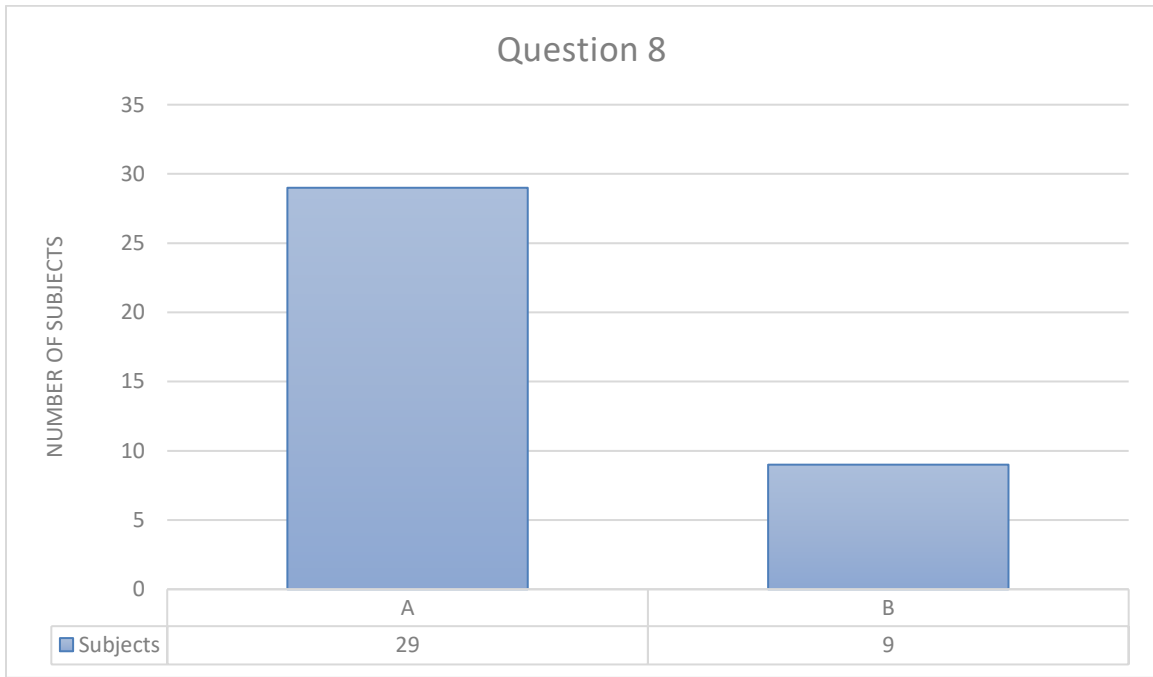
Figure 24. Answers for Question 8

This is the biggest difference between both tools. Its approach to compare different resources for the same variant were appreciated by the surveyed individuals.

Question	In order to take business decisions about a process, do you prefer to have aggregated information that shows a trend, detailed information that doesn't show trends, or both?
Answers	(A) Aggregated information (B) Detailed information (C) Both
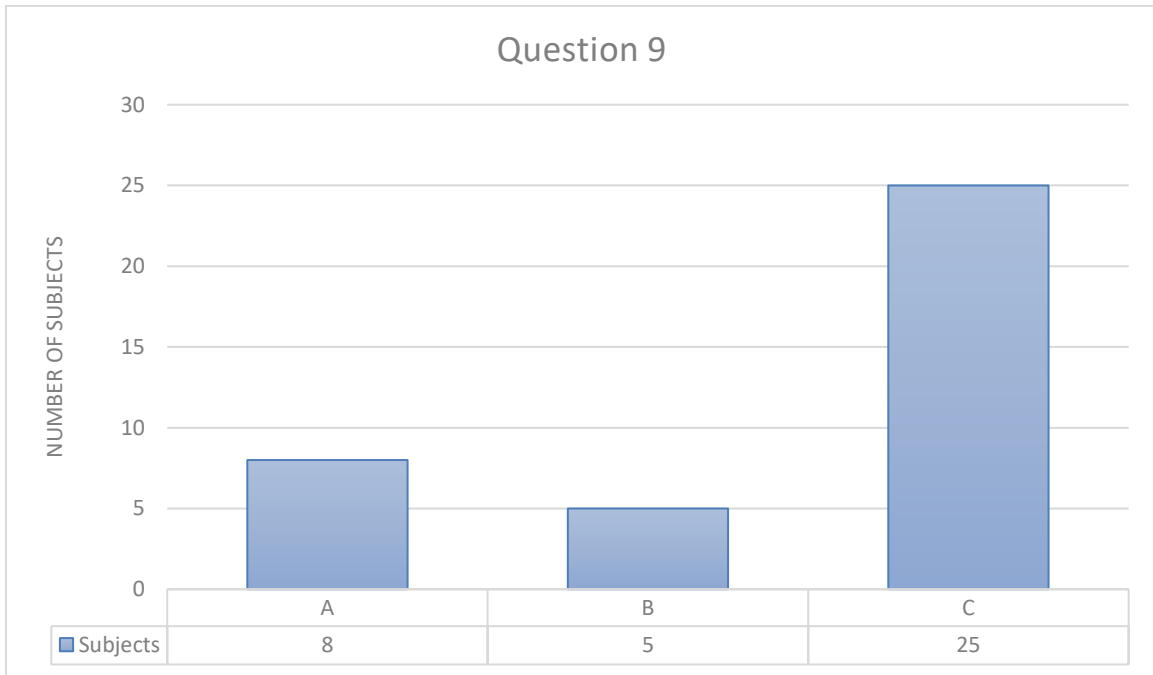
Figure 25. Answers for Question 9

This question was put in place in order to verify which approach is preferred by the surveyed individuals. While Holograph prefers to show aggregated information, the data that Disco presents tends to be much more detailed and give the ability to dig into specific statistics. The subjects overwhelmingly chose that they prefer both approaches, which means that an approach that combines the results that both Disco and Holograph is preferred. The implications of this will be discussed later.

6.4 HANDOFF IMPACT

The final section was reserved for the analysis of the impact of handoffs in a business process. Just as the previous version, questions were asked about the preference between Holograph and Disco through screenshots presented in the questionnaire, as well as a question about preferences in the display of information.

Question	Which analysis do you find better to analyze the handoff between different resources across a business process?
Answers	(A) Visualization 1 (B) Visualization 2
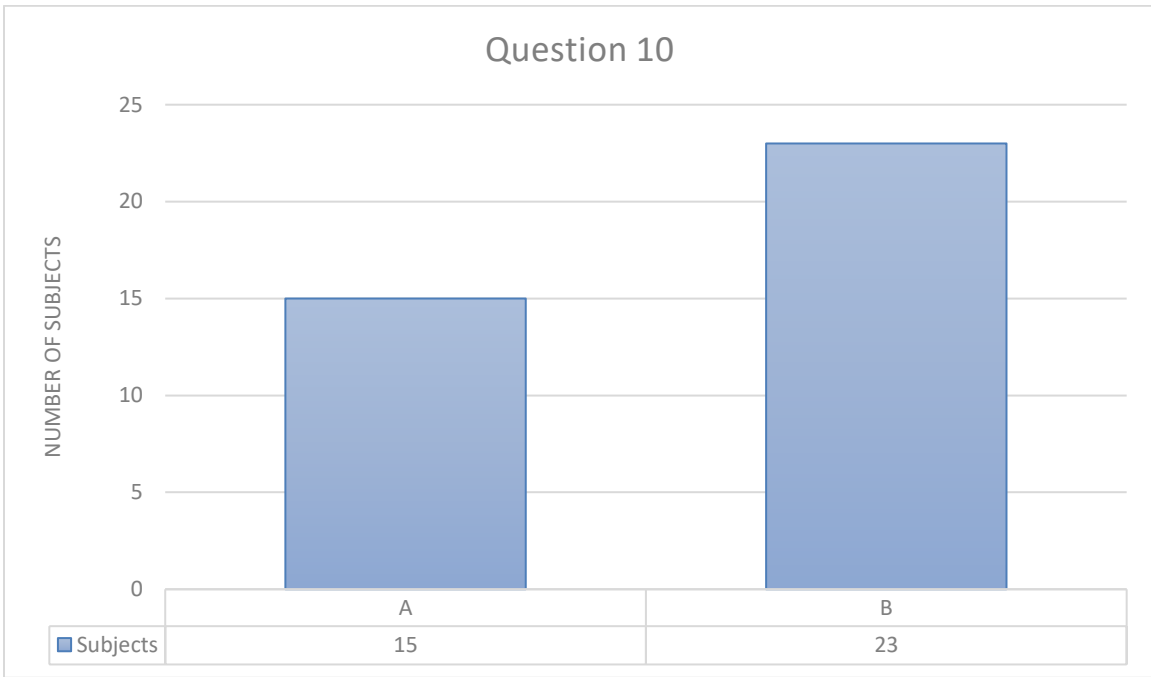
Figure 26. Answers for Question 10

Handoff information is presented in Holograph just as other views, with a bar graph plus information at the bottom regarding the involved resources, average time, total time, and number of occurrences. Disco presents a directly-follows graph that visually displays how the resources interact between each other. The surveyed people preferred the appealing visual way that Disco uses for this analysis, edging Holograph by eight votes.

Question	Which analysis gives you more information in order to take a decision about handoffs within a business process, either by changing the resources that take part in a handoff or by restructuring the process to avoid handoffs between specific activities?
Answers	(A) Visualization 1 (B) Visualization 2
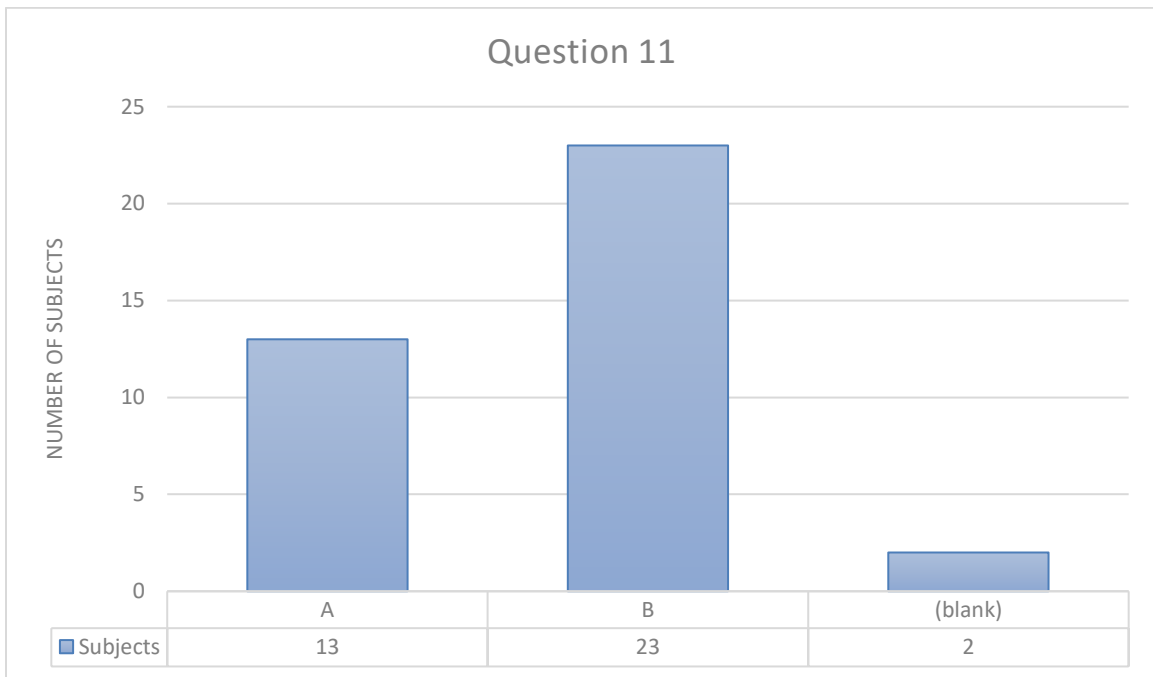
Figure 27. Answers for Question 11

This is the second question regarding the confidence that each tool projects. This is where Disco's elaborate map is useful towards projecting an image of confidence and trust. Over 60% of the subjects chose Disco. Also, it is important to note that this is the only question that had blank answers.

Question	If you want to analyze the impact that a specific handoff has on a business process, do you find more useful to know the amount of times that a handoff occurs, the average execution time that the handoff, or both?
Answers	(A) Amount of times (B) Execution time (C) Both
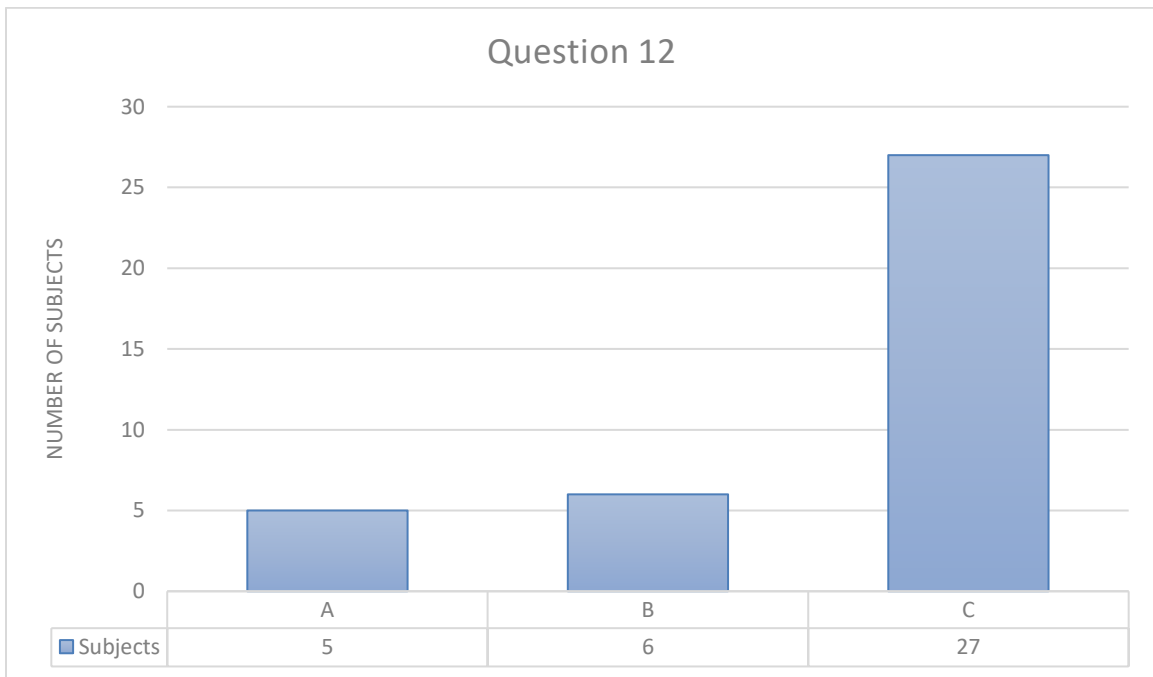
Figure 28. Answers for Question 12

As done for the previous section, this question was asked to validate what is being displayed by Holograph. An overwhelming majority preferred to have both the amount of handoffs occurred and average execution time. It is strange, then, that while Holograph displays this, it lost the dispute against Disco. This will be discussed later as well.

6.5 TOTAL TIME

A measurement of the total time that the subjects took to finish the questionnaire was also done. As expected, the individuals took approximately 15 minutes to finish it. This was the goal set for the questionnaire, since its primary aim was to be as easy as possible for the subjects. The longest time was 27 minutes.
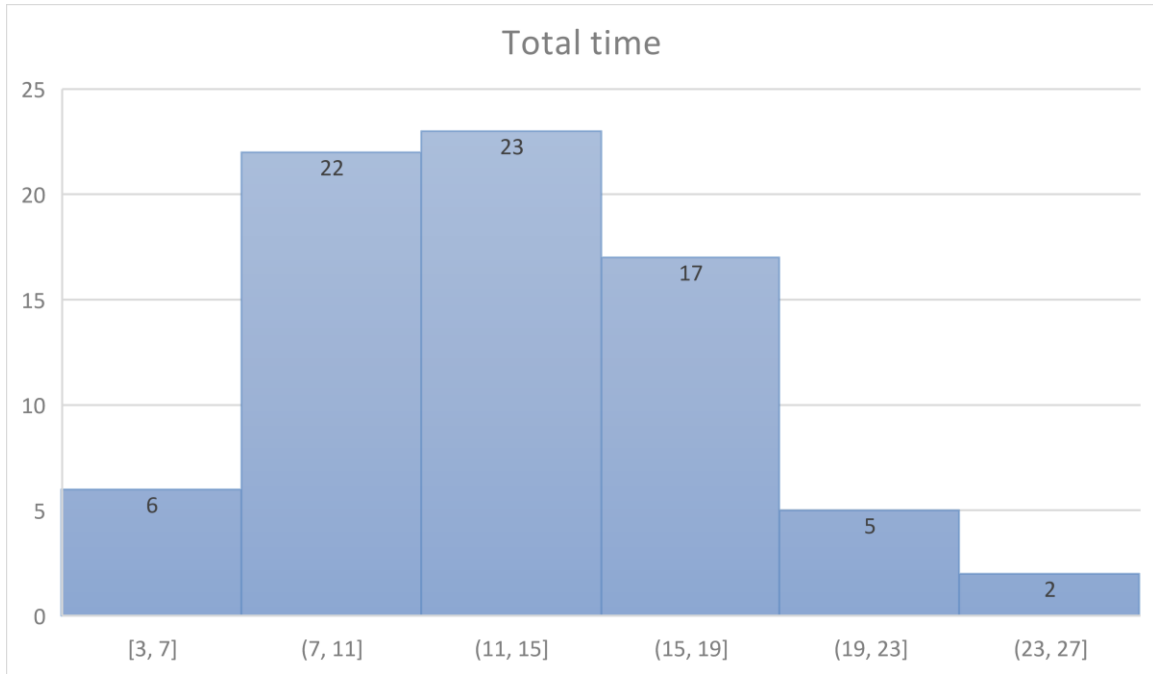
Figure 29. Total time the subjects took to answer, in minutes

It is important to note that six people took between 3 and 7 minutes. This is probably too short to answer the questionnaire in a conscious way, so hopefully these results haven't affected that much the overall numbers.

6.6   OTHER ANALYSES

As previously said, the questionnaires had two different versions. One featured Holograph's screenshots first (called Group 1 from this moment onwards), and the second one had Holograph's screenshots second (called Group 2). This was done in order to minimize the impact that the order of presentation of the tools could have on the overall results. In order to verify that this has been successful, several analyses on the results by both groups was done.

6.6.1   T-Test

The number of surveyed individuals that chose Holograph in both groups were added, and the values were put under a t-test analysis in order to determine if the groups are the same between each other regardless of the order of the visualizations.

Question	Group 1	Group 2
4	12	10
5	7	5
6	10	6
7	13	13
8	14	15
10	6	9
11	5	8

This resulted in the following numbers:

Parameter	Result
t-value	0.07426
p-value threshold	0.05
p-value	0.471013

Because the p-value is a lot higher than 0.05, we cannot say that the two different groups are either equal or different. This high p-value can be attributed to the low amount of questions featured in the survey. Only a total of seven questions were affected by the different order of the visualizations.

As a result of this, a more manual, specialized analysis was done in order to delve deeper into the differences between both groups and determine what is the result of having two different questionnaires.

6.6.2 Manual analysis

The data obtained through the questionnaire is presented below. The labeling of the groups is the same as the one used in the previous section. Since both groups had 19 answers unless otherwise noted, then only the Holograph percentages have been presented here.

Question	Group	Total answers	Percentage	Difference
4	1	12	63.16%	10.53%
	2	10	52.63%	
5	1	7	36.84%	10.52%
	2	5	26.32%	
6	1	10	52.63%	21.05%
	2	6	31.58%	
7	1	13	68.42%	0%
	2	13	68.42%	
8	1	14	73.68%	-9.65%
	2	15	83.33%	
10	1	6	31.57%	-15.8%
	2	9	47.37%	
11	1*	5	29.41%	-12.7%
	2	8	42.11%	

*Only 17 answers were valid, two were blank.

Measurement	Value	On question
Average difference	0.56%	-
Highest difference	21.05%	6
Lowest difference	0%	7

Measurement	Average for Group 1	Average for Group 2
Q4-6	50.88%	36.84%

61

Q7-8	71.05%	75.88%
Q10-11	30.49%	44.74%
Overall	**50.81%**	**50.25%**

From these results we can draw several conclusions. First, we can see a variation between the results according to the order in which the screenshots were presented. However, they do not follow a logical structure. For the first three questions that asked for comparisons, we can see a positive difference between the two variations. This means that more people chose Holograph when the screenshot was placed first. Then, the second group displays closer numbers between each question, with question 7 displaying a 0% difference between both groups. Finally, for questions 10 and 11, we see a negative difference, signaling that Holograph was more chosen on the questionnaire that had its screenshot in second place.

This means that, out of seven questions, Holograph was favored in three by being shown first, on three when it was shown second, and on one it was absolutely irrelevant. When we analyze the average of responses for Holograph, we can see that the numbers are almost the same, meaning that the respondents chose Holograph for 50% of the questions on average on both versions. Since there is not a structured tendency towards one or the other, and the preferences of the surveyed people seems to be more related to the groups of questions rather than the order in which the images are shown, we can conclude that the order could have been a factor that changed the answers by a small margin, but it wasn't big enough to be considered significant or relevant for the study.

Finally, a small analysis was made regarding the amount of times people chose Holograph over Disco.
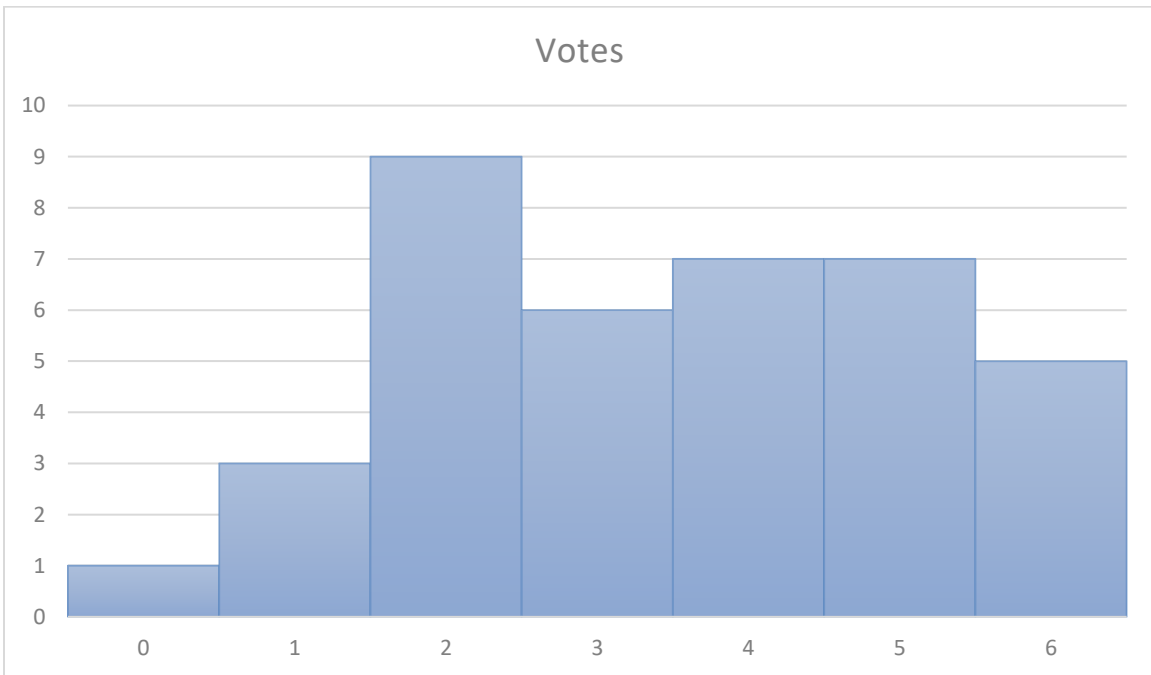
Figure 30. Number of votes for Holograph within the same questionnaire

We can see that the scores are pretty evenly laid out. Exactly half of the participants chose Holograph for four or more of the questions, while five of them preferred Holograph for almost all questions. One of the surveyed individuals who chose Holograph for six questions didn't answer question 11, so that data point is incomplete. Only one person didn't choose Holograph on any of the available questions. This corroborates the previously seen results, where Holograph is preferred mainly for some analysis but not others. More about this will be explained in the Discussion section.

6.7 DISCUSSION

The overall result that we can get from the questionnaire answers is that Holograph is still not in the development point that it should be to be considered a significant improvement over current tools in the market. However, even with a small development process, we can see the validation of several of the ideas and visualizations that were proposed in the new approach. There are several very encouraging results to be extracted.

Disco does not have a native way of presenting resource impact, performance, or other similar information. It is a tool primarily aimed at analyzing activities in an event log and infer a process map out of it. Unlike Minit, Disco does not have a social network map or something similar to make analysis of resources easier.

There is a tab under statistics that presents some basic information of the resource, such as relative frequency, median duration, among others. However, this is done on a log level, not on a variant level (or activity level, in some cases), as Holograph tries to do.

In order to achieve results that are more aimed towards the analysis of resources, the analyst has to use a workaround. When importing a log into Disco, the tool lets you select which field needs to be assigned to activities. By selecting the resource column as an activity, Disco presents the same in-depth analysis with resource information.

The biggest benefit of this information is the extraction of a map that shows how each resource is connected to others (in other words, the handoff). Since Disco natively provides information about how different resources interact with each other timewise, this is already a valuable source of information. This could be the beginning of an analysis about the impact of resources in a process.

However, one area that Disco lacks is the analysis of resources in a more specific way. Of course, Disco was not created for this kind of analysis in mind, so it makes sense that it is not suited for it. For example, there is no way of analyzing how the time of execution of a process changes from before a resource is involved in it versus after it has done its work. It is not possible to see how good or bad a resource is depending on where the resource is located within the process. These are some of the limitations that the proposed approach addresses.

6.7.1 General opinion

By providing users with ways of analyzing resources impact on flows, variants, and handoffs, Holograph aims to cover the areas that are not fully addressed by current implementations. This is why results for the first three questions are encouraging, since

they present a mostly positive result regarding the need of analyzing the impact of resources.

A total of 55 percent of the surveyed individuals agreed that the human resources can impact the total time in which a task is executed. This means that, since the time that an activity takes is partly influenced by the individual itself, the subjects are interested in analysis that would yield the impact of said individual in a business process. This question validates the different options for visualizing individual impact in an event log that we can find in Holograph.

However, we also know that the time an activity takes is also affected by the handoff. This means that the output of previous activities directly affects the execution time of subsequent activities. For this, question 2 and 3 were asked.

Question 2, which dealt with the impact of handoffs in the execution of business processes, presented a more polarizing result than the first one. A total of 38% of the surveyed people think that handoffs do affect the outcome of an activity, while the same amount of people are either fully or partially convinced that handoff is not an important factor. Finally, 18% of the subjects weren't sure if handoffs really affect the execution time of an activity, leading to the highest percentage of neutral answers out of the three questions that had a scale answer possibility.

The primary reason for this polarization of results is perhaps the wording of the question. The word "handoff" is tricky to translate to Spanish, so the word "relevo" might not mean exactly the same as it was intended. This led to unfamiliarity with the word and the whole question, which in turn yielded a high amount of indecision. This is why a total of 71% of the candidates answered with options B, C or D.

Another reason for this seemingly strange result is that, even though the subjects are highly trained in their areas, their knowledge of business process management is still growing. There is a chance that they don't fully grasp the impact that a handoff has on an activity. Instead, they mostly correlate the active time with the resource that executes it. This also explains the high results for question 1 regarding the impact of resources in an activity.

Finally, one of the indicators that address the curious results of question 2, is the results of question 3. This question displays more logical results and confirm that the answers for question 2 should have been more approximated towards a positive answer.

Question 3 shows one of the most dominant results of the whole questionnaire, with 92% of the subjects answering positively, and 77% of them being completely sure about the need of analyzing groups of resources apart from individual impact. This result also validates the numerous analysis that Holograph proposes for groups of resources, and the need for tools that display this information.

Regardless of the peculiar results of question 2, the first section as a whole is positive towards the existence of an approach to get information about resource impact. With highly positive results for questions 1 and 3, and satisfactory answers on question 2, we can conclude that resource impact, both individual and in groups, is an area that needs

to be addressed and that it is recognized by the community as one of the biggest factors in the execution of a business process.

6.7.2 Resource performance

In order to analyze the individual impact of a resource in a process, the performance of a resource should be the starting point. It is important to compare the performance of different individuals in order to start to assess possible ways of improving the overall process.

This is why Holograph's first analysis was the comparison of active times according to activity. There is the possibility to compare them by using different graphs and to change the analysis between total time and average time. Disco also has resource performance analysis natively (no need to input the resource parameter as an activity).

Question 4 of the questionnaire asked the subjects to compare the visualization that Holograph provides vs the analysis that Disco displays when selecting resource performance. The comparison needed to be made on the basis of which tool was better for analyzing the performance of resources across different activities.

A total of 55% of the individuals favored Holograph's bar graph over Disco's list of information. Since Holograph provides a very easy way of changing the activity and visualization aspects of the graph, it is logical that Holograph was favored in this particular question.

However, question 5 asked about which tool provides more detailed information about the performance of a resource and how it compares to other resources, and here Disco's analysis shined. A total of 68% preferred Disco's list of resources with different durations rather than Holograph's charts.

Finally. Question 6 asked about the confidence that the subject would put on each tool in order to make a business decision about a process. Here, Disco again edged Holograph with 55% of the total votes.

Although Holograph displayed healthy numbers for the fourth question, Disco's victory on questions 5 and 6 shows flaws in Holograph's way of displaying information to the user. The analysis they do are somewhat similar, but Holograph, for this particular screen, prefers to display synthesized information rather than a list of data. This did not go down well with the surveyed students, which seem to prefer a bigger amount of information at their disposal. Apart from that, it seems like the average and total time by resource is not enough for Holograph, and additional parameters such as relative frequency and duration range need to be displayed to provide a more complete set of data.

As the limitations section addresses, one of the biggest disadvantages that Holograph presented here is that the question has been asked on paper. There are several different visualization options in Holograph which could have tipped the scale a little bit more towards Holograph, but the fact that the users need more information in order to find this analysis a little bit more useful is a great piece of information that should be applied to future development of Holograph, especially since the goal of it is to aid in the decision making of companies. However, as it stands, the initial visualization for it is shaping up to be really useful towards the overall resource analysis.

65

6.7.3   Impact of resources in flows

After analyzing the performance, one of the most important aspects that impact a process is how the involvement of a process affects a specific flow. If a resource being involved affects either positively or negatively the execution, then additional studies should be made to figure out what are the underlying causes behind it.

With this information in mind, Holograph has a specific analysis that deals exclusively with this impact. The user can pick a variant and an activity within the variant, and Holograph will display information about how much time did a process take before reaching the resource, and then how much time it took after. Even though it is not set in stone, the fluctuations between these numbers should start to be an indicator about the need of further analysis.

Disco has a specific tool for this in the form of a Cases tab. This displays a list of variants and then a list of cases for the specific variant. Upon selection of a cases, a sequence is shown in which we can see the activities that took place, the resource involved, the active time, and the waiting time between activities.

This section is where Holograph's analysis has really shined. Question 7, just as question 4, asked about the preference of the individual regarding which visualization is better suited for the specific analysis. In this case, the analysis was the impact of a resource in a process. In this question, a total of 68% approved of Holograph's way of displaying information. The lack of a specialized tool for analyzing this kind of information affects its performance among the individuals.

If question 7 wasn't enough to corroborate that the proposed approach is really helpful for this kind of analysis, then the next question will end all doubts. Question 8 asked about which visualization is better for comparing two different resources executing a case within a variant. Since Disco does not offer something similar, because it was not designed with that in mind, Holograph got a total of 76% of votes, the highest percentage out of all the questions that involved selecting among two different visualizations.

These two questions are extremely encouraging for the proposed approach and validate that the analysis that has been described and constructed is really valuable for companies. By presenting a way of directly getting the impact of a resource on the active time of next activities and compare them to the performance of previous actions, one can start to infer some situations in a business process. The researched individuals seem to agree that this is a useful analysis, and that what is currently available is not enough to draw meaningful conclusions.

Still, it is peculiar to see that 24% seem to feel like Disco is better aimed at comparing the impact of two different resources on a specific variant. Through Disco, this information could be obtained, albeit in a manual way. That one out of four decided that Disco was still a better tool for the job is a little intriguing, but shows how well trusted Disco is in the community, and why their interface and tools are respected by analysts throughout the world.

To finish this section, one more question was added, but it didn't involve directly the visualizations that Disco and Holograph present to users. Question 9 is different from

the previous questions, asking for the preference of the individual regarding the information that should be shown in a tool like this. It asks about what information does the individual needs in order to take decisions about a process. The choices are to have detailed information, aggregated information, or both.

A total of 65% of the subjects responded that they would prefer to have both. This explains why some people still feel that Disco's approach of displaying all information is more trustworthy than Holograph's approach of aggregating the information that is available. To 79% of the individuals, having detailed information is still a valuable source of information, so having only an aggregated view is not enough. Even though Holograph's approach was validated by the subjects on this section, the fact that they still express their desire for detailed information is a lesson that should be learned and should be kept in mind for future developments of this approach and tool.

6.7.4 Handoff impact

One of the points where Holograph was touted to be vastly superior to what Disco presents, but ultimately fell short, was the handoff analysis. This is one of the biggest aspects of analyzing resource performance, so getting a good analysis of what can be improved here is of paramount importance.

Question 10, as all of the first questions in the comparison sections, dealt with the preference of the subject regarding the visualizations that Disco and Holograph display. In the case of Holograph, the group by activity section was shown, displaying information of two different activities and all the involved resources in said activities. For Disco, its trademark process map is shown, with the previously discussed workaround for getting this graph for resources instead of activities.

Holograph languished behind Disco's superb graph, with 60% of the subjects expressing a preference towards the process map. Even though this is one of the screens that Holograph presents with both detailed and aggregated information, it was not enough to grab the attention of the individuals.

For question 11, it was requested again to select which visualization would be more beneficial in a business where a decision has to be made regarding the handoffs of a process. As previously, 60% of the subjects chose Disco's map as a better tool for taking decisions. Curiously, this is the only question of the questionnaire to register two blank answers. Since there were the same number of votes for Disco in both questions, and 66% of the people selected the same visualization for both questions 10 and 11, we could have inferred that at least one of those votes would have been for Holograph and the other either for Holograph or for Disco. However, even with this proposition, Holograph's visualization just wasn't enough to compete against Disco.

This can be attributed to several factors. It is undeniable that Disco's process discovery function, and the output in the form of a process map is top-notch and is hard to compete against with normal graphs, like column or pie graphs. In order to provide something that is better than Disco and that is useful to researchers, something more innovative has to be proposed in the visualization part.

The information shown in Holograph is valuable (as will be seen in question 12), and the potential for it is very high, since a small subset of the proposed approach was actually implemented, but the visualization of this information has to be done in a better way. Clearly, a bar graph is not enough.

Another aspect that could have affected Holograph negatively is the choice that was made for the screenshots of the previous two questions. While different types of visualization have been shown for the Disco part of the questions, Holograph presents natively a bar graph in five of its six analyses, with only one of them letting the user change the type of visualization.

After seeing three questions with almost identical visualizations for Holograph, the surveyed individual might have either grown tired of the bar visualization or felt like Holograph wasn't presenting any new information compared to Disco. Seeing something as useful as Disco's process map alongside the third bar graph in a row presented by Holograph might have make the proposed implementation look worse than it should.

In order to address this, another type of graph should have been shown for the previous two analyses. However, this might have felt like a patch instead of addressing fully the inconvenience that Holograph's bar graph is not enough to entice researchers to draw conclusions with this analysis.

The last question dealt with the information that the surveyed people would like to see in a tool like this. The options were to present amount of times of handoff information, average execution time, or both. The vast majority of individuals preferred to have both piece of information in an analysis, with 71% of the individuals selecting the option "Both."

These results are encouraging for Holograph, since it shows total time, average time, and number of executions for a specific handoff. Disco is also displaying this information, thanks to the tabs frequency and performance that it incorporates into the process map visualization. Clearly, this is the way to go and suggests that a visualization that puts its focus on these numbers is something that the surveyed individuals would prefer instead of showing them either in a list, as is done by Holograph, or not at the same time, as is done in Disco.

Overall, we can conclude that the evaluation was a positive experiment. This is not from a results point of view since the flaws of Holograph were made apparent by the surveyed individuals. However, the knowledge that was acquired through this questionnaire is extremely valuable in order to decide what to do with future developments, not only regarding Holograph, but also for other people who would be interested to implement a tool according to the proposed approach. Also, for people who want to take this further and add more to the proposal, this information would be very useful in order to determine priorities and construct a feature map.

Nevertheless, it is impossible to hide the fact that Holograph still needs refinement in order to be the tool that is envisioned. Although several aspects are encouraging, such as the positive reception of the flow analysis, right choice of approach regarding information being displayed, and decent results on the active time visualization, there is

still a lot of work to do, especially in the group impact section. This is the weakest part of Holograph according to the results of the questionnaire, and with it being one of the pillars of the new approach, it has to be refined further in order to give users a visualization that they can trust, gives valuable information, and is helpful towards decision making in a company.

6.8   LIMITATIONS OF THE EVALUATION

Even though the experiment yielded interesting results that are going to be useful for future developments of Holograph, there were some limitations that made the evaluation a little less effective than it could have been, either in the form of inconveniences to the surveyed people, or lack of an opportunity of getting more information out of them.

The first problem that was even expressed by them is that the questionnaire was done on paper printed on both sides. The main issue with this was that, in order to compare both visualizations, the subjects had to flip back and forth between the pages in order to compare between both screenshots. This made it more difficult for them to answer the questions and could have been influential in a minor way to the results that were selected.

The cause for this was that the experiment was done in a classroom with 38 individuals that might not have a computer to do this online. It was of paramount importance for the experiment that it was conducted in a partially controlled environment, which is why the questionnaire wasn't sent to a wider audience in different environments, knowledge levels, or locations. By having all individuals under the same conditions working at the same time, any variations caused by previously said factors was eliminated, and the results of the experiment, in turn, were more robust and trustworthy.

Another problem with the questionnaire was already discussed. Holograph, on this first development version, has six analyses. Out of them, five present their results with bar graphs. This factor wasn't well received by the surveyed individuals, who saw three bar charts of Holograph versus the detailed screens that Disco proposes.

Although Holograph displays different graphics in some screens, this was a result of Holograph not having enough customization regarding visualization possibilities. Regardless, some more planning could have been done in order to show users different graphs so that the difference between screens was more apparent.

Since the experiment was conducted abroad in an environment that did not involve a computer, then the users were not able to use the tool, as they could do with Disco for previous assignments of their course. This meant that all users were already familiar with Disco, but rather unfamiliar with Holograph. Knowing that Disco is a robust tool and that they have seen it in action already could have made them more biased towards Disco.

Ideally, the users should have also the chance of using Holograph before making a questionnaire, or, at least, given a small introduction about it and what is its functionality. This lack of previous introduction versus the familiarity with Disco by the subjects could have tilted the scale towards Holograph in some cases, especially because some screens are able to be rendered with different filters and type charts, something that is not apparent in the screenshots.

All of the sections of the questionnaire had an explanation of what the screen is displaying and how can it be customized (where applicable). However, reading about it versus using the tool for sure made a difference in the perception that the subjects had on both tools.

Finally, another limitation of the experiment is that additional information apart from the questionnaire was not able to be extracted. Other than a couple of comments manifested by students to professor Dumas after completing the questionnaire, there was no way of getting additional input, such as their reasoning behind some choices, what led them to choose one tool over the other, what could have helped them to perform the questionnaire better, among others. This is why most of the assumptions about the reason behind some choices can only be inferred based on the results and some identified issues in the questionnaire. However, there is no certain way of corroborating these assumptions.

7 CONCLUSION

In this thesis, a new approach for assessing the impact of resource assignment in a business process was proposed. This approach tries to overcome the limitations of current proposals and commercial tools, by building a method from the ground up specifically aimed at analyzing the impact of resource intervention in a process. Several analyses are included within the proposal, such as analysis of active time by activity, analysis of the impact within different variants in the event log, analysis of the impact of a resource on the activities after a resource is involved versus the performance before (called by flows), analysis based on the pitcher analogy, and analysis of the impact of groups of resources among each other (called handoff analysis).

Based on said approach, a web application called Holograph was developed in order to validate that the different kinds of analysis are useful to companies and aid in the decision making regarding business processes. This tool, composed of a front-end and a back-end, aims to show how an implementation of the approach would work, and how the guidelines proposed graphically give valuable information about the event log. The tool was done with state-of-the-art technology and with an architecture in mind aimed at providing extensibility for future projects and additions to the current implementation.

In order to validate the new method, a questionnaire was done to managers experienced in process mining concepts and tools. This set of questions had the purpose of getting more information about the three dimensions that were defined as part of the experimental design: correctness, confidence, and usability. The questionnaire was done comparing visualizations of Holograph and Disco, an established process mining tool.

Although the results were mixed, with Disco being vastly preferred in some areas and Holograph in others, it was determined that the concepts introduced in the approach are beneficial and of interest of the process mining community. Also, some of the results for Holograph itself, especially when dealing with flow analysis, were very satisfactory and set a great foundation for future work. Likewise, the bad results in group impact were a great insight in how to make Holograph a better tool for the job.

There is a lot of work that could be done after this proposal. First, there is a lot of analysis done for the impact of single resources, with active time, variant, flow, and individual analyses already proposed and implemented. For handoff analysis, for now it is limited to being done by resource and by activity, and only in pairs. This is the first obvious area of improvement. More emphasis should be put in the group analysis in order to gather more information out of more resources simultaneously.

Also, through the questionnaire it was determined that, while Holograph is working with a great set of data and generating useful information, the visualization aspect of it is lacking. Part of this flaw is the lack of different visualization options available for the user. Either a new library for charts should be implemented, or the ones that are currently available through NGX-Charts should be utilized in a way that makes it more appealing and useful for the analyst.

Finally, there were also some expressed concerns about the lack of detailed information in Holograph. While the philosophy of it is to show aggregated information

that gives at-a-glance data for the analyst to work with, some of the surveyed individuals preferred the detailed approach that tools like Disco and Minit have adopted. A compromise could be made between both approaches in order to make Holograph both a tool that shows generated information that is useful and leads to the discovery of new trends, while also showing the detailed information from where these aggregates come from.

8 REFERENCES

Bolt, A., Leoni, M. d., & van der Aalst, W. M. (2016). A Visual Approach to Spot Statistically-Significant Differences in Event Logs Based on Process Metrics. *Advanced Information Systems Engineering*, 151-166.

Bolt, A., van der Aalst, W. M., & de Leoni, M. (2017). Finding Process Variants in Event Logs. *On the Move to Meaningful Internet Systems*, 70-77.

Cambridge Dictionary. (s.f.). *KPI Definition.* Cambridge: Cambridge University Press.

Dumas, M. (2016). *Process Monitoring.*

Dumas, M., & García-Bañuelos, L. (2015). Process Mining Reloaded: Event Structures as a Unified Representation of Process Models and Event Logs. *Application and Theory of Petri Nets and Concurrency*, 43-58.

Dumas, M., Rosa, M. L., Mendling, J., & Reijers, H. A. (2018). *Fundamentals of Business Process Management, 2nd Edition.* Springer.

Hompes, B. F., Maaradji, A., Rosa, M. L., Dumas., M., Buijs, J. C., & Aalst, W. M. (2017). Discovering Causal Factors Explaining Business Process Performance. *Advanced Information Systems Engineering: 29th International Conference, CAiSE 2017, Proceedings (Lecture Notes in Computer Science, Volume 10253)* (págs. 177-192). Essen: Springer.

La Rosa, M., Dumas, M., Uba, R., & Dijkman, R. (2013). Business Process Model Merging: An Approach to Buiness Process Consolidation. *ACM Transactions on Software Engineering and Metholodogy, 22*(2).

Liu, L., & Ozsu, M. T. (2009). *Encyclopedia of Database Systems.* Springer US.

Major League Baseball. (2018). *Positions - Glossary.* Recuperado el 31 de July de 2018, de MLB: http://m.mlb.com/glossary/positions

Nguyen, H., Dumas, M., La Rosa, M., Maggi, F. M., & Suriadi, S. (2016). Business Process Deviance Mining.

Pai Ballambettu, N., Agumbe Suresh, M., & Chandra Bose, R. P. (2017). Analyzing Process Variants to Understand Differences in Key Performance Indices. *Advanced Information Systems Engineering*, 298-313.

Parmenter, D. (2015). *Key Performance Indicators: Developing, Implementing, and Using Winning KPIs.* John Wiley & Sons.

Parthasarathy, S., Zaki, M. J., Ogihara, M., & Dwarkadas, S. (1999). Incremental and Interactive Sequence Mining.

Pika, A., Leyer, M., Wynn, M. T., Fidge, C. J., Hofstede, A. H., & Aalst, W. M. (2017). Mining Resource Profiles from Event Logs. *ACM Transactions on Management Information Systems, 8*(1), 1-30.

Pika, A., Wynn, M., Fidge, C., Hofstede, A. t., Leyer, M., & Aalst, W. v. (2014). An Extensible Framework for Analysing Resource Behaviour Using Event Logs. *dvanced Information Systems Engineering. CAiSE 2014. Lecture Notes in Computer Science, vol 8484.* (págs. 564-579). Springer.

Rozinat, A., & van der Aalst, W. (26 de June de 2006). Conformance checking of processes based on monitoring real behavior. *Elsevier*, 64-95.

Springer London. (2010). *Key Performance Indicators.* London: Springer.

Stojadinović, Z., Marinković, D., & Ivković, B. (2014). Human Resource Performance Measurement Framework for Construction Projects and Companies. *Technical Gazette, 21*(1), 69-78.

Suriadi, S., M. T. Wynn, C. O., & Dijk, N. J. (2013). Understanding Process Behaviours in a Large Insurance Company in Australia: A Case Study. *Advanced Information Systems Engineering*, 449-464.

Swinnen, J., Depaire, B., Jans, M. J., & Vanhoof, K. (2011). A Process Deviation Analysis - A Case Study. *Business Process Management Workshops*, 87-98.

van Beest, N. R., Dumas, M., Garcia-Banuelos, L., & La Rosa, M. (2015). Log Delta Analysis: Interpretable Differencing of Business Process Event Logs. *International Conference on Business Process Management*.

Wynn, M., Poppe, E., Xu, J., Hofstede, A. t., Brown, R., Pini, A., & Aalst, W. v. (2 de May de 2017). ProcessProfiler3D: A visualisation framework for log-based process performance comparison. *Elsevier*, 93-108.

Zaki, M. J. (2000). Sequence Mining in Categorical Domains: Incorporating Constraints.

APPENDIX

I.   QUESTIONNAIRE (ENGLISH VERSION, HOLOGRAPH SECOND)

The aim of this questionnaire is to gather data about the usefulness, perceived ease of use, and effectiveness of different visualization techniques used for analysis the impact a resource has on the execution of a business process.

When answering the questions, please bear in mind the following:

- This questionnaire is completely anonymous and will be used in an aggregated way with other answers.
- Rate the tools based on the usefulness of the information they show, not on their appearance.
- Even though you might want to choose two answers for the same question, please only select **one** answer.

Part 1

- Q1: In general, the average cycle time of a process is determined primarily by the execution time of the activities in the process.

(A) Completely agree / (B) Partially agree / (C) Neither agree nor disagree / (D) Partially disagree / (E) Totally Disagree / (F) I don't know

- Q2: In general, the average cycle time of a process is determined primarily by the time the handoffs between the involved resources take.

(A) Completely agree / (B) Partially agree / (C) Neither agree nor disagree / (D) Partially disagree / (E) Totally Disagree / (F) I don't know
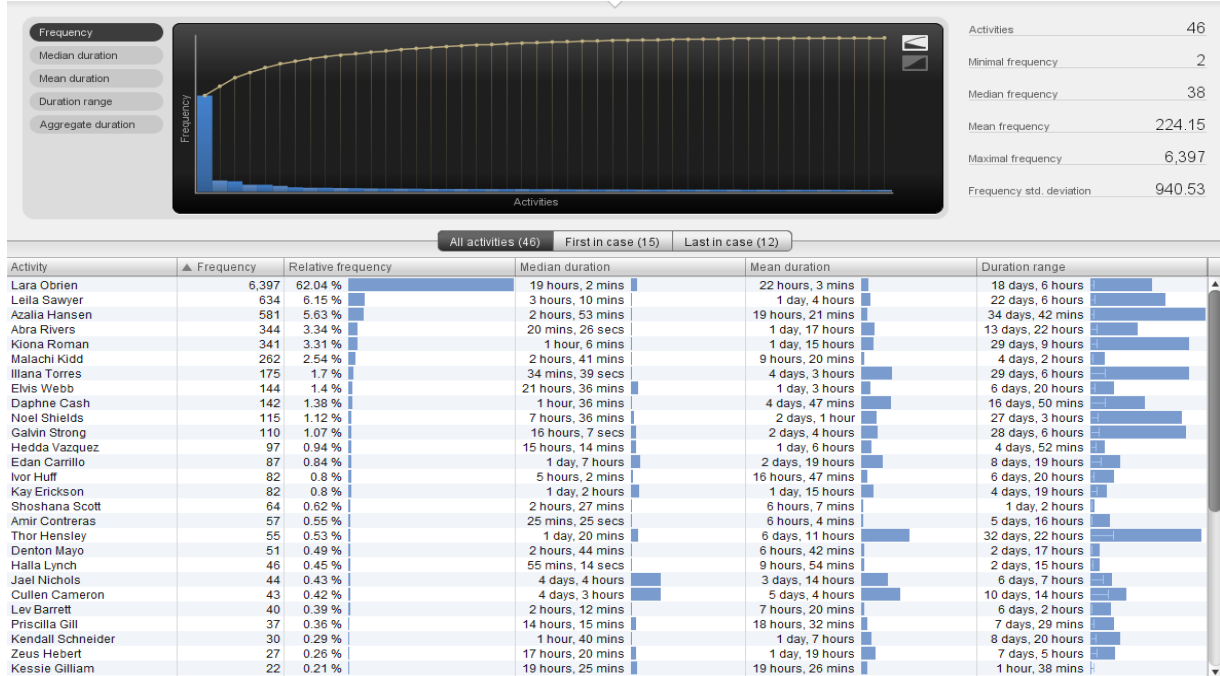
- Q3: Analyzing the impact of a group (like a department or team) is as important as analyzing individual impact on a business process.

(A) Completely agree / (B) Partially agree / (C) Neither agree nor disagree / (D) Partially disagree / (E) Totally Disagree / (F) I don't know
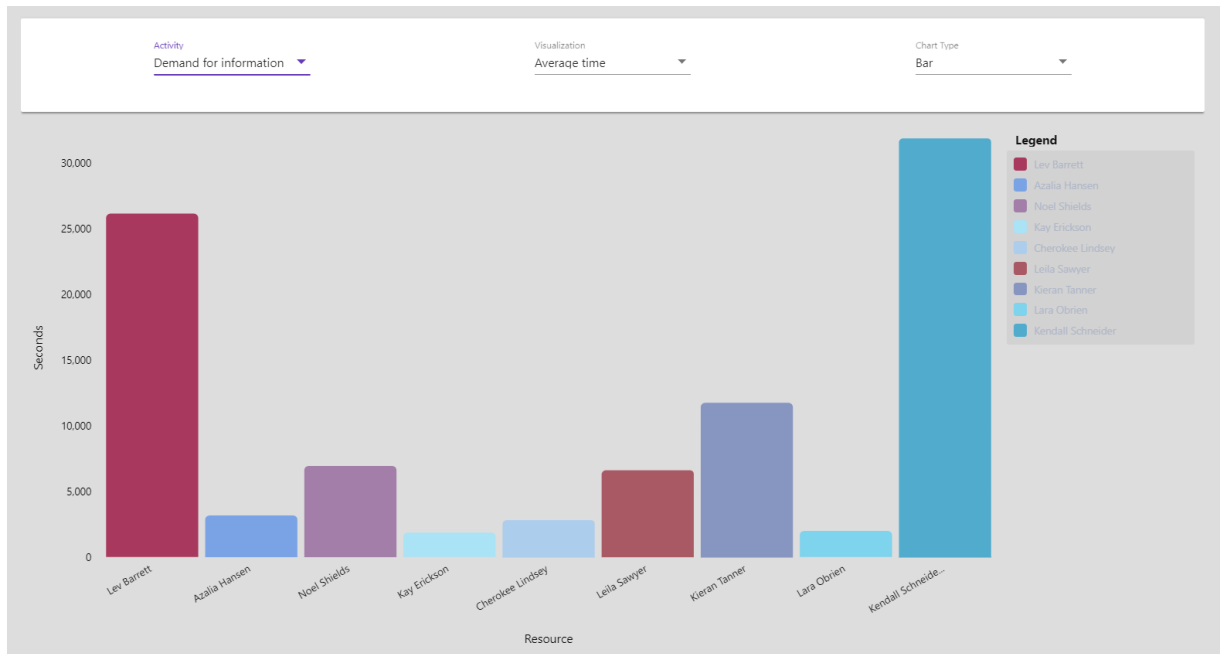
Part 2

Consider these two screenshots.

Tool 1 (Ignore the fact that resources are referred to as "Activity." These are all resources):

Frequency
Median duration
Mean duration
Duration range
Aggregate duration

Activities	46
Minimal frequency	2
Median frequency	38
Mean frequency	224.15
Maximal frequency	6,397
Frequency std. deviation	940.53

All activities (46) | First in case (15) | Last in case (12)

Activity	Frequency	Relative frequency	Median duration	Mean duration	Duration range
Lara Obrien	6,397	62.04 %	19 hours, 2 mins	22 hours, 3 mins	18 days, 6 hours
Leila Sawyer	634	6.15 %	3 hours, 10 mins	1 day, 4 hours	22 days, 6 hours
Azalia Hansen	581	5.63 %	2 hours, 53 mins	19 hours, 21 mins	34 days, 42 mins
Abra Rivers	344	3.34 %	20 mins, 26 secs	1 day, 17 hours	13 days, 22 hours
Kiona Roman	341	3.31 %	1 hour, 6 mins	1 day, 15 hours	29 days, 9 hours
Malachi Kidd	262	2.54 %	2 hours, 41 mins	9 hours, 20 mins	4 days, 2 hours
Illana Torres	175	1.7 %	34 mins, 39 secs	4 days, 3 hours	29 days, 6 hours
Elvis Webb	144	1.4 %	21 hours, 36 mins	1 day, 3 hours	6 days, 20 hours
Daphne Cash	142	1.38 %	1 hour, 36 mins	4 days, 47 mins	16 days, 50 mins
Noel Shields	115	1.12 %	7 hours, 36 mins	2 days, 1 hour	27 days, 3 hours
Galvin Strong	110	1.07 %	16 hours, 7 secs	2 days, 4 hours	28 days, 6 hours
Hedda Vazquez	97	0.94 %	15 hours, 14 mins	1 day, 6 hours	4 days, 52 mins
Edan Carrillo	87	0.84 %	1 day, 7 hours	2 days, 19 hours	8 days, 19 hours
Ivor Huff	82	0.8 %	5 hours, 2 mins	16 hours, 47 mins	6 days, 20 hours
Kay Erickson	82	0.8 %	1 day, 2 hours	1 day, 15 hours	4 days, 19 hours
Shoshana Scott	64	0.62 %	2 hours, 27 mins	6 hours, 7 mins	1 day, 2 hours
Amir Contreras	57	0.55 %	25 mins, 25 secs	6 hours, 4 mins	5 days, 16 hours
Thor Hensley	55	0.53 %	1 day, 20 mins	6 days, 11 hours	32 days, 22 hours
Denton Mayo	51	0.49 %	2 hours, 44 mins	6 hours, 42 mins	2 days, 17 hours
Halla Lynch	46	0.45 %	55 mins, 14 secs	9 hours, 54 mins	2 days, 15 hours
Jael Nichols	44	0.43 %	4 days, 4 hours	3 days, 14 hours	6 days, 7 hours
Cullen Cameron	43	0.42 %	4 days, 3 hours	5 days, 4 hours	10 days, 14 hours
Lev Barrett	40	0.39 %	2 hours, 12 mins	7 hours, 20 mins	6 days, 2 hours
Priscilla Gill	37	0.36 %	14 hours, 15 mins	18 hours, 32 mins	7 days, 29 mins
Kendall Schneider	30	0.29 %	1 hour, 40 mins	1 day, 7 hours	8 days, 20 hours
Zeus Hebert	27	0.26 %	17 hours, 20 mins	1 day, 19 hours	7 days, 5 hours
Kessie Gilliam	22	0.21 %	19 hours, 25 mins	19 hours, 26 mins	1 hour, 38 mins

Tool 2:

Activity: Demand for information
Visualization: Average time
Chart Type: Bar

Legend
Lev Barrett
Azalia Hansen
Noel Shields
Kay Erickson
Cherokee Lindsey
Leila Sawyer
Kieran Tanner
Lara Obrien
Kendall Schneider

(Y axis: Seconds — 0, 5,000, 10,000, 15,000, 20,000, 25,000, 30,000; X axis: Resource — Lev Barrett, Azalia Hansen, Noel Shields, Kay Erickson, Cherokee Lindsey, Leila Sawyer, Kieran Tanner, Lara Obrien, Kendall Schneide...)

These are two different tools for visualizing information regarding the resource performance. The first one shows aggregate information about all resources on all activities of the event log. The second one lets you choose a specific activity and shows the average

time a resource takes to execute it, while comparing it to other resources that also execute the same activity through a graph.

- Q4: Which analysis do you find better to analyze the performance of resources across different activities within a process?

(A) Tool 1 / (B) Tool 2

- Q5: Which analysis gives more targeted information about how the performance of a resource compares to other resources?
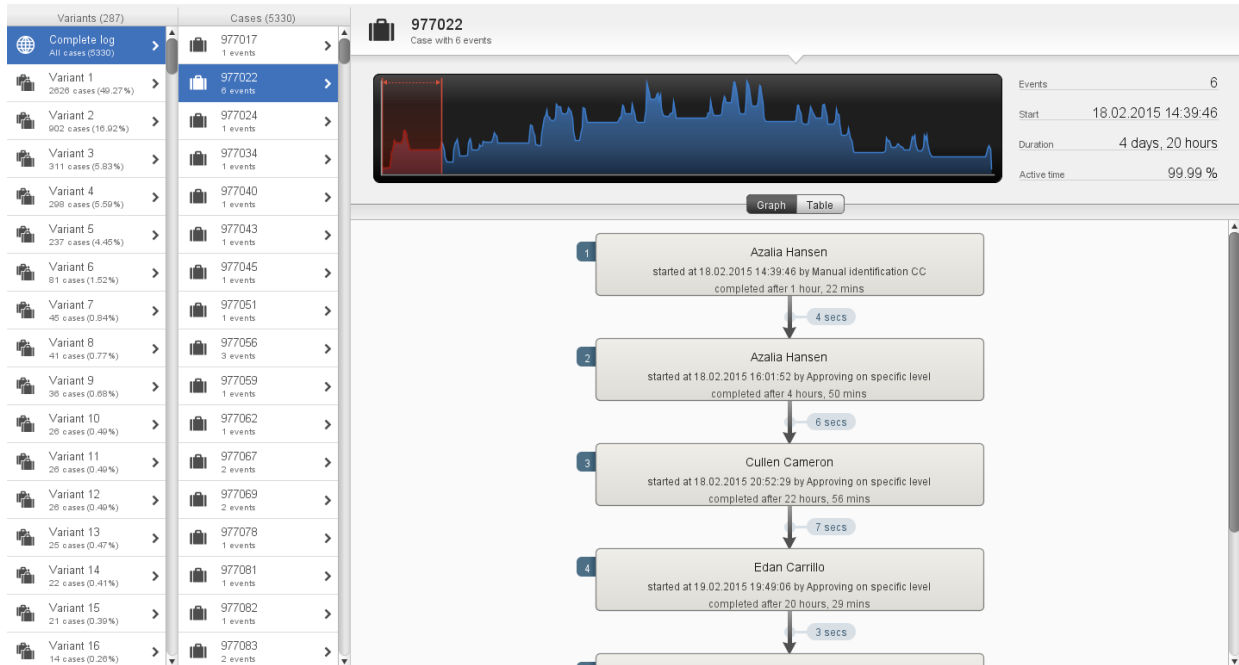
(A) Tool 1 / (B) Tool 2

- Q6: Which analysis would help you more in taking a decision about the business process, either by changing the resources that execute tasks or by changing the task itself to make it easier?

(A) Tool 1 / (B) Tool 2

Part 3

Consider these two screenshots:

Tool 1:

Tool 2:
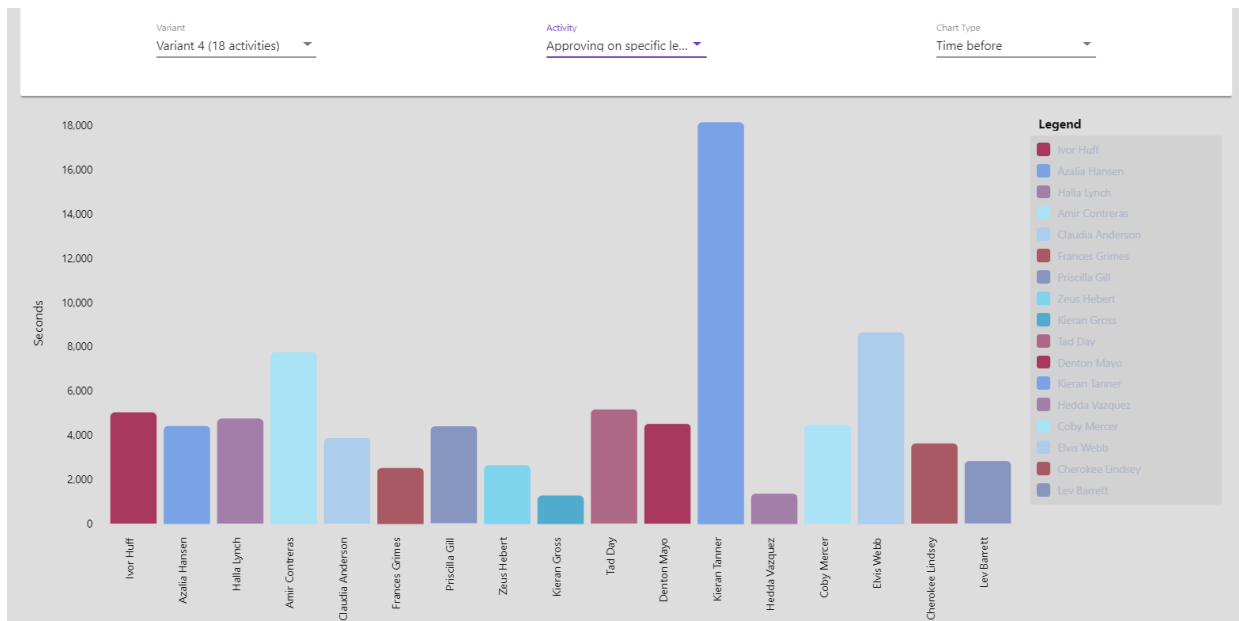
The first tool gives you information about a specific executed case. All of the events that occur within that case are shown with their start time and completion duration. All cases within a variant are shown, so you can choose to analyze each case within the same variant individually. The second tool analyzes the impact of a resource on a variant. It aggregates the variant data and lets you see the average time taken for all activities before and after (toggled by a drop down) they reach the resource. For example, if the variant has 5 activities and you choose activity #3, the tool lets you see the average execution time of activities 1

and 2, and then of activities 4 and 5. These times are shown in a graph along with other resources that execute the variant.

- Q7: Which analysis do you find better to analyze the impact of a resource across a business process?

(A) Tool 1 / (B) Tool 2

- Q8: Which tool do you think is better for comparing between two or more resources executing a variant?
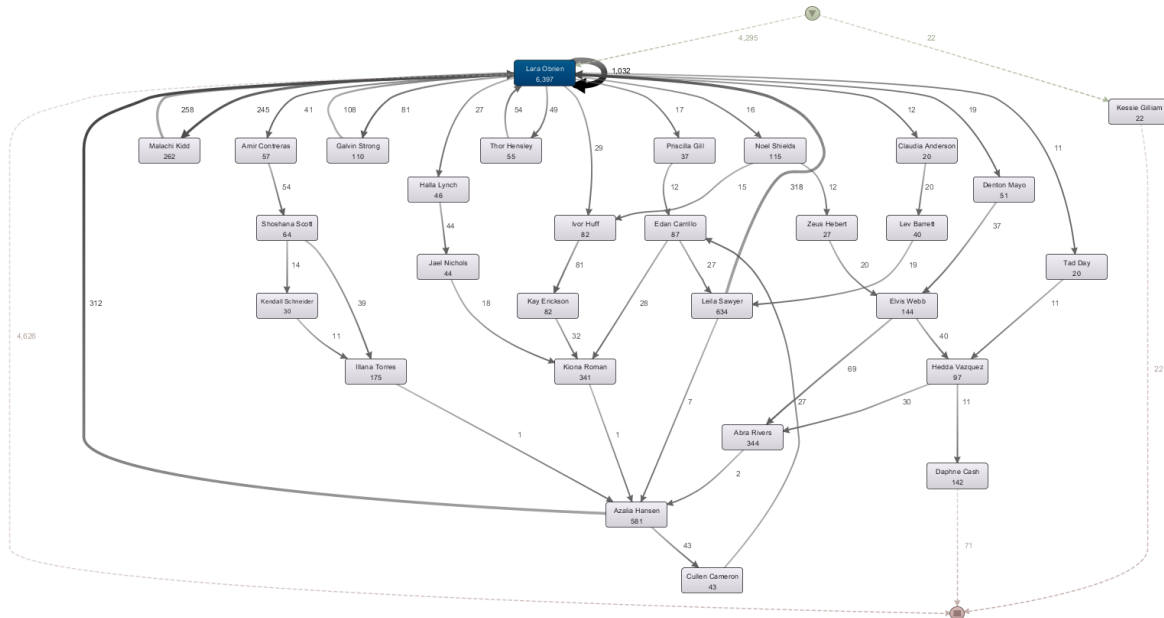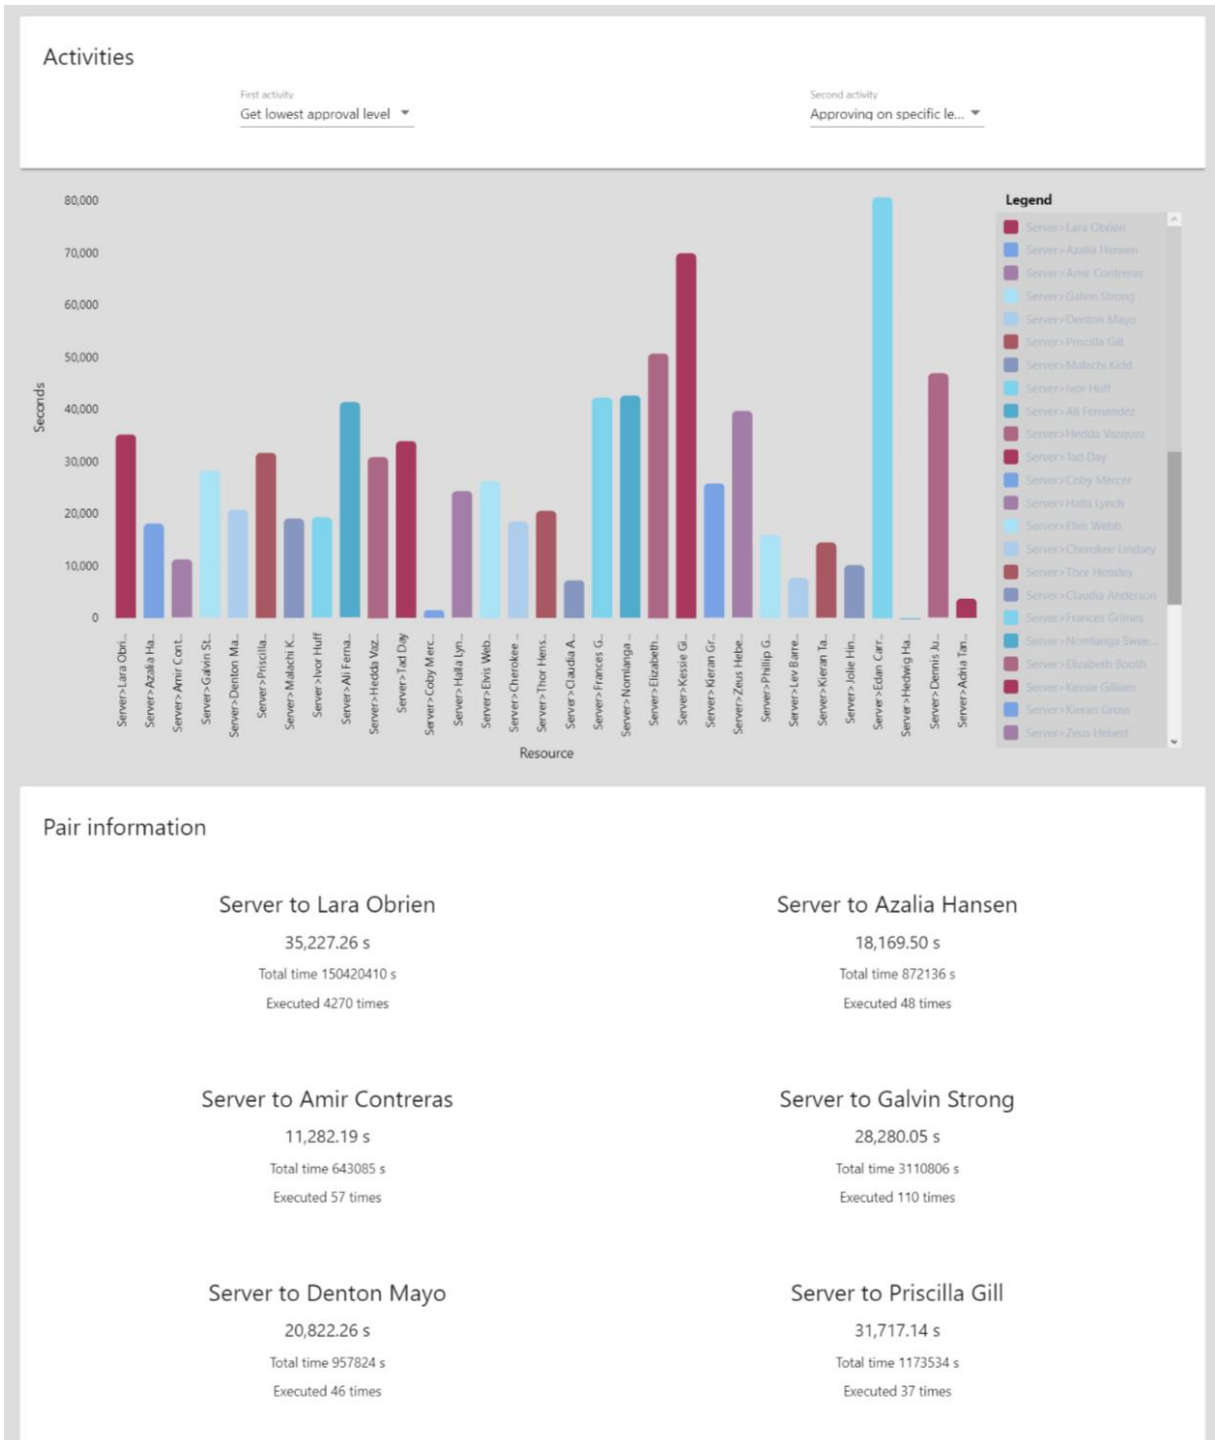
(A) Tool 1 / (B) Tool 2

- Q9: In order to take business decisions about a process, do you prefer to have aggregated information that shows a trend, detailed information that doesn't show trends, or both?

(A) Aggregated information / (B) Detailed information / (C) Both

Part 4

Consider these two screenshots:

Activities

First activity
Get lowest approval level ▼

Second activity
Approving on specific le... ▼

Legend
- Server>Lara Obrien
- Server>Azalia Hansen
- Server>Amir Contreras
- Server>Galvin Strong
- Server>Denton Mayo
- Server>Priscilla Gill
- Server>Malachi Kidd
- Server>Ivor Huff
- Server>Ali Fernandez
- Server>Hedda Vazquez
- Server>Tad Day
- Server>Coby Mercer
- Server>Halla Lynch
- Server>Elvis Webb
- Server>Cherokee Lindsey
- Server>Thor Hensley
- Server>Claudia Anderson
- Server>Frances Grimes
- Server>Nomlanga Swee...
- Server>Elizabeth Booth
- Server>Kessie Gilliam
- Server>Kimani Grost
- Server>Zeus Hebert

Pair information

Server to Lara Obrien	Server to Azalia Hansen
35,227.26 s	18,169.50 s
Total time 150420410 s	Total time 872136 s
Executed 4270 times	Executed 48 times

Server to Amir Contreras	Server to Galvin Strong
11,282.19 s	28,280.05 s
Total time 643085 s	Total time 3110806 s
Executed 57 times	Executed 110 times

Server to Denton Mayo	Server to Priscilla Gill
20,822.26 s	31,717.14 s
Total time 957824 s	Total time 1173534 s
Executed 46 times	Executed 37 times

The first tool shows a map in which we can see the handoff between different resources. Also, information about how many times a handoff occurs is shown. The second tool lets you analyze the handoff between resources by selecting two activities. The tool shows the number of times, average time, and total time that it takes for the two resources to do an activity together and compares it in a graph to handoffs for same activities but with different resources.

80

- Q10: Which analysis do you find better to analyze the handoff between different resources across a business process?

(A) Tool 1 / (B) Tool 2

- Q11: Which analysis gives you more information in order to take a decision about handoffs within a business process, either by changing the resources that take part in a handoff or by restructuring the process to avoid handoffs between specific activities?

(A) Tool 1 / (B) Tool 2

- Q12: If you want to analyze the impact that a specific handoff has on a business process, do you find more useful to know the amount of times that a handoff occurs, the average execution time that the handoff, or both?

(A) Amount of times / (B) Execution time / (C) Both

Time of questionnaire completion: _____

Thank you for your participation!

II. LICENSE
Non-exclusive licence to reproduce thesis and make thesis public

I, **Luis Gerson Noboa Martillo**

1. herewith grant the University of Tartu a free permit (non-exclusive licence) to:
    1.1. reproduce, for the purpose of preservation and making available to the public, including for addition to the DSpace digital archives until expiry of the term of validity of the copyright, and
    1.2. make available to the public via the web environment of the University of Tartu, including via the DSpace digital archives until expiry of the term of validity of the copyright,

of my thesis
Holograph: A Tool for Assessing the Impact of Resource Assignment on Business Process Performance Based on Event Logs,
supervised by Marlon Dumas.

2. I am aware of the fact that the author retains these rights.
3. I certify that granting the non-exclusive licence does not infringe the intellectual property rights or rights arising from the Personal Data Protection Act.

Tartu, **09.08.2018**