Handling of Building Permit Applications in The Netherlands: A Multi-Dimensional Analysis

P.M. Dixit[1,2], B.F.A. Hompes[1,2], N. Tax[1,2], S.J. van Zelst[1]

[1] Department of Mathematics and Computer Science
Eindhoven University of Technology, Eindhoven, The Netherlands
{p.m.dixit,b.f.a.hompes,n.tax,s.j.v.zelst}@tue.nl
[2] Philips Research, Eindhoven, The Netherlands
{prabhakar.dixit,niek.tax}@philips.com

Abstract. As part of the BPI Challenge 2015, five event logs originating from five different Dutch municipalities have been analyzed. The process captured by these event logs concerns the application process of building permit requests. The event logs have been analyzed and compared on the organizational dimension, the performance dimension, and the control-flow dimension. With respect to the organizational dimension we found existing collaborations and re-allocation of resources, and the effect of this on the executed activities. From a control-flow point of view we discuss the effect of a central registration system, i.e. the OLO system, being used. Our performance analysis shows the average throughput times to differ significantly between the different municipalities. Furthermore we found cases in which the OLO system seems to be used to have a significantly higher throughput time than cases in which this system is not used.

Keywords: BPI challenge 2015, process minning, process discovery

1 Introduction

Building permit applications in the Netherlands are handled by a designated department of the municipality governing the intended building location. The process of handling such building permits should theoretically be identical throughout different municipalities, however, slight differences may occur. Differences might for example be caused by changes of rules and regulations that are adopted at different points in time by the municipalities.

For the BPI Challenge 2015[3], five event logs have been made available where each event log contains activities of a single municipality. The activities in the logs are related to the permit application process and are logged by their IT systems [1]. In this paper we present a thorough analysis of the event data as well as the corresponding findings. The event data has been analyzed on three different dimensions being the organizational dimension, the performance dimension, and the control-flow dimension. The organizational structure analysis

[3] http://www.win.tue.nl/bpi/2015/challenge

forms the basis of this report. The control-flow and performance analysis will take the concept drifts found in the organizational structure as a starting point. The analysis is supported by existing and newly developed plug-ins in release 6.5 of the process mining toolkit *ProM*[4] [2], developed at the Eindhoven University of Technology. Plugins developed for the analyses described in this report are available in the BPIC2015 package[5] on the ProM SVN repository.

The remainder of this paper is structured as follows. In Section 2 we present the main findings of a exploratory high-level analysis of the five data sets. Section 3 discusses the analysis on an organizational level. Section 4 discusses the impact of changes in organizational structure on control flow. The impact of organizational structure changes on performance is discussed in Section 5. The remainder of this paper assumes the reader to have a basic understanding of general concepts within the field of process mining, e.g., the concepts of *event logs*, *traces*, *events*, etc. We refer to [3] as a reference and a complete overview of the field.

2 Exploratory Analysis

2.1 Global Overview

The available event data consists of all building permit applications submitted to five Dutch municipalities over a period of approximately four years, starting at late 2010 and ending early 2015. The traces in the event logs contain information regarding the main application process as well as objection and appeal procedures in various stages. A combined visualization of the five event logs, using a dotted chart visualization [4], is shown in Figure 1.

Each dot in the visualization corresponds to an event[6]. The coloring of the dots is based on the municipality in which the event was executed. The vertical axis describes the trace identifier which is sorted in an ascending fashion *per municipality*. The horizontal axis describes the timestamp of an event.

The execution of traces seems to be relatively constant throughout time. Using the intensity chart of the dotted chart visualization, i.e., the chart in the lower pane, we identify a minor number of events being logged on times that do not agree well with the overall period of logging. The events are logged around the beginning of 2010 and possibly hint to noise, i.e. inappropriate logging.

Several meta-data are available on both trace and event level, of which the most prominent frequent elements are presented here in an informal fashion:

- *Trace level*
  - Cace identifier.
  - Indication of costs associated to the Cace.
  - Indication of the last *phase* executed for the given cace.

[4] ProM 6.5 can be downloaded from `http://www.promtools.org`

[5] `https://svn.win.tue.nl/repos/prom/Packages/BPIC2015/`

[6] Not all events are visible, within rendering sampling of events is used for performance purposes.

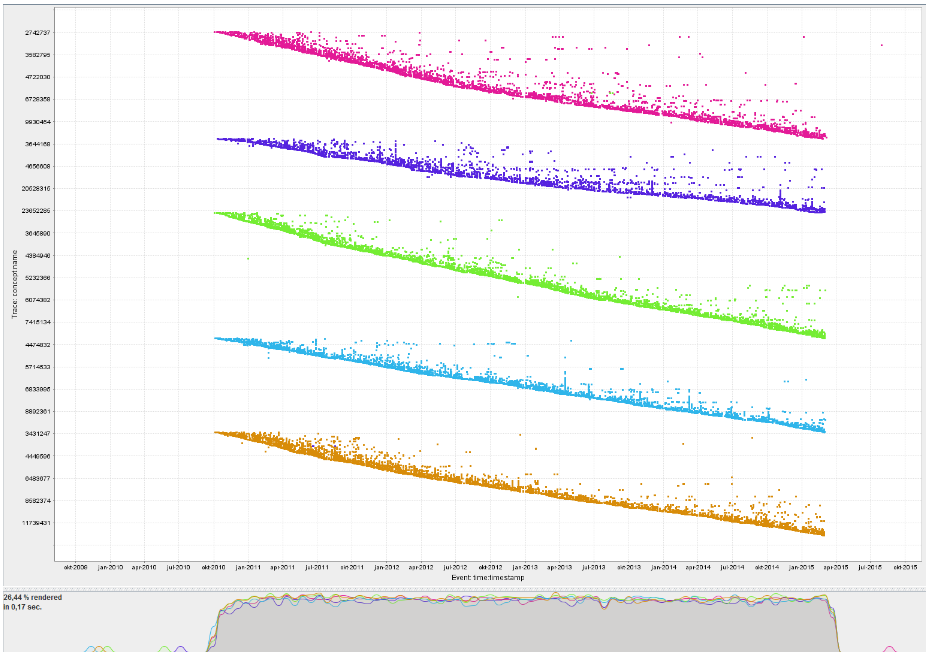

Fig. 1: Dotted chart visualization of the five event logs. Municipalities are sorted ascending in a top down fashion, i.e., municipality 1 is purple, municipality 2 is dark blue...

  - Status of the case.
  - Indication of inclusion of *sub-cases* within the given cace.
  - Resource (actor) responsible for the case.
  - Type of permit.
- *Event level*
  - Activity identifiers, both an activity code as well as a description.
  - Resource executing the event.
  - Resource monitoring the event.
  - Time stamp of execution of event.
  - Planned time of execution of the event.
  - Due date of execution of the event.

Table 1 gives a global overview of the data characteristics per municipality. We identify municipality 3 to handle a relatively high amount of cases whereas municipality 2 seems to handles a relatively low amount of cases. The tables have turned with respect to the average number of events executed per case, i.e., municipality 2 executes 53 events per case on average whereas municipality 3 executes only 42 events per case. More interestingly, although municipality 2 is involved in the least number of cases, it in turn executed the highest number of different types of activities (*event classes*). A possible explanation for the

Table 1: Data characteristics per municipality.

	Cases	Events	Event Classes	Avg. Events per Case	No. of executing resources
Municipality 1	1.199	52.217	398	44	23
Municipality 2	832	44.354	410	53	11
Municipality 3	1.409	59.681	383	42	14
Municipality 4	1.053	47.293	356	45	10
Municipality 5	1.156	59.083	389	51	22
Total	5649	262.628	500	47	72

differences in number of event classes could be differences and/or deviations in the process of handling building permits throughout different municipalities. It could also be the case that the granularity of event logging differs between municipality, i.e. two related smaller tasks might be logged as two low-level events by one municipality while another municipality logs those two as one more high-level event. The high number of average events per case and the high number of event classes in combination with the low number of cases could be an indication that municipality 2 logs on a more fin-grained granularity.

A total of 500 distinct event classes can be identified throughout the five municipality event logs. Not all event classes however are present in each municipality's individual event log. Figure 2a depicts the distribution of event classes shared between municipalities. Out of the 500 event classes, around 77% are executed within at least three different municipalities. The remaining 23% are executed in at most two different municipalities. Interestingly, a total of 15% of the event classes, i.e 74 distinct event classes, are only executed uniquely in one municipality. Figure 2b depicts the distribution of unique event classes over the different municipalities. Municipality 2 has a total share of 38% of all unique cases, explaining the high value of different event classes within the correspond-

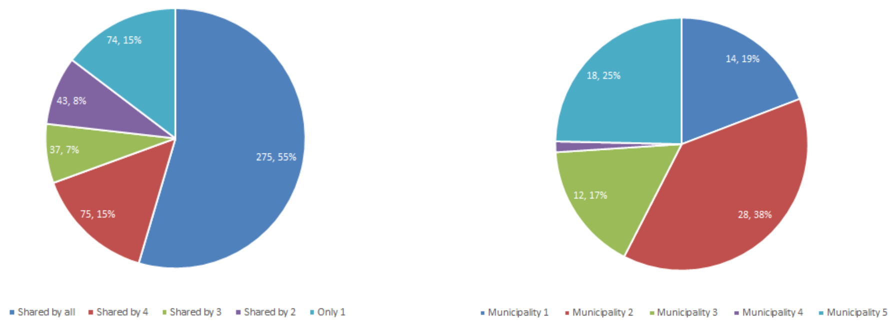

(a) (b)

Fig. 2: Distribution of event classes shared between municipalities and distribution of unique event classes over different municipalities

ing event log (410). Municipality 4 only executed 1 unique event class. Again this is in line with the number of event classes present within the corresponding event log (356).

Different types of permits can be identified in the event data, e.g. construction permits, environmental permits, demolition etc. Each permit application in the log is labeled with a set of permit types. Hence, a permit application might fall into multiple types. For example, a permit involving construction work might also involve some demolition work, in which case the application is labeled both as of type building permit and as of type demolition permit. Figure 3 lists the distribution of permit types for the five municipalities. The stacked bar chart shows that 40% to 50% of all applications are labeled solely as a construction permit application. In addition, the third, the eight, and the tenth most frequent set of permit application types also includes construction permit as one of its types.

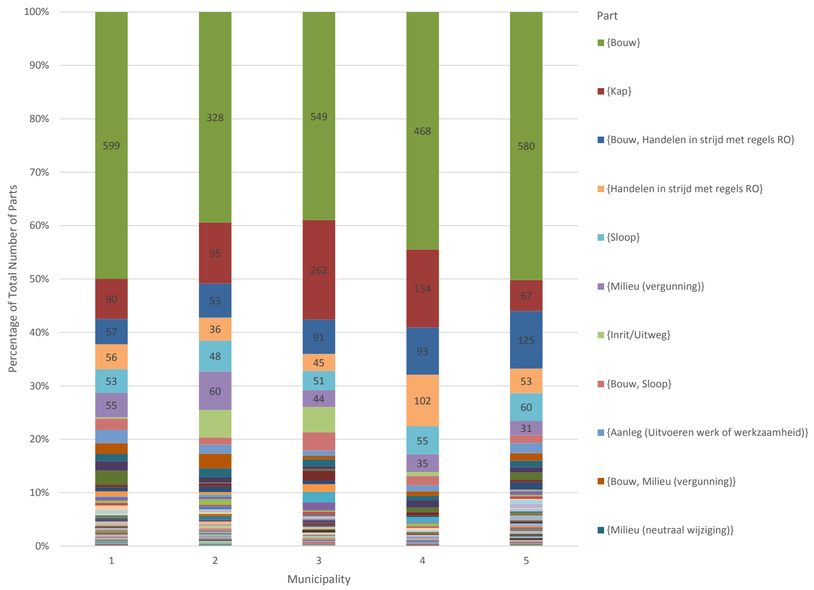

Fig. 3: Distribution of cases over the permit types per municipality

2.2 Log quality

The following sections discuss several types of logging inconsistencies that were found in the event logs, which have consequences for the analyses described in the succeeding parts of this paper.

Case Identification A sanity check performed on the five event logs shows that there is a slight overlap in trace identifiers across municipalities[7]. There are two identifiers that occur in two municipalities. Trace id *6038724* occurs in both municipality 1 and 3. Trace id *4020737* occurs in municipality 2 and 5. In both cases the occurrence of typical events indicate that these traces in fact describe two distinct cases, rather than cases shared across municipalities. Also, in both of the cases other trace attributes such as the type of the permit differ significantly. For this reason we do not assume these traces to be a special case shared amongst different municipalities, and therefor, treat them as individual cases.

Time Granularity in Event Logging When inspecting events within the data, it becomes clear that there is a large inconsistency in the level of granularity of the associated timestamps. Many events have timestamp 00:00:00. Although not impossible, it seems unlikely that the corresponding events were actually executed at this time. Therefore, we assume that event with such a timestamp have a granularity at day level. In the succeeding parts of this paper we refer to such timestamps as *coarse-grained timestamps* and to event with such a timestamp as *coarse-grained events*. Table 2 shows an overview of the timestamp granularity of the events per municipality.

To assess the distribution of fine-grained and coarse-grained events, for each trace in each event log we have computed the relative occurrence of fine-grained events. A relative occurrence of 0 means that there is no event having fine-grained time granularity whereas a relative occurrence of 1 means that all events in the trace have a fine-grained time granularity. Using the figures related to relative occurrence of fine-grained events, we have computed a kernel density plot which is depicted in Figure 4.

Interestingly, the kernel density plot shows that there are roughly two groups of traces distinguishable within the data. One group of traces seems to have a majority of coarse-grained timestamps, whereas another group of traces seems to have a majority of fine-grained timestamps.

Table 2: Number of events per timestamp granularity per municipality

	Coarse-grained timestamp	Fine-grained timestamp	Ratio
Municipality 1	23,992	28,225	0.541
Municipality 2	19,453	24,901	0.561
Municipality 3	25,440	34,241	0.574
Municipality 4	22,813	24,480	0.518
Municipality 5	28,955	30,128	0.510
Total	120,653	141,975	0.541

[7] This is in fact visible in Figure 1.

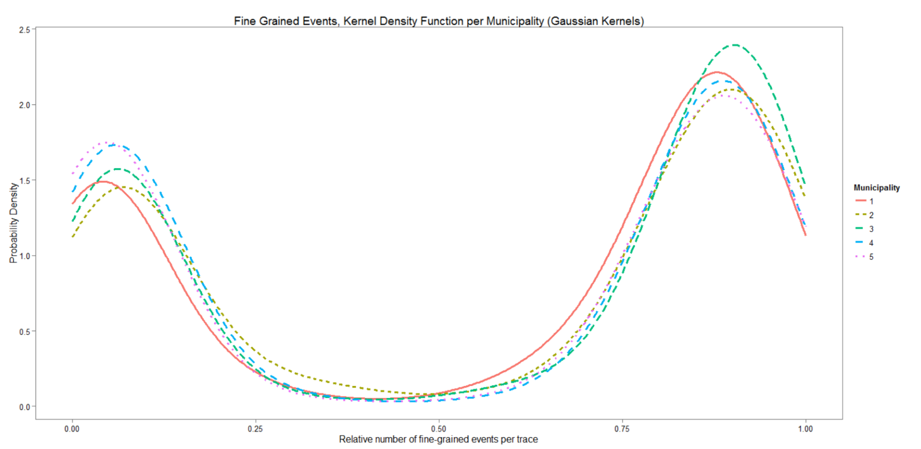

Fig. 4: Kernel density plot (based on Gaussian kernels) of the relative occurrence of fine-grained events within traces

Projection of the event log onto those events that have a fine-grained timestamp, i.e. removing all events having a coarse-grained timestamp, shows that within the aforementioned distribution there is an underlying temporal dimension. Consider Figure 5 in which we have depicted this filtered log using a dotted chart visualization. The chart clearly shows that from around mid 2013 logging of events became significantly more coarse-grained. Note that this is also reflected by the intensity chart within the dotted chart visualization. The height of the intensity chart seems to be inconsistent with the figures presented in Table 2, this is however due to the fact that the chart uses a logarithmic scale.

Using Figure 4 we decided to project the event data onto traces that have at least 50% fine-grained timestamp based events. A dotted chart visualization of the resulting projected event data is depicted in Figure 6 Interestingly we note that due to the decrease in accurate logging starting from mid 2013, all traces starting around that period seem to have less than 50% of fine-grained timestamp based events. When regarding the time-frame of the total dataset, this time point is in-line with the distribution presented in Figure 4, as it suggest a slightly higher probability of traces having more than 50% fine-grained events.

Table 3 shows the statistics of coarse-grained and fine-grained events per municipality in the time range from July 2013 to the end of the log. July 2013 is the time from which the municipalities stop frequently recording traces consisting of more than 50% fine-grained events. The table clearly indicates that after June 2013, all municipalities generally seem to log their activity at day level.

Event ordering We have found that the events in all traces are ordered based on their timestamp. As a result, per trace, events that have a coarse timestamp granularity all are placed before events having a fine-grained timestamp, happening on the same day. The events that have coarse timestamp granularity should in reality probably be intertwined with those that have fine-grained timing information. This is supported by the last three digits of the activity codes that hint on the order in which activities are executed. We can only safely assume the order of those events for which we do have a fine-grained timestamp.

Additionally, we computed the average number of traces per day for each municipality. The averages and standard deviations are presented in Table 4.

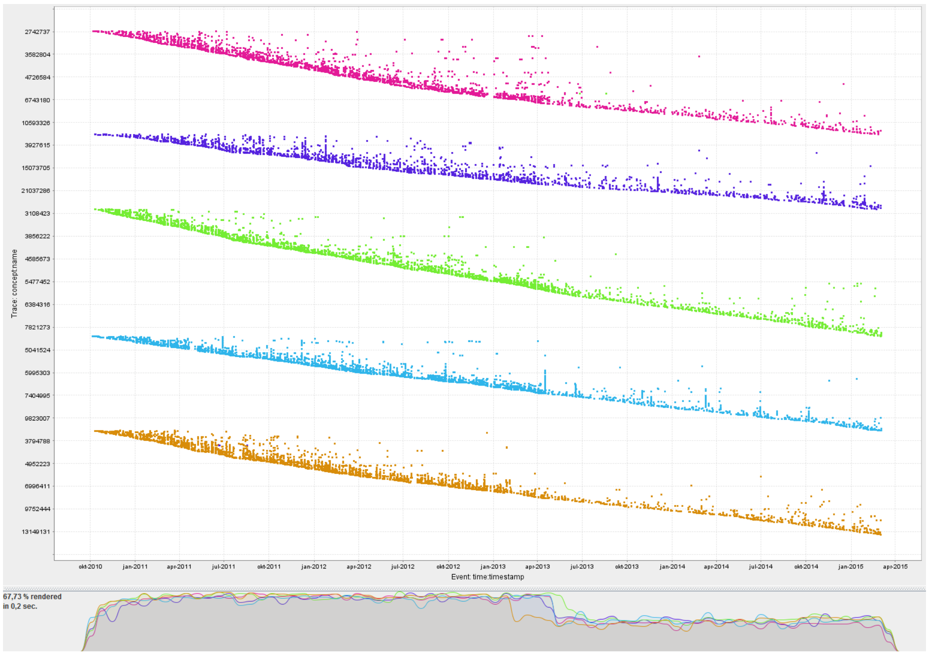

Fig. 5: Dotted chart visualization of the five event logs projected on fine-grained events. Municipalities are sorted ascending in a top down fashion, i.e., municipality 1 is purple, municipality 2 is dark blue...

Table 3: Number of events per timestamp granularity per municipality that where logged after June 2013

	Coarse-grained timestamp	Fine-grained timestamp	Ratio	Ratio (whole log)
Municipality 1	17,910	733	0.039	0.541
Municipality 2	14,956	1,136	0.071	0.561
Municipality 3	19,849	1,257	0.060	0.574
Municipality 4	17,744	1,149	0.061	0.518
Municipality 5	19,230	1,011	0.050	0.510
Total	89,689	5,286	0.059	0.541

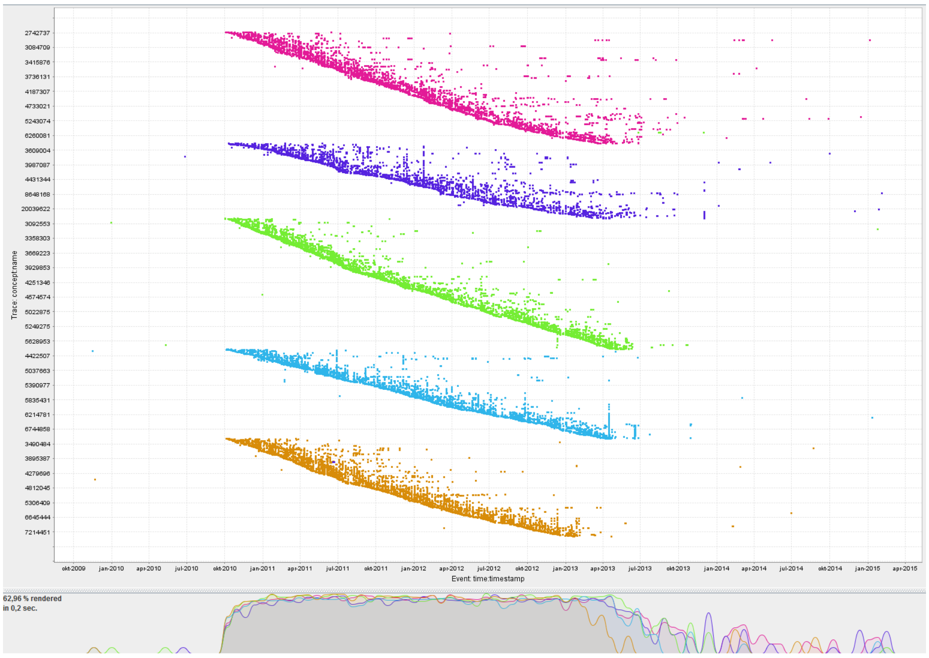

Fig. 6: Dotted chart visualization of the five event logs projected on traces with at least 50% fine-grained events. Municipalities are sorted ascending in a top down fashion, i.e., municipality 1 is purple, municipality 2 is dark blue...

As the figures in the table clearly indicate, multiple cases have events executed on the same day. Besides false ordering of events due to timing information, we also noticed irregularities with respect to semantics. For example, some traces show that letters have been sent to applicants before the application was received. Based on the combination of both a high ratio of coarse-grained events and semantically unexpected orderings of events which are unlikely to be correct, we regard the as-is control-flow of the event logs to be untrustworthy.

Table 4: Average number of events per day per municipality.

	Average	Std. Dev.
Municipality 1	7.25	2.60
Municipality 2	9.22	4.68
Municipality 3	7.80	2.81
Municipality 4	7.42	2.63
Municipality 5	9.06	3.19

2.3 Open Versus Closed Cases

The cases in the event logs have a case attribute *caseStatus*, which indicates whether or not the handling of the case has finished (status 'G' vs 'O' from the Dutch words 'Gesloten' (closed) and 'Open' (opened)). The distribution of case status over time is shown in Figure 7. There are many cases that seem to have ended (as no events have been recorded for quite some time) but still have status O. Especially municipality 1 and 3 seem to have a lot of those cases. We cannot know for sure whether these cases are in fact opened or closed or whether we should make additional assumptions. Thus, we restrict our analyses to those cases actually marked as closed for the analyses that could be affected otherwise.

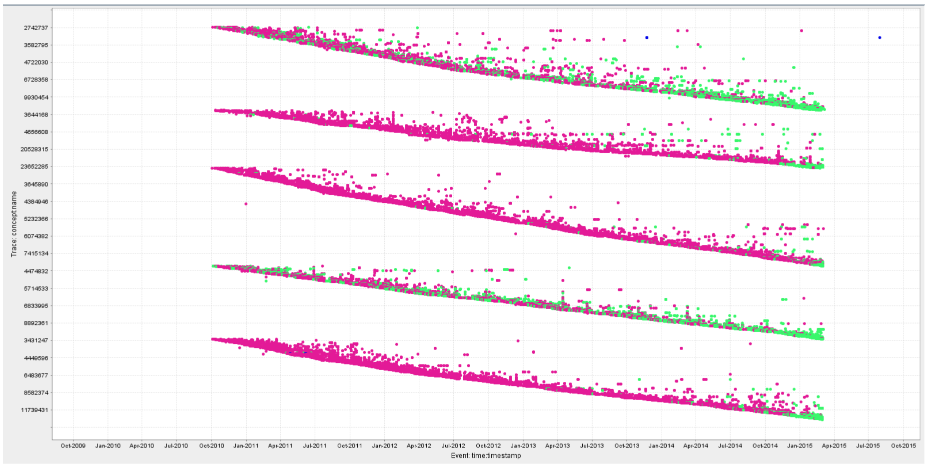

Fig. 7: Dotted chart visualization of the logged cases status. Green dots represent events of opened cases while red dots represent events of cases marked closed.

3 Organizational Structure

Three data attributes in the event log represent the resources involved in a case: *Resource* and *monitoringResource*, which are both event attributes, and *Responsible_actor*, which is a case attribute. This section will discuss the collaborations between the five municipalities for the different types of resources.

The analyses performed in the following three subsections share a common methodology: Based on the merged log containing the cases of all five municipalities, a C4.5 decision tree [5] is learned to predict the municipality based on one of the three resource attributes and the date attribute. Resources that provide work for multiple municipalities indicate collaboration or movement between municipalities. For resources that are performing work for a single municipality, solely the resource attribute will be a predictor for the municipality target variable. For resources that are performing work for multiple resources, one or more splits will be made by the decision tree learner on the timestamp variable. The value of the first timestamp split provides insight in when two municipalities started collaborating.

More in-depth resource-specific analysis is performed on a set of selected resources. We explore which resources changed location, and which collaborations between municipalities happened.

3.1 Decision Tree Analysis

Resource Using the decision tree mined from the resource and the timestamp attribute, the municipality can be predicted with an accuracy of 84,76%, up from a prior of 24%. All but six resources perform work for only a single municipality. Figure 8 highlights parts of the tree that show the resources that work for multiple municipalities. The decision tree shows that resources *560530*, *560532* and *560598* have performed work for both municipality 2 and 5, while resources *560752*, *560849*, performed work for municipalities 4 and 5. Resource *6* has performed work for municipality 1, 3 and 4.

Responsible Actor From the responsible actor and the timestamp attributes the municipality can almost perfectly be predicted (99.92%) by the decision tree, shown in Figure 9a. Only a single responsible actor, resource number *569598*, is the responsible actor in more than one municipality, namely municipalities 2 and 5.

Monitoring Resource From the monitoring resource and timestamp attributes, the prediction accuracy is poor with 22.72%, which is comparable to the prior. Figure 9b shows this decision tree. This indicates that monitoring resources are shared between municipalities and no single monitoring resource works exclusively for a single municipality.

3.2 Movement of Resources

Based on the decision trees analyses we can conclude that six resources work for more than one municipality. Only one resource works for more than one municipality as responsible actor. Monitoring resources all seem to be shared between municipalities, therefore there are no monitoring resources of particular interest that were selected for further analysis. Resource *6* is only very occasionally active in three municipalities over the four years, indicating (s)he either performs a very specialized task or only works in this process when necessary. Therefor, we focus our attention to the remaining five resources. Their movement is summarized in Table 5.

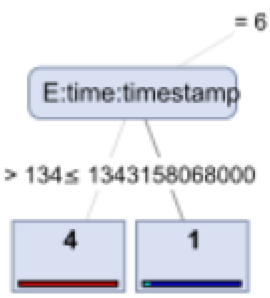

(a) Resource 6

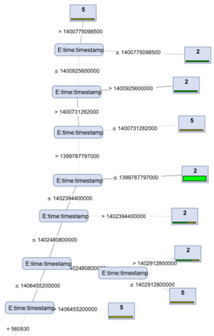

(b) Resource 560530

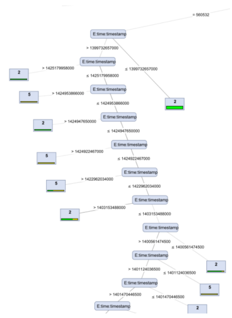

(c) Resource 560532

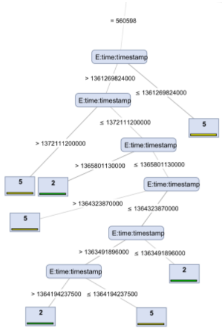

(d) Resource 560598

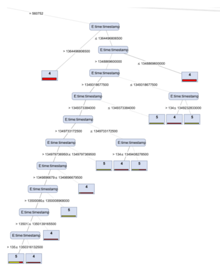

(e) Resource 560752

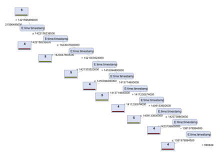

(f) Resource 560849

Fig. 8: There are six resources performing work for multiple municipalities.

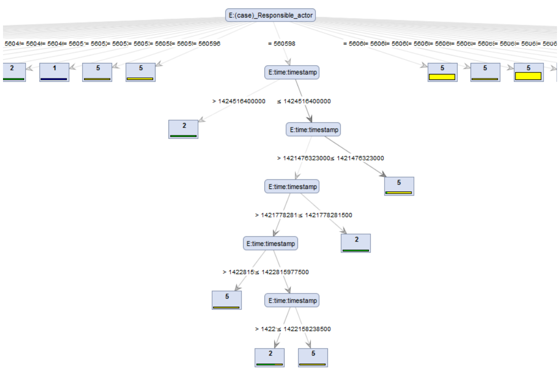

(a) Resource 569598 is the only resource responsible for events of multiple municipalities.

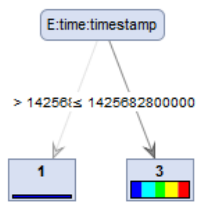

(b) The monitoring resource does not seem to have any predictive value for municipality.

Fig. 9: The responsible and monitoring resources that work for multiple municipalities.

Table 5: Movement of resources. Six resources work across municipalities.

Resource	From-municipality	To-municipality	Duration
6	1 and 3	4	Q4 2010 (1 and 3) and Q1 2014 (4)
560530	2	5	Q2 2014 - Current
560532	2	5	Q2 2014 - Current
560598	5	2	Q1 2013 - Q2 2013
560752	4	5	Q4 2012 - Q1 2013
560849	4	5	Q4 2013 - Current

Short-term Movement of Resources Figure 10 shows the timelines for resources *560598* and *560752*. They show the number of events performed per municipality over time. We can see that these resources temporarily worked in two municipalities at the same time.

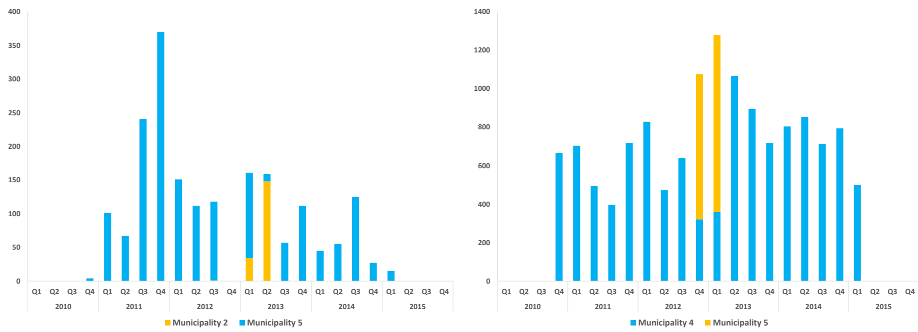

(a) Resource 560598 (b) Resource 560752

Fig. 10: Number of events performed per municipality over time. Both resources temporarily worked in two municipalities at the same time.

Resource *560598* originally only works at municipality 5, however (s)he works for municipality 2 for a brief period in Q1 and Q2 of 2013. From Figure 11a we can see that resource *560598*, who works on very specialized activities, fills in during the period when resource *560458* is away at municipality 2.

Similarly, resource *560752* originally worked only at municipality 4, but also worked for a brief period (from Q4 2012 to Q1 2013) at municipality 5. The reason can be seen in Figure 11b. Resource *1254625* was heavily involved in performing activities before Q4 2012. However, between Q4 2012 and Q1 2013 (s)he performed very few activities. As a result, *560752* seems to have joined municipality *5* to fill in. From Q2 2013 a new resource *8492512* started working at this municipality, thereby relieving *560752* of his/her duties at this municipality. Interestingly we can see that during the time *8492512* started working, resource *1254625* helps perform part of the activities.

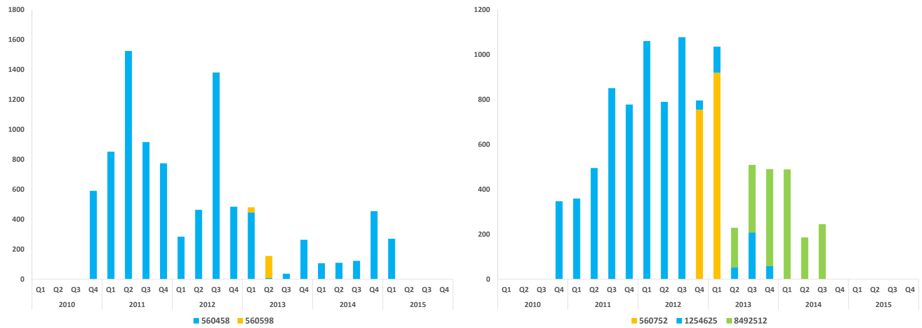

(a) Resources *560458* and *560598* at municipality *2*.

(b) Resources *560752*, *1254625*, and *8492512* at municipality *5*.

Fig. 11: The number of events over time performed by different resources.

From the above two examples, it is evident that there is a big collaboration between municipalities wherein resources having similar profiles are flexible enough to move across locations and municipalities.

Long-term Movement of Resources In Section 3.2, we looked at two of the five resources who worked at multiple municipalities. Both these resources worked for a brief period at a different municipality as a *replacement* of some other resource. Also, both the cases occurred some time ago in history and both the resources returned to their original municipalities after a while. Figure 12 shows the number of events performed over time of the remaining 3 resources who moved across municipalities. All these three resources started doing work for municipality 5 and continue to do so up until the end of the available data.

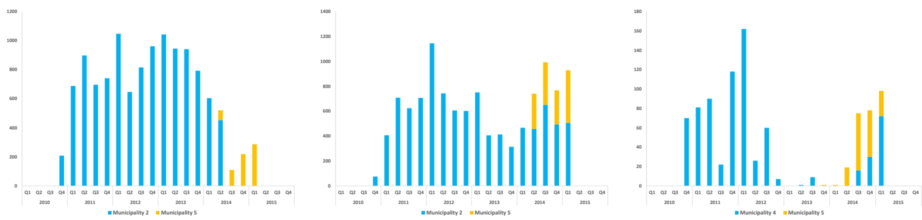

(a) Resource 560530

(b) Resource 560532

(c) Resource 560849

Fig. 12: Number of events over time for resources which started doing work for municipality 5 and continue to do so.

This recent movement of resources corresponds to question 3 of the competition: *The employees of two of the five municipalities have physically moved into* *the same location recently. Did this lead to a change in the processes and if so, what is different?*

In lieu of answering said question, we start with an analysis of the movement of the three employees to the new municipality (5). From Figure 12 we can see that almost all three resources who moved to municipality 5 registered activities at municipality 5 from May 2014 onwards. The process might depend on the *type* of the permit, therefore we zoom in to a a specific type of permit application. We've shown that more than half of the cases are of type *Bouw*. We now concentrate only on *Bouw* applications, as it gives us an homogeneous and large enough dataset. Also, as there would be many occurrences of change in processes in the past due to other reasons, such as change in protocols, legal requirements etc., we consider a time window of 10 months before and after May 2014. This results in logs of 89 and 115 cases for *Bouw* cases before and after May 2014 respectively.

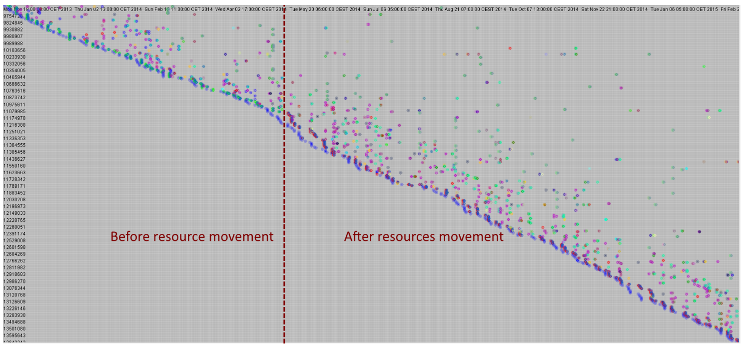

Fig. 13: Dotted chart showing concept drift for after movement of new resources to municipality 5 for cases with *Bouw* parts.

We start our comparison analysis with the dotted chart diagram Figure 13. This is a zoomed-in version of dotted chart to clearly visualize the differences in the process. The vertical dotted line distinguishes all the cases which occurred before and after the movement of new resources in the municipality. Each *dot* in the diagram represents an event. The horizontal axis represent the individual cases and the vertical axis represent time. All the cases start with an activity of the same event class - *register submission date request* (blue dot). However, for cases which had started before the resource movement, the first activity is immediately followed by the green activity, and then the next activities are spread over a longer period of time. On the other hand, for cases after resource movement, the first activity is immediately followed by red, blue and purple activities. The new resources were the ones who mostly performed these activities, which hints that their introduction caused this change in process flow.

3.3 Resource Collaboration

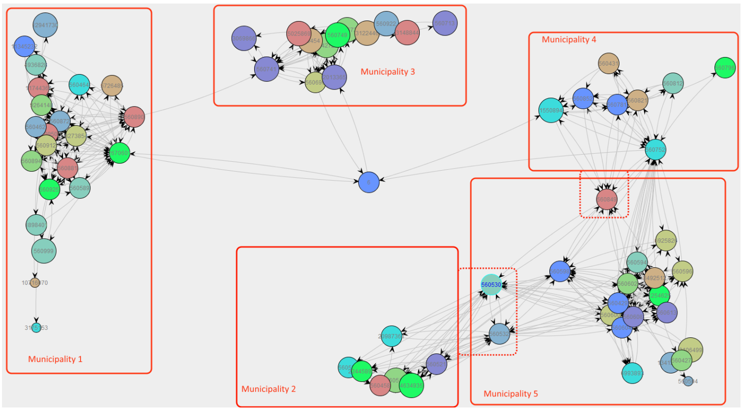

Fig. 14: Social network analysis of resources of all five municipalities representing handover of work, also showing the movement across municipalities.

Figure 14 shows the handover of work diagram for all the resources across all the municipalities. Each municipality is distinguished by a solid red box. Resource *6* which is at the centre of the figure has registered activities at three municipalities. As explained earlier, resource 6 worked for small time intervals (∼2 months) on very few activities at all 3 municipalities. Municipalities 3 and 1 collaborate the least amongst all the five municipalities. Other than resource *6*, the only other time when these municipalities were involved in collaboration was when resource *560890* from municipality 1 worked at municipality 3 for a brief period of time. This explains the outgoing edge from resource *560890* of municipality 1 to municipality 3. Municipality 5 is heavily involved in collaborations, as evident from the highly intertwined web of connections between its resources across other municipalities. In Table 5, all other resource movements have municipality 5 involved, and all these resource movements are evident from Figure 14. Also, as discussed in Section 3.2, the three resources which previously belonged to two different municipalities (shown by red dotted boxes currently work at municipality 5. Also, upon investigating the betweenness analysis of the handover of work network, it is evident that there is a high degree of collaboration *within* each municipality. This follows from the fact that most of the resources at the municipality are involved in multiple activities within the municipality. The *outliers* of the network involved in this case perform the handover duties at a much lower frequency than other resources within the municipality, and hence are on the edge of the network with respect to each municipality. The working together social network results in a similar graph and also shows groups of resources spread in 5 clusters corresponding to 5 municipalities.

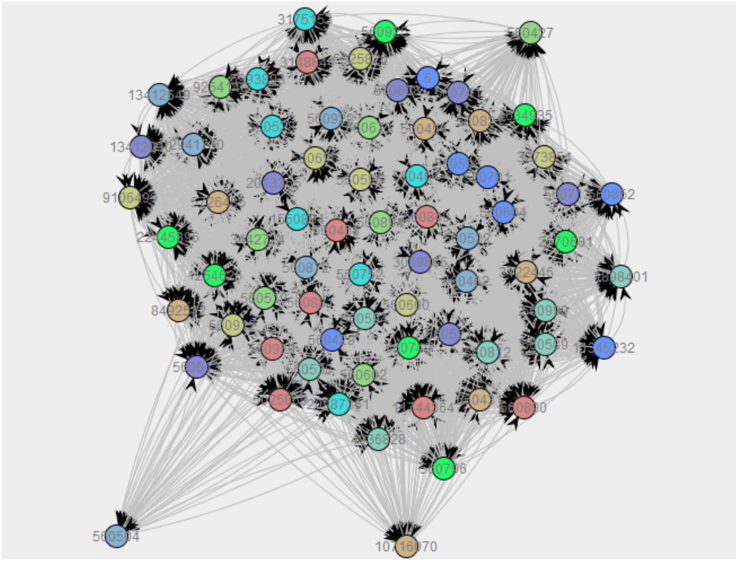

Fig. 15: Social network analysis for similar tasks for all the resources across all five municipalities.

Figure 15 shows the distribution of similar tasks for all resources across all municipalities. We see a big cluster containing almost all the resources in the middle. There are three resources that are not a part of this cluster, and these are the resources who have worked on only one case each, performing between 2 to 10 specialized activities in total. From Figure 15 we can conclude that all the resources work on very similar tasks and each resource is capable of performing other resource's task within the municipality as well as across municipalities. This is also reflected by the movement of resources across municipalities.

3.4 Resource Roles

Subprocesses In order to see what the roles of people involved in various stages of the process are we have mapped the number of events performed per subprocess per resource in Figure 17. Note that the main process (HOOFD) is not included in these graphs as the number of events performed for it is much higher. We can see some resources having similar preferred or assigned tasks. However, mostly resources seem to be flexible, i.e. there don't seem to be resources that only perform activities in a certain phase or subprocess. This is true for all five municipalities.

Responsibilities Besides looking into the roles of resources with respect to their activity in the various subprocesses we can also consider resource roles as categories of responsibilities. As explained in Section 2, there are three types of responsibilities a resource can have: executing resource, monitoring resource, and/or responsible actor. After analyzing the data we have found that there are in fact only three combinations of these responsibilities occurring in the logs. These are: only executing resource, executing and monitoring resource, and all three. For example, there are no resources that are solely monitoring resource. As such, we now consider these three combinations as roles. Table 6 lists the distribution of resources over the three groups. We can see that the municipalities all have a similar distribution, except municipality 5. In this municipality relatively more resources are just executing resource.

Table 6: Number of resources per group per municipality.

Municipality	just resource	resource and monitoring resource	all three roles
1	3 (13,04%)	4 (17,39%)	16 (69,57%)
2	2 (18,18%)	1 (9,09%)	8 (72,73%)
3	0	3 (23,08%)	10 (72,73%)
4	1 (11,11%)	1 (11,11%)	7 (77,78%)
5	6 (37,5%)	4 (25%)	6 (37,5%)

Figure 16 lists a percentage-wise overview of the events performed by the three roles per municipality. We can see that the resources that only execute tasks mainly perform work in municipality 1 and 2, relatively. In these two municipalities, most of the work is still performed by resources that are both monitoring and responsible actors. In municipality 3,4 and 5 we note that most activities are performed by the other two groups. Thus, if the groups identified indeed correspond to the hierarchy within the municipalities, this hints on a *flat hierarchical structure*.

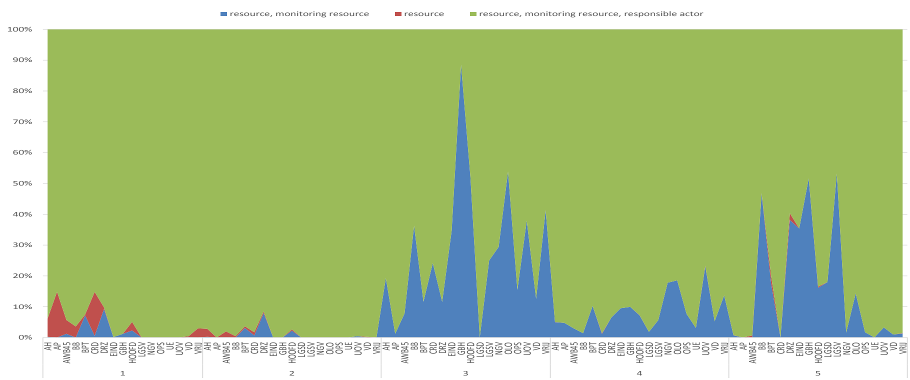

Fig. 16: Distribution of events over organizational groups per process per municipality. The graph hints on a flat hierarchical structure.

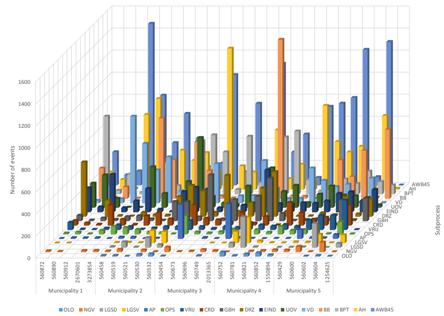

(a) Top 5 resources per municipality

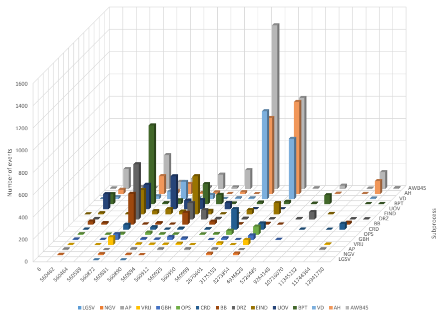

(b) Municipality 1

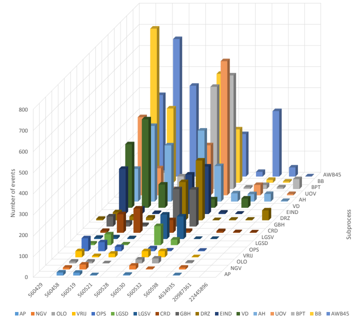

(c) Municipality 2

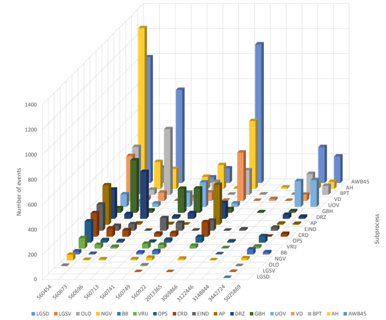

(d) Municipality 3

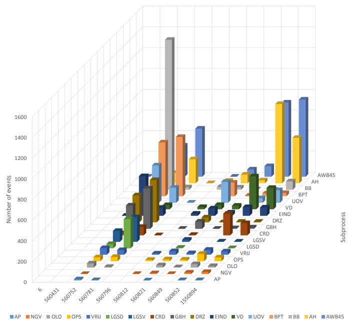

(e) Municipality 4

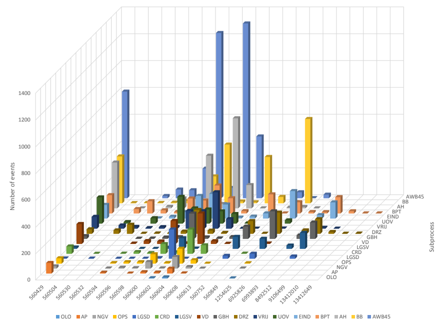

(f) Municipality 5

Fig. 17: Number of events per resource per subprocess. Subprocesses are sorted on requency. Resources seem to be flexible and no clear roles can be distinguished.

4 Control-flow Evaluation

4.1 Resource Re-Allocation

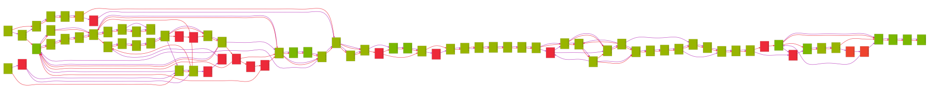

Fig. 18: Differences in process flow before and after the movement of resources to municipality 5. The activities in green are the **common** activities in the processes before and after the movement of resources. Activities in red correspond to the activities which were performed before the movement of resources and discontinued after or vice versa.

Using the insights from dotted chart analysis we see that there is a clear difference in process wherein some activities were performed before are not performed after the movement of resources, and vice-versa. We used a 10 month period before and after resource movement to analyse the impact of the movement of resources. This limits the impact (interference) of other process modifications due to e.g. regulatory changes. Figure 18 distinguishes the two processes by plotting the the common and uncommon activities for both logs (before & after resource movement) on top of each other. The green activities correspond to the activities common in both processes, whereas the red activities are the activities which were mostly performed in either one of the processes (before or after resource movement). The red edge between activities corresponds to the process - after the movement of resources, and the pink edge between activities corresponds to the process - before the movement of resources. The overall structure of the process is quite similar (before and after the movement of resources), as indicated by the pairs in which red and pink edges appear in the figure. However there are some activities (in red) which occur only before (or only after) the movement of resources.

Table 7 shows the differences in activities which occur only before (or after) the resources movement. There are 9 activities which occurred regularly before the movement of resources to the municipality. However these activities did not occur (or occurred very few times) after the movement of resources. There are 7 activities which never occurred before May 2014, and started occurring after May 2014. The resources responsible for these activities are indeed two out of the three new resources (*560530 and 560532*). This strongly suggests that these resources have brought in some new activities which are now being performed at this municipality (replacing some of the previous activities). The third resource *560849* is not responsible for any of the *new* activities introduced in the municipality. This can be attributed to the fact that this resource performs very few activities (±15 per month), and seems responsible only for certain specialized activities which are not related to *Bouw* cases.

Table 7: ***Emphasized activities*** are the activities which occur in a considerably higher frequency before (or after resp.) the resource movement.

	After movement (post May'14)		Before movement (pre May'14)	
Activity Name (Concept Name)	Relative freq (Abs. freq)	Case status	Relative freq (Abs. freq)	Case status
applicant is stakeholder (03_GBH_005)	0.02 (2)	G(2) O(0)	***1.0 (89)***	***G(88) O(1)***
MER present in supplement (01_HOOFD_193)			***0.21 (19)***	***G(18) O(1)***
inform BAG administrator (01_HOOFD_050)	***0.29 (33)***	***G(4) O(29)***		
investigate BAG objects (01_HOOFD_055)	***0.17 (20)***	***G(2) O(18)***		
terminate on request (05_EIND_010)	0.03 (3)	G(3) O(0)	***1.0 (89)***	***G(88) O(1)***
calculate provisional charges (16_LGSV_010)	0.01 (2)	G(2) O(0)	***1.0 (89)***	***G(88) O(1)***
publish (01_HOOFD_090)	***0.98 (113)***	***G(70) O(43)***	0.02 (2)	G(2) O(0)
create publication document (01_HOOFD_100_0)	***0.40 (47)***	***G(11) O(36)***		
registration date publication (01_HOOFD_101)	***0.4 (46)***	***G(10) O(36)***		
create subcases content (01_HOOFD_250_0)	***0.79 (91)***	***G(61) O(30)***		
assessment of content completed (01_HOOFD_370)	0.01 (1)	G(1) O(0)	***0.85 (76)***	***G(74) O(2)***
calculate final charges			***0.76 (68)***	***G(68) O(0)***
create monitoring case oversight (01_HOOFD_532_0)	***0.72 (83)***	***G(57) O(26)***		
registration date publication (01_HOOFD_101b)	0.01 (1)	G(1) O(0)	***0.42 (37)***	***G(37) O(0)***
read publication date field (01_HOOFD_809)	0.03 (4)	G(4) O(0)	***0.45 (40)***	***G(40) O(0)***
read publication date field (01_HOOFD_809s)	0.02 (3)	G(3) O(0)	***0.37 (33)***	***G(33) O(0)***

4.2 Omgevingsloket Online

To assess the differences in behavior between municipalities without taking into account the sequential ordering of events, we calculated the set of activities that occur within each municipality and compared the municipalities in a pairwise manner. The most notable difference in the resulting activity set differences is found in the degree in which the online portal *Omgevingsloket online* (often abbreviated to OLO) is used in the process. Table 8 shows the occurrence of OLO events per municipality and their frequency.

Table 8: Usage of the OLO portal in the five municipalities.

Municipality	OLO event class
1	794x OLO messaging active
1	240x reception through OLO
2	590x OLO messaging active
2	148x reception through OLO
2	34x received OLO documents
2	16x application submitted through OLO
2	12x send message OLO-status in progress
2	5x send message OLO status additions required
3	938x OLO messaging active
3	295x reception through OLO
3	4x application submitted through OLO
3	2x send message OLO-status in progress
3	1x send message OLO status additions required
4	722x OLO messaging active
4	176x reception through OLO
4	101x received OLO documents
4	35x application submitted through OLO
4	23x send message OLO-status in progress
4	10x send message OLO status additions required
5	779x OLO messaging active
5	253x reception through OLO
5	46x received OLO documents
5	5x application submitted through OLO
5	3x send message OLO-status in progress
5	1x request advice through OLO
5	1x send message OLO status additions required
5	1x send message OLO status decision

The most prominent OLO event class in each municipality is *OLO messaging active*, followed by *reception through OLO*. Figure 19 shows the occurrence of OLO events over time per municipality. Notable in this figure is the gap of OLO events that is observed for each municipality that starts at the beginning of July 2011 At each municipality there is a gap of OLO events that starts at the beginning of July 2011 and ends between November 2011 and January 2012 depending on the municipality. The start of this gap coincides with the release of a new version of the Omgevingsloket Online was on the first of July 2011[8]. The gap ends at January 16th of 2012, at which municipality 4 starts logging OLO activities again as the last of the five municipalities. The *OLO messaging active* event class occurs for the first time in the log after the gap, and occurs in 3682 of the 4096 traces of all municipalities combined that start after the gap (after January 16th of 2012). The second most frequent OLO event class, *reception through OLO*, occurs in 1065 of the 1209 traces that ended before the gap in OLO events, but never occurs after this gap. It seems to be the case that the event class *reception through OLO* was replaced by event class *OLO messaging active* in a new version of OLO, released in July 2011.

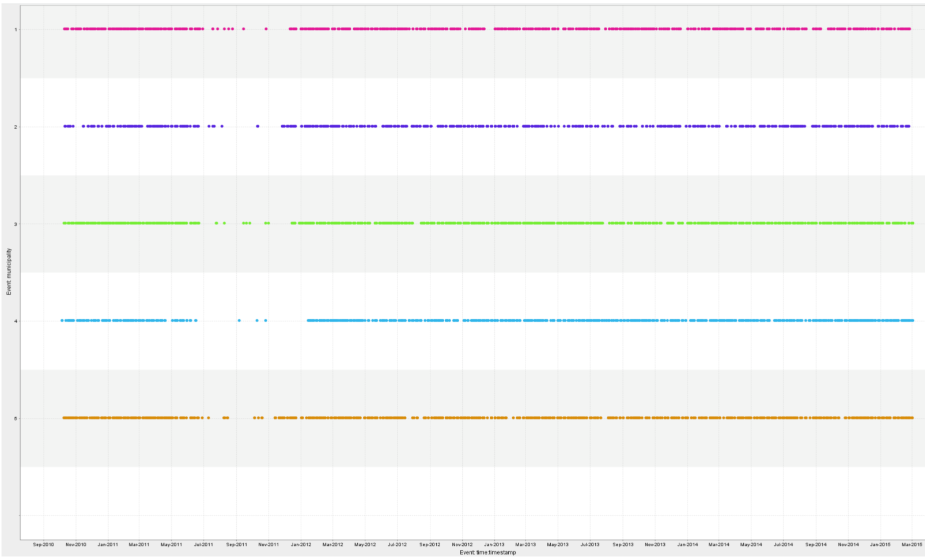

Fig. 19: The occurrence of OLO events over time per municipality

Since *OLO messaging active* and *reception through OLO* seem to be executed by default, analyzing the occurrence OLO activities other than those two gives an indication of the degree in which the OLO system is used in each municipality. Table 8 shows that municipality 1 performs no other activities with the OLO system and municipality 3 performs very little activities with the OLO system. Municipality 4 performs the most activities with the OLO system.
Figure 20 shows the event class of OLO events per municipality over time when *OLO messaging active* and *reception through OLO* are not considered. It is

[8]http://www.wabobank.nl/nieuwe-release-omgevingsloket-online-beschikbaar-op-1-juli

noticeable that the OLO events start occurring after the new OLO system was suspected to be introduced. A small number of events occurred shortly after July 2011, which coincides with the start of the gap of OLO events identified in Figure 19 and then start occurring again after the gap of OLO events that was identified in Figure 19. From figures Figure 19 and 20 combined we can conclude that the only OLO events that occurred before July 2011 is *reception through OLO*, all other OLO events seemed to have been enabled by a new OLO system in July 2011.

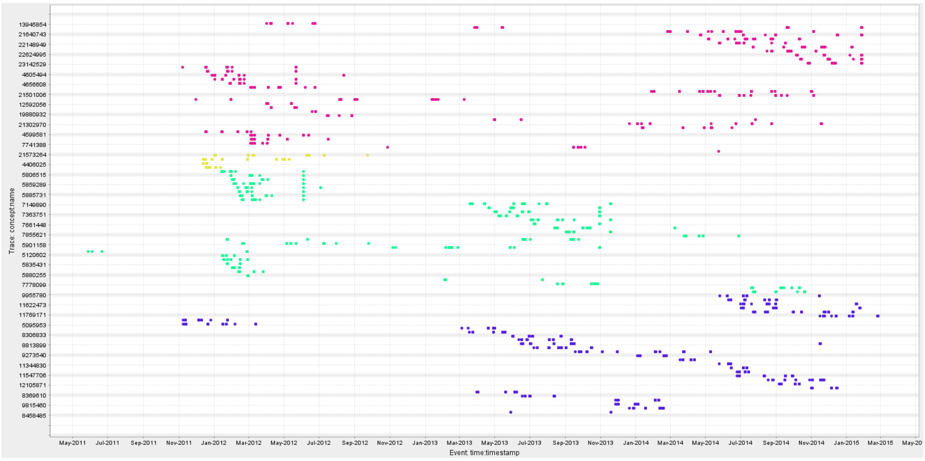

Fig. 20: The occurrence of OLO events without *OLO messaging active* and *reception through OLO* over time per municipality. Pink represents municipality 2, yellow represents municipality 3, light blue represents municipality 4 and dark blue represents municipality 5

5 Performance Evaluation

To assess the general performance of the five municipalities we have measured the throughput time in the order of days of each *completed trace* within the five event logs. The basic aggregate throughput time statistics of each municipality are depicted in Table 9. The kernel density function for each municipality, based upon the same underlying data, is depicted in Figure 21.

5.1 Municipality Differences

We clearly identify municipality 3 having the, on average, lowest throughput time, followed by municipality 1. The difference between the two municipalities, in terms of average, exceeds 30 days. Even when we compare the median

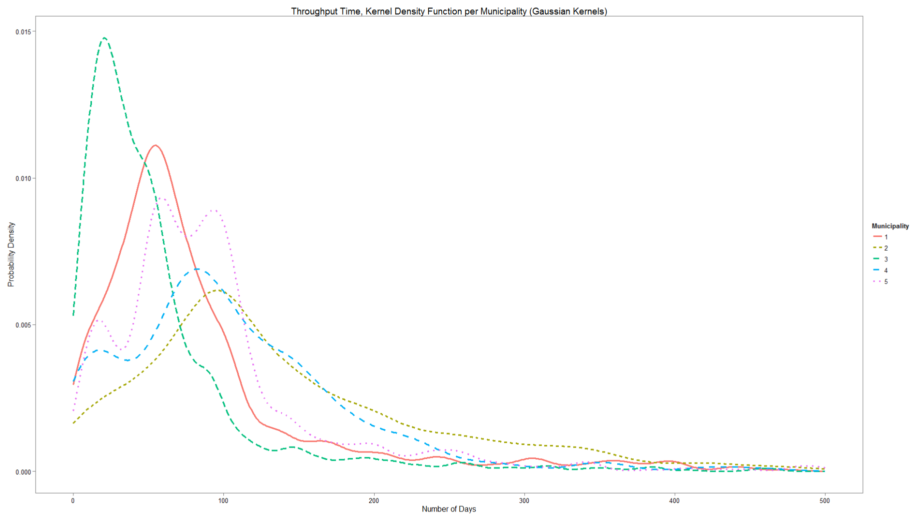

Fig. 21: Kernel density function per municipality, using Gaussian kernels

throughput time of the two municipalities, the difference exceeds 20 days. Thus, municipality 3 seems to be greatly outperforming the other municipalities. Interestingly, municipality 3 is also the municipality having the largest number of cases (Table 1) Municipality 2 seems to be the least performing municipality in terms of throughput time. Again we identify the relation with the number of cases present in the municipality, as municipality 2 has the least number of cases within its corresponding event log.

To assess the significance of the differences in throughput time for the given municipalities we have applied the Mann-Whitney U test [6] to the throughput times. The test was performed using a confidence interval of 0.99, i.e. $\alpha = 0.01$. The results of the test are depicted in Table 10. Each p-value depicted in the table is smaller than the

We additionally computed the average throughput time per permit type, per municipality. Within this analysis we identify a very interesting phenomenon. After filtering out types that occur rarely, for each municipality, the top-5 types having the *longest* throughput time are the same:

Table 9: Aggregate throughput statistics in days per municipality.

	Avg. Throughput	Median Throughput	Std. Dev. Throughput
Municipality 1	95.85	62.5	124.91
Municipality 2	159.3	115	150.38
Municipality 3	62.28	39	97.42
Municipality 4	110.6	92.5	96.43
Municipality 5	100.61	79	107.24

Table 10: p-values of the throughput times of the five municipalities (M1 .. M5) based on the Mann-Whitney U test [6] with $\alpha = 0.01$.

	M1	M2	M3	M4	M5
M1		$< 2.2 \cdot 10^{-16}$	$< 2.2 \cdot 10^{-16}$	$1.228 \cdot 10^{-12}$	$1.685 \cdot 10^{-07}$
M2			$< 2.2 \cdot 10^{-16}$	$8.881 \cdot 10^{-12}$	$< 2.2 \cdot 10^{-16}$
M3				$< 2.2 \cdot 10^{-16}$	$< 2.2 \cdot 10^{-16}$
M4					$5.53 \cdot 10^{-05}$
M5					

- {Milieu (Vergunning)}
- {Bouw, Handelen in strijd met regels RO}
- {Kap}
- {Bouw}
- {Handelen in strijd met regels RO}
- {Sloop}

The type *Milieu (Vergunning)* has the longest throughput time for all municipalities. Looking at the individual types, we identify a comparable trend w.r.t. the overall throughput times. For example, the average throughput time for the top-5 types of municipality 3 are all lower compared to those of municipality 2. The overall differences seem not to be explainable in terms of the types associated to the cases being executed within the municipalities.

We also inspected the average number of resources involved in handling a case. The average number of resource per case are remarkably close:

- Municipality 1: 2.65
- Municipality 2: 2.49
- Municipality 3: 2.47
- Municipality 4: 2.58
- Municipality 5: 2.81

Although municipality 3 does have the lowest average number of resources involved in handling a case, the figure is very close to municipality 2. Hence the average number of resources per case does not seem to be a good indicator for the throughput time.

To assess the actual indicator, i.e., explaining why municipality 3 is significantly faster, a thorough analysis involving control-flow should be conducted.

Table 11: p-values of the throughput times of OLO and Non-OLO cases based on the Mann-Whitney U test [6] with $\alpha = 0.01$.

	p-value	# of OLO cases	avg. throughput OLO	avg. throughput non-OLO
M1	-	0	-	95.8
M2	0.4036	33	173.8	158.8
M3	0.636	4	110.5	62.1
M4	0.04262	31	142.5	108.6
M5	0.8309	30	105.5	100.5
Combined	$1.281 \cdot 10^{-6}$	98	140.4	98.7

For example, bottleneck analysis can highlight the inefficiencies within the processes of the municipalities and might therefor provide insights into the potential causes for high or low throughput times. However, due to the data quality problems as presented in Section 2, we refrain from performing such control-flow oriented, analysis.

5.2 OLO vs. Non-OLO Differences

In this section we use the Mann-Whitney U test to determine the difference in throughput times between cases in which the OLO system is used and cases in which the OLO system is not used. Here, a case will be regarded as a case in which the OLO system is used when it contains at least one of the OLO activities listed in Table 8 other than the activities *OLO messaging active* and *reception through OLO*, which are by default present in each case. Only complete cases are considered, as throughput time cannot be determined for incomplete cases. The log consists of 98 cases that contains OLO activity other than *OLO messaging active* or *reception through OLO*. By applying the Mann-Whitney U test we found that cases with no OLO activity are significantly faster than cases with OLO activity, with a p-value of $1.281 \cdot 10^{-6}$. The average throughput time of cases with OLO activity is 140.4 days, while the average throughput time of non-OLO activity is 98.7 days. When we look at each municipality individually, we do not find a significant difference in throughput time between OLO and non-OLO cases because the number of OLO cases per municipality is small, but we do see that non-OLO cases are faster on average than OLO cases for all municipalities.

6 Conclusion

In this report we present our findings of the analysis of five event logs containing data related to building permit application requests, as part of the BPI Challenge 2015. We found several inconsistencies in the event data, most prominently the inaccurate logging of events. Although data inconsistencies are a common problem in real event logs, the predominant presence of inconsistent data logging prohibited the use of more advanced, control-flow-based analysis techniques including the state of the art in process mining techniques.

The analysis covers three dimensions: the organizational dimension, the performance dimension and the control-flow dimension. With respect to the organizational dimension, we identified collaborations between municipalities 1 and 4, 4 and 5, and municipalities 2 and 5. With regard to the different sub-processes, we identified that some resources share a similar profile in preferred/assigned tasks. However, there don't seem to be any clear roles of resources with regard to their involvement in the different activities or sub-processes.

Control-flow analysis focused on the effect of resource re-allocation. We identified a solid effect of resource re-allocation on the activities being performed within the municipality.

From a performance point of view we found the throughput times of all municipalities to be significantly different from each other. Based on the significance of the pairwise differences and the average throughput times we can rank the municipalities in the order 3,1,5,4,2, from lowest to highest average throughput time.

We identified the fact that some of the municipalities use a central system, abbreviated as OLO. We identified a gap in the use of OLO in the beginning of the logged time range, after which a sudden change in the type of activity executed in the beginning was noticeable. The aforementioned observation hints at some temporal down-time of the OLO-system, possibly related to an update of the system, potentially related to a change in regulations.

The frequency of use of the OLO system differs throughout the municipalities. Interestingly, municipality 3, which has the lowest average throughput time, does not seem to be using the OLO system frequently. All municipalities combined we found the OLO-based cases to be significantly slower than the non-OLO cases. Per municipality individually this effect was not found to be significant, possibly due to the number of OLO-based cases being rather low.

We did not find any control-flow based indicator explaining the (significant) differences in terms of average throughput time among the municipalities. As motivated, a bottleneck analysis could potentially highlight inefficiencies within the processes and might thereby provide causes for high or low throughput times of the municipalities. Due to data quality problems, we refrained from performing this type of analysis. Identified data inconsistencies forms a basis for an improvement in logging quality necessary to enable such analysis in the future.

References

1. van Dongen, B.F.: BPI Challenge 2015 (2015) 1
2. Verbeek, H.M.W., Buijs, J.C.A.M., van Dongen, B.F., van der Aalst, W.M.P.: ProM 6: The process mining toolkit. In: Proc. of BPM Demonstration Track 2010. Volume 615., CEUR-WS.org (2010) 34–39 2
3. van der Aalst, W.M.P.: Process Mining - Discovery, Conformance and Enhancement of Business Processes. Springer (2011) 2
4. Song, M., van der Aalst, W.M.P.: Supporting process mining by showing events at a glance. In: 7th Annual Workshop on Information Technologies and Systems. (2007) 139–145 2
5. Quinlan, J.R.: C4.5: programs for machine learning. Elsevier (1993) 10
6. Mann, H.B., Whitney, D.R.: On a test of whether one of two random variables is stochastically larger than the other. The annals of mathematical statistics (1947) 50–60 25, 26, 27