

Discovering Social Networks from Event Logs

WIL M. P. VAN DER AALST¹, HAJO A. REIJERS¹ &
MINSEOK SONG^{1,2}

¹*Department of Technology Management, Eindhoven University of Technology, 513, NL-5600 MB, Eindhoven, The Netherlands (E-mail: w.m.p.v.d.aalst@tm.tue.nl);* ²*Department of Industrial Engineering, Pohang University of Science and Technology, San 31 Hyoja-Dong, Nam-gu, Pohang, 790-784, South Korea*

Abstract. Process mining techniques allow for the discovery of knowledge based on so-called “event logs”, i.e., a log recording the execution of activities in some business process. Many information systems provide such logs, e.g., most WFM, ERP, CRM, SCM, and B2B systems record transactions in a systematic way. Process mining techniques typically focus on performance and control-flow issues. However, event logs typically also log the *performer*, e.g., the person initiating or completing some activity. This paper focuses on mining social networks using this information. For example, it is possible to build a social network based on the hand-over of work from one performer to the next. By combining concepts from workflow management and social network analysis, it is possible to discover and analyze social networks. This paper defines metrics, presents a tool, and applies these to a real event log within the setting of a large Dutch organization.

Key words: business process management, data mining, Petri nets, process mining, social network analysis, workflow management

1. Introduction

This paper builds on concepts from *business process* management (workflow management in particular) and sociometry (social network analysis in particular).

Business process management is concerned with *process-aware information systems*, i.e., systems supporting the design, analysis, and enactment of operational business processes. Typical examples of such process-aware systems are workflow management systems where the process is driven by an explicit process model (Jablonski and Bussler, 1996; Leymann and Roller, 1999; Aalst and Hee, 2002). However, in many other process-aware information systems the process model is less explicit and users can deviate from the “normal flow”, i.e., these systems allow for more flexibility.

Sociometry, also referred to as sociography, refers to methods presenting data on interpersonal relationships in graph or matrix form (Burt and Minor, 1983; Scott, 1992; Wasserman and Faust, 1994). The term sociometry was coined by Jacob Levy Moreno who conducted the first long-range sociometric study from 1932 to 1938 at the New York State Training School for

Girls in Hudson, New York (Moreno, 1934). As part of this study, Moreno used sociometric techniques to assign residents to various residential cottages. He found that assignments on the basis of sociometry substantially reduced the number of runaways from the facility. Many more sociometric studies have been conducted since then by Moreno and others. In most applications of sociometry, the assessment is based on surveys (also referred to as sociometric tests). With the availability of more electronic data, new ways of gathering data are enabled (Feldman, 1987). By analyzing the history of a user's e-mail interactions, personal networks can be extracted. One of the first social-networked tools developed for this purpose is ContactMap (Nardi et al., 2002). BuddyGraph (<http://www.buddygraph.com>) and MetaSight (<http://www.metasight.co.uk>) are other examples. By using logs on e-mail traffic as a starting point, meaningful organizational patterns can be distinguished (see e.g., Ogata et al., 2001; Nardi et al., 2002; Begole et al., 2002; Farnham et al., 2004a, b; Fisher and Dourish, 2004). Similarly, information on the Web can be used for the analysis of social networks (Culotta et al., 2004). For example, Usenet data has been used to characterize the "authority" of individuals based on posting patterns (Smith, 1999).

For the analysis of social networks around *business processes* such approaches are less useful, since they are based on unstructured information. For example, when analyzing e-mail it is difficult, but also crucial, to distinguish between e-mails corresponding to particular activities within a business process (e.g., the decision with respect to a loan request) and e-mails representing less relevant operational details (e.g., scheduling a meeting). Fortunately, many enterprise information systems store relevant events in a more structured form. For example, workflow management systems typically register the enabling, start and completion of activities (Jablonski and Bussler, 1996; Leymann and Roller, 1999; Fischer, 2001; Aalst and Hee, 2002). ERP systems like SAP log all transactions, e.g., users filling out forms, changing documents, etc. Business-to-business (B2B) systems log the exchange of messages with other parties. Call center packages but also general-purpose CRM systems log interactions with customers. These examples show that many systems have some kind of *event log* often referred to as "history", "audit trail", "transaction file", etc. (Agrawal et al., 1998; Grigori et al., 2001; Sayal et al., 2002; Aalst et al., 2003).

When people are involved in events, logs will typically contain information on the person executing or initiating the *event*. We only consider events both referring to an *activity* and a *case* (Aalst et al., 2003). The case (also named process instance) is the "thing" which is being handled, e.g., a customer order, a job application, an insurance claim, a building permit, etc. The activity (also named task, operation, action, or work-item) is some operation on the case, e.g., "contact customer". An event may be denoted by (c, a, p) where c is the case, a is the activity, and p is the person. Events are ordered in

time allowing the inference of causal relations between activities and the corresponding social interaction. For example, if (c, a_1, p_1) is directly followed by (c, a_2, p_2) , there is some handover of work from p_1 to p_2 (note that both events refer to the same case). If this pattern (i.e., there is some handover of work from p_1 to p_2) occurs frequently but there is never a handover of work from p_1 to p_3 although p_2 and p_3 have identical roles in the organization, then this may indicate that the relation between p_1 and p_2 is stronger than the relation between p_1 and p_3 . Using such information it is possible to build a social network expressed in terms of a graph (“sociogram”) or matrix.

Social Network Analysis (SNA) refers to the collection of methods, techniques and tools in sociometry aiming at the analysis of social networks (Burt and Minor, 1983; Scott, 1992; Wasserman and Faust, 1994). There is an abundance of tools allowing for the visualization of such networks and their analysis. A social network may be dense or not, the “social distances” between individuals may be short or long, etc. An individual may be a so-called “star” (directly linked to many other individuals) or an “isolate” (not linked to others). However, also more subtle notions are possible, e.g., an individual who is only linked to people having many relationships is considered to be a more powerful node in the network than an individual having many connections to less connected individuals.

The work presented in this paper applies the results from sociometry, and SNA in particular, to events logs in today’s enterprise information systems. The main challenge is to derive social networks from this type of data. This paper presents the approach, the various metrics that can be used to build a social network, our tool *MiSoN* (Mining Social Networks),¹ and a case study. The paper extends the results presented in (Aalst and Song, 2004) by providing concrete metrics and demonstrating these using a case study.

The case study, used to demonstrate the applicability of our approach, has been conducted within a Dutch national public works department employing about 1000 civil servants. Based on the particular process of handling invoices, we constructed several social networks using the various metrics defined in this paper. As a starting point we used an event log with about 5000 cases and more than 33,000 events. One of the metrics we applied was the handover of work mentioned before, i.e., the strength of the link from person p_1 to person p_2 is calculated based on the number of times one activity by p_1 is followed by an activity by p_2 for the case (e.g., an invoice). The resulting social networks were analyzed using a variety of SNA techniques and the outcomes were discussed with the management of the Dutch national public works department.

The paper is organized as follows. Section 2 introduces the concept of process mining. Section 3 focuses on the mining of organizational relations, introducing concepts from SNA but also showing which relations can be derived from event logs. Section 4 defines the metrics we propose for mining

organizational relations. We propose metrics based on (possible) causality, metrics based on joint cases, metrics based on joint activities, and metrics based on special event types (e.g., delegation). Then we present our tool MiSoN. Section 6 discusses the case study. Section 7 presents related work. Finally, Section 8 concludes the paper.

2. Process mining: An overview

The goal of process mining is to extract information about processes from transaction logs (Aalst et al., 2003). We assume that it is possible to record events such that (i) each event refers to an *activity* (i.e., a well-defined step in the process), (ii) each event refers to a case (i.e., a process instance), (iii) each event refers to a *performer* (the person executing or initiating the activity), and (iv) events are totally ordered. An *event log* is a collection of events. Any kind of transactional information system, e.g., ERP, CRM, or workflow management systems, will offer this information in some form (Jablonski and Bussler, 1996; Leymann and Roller, 1999; Fischer, 2001; Aalst and Hee, 2002).

To clarify notions such as *activity*, *case*, *performer*, and *event log*, let us consider some examples.

- Consider the *cancer treatment process* in a hospital. Each case in this process refers to a patient having cancer. Examples of activities that may take place are visits of the patient to the hospital for chemotherapy, consultation of a specialist, a bone marrow transplant, etc. The performers are all kinds of health-care professionals, e.g., doctors, specialists, nurses, surgeons, oncologists, radiologists, etc. An event may be the administration of chemotherapy to a patient by a nurse at a given point in time. The event log for the cancer treatment process will contain all events relevant for this process.
- Another example is the *processing of job applications*. Each case refers to a job application. Possible activities are confirming the application, inviting the applicant for an interview, reimbursing the travel costs, making a decision, doing a medical test, etc. The performers in the process include the personnel officers, department managers, doctors, etc. An example of an event is the invitation of an applicant for an interview by a personnel officer. The event log for the job application process will contain more such events ordered in time.
- The last example comes from the scientific domain. Consider the *reviewing process* for a journal. The cases are papers that are submitted to the journal. Examples of activities are the reviewing of a paper, the acceptance of a paper, and the notification of authors. Among the performers in this process are reviewers, area editors, editors in chief, authors, and editorial assistants. An example of an event is the notification of the corresponding author by an editorial assistant.

A more abstract example of an event log is shown in Table I. Each row refers to a single event, e.g., the last row refers to the execution of activity D by Pete for the 4th case.

Many information systems offer the information shown in Table I in some form. Any workflow management system will provide the information shown in Table I (and more). Note that we do not assume the presence of a workflow management system. The only assumption we make, is that it is possible to collect logs with event data. For example, ERP systems, CRM systems, HRM systems, hospital information systems, reviewing systems, etc. offer this information in some form. Nevertheless, it may take some efforts to extract this information in the right form.

Event logs, such as the one shown in Table I, can be used to construct models that explain some aspect of the behavior registered. The term *process mining* refers to methods for distilling a structured process description from a set of real executions (Agrawal et al., 1998; Grigori et al., 2001; Sayal et al., 2002; Aalst et al., 2003). The term “structured process description” may be interpreted in various ways, ranging from a control-flow model expressed in terms of a classical Petri net to a model incorporating organizational, temporal, informational, and social aspects. In Section 7, references to the state-of-the-art using these interpretations are given. In this paper, however, we focus on the social aspect of mining event logs.

Table I. An event log

Case identifier	Activity identifier	Performer
Case 1	Activity A	John
Case 2	Activity A	John
Case 3	Activity A	Sue
Case 3	Activity B	Carol
Case 1	Activity B	Mike
Case 1	Activity C	John
Case 2	Activity C	Mike
Case 4	Activity A	Sue
Case 2	Activity B	John
Case 2	Activity D	Pete
Case 5	Activity A	Sue
Case 4	Activity C	Carol
Case 1	Activity D	Pete
Case 3	Activity C	Sue
Case 3	Activity D	Pete
Case 4	Activity B	Sue
Case 5	Activity E	Clare
Case 5	Activity D	Clare
Case 4	Activity D	Pete

2.1. DISCOVERING SOCIAL NETWORKS

When distilling a process model from an event log, the focus is on the various process activities and their dependencies. When deriving roles and other organizational entities, the focus is on the relation between people or groups of people and the process. Another perspective is to focus on the relations among individuals (or groups of individuals) acting in the process, in other words: the social network. Consider for example the event log of Table I. Although Carol and Mike can execute the same activities (B and C), Mike is always working with John (cases 1 and 2) and Carol is always working with Sue (cases 3 and 4). Probably Carol and Mike have the same role but based on the small sample shown in Table I it seems that John is not working with Carol and Sue is not working with Carol.² These examples show that an event log can be used to derive relations between performers of activities, thus resulting in a sociogram. For example, it is possible to generate a sociogram based on the transfers of work from one individual to another as is shown in Figure 1. Each node represents one of the six performers and each arc represents that there has been a transfer of work from one individual to another. The definition of “transfer of work from A to B” is based on whether for the same case an activity executed by A is directly followed by an activity executed by B. For example, both in case 1 and 2 there is a transfer from John to Mike. Figure 1 does not show frequencies. However, for analysis purposes these frequencies can be added. The arc from John to Mike would then have weight 2. Typically, we do not use absolute frequencies but weighted frequencies to get relative values between 0 and 1. Figure 1 shows that work is transferred to Pete but not vice versa. Mike only interacts with John, Carol only interacts with Sue. Clare is the only person transferring work to herself.

For a simple network with just a few cases and performers the results may seem trivial. However, for larger organizations with many cases it may be possible to discover interesting structures. Sociograms as shown in Figure 1 can be used as input for SNA tools that can visualize the network in various ways, compute metrics like the density of the network, analyze the role of an individual in the network (for example the “centrality” or “power” of a performer), and identify cliques (groups of connected individuals). Section 3




Figure 1. The sociogram based on the event log shown in Table I.

will discuss this aspect in more detail and Section 4 will provide concrete metrics to derive sociograms from event logs.

2.2. OTHER TYPES OF MINING

Table I contains the *minimal information* we assume to be present in an event log. Using the information one can also discover other models (i.e., not just sociograms). For example, we have developed techniques and tools to discover the process model. Figure 2 shows the resulting Petri net model after applying our α -algorithm (Aalst et al., 2004) to Table I. The model shows that the process always starts with *A* and ends with *D*. In between these tasks either *B* and *C* are executed or *E* alone. *B* and *C* are concurrent, i.e., they can be executed in any order. Given the focus of this paper, we will not elaborate further on process discovery. See Section 7 for pointers to related work.

In many applications, the event log contains a *time stamp* for each event and this information can be used to extract additional causality information. In addition, a typical log also contains information about the *type of event*, e.g., a start event (a person selecting an activity from a worklist), a complete event (the completion of a activity), a withdraw event (a scheduled activity is removed), etc. Moreover, we are also interested in the relation between attributes of the case and the actual route taken by a particular case or allocation of work to workers. For example, when handling traffic violations: Is the make of a car relevant for the routing of the corresponding traffic violation? (e.g., People driving a Ferrari always pay their fines in time.) Another example directly related to SNA would be to see whether the sociograms for different types of cases (e.g., private and corporate customers) differ.

The presence of timing information and information on cases/activities allows for more advanced forms of process mining, e.g., methods trying to explain the performance indicators like flow times in term of the attributes/performers of cases. Another interesting application of process mining is




Figure 2. A process model based on the event log shown in Table I discovered by the α -algorithm.

fraud detection, i.e., detecting suspicious patterns that may indicate security violations (cf. four eyes principle (Aalst and Hee, 2002)).

2.3. COMPLETENESS AND NOISE

For this simple example treated so far (i.e., Table I), it is quite easy to generate the process model shown in Figure 2 or the sociogram shown in Figure 1. For more realistic situations there are however a number of complicating factors:

- *Completeness.* For larger workflow models and models exhibiting alternative and parallel routing, the workflow log will typically not contain all possible routes. Consider 10 activities which can be executed in parallel. The total number of interleavings is $10! = 3628800$. It is not realistic that each interleaving is present in a log. Moreover, certain paths through the process model may have a low probability and therefore remain undetected. Similar remarks hold for the organizational model and social network. For example, a person has a role but just by coincidence did not execute some or all activities corresponding to that role. Another example is that two individuals work together frequently but during the data collection period one of them was on a sabbatical leave. As a result the log is not complete in the sense that it captures possible and/or typical behavior.
- *Noise.* Parts of the log may be incorrect, incomplete, or refer to exceptions. Events can be logged incorrectly because of human or technical errors. Events can be missing in the log if some of the activities are manual or handled by another system/organizational unit. Events can also refer to rare or undesired events. Consider for example the workflow in a hospital. If due to time pressure the order of two events (e.g., make X-ray and remove drain) is reversed, this does not imply that this would be part of the regular medical protocol and should be supported by the hospital's workflow system. Also two causally unrelated events (e.g., take blood sample and death of patient) may happen next to each other without implying a causal relation (i.e., taking a sample did not result in the death of the patient; it was sheer coincidence). Clearly, exceptions which are recorded only once should not automatically become part of the regular workflow.

2.4. LEGAL ISSUES AND ETHICAL ISSUES

To conclude this section, we point out legal issues relevant when mining event logs. Clearly, event logs can be used to systematically measure the performance of employees. The legislation with respect to issues such as privacy and protection of personal data differs from country to country. For example,

Dutch companies are bound by the Personal Data Protection Act (Wet Bescherming Persoonsgegevens) which is based on a directive from the European Union. The practical implications of this for the Dutch situation are described in (CBP n.d.; Hulsman and Ippel, 1994; Sauerwein and Linnemann, 2001). Event logs are not restricted by these laws as long as the information in the log cannot be traced back to individuals. If information in the log can be traced back to a specific employee, it is important that the employee is aware of the fact that her/his activities are logged and the fact that this logging is used to monitor her/his performance. Note that in a log we can deliberately abstract from information about the workers executing activities and still mine the process, organizational, and social structures (simply hide identities).³ Therefore, it is possible to avoid collecting information on the productivity of individual workers and legislation such as the Personal Data Protection Act does not apply. Nevertheless, the logs of most workflow systems contain information about individual workers, and therefore, this issue should be considered carefully. Moreover, to use social network analysis as an operational tool to improve work processes, employees should approve and it is vital not to misuse the information gathered.

3. Mining organizational relations

In the previous section, we provided an overview of process mining. In this section, we focus on the main topic of this paper: mining organizational relations as described in Section 2.1. The goal is to generate a sociogram that can be used as input for standard software in the SNA (Social Network Analysis) domain.

The motivation for doing this is twofold. First of all, existing systems record information about human activity. This information can be structured in the form of a sociogram. Second, there is a wide variety of mature techniques and tools to analyze such sociograms. Therefore, it is both interesting and feasible to use this as a starting point for investigating the social context of work processes. A better understanding of this social context may reveal a mis-alignment between the information system and its users and may provide insights that can be used to increase the efficiency and effectiveness of processes and organizations.

In this section, we first introduce the fundamentals of SNA and then focus on the question how to derive sociograms from event logs.

3.1. SOCIAL NETWORK ANALYSIS

Applications of SNA range from the analysis of small social networks to large networks. For example, the tool InFlow (<http://www.orgnet.com>) has been used to analyze terrorist network surrounding the September 11th 2001

events. However, such tools could also be used to analyze the social network in a classroom. In literature, researchers distinguish between *sociocentric* (whole) and *egocentric* (personal) approaches. Sociocentric approaches consider interactions within a defined group and consider the group as a whole. Egocentric approaches consider the network of an individual, e.g., relations among the friends of a given person. From a mathematical point of view both approaches are quite similar. In both cases the starting point for analysis is a graph where nodes represent people and the arcs/edges represent relations. Although this information can also be represented as a matrix, we use the graph notation. The graph can be undirected or directed, e.g., A may like B but not vice versa. Moreover, the relations may be binary (they are there or not) or weighted (e.g., “+” or “-”, or a real number). The weight is used to qualify the relation. The resulting graph is named a *sociogram* (Scott, 1992; Wasserman and Faust, 1994). Note that Figure 1 shows an example of sociogram with directed links which are not weighted.

In a mathematical sense a sociogram is a graph (P, R) where P is the set of individuals (in the context of process mining referred to as performers) and $R \subseteq P \times P$ (Scott, 1992; Wasserman and Faust, 1994). If the graph is undirected, R is symmetric. If the graph is weighted, there is an additional function W assigning a value to all elements of R . When looking at the graph as a *whole* there are notions like *density*, i.e., the number of elements in R divided by the maximal number of elements. For example the density of the graph shown in Figure 1 is $8/(6 * 6) = 0.22$. Another metric based on weighted graphs is the maximal geodesic distance in a graph. The geodesic distance of two nodes is the distance of the shortest path in the graph based on R and W .

When looking at one specific individual (i.e., a node in the graph), many notions can be defined (Scott, 1992; Wasserman and Faust, 1994). If all other individuals are in short distance to a given node and all geodesic paths (i.e., shortest path in the graph) visit this node, clearly the node is very central (like a spider in the web). There are different metrics for this intuitive notion of *centrality*. The Bavelas–Leavitt index of centrality is a well-known example that is based on the geodesic paths in the graph (Bavelas, 1948). Let i be an individual (i.e., $i \in P$) and $D_{j,k}$ the geodesic distance from an individual j to an individual k . The Bavelas–Leavitt index of centrality is defined as $BL(i) = (\sum_{j,k} D_{j,k}) / (\sum_{j,k} D_{j,i} + D_{i,k})$. Note that the index divides the sum of all geodesic distances by the sum of all geodesic distances from and to a given resource. Other related metrics are *closeness* (1 divided by the sum of all geodesic distances to a given resource) and *betweenness* (a ratio based on the number of geodesic paths visiting a given node) (Freeman, 1977; Freeman, 1979; Burt and Minor, 1983; Scott, 1992; Wasserman and Faust, 1994). Other notions include the *emission* of a resource (i.e., $\sum_j W_{i,j}$), the *reception* of a resource (i.e., $\sum_j W_{j,i}$), and the *determination* degree (i.e., $\sum_j W_{j,i} - W_{i,j}$) (Burt and Minor, 1983; Scott, 1992; Wasserman and Faust, 1994). Another

interesting metric is the *sociometric status* which is determined by the sum of input and output relations, i.e., $\sum_j D_{j,i} + D_{i,j}$. All metrics can be normalized by taking the size of the social network into account (e.g., divide by the number of resources). Using these metrics and a visual representation of the network one can analyze various aspects of the social structure of an organization. For example, one can search for densely connected clusters of resources and structural holes (i.e., areas with few connections), cf. (Burt and Minor, 1983; Scott, 1992; Wasserman and Faust, 1994).

Let us apply some of these notions to the sociogram shown in Figure 1 where the arcs indicate (unweighted) frequencies. The sociometric status of Clare is 3 (if we include self-links, otherwise 1), the sociometric status of Pete is 2, the emission of John is 2, the emission of Pete is 0, the reception of Pete is 2, the reception of Sue is 1, the determination degree of Mike is 0, etc. The Bavelas–Leavitt index of centrality of John is 4.33 while the same index for Sue is 3.25. The numbers are unweighted and in most cases these are made relative to allow for easy comparison. Tools like AGNA, Egonet, InFlow, CliqueFinder, MetaSight, NetForm, NetMiner, NetVis, StOCNET, UCINET, and Visone are just some of the many SNA tools available. For more information on SNA we refer to (Burt and Minor, 1983; Bernard et al., 1990; Scott, 1992; Wasserman and Faust, 1994).

3.2. DERIVING RELATIONS FROM EVENT LOGS

After showing the potential of SNA and the availability of techniques and tools, the main question is: *How to derive meaningful sociograms from event logs?* To address this question we identify different *metrics*. Each metric assigns a weight $W_{i,j}$ to the relationship between individuals i and j . If $W_{i,j}$ is above a certain threshold τ , it will be included in R (i.e., $(i, j) \in R$ if and only if $W_{i,j} > \tau$ for any $i, j \in P$). This way we get a weighted graph (P, R, W) that can be used by tools such as AGNA and NetMiner. In this paper, we will focus on four types of metrics that can be derived from event logs: (1) metrics based on (possible) causality, (2) metrics based on joint cases, (3) metrics based on joint activities, and (4) metrics based on special event types.

Metrics based on (possible) causality monitor for individual cases how work moves among performers. One of the examples of such a metric is *handover of work*. Within a case (i.e., process instance) there is a handover of work from individual i to individual j if there are two subsequent activities where the first is completed by i and the second by j . This notion can be refined in various ways. For example, knowledge of the process structure can be used to detect whether there is really a causal dependency between both activities. It is also possible to not only consider direct succession but also indirect succession using a “causality fall factor” β , i.e., if there are three activities in-between an activity completed by i and an activity completed by j , the causality fall factor

is β^3 . A related metric is *subcontracting* where the main idea is to count the number of times individual j executed an activity in-between two activities executed by individual i . This may indicate that work was subcontracted from i to j . Again all kinds of refinements are possible.

Metrics based on joint cases ignore causal dependencies but simply count how frequently two individuals are performing activities for the same case. If individuals work together on cases, they will have a stronger relation than individuals rarely working together.

Metrics based on joint activities do not consider how individuals work together on shared cases but focus on the activities they perform. The assumption here is that people doing similar things have stronger relations than people doing completely different things. Each individual has a “profile” based on how frequent they conduct specific activities. There are many ways to measure the “distance” between two profiles thus enabling many metrics.

Metrics based on special event types consider the type of event. Thus far we assumed that events correspond to the execution of activities. However, there are also events like reassigning an activity from one individual to another. For example, if i frequently delegates work to j but not vice versa it is likely that i is in a hierarchical relation with j . From an SNA point of view these observations are particularly interesting since they represent explicit power relations.

The sociogram shown in Figure 1 is based on the causality metric hand-over of work. In the next section, we will define the metrics in more detail.

4. Metrics

In this section, we define the metrics we have developed to establish relationships between individuals from event logs. We address all four types introduced in Section 3.2. Recall that each metric will assign a weight $W_{i,j}$ to the relationship between individuals i and j . Before we define the various metrics in detail, we introduce a convenient notation for event logs.

Definition 4.1 (Event log). Let A be a set of activities (i.e., atomic workflow/process objects, also referred to as tasks) and P a set of performers (i.e., resources, individuals, or workers). $E = A \times P$ is the set of (possible) events, i.e., combinations of an activity and a performer (e.g., (a, p) denotes the execution of activity a by performer p). $C = E^*$ is the set of possible event sequences (traces describing a case). $L \in \mathcal{B}(C)$ is an *event log*. Note that $\mathcal{B}(C)$ is the set of all bags (multi-sets) over C .

Note that this definition of an event slightly differs from the informal notions used before. First of all, we abstract from additional information such as time stamps, data, etc. Secondly, we do not consider the ordering of events corresponding to different cases. For convenience, we define two operations on events: $\pi_a(e) = a$ and $\pi_p(e) = p$ for some event $e = (a, p)$.

4.1. METRICS BASED ON (POSSIBLE) CAUSALITY

Metrics based on causality take into account both handover of work and subcontracting. The basic idea is that performers are related if there is a causal relation through the passing of a case from one performer to another. For both situations, three kinds of refinements are applied. First of all, one can differentiate with respect to the degree of causality, e.g., the length of handover. It means that we can consider not only direct succession but also indirect succession. Second, we can ignore multiple transfers within one instance or not. Third, we can consider arbitrary transfers of work or only consider those where there is a causal dependency (for the latter we need to know or be able to derive the process model). Based on these refinements, we derive $2^3 = 8$ variants for both the handover of work and subcontracting metrics. These variants are all based on the same event log. Before defining the metrics, some of the basic notions that can be applied to a single case $c = (c_0, c_1, \dots)$ are specified.

Definition 4.2. (\triangleright , \triangleright_c) Let L be a log. Assume that \longrightarrow denotes some causality relation derived from the process model. For $a_1, a_2 \in A$, $p_1, p_2 \in P$, $c = (c_0, c_1, \dots) \in L$, and $n \in \mathbb{N}$:

$$\begin{aligned}
- p_1 \triangleright_c^n p_2 &= \exists_{0 \leq i < |c| - n} \pi_p(c_i) = p_1 \wedge \pi_p(c_{i+n}) = p_2 \\
- |p_1 \triangleright_c^n p_2| &= \sum_{0 \leq i < |c| - n} \begin{cases} 1 & \text{if } \pi_p(c_i) = p_1 \wedge \pi_p(c_{i+n}) = p_2 \\ 0 & \text{otherwise} \end{cases} \\
- p_1 \triangleright_c^n p_2 &= \exists_{0 \leq i < |c| - n} \pi_p(c_i) = p_1 \wedge \pi_p(c_{i+n}) = p_2 \wedge \pi_a(c_i) \rightarrow \pi_a(c_{i+n}) \\
- |p_1 \triangleright_c^n p_2| &= \sum_{0 \leq i < |c| - n} \begin{cases} 1 & \text{if } \pi_p(c_i) = p_1 \wedge \pi_p(c_{i+n}) = p_2 \wedge \pi_a(c_i) \rightarrow \pi_a(c_{i+n}) \\ 0 & \text{otherwise} \end{cases}
\end{aligned}$$

$p_1 \triangleright_c^n p_2$ denotes the function which returns *true* if within the context of case c performers p_1 and p_2 both executed some activity such that the distance between these two activities is n . For example, for case 1 shown in Table I, $John \triangleright_c^1 Mike$ equals 1 (i.e., true) and $John \triangleright_c^3 Pete$ equals 1 (i.e., true). In this definition, if the value of n equals 1, it refers to direct succession. If n is greater than 1, it refers to indirect succession. However, it ignores both multiple transfers within one instance and causal dependencies. $|p_1 \triangleright_c^n p_2|$ denotes the function which returns the number of times $p_1 \triangleright_c^n p_2$ in the case c . In other words, it considers multiple transfers within one instance. $p_1 \triangleright_c^n p_2$ and $|p_1 \triangleright_c^n p_2|$ are similar to $p_1 \triangleright_c^n p_2$ and $|p_1 \triangleright_c^n p_2|$ but in addition they take into account whether there is a real causal dependency. For example, consider case 1 shown in Table I. The order of events is: A (John), B (Mike), C (John), and D (Pete). If we calculate the relationships among activity B, C, and D, $Mike \triangleright_c^1 John$ equals 1 and $Mike \triangleright_c^1 Pete$ equals 0. However, $Mike \triangleright_c^1 John$ equals 0, i.e., although an activity conducted by Mike is followed by an

activity conducted by John there is not a causal dependency between B and C because both activities are in parallel. However, there is causal dependency between activity B and D (see Figure 2) and, therefore, $Mike \succeq_c^2 Pete$ equals 1. The information on causality can be added if the process model is known. If necessary, this information can also be derived from the log by using for example the α -algorithm (Aalst et al., 2004).

Using such relations, we define handover of work metrics. Based on three kinds of refinements mentioned before, eight variants are derived as follows.

Definition 4.3 (Handover of work metrics). Let L be a log. For $p_1, p_2 \in P$ and some β ($0 < \beta < 1$):

$$\begin{aligned}
- p_1 \triangleright_L p_2 &= \left(\sum_{c \in L} |p_1 \triangleright_c^1 p_2| \right) / \left(\sum_{c \in L} |c| - 1 \right) \\
- p_1 \dot{\triangleright}_L p_2 &= \left(\sum_{c \in L \wedge p_1 \triangleright_c^1 p_2} 1 \right) / |L| \\
- p_1 \triangleright_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n < |c|} \beta^{n-1} |p_1 \triangleright_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{1 \leq n < |c|} \beta^{n-1} (|c| - n) \right) \\
- p_1 \dot{\triangleright}_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n < |c| \wedge p_1 \triangleright_c^n p_2} \beta^{n-1} \right) / \left(\sum_{c \in L} \sum_{1 \leq n < |c|} \beta^{n-1} \right) \\
- p_1 \triangleright_L p_2 &= \left(\sum_{c \in L} |p_1 \succeq_c^1 p_2| \right) / \left(\sum_{c \in L} |c| - 1 \right) \\
- p_1 \dot{\triangleright}_L p_2 &= \left(\sum_{c \in L \wedge p_1 \succeq_c^1 p_2} 1 \right) / |L| \\
- p_1 \triangleright_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n < |c|} \beta^{n-1} |p_1 \succeq_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{1 \leq n < |c|} \beta^{n-1} (|c| - n) \right) \\
- p_1 \dot{\triangleright}_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n < |c| \wedge p_1 \succeq_c^n p_2} \beta^{n-1} \right) / \left(\sum_{c \in L} \sum_{1 \leq n < |c|} \beta^{n-1} \right)
\end{aligned}$$

$p_1 \triangleright_L p_2$ means dividing the total number of direct successions from p_1 to p_2 in a process log by the maximum number of possible direct successions in the log. $p_1 \dot{\triangleright}_L p_2$ ignores multiple transfers within one instance (i.e., case). For example, in Table I, $John \triangleright_L Mike$ equals 2/14 and $John \dot{\triangleright}_L Mike$ equals 2/5. Note that metric \triangleright_L defines a weight function W , i.e., $p_1 \triangleright_L p_2 = W_{p_1, p_2}$ is the weight of the link from p_1 to p_2 in the corresponding sociogram. As indicated before, a threshold may be used to remove links from the sociogram.

$p_1 \triangleright_L^\beta p_2$ and $p_1 \dot{\triangleright}_L^\beta p_2$ deal with indirect succession by introducing a ‘‘causality fall factor’’ β in this notation. If within the context of a case there are n

Table II. Handover of work metrics according to the causality fall factor β

Beta	$John \triangleright_L^\beta Mike$	$John \triangleright_L^\beta John$	$John \triangleright_L^\beta pete$
0.1	0.3116 (1/3.21)	0.0312 (0.1/3.21)	0.0031 (0.01/3.21)
0.5	0.2352 (1/4.25)	0.1176 (0.5/4.25)	0.0588 (0.25/4.25)
0.9	0.1783 (1/5.61)	0.1604 (0.9/5.61)	0.1444 (0.81/5.61)

events in-between two performers, the causality fall factor is β^n . $p_1 \triangleright_L^\beta p_2$ considers all possible successions, while $p_1 \dot{\triangleright}_L^\beta p_2$ ignores multiple transfers within one case. For example, in Table II, if β equals 0.5, then $John \triangleright_L Pete$ equals 2.5/19.5 and $John \dot{\triangleright}_L Pete$ equals 2.5/8.5. If we use a β close to 1, the effect of the distance between performers decreased. For example, suppose that only case 1 exists in Table I, we calculate the handover of metrics from John in Activity A to Mike, John in Activity B , and Pete, according to various values of β . Table II shows the results. If the value β increases in value, the variance of resulting values decreases.

The remaining four metrics $p_1 \triangleright_L p_2$, $p_1 \dot{\triangleright}_L p_2$, $p_1 \underline{\triangleright}_L p_2$, and $p_1 \underline{\dot{\triangleright}}_L p_2$ are similar to the previous four kinds of metrics, but take into account real causal dependencies. For example, $p_1 \underline{\triangleright}_L p_2$ means that the total number of direct successions from p_1 to p_2 in a log is divided by the maximum number of possible direct successions in the log when p_1 and p_2 are causally related. Note that each of the eighth metrics defines a different weight function W and implicitly another sociogram.

From above definitions, we derive general formulations of the metrics. The eight metrics mentioned can be merged into the following four metrics.

Definition 4.4 (General forms of handover of work metrics). Let L be a log. For $p_1, p_2 \in P$, some β ($0 < \beta \leq 1$) and $k \in \mathbb{N}$.

$$\begin{aligned}
-p_1 \triangleright_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k)} \beta^{n-1} |p_1 \triangleright_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k)} \beta^{n-1} (|c|-n) \right) \\
-p_1 \dot{\triangleright}_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k) \wedge p_1 \triangleright_c^n p_2} \beta^{n-1} \right) / \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k)} \beta^{n-1} \right) \\
-p_1 \underline{\triangleright}_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k)} \beta^{n-1} |p_1 \underline{\triangleright}_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k)} \beta^{n-1} (|c|-n) \right) \\
-p_1 \underline{\dot{\triangleright}}_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k) \wedge p_1 \underline{\triangleright}_c^n p_2} \beta^{n-1} \right) / \left(\sum_{c \in L} \sum_{1 \leq n \leq \min(|c|-1,k)} \beta^{n-1} \right)
\end{aligned}$$

In these alternative formulations, we introduce a “calculation depth factor” k . When we calculate metrics, k specifies maximum degree of causality. For example, if k equals 3, it considers the case of direct succession, one event in between two performers, and two events in-between two performers. Note that if $\beta=1$, $k=1$, then $p_1 \triangleright_L^{1,1} p_2 = p_1 \triangleright_L p_2$, and if $k > \max(|c|)$, then $p_1 \triangleright_L^{\beta,k} p_2 = p_1 \triangleright_L^\beta p_2$. This rule is also applied to the other three metrics. Further, when we calculate the metrics, a suitable value for k is important for the efficiency of calculation. Logs are typically very large. Therefore considering all possible successions may be inefficient.

After defining metrics for handover of work we now consider another class of metrics based on (possible) causality: *subcontracting metrics*. In the case of subcontracting, the three refinements mentioned before can also be applied. However the concept of direct and indirect succession is changed. Direct succession means there is only one activity in-between two activities executed by one performer. While indirect succession means, there are multiple activities in-between two activities executed by one performer. We also introduce causality fall factor β for indirect succession. For example, assume that there are four activities. Both the first and the fourth activity are executed by a performer i , while the second and third activity are executed by performer j and k respectively. In this situation, we can derive two relations which are from a performer i to a performer j and from a performer i to a performer k . Again we use a causality fall factor β . The second and third refinements are the same as for handover of work. Before defining metrics, the basic notions applied to a single case $c=(c_0, c_1, \dots)$ are specified.

Definition 4.5. (\diamond , \diamondsuit) Let L be a log. Assume that \rightarrow denotes some causality relation. In the context of L and \rightarrow , we define a number of relations. For $a_1, a_2 \in A, p_1, p_2 \in P, c = (c_0, c_1, \dots) \in L, |c| > 2, n \in \mathbb{N}$, and $n > 1$:

$$-p_1 \diamond_c^n p_2 = \exists_{0 \leq i < j < i+n < |c|} \pi_p(c_i) = p_1 \wedge \pi_p(c_j) = p_2 \wedge \pi_p(c_{i+n}) = p_1$$

$$-|p_1 \diamond_c^n p_2| = \sum_{0 \leq i < |c| - n} \sum_{i < j < i+n} \begin{cases} 1 & \text{if } \pi_p(c_i) = p_1 \wedge \pi_p(c_j) = p_2 \wedge \pi_p(c_{i+n}) = p_1 \\ 0 & \text{otherwise} \end{cases}$$

$$-p_1 \diamondsuit_c^n p_2 = \exists_{0 < j < i+n < |c|} \pi_p(c_i) = p_1 \wedge \pi_p(c_j) = p_2 \wedge \pi_p(c_{i+n}) = p_1 \wedge \pi_a(c_i) \rightarrow \pi_a(c_j) \rightarrow \pi_a(c_{i+n})$$

$$-|p_1 \diamondsuit_c^n p_2| = \sum_{0 \leq i < |c| - n} \sum_{i < j < i+n} \begin{cases} 1 & \text{if } \pi_p(c_i) = p_1 \wedge \pi_p(c_j) = p_2 \wedge \pi_p(c_{i+n}) = p_1 \\ & \wedge \pi_a(c_i) \rightarrow \pi_a(c_j) \rightarrow \pi_a(c_{i+n}) \\ 0 & \text{otherwise} \end{cases}$$

$p_1 \diamond_c^n p_2$ denotes the function which returns *true* if performer p_2 executed an activity in-between two activities executed by performer p_1 and distance between these two activities executed by performer p_1 is n . For example, for case 1 shown in Table I, $John \diamondsuit_c^2 Mike$ equals 1. However, it ignores both

multiple transfers within one instance and causal dependencies. $|p_1 \diamond_c^n p_2|$ denotes the function which returns the number of times $p_1 \diamond_c^n p_2$ in the case c . In other words, it considers multiple transfers within one instance. $p_1 \diamond_c^n p_2$ and $|p_1 \diamond_c^n p_2|$ are similar to $p_1 \diamond_c^n p_2$ and $|p_1 \diamond_c^n p_2|$ but in addition they take into account whether there is a real causal dependency. For example, consider case 1 shown in Table I. $John \diamond_c^n Mike$ equals 0, because activity B and C do not have a causal dependency.

Using such relations, we define subcontracting metrics. Again eight variants are identified.

Definition 4.6 (In-between metrics). Let L be a log. For $p_1, p_2 \in P$, $c = (c_0, c_1, \dots) \in L$, $|c| > 2$, and some β ($0 < \beta < 1$):

$$\begin{aligned}
 - p_1 \diamond_L p_2 &= \left(\sum_{c \in L} |p_1 \diamond_c^2 p_2| \right) \left(\sum_{c \in L} (|c| - 2) \right) \\
 - p_1 \hat{\diamond}_L p_2 &= \left(\sum_{c \in L \wedge p_1 \diamond_c^2 p_2} 1 \right) / |L| \\
 - p_1 \diamond_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n < |c|} \beta^{n-2} |p_1 \diamond_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{2 \leq n < |c|} \beta^{n-2} (|c| - n)(n - 1) \right) \\
 - p_1 \hat{\diamond}_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n < |c| \wedge p_1 \diamond_c^n p_2} \beta^{n-2} \right) / \left(\sum_{c \in L} \sum_{2 \leq n < |c|} \beta^{n-2} \right) \\
 - p_1 \underline{\diamond}_L p_2 &= \left(\sum_{c \in L} |p_1 \diamond_c^2 p_2| \right) / \left(\sum_{c \in L} (|c| - 2) \right) \\
 - p_1 \hat{\underline{\diamond}}_L p_2 &= \left(\sum_{c \in L \wedge p_1 \diamond_c^2 p_2} 1 \right) / |L| \\
 - p_1 \underline{\diamond}_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n < |c|} \beta^{n-2} |p_1 \diamond_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{2 \leq n < |c|} \beta^{n-2} (|c| - n)(n - 1) \right) \\
 - p_1 \hat{\underline{\diamond}}_L^\beta p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n < |c| \wedge p_1 \diamond_c^n p_2} \beta^{n-2} \right) / \left(\sum_{c \in L} \sum_{2 \leq n < |c|} \beta^{n-2} \right)
 \end{aligned}$$

$p_1 \diamond_L p_2$ means dividing the total number of direct subcontracting occurrences between p_1 and p_2 in a process log by the maximum number of possible direct subcontracting occurrences in the log. $p_1 \hat{\diamond}_L p_2$ ignores multiple subcontracting occurrences within one instance (i.e., case). For example, in Table I, $John \diamond_L Mike$ equals 2/9 and $John \hat{\diamond}_L Mike$ equals 2/5. $p_1 \diamond_L^\beta p_2$ and $p_1 \hat{\diamond}_L^\beta p_2$ deal with the situation where the distance between these two activities executed by performer p_1 is greater than 2. Again we introduce a ‘‘causality fall factor’’ β in a fashion similar to the handover of work metrics. If within the context of a case there are n events in-between two activities executed by the same performer, the causality fall factor is β^n . $p_1 \diamond_L^\beta p_2$ considers all possible subcontracting occurrences, while $p_1 \hat{\diamond}_L^\beta p_2$ ignores multiple subcontracting within one case. For example, in Table II, if β equals 0.5, then $John \diamond_L Mike$ equals 2/13 and $John \hat{\diamond}_L Mike$ equals 2/7. Again $p_1 \underline{\diamond}_L p_2, p_1 \hat{\underline{\diamond}}_L p_2, p_1 \underline{\diamond}_L^\beta p_2,$ and $p_1 \hat{\underline{\diamond}}_L^\beta p_2$ are similar but take into account real causal dependencies. For

example, $p_1 \diamond_L p_2$ means that the total number of direct subcontracting from p_1 to p_2 in a process log is divided by the maximum number of possible direct subcontracting in the log when p_1 and p_2 are causally related.

As before we can derive more general formulations for the metrics. The eight metrics mentioned above can be merged into four metrics as shown in the following definition.

Definition 4.7 (General forms of in-between metrics). Let L be a log. For $p_1, p_2 \in P$, some β ($0 < \beta \leq 1$) and $k \in \mathbb{N}$ ($k > 1$)

$$\begin{aligned} -p_1 \diamond_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k)} \beta^{n-2} |p_1 \diamond_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k)} \beta^{n-2} (|c|-n)(n-1) \right) \\ -p_1 \dot{\diamond}_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k) \wedge p_1 \diamond_c^n p_2} \beta^{n-2} \right) / \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k)} \beta^{n-2} \right) \\ -p_1 \underline{\diamond}_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k)} \beta^{n-2} |p_1 \underline{\diamond}_c^n p_2| \right) / \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k)} \beta^{n-2} (|c|-n)(n-1) \right) \\ -p_1 \dot{\underline{\diamond}}_L^{\beta,k} p_2 &= \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k) \wedge p_1 \underline{\diamond}_c^n p_2} \beta^{n-2} \right) / \left(\sum_{c \in L} \sum_{2 \leq n \leq \min(|c|-1, k)} \beta^{n-2} \right) \end{aligned}$$

Again we also introduce a ‘‘calculation depth factor’’ k . When calculating the metrics, k specifies maximum distance between two activities executed by one performer. For example, if k equals 3, it considers the case of one activity in between two activities executed by one performer and two activities in between two activities executed by one performer. Note that if $\beta = 1$, $k = 2$, then $p_1 \diamond_L^{1,2} p_2 = p_1 \diamond_L p_2$, and if $k > \max(|c|)$, then $p_1 \diamond_L^{\beta,k} p_2 = p_1 \diamond_L^\beta p_2$.

4.2. METRICS BASED ON JOINT CASES

For this type of metrics we ignore causal dependencies and simply count how often two individuals are performing activities for the same case.

Definition 4.8 (Working together metrics). Let L be a log. For $p_1, p_2 \in P$: $p_1 \bowtie_L p_2 = \sum_{c \in L} p_1 \bowtie_c p_2 / \sum_{c \in L} g(c, p_1)$ if $\sum_{c \in L} g(c, p_1) \neq 0$, otherwise $p_1 \bowtie_L p_2 = 0$, where for $c = (c_0, c_1, \dots) \in L$: $p_1 \bowtie_c p_2 = 1$ if $\exists_{0 \leq i, j < |c| \wedge i \neq j} \pi_p(c_i) = p_1 \wedge \pi_p(c_j) = p_2$, otherwise $p_1 \bowtie_c p_2 = 0$: $g(c, p_1) = 1$ if $\exists_{0 \leq i < |c|} \pi_p(c_i) = p_1$, otherwise $g(c, p_1) = 0$

Note that in this definition we divide the number of joint cases by the number of cases in which p_1 appeared. It is important to use a relative notation. For example, suppose that p_1 participates in three cases, p_2 participates in six cases, and they work together three times. In this situation, p_1 always work together with p_2 , but p_2 does not. Thus, the value for $p_1 \bowtie_L p_2$

has to be larger than the value for $p_2 \bowtie_L p_1$. Let us apply this metric to analyze the relationship between John and Pete based in the log shown in Table I. In the log, John appeared in two cases, Pete in four cases, and they work together on two cases. Thus, $John \bowtie_L Pete = 2/2$ and $Pete \bowtie_L John = 2/4$.

Moreover, alternative metrics can be composed by taking the distance between activities into account, e.g., use variants like $(p_1 \triangleright_L^\beta p_2 + p_2 \triangleright_L^\beta p_1)/2$ or $(p_1 \triangleright_L^\beta p_2 + p_2 \triangleright_L^\beta p_1)/2$.

4.3. METRICS BASED ON JOINT ACTIVITIES

To calculate the metrics based on joint activities, first we make a “profile” based on how frequent individuals conduct specific activities. In this paper, we use a *performer by activity matrix* to represent these profiles. This matrix simply records how frequent each performer executes specific activities.

Definition 4.9. (Δ) Let L be a log. For $p_1 \in P$, $a_1 \in A$, and $c = (c_0, c_1, \dots) \in L$:

$$-p_1 \Delta_c a_1 = \sum_{0 \leq i < |c|} \begin{cases} 1 & \text{if } \pi_a(c_i) = a_1 \wedge \pi_p(c_i) = p_1 \\ 0 & \text{otherwise} \end{cases}$$

$$-p_1 \Delta_L a_1 = \sum_{c \in L} p_1 \Delta_c a_1$$

Note that Δ defines a matrix with rows P and columns A . Table III shows the performer by activity matrix derived from Table I.

After creating the matrix, we measure the distance between two performers by comparing the corresponding row vectors. A simple distance measure is *Minkowski distance* which can be seen as a generalization of the Euclidean distance. But the Minkowski distance only gives good results if performers execute comparable volumes of work. Therefore, we also use the *Hamming distance* which does not consider the absolute frequency but only whether it is 0 or not. Another metric is *Pearson’s correlation coefficient* which is frequently used to find the relationship among cases.

Table III. The performer by activity matrix

Performer	Activity A	Activity B	Activity C	Activity D	Activity E
John	2	1	1	0	0
Sue	3	1	1	0	0
Mike	0	1	1	0	0
Carol	0	1	1	0	0
Pete	0	0	0	4	0
Clare	0	0	0	1	1

Definition 4.10. $(\Delta_L^{MD,n}, \Delta_L^{HD}, \Delta_L^{PC})$ Let L be a log and Δ_L be a performer by activity matrix. For $p_1, p_2 \in P, n \in \{1,2,3,\dots\}$:

$$\begin{aligned}
 -p_1 \Delta_L^{MD,n} p_2 &= \left(\sum_{a \in A} |(p_1 \Delta_L a) - (p_2 \Delta_L a)|^n \right)^{1/n} \\
 -p_1 \Delta_L^{HD} p_2 &= \left(\sum_{a \in A} \delta(p_1 \Delta_L a, p_2 \Delta_L a) \right) / |A| \\
 \text{where } \delta(x, y) &= \begin{cases} 0 & \text{if } (x > 0 \wedge y > 0) \vee (x = y = 0) \\ 1 & \text{otherwise} \end{cases} \\
 -p_1 \Delta_L^{PC} p_2 &= \frac{\sum_{a \in A} ((p_1 \Delta_L a) - \bar{X})(p_2 \Delta_L a - \bar{Y})}{\sqrt{\sum_{a \in A} ((p_1 \Delta_L a) - \bar{X}) \sum_{a \in A} ((p_2 \Delta_L a) - \bar{Y})}} \\
 \text{where } \bar{X} &= \sum_{a \in A} (p_1 \Delta_L a) / |A|, \bar{Y} = \sum_{a \in A} (p_2 \Delta_L a) / |A|
 \end{aligned}$$

The Minkowski distance $\Delta_L^{MD,n}$ has a parameter n : $n = 1$ is the Rectilinear distance also referred to as Manhattan distance, $n = 2$ is the Euclidean distance, and for large values of n the metric approximates the Chebyshev distance. The Hamming distance Δ_L^{HD} does not have a parameter but could be extended with some threshold value. In the case of Pearson’s correlation coefficient, the result ranges from $+1$ to -1 . A correlation of $+1$ means that there is a perfect positive linear relationship between variables. A correlation of -1 means that there is a maximal negative linear relationship between variables. In other words, if the distance between performers is small, the correlation is closer to 1 , if it is large, the correlation is closer to -1 .

To illustrate the limitations of simple metrics like the Minkowski distance we consider Table III. Clearly, from an intuitive point of view the distance between Sue and Carol should be smaller than the distance between Carol and Clare because Carol and Clare have no activities in common. The Minkowski distance ($n = 1$) between Sue and Carol equals 3 and the distance between Carol and Clare equals 4. However, if Sue would have executed activity B and activity C also three times, the distance between Sue and Carol would be 7 and thus incorrectly suggest that Carol is closer to Clare than to Sue. The Hamming distance is more robust and would indicate in both cases that Carol is closer to Sue: *Sue* Δ_L^{HD} *Carol* equals $1/5$ and *Carol* Δ_L^{HD} *Clare* equals $4/5$. If we calculate the Pearson’s correlation coefficient, *Sue* Δ_L^{PC} *Carol* equals 0.2182 and *Carol* Δ_L^{PC} *Clare* equals -0.6667 . Hence, the result of Pearson’s correlation leads to the same conclusion as the Hamming distance.

Note that if the volume of work varies significantly, the metrics are not suitable. For example, it is difficult to compare the profile of a part-time worker with a full-time worker. Thus, in some cases we first apply the $\log_k (X + 1)$ function on the values of the performer by activity matrix, i.e., use a logarithmic scale for Δ_L . Note that we need to add “ $+ 1$ ” to avoid negative values.

4.4. METRICS BASED ON SPECIAL EVENT TYPES

The types of metrics mentioned in previous subsections do not consider event types. They more or less assume that all events correspond to the completion of an activity. But events can contain various event types such as *schedule*, *assign*, *withdraw*, *reassign*, *start*, *suspend*, *resume*, *pi_abort*, *ate_abort*, *complete*, *autoskip*, *manualskip*, and *unknown*.⁴ For example, *schedule* refers to the enabling of a task for a specific case, *assign* refers to the allocation of such an enabled task to a user, *start* refers to the actual start of a task, and *complete* refers to the completion of a task. Event types such as *withdraw*, *reassign*, *suspend*, *resume*, *pi_abort*, and *ate_abort* may refer to exceptions which are interesting from the viewpoint of SNA. See (Aalst et al., 2003) for some more information on the various event types.

In this subsection, we take into account metrics based on special event types. In particular, we concentrate on the *reassign* event type. To define metrics based on special event types, we suppose that log lines have an event type. For convenience, we define an operation on events: $\pi_{et}(e) = \text{event type}$ for some event $e = (a, p)$. Note that Definition 4.1 could be extended to capture event types such as used by commercial systems. In the next section we define an XML format to capture this information.

Before defining metrics, the basic notations used for a single case $c = (c_0, c_1, \dots)$ are specified as follows.

Definition 4.11 (*follow*, ∇). Let L be a log. For $p_1, p_2 \in P$, $c = (c_0, c_1, \dots) \in L$, and some event type *event type*:

$$\begin{aligned}
 - \text{follow}(c, i, j) &= \pi_a(c_i) = \pi_a(c_j) \wedge \forall_{i < k < j} \pi_a(c_k) \neq \pi_a(c_i), \text{ for } 0 \leq i < j < |c| \\
 - p_1 \nabla_c^{\text{event type}} p_2 &= \begin{cases} \exists_{0 \leq i < j < |c|} \text{follow}(c, i, j) \wedge \pi_p(c_i) = p_1 \wedge \\ \pi_{et}(c_i) = \text{event type} \wedge \pi_p(c_j) = p_2 \end{cases} \\
 - |p_1 \nabla_c^{\text{event type}} p_2| &= \sum_{0 \leq i < |c|} \begin{cases} 1 & \text{if } \exists_{i < j < |c|} \text{follow}(c, i, j) \wedge \pi_p(c_i) = p_1 \\ & \wedge \pi_{et}(c_i) = \text{event type} \wedge \pi_p(c_j) = p_2 \\ 0 & \text{otherwise} \end{cases}
 \end{aligned}$$

In a log, there may be several events that correspond to the same activity. If the activity a is reassigned from a performer p_1 to a performer p_2 , we can find two events c_i and c_j such that $c_i = (a, p_1)$, $\pi_{et}(c_i) = \text{'reassign'}$, $c_j = (a, p_2)$, and $\pi_{et}(c_j)$ is some event type. Thus, we need *follow* to find a next event which is related to c_i . $p_1 \nabla_c^{\text{event type}} p_2$ denotes the function which returns true if within the context of the case c performers p_1 and p_2 both executed the same activity and p_1 was responsible for a specific type of event and p_2 is the first performer of some event for the same activity. $|p_1 \nabla_c^{\text{event type}} p_2|$ denotes the function which returns the number of times $p_1 \nabla_c^{\text{event type}} p_2$ in the case c . Using such relations, we define reassignment metrics. Recall that *reassign* is a special event type corresponding to the delegation from one performer to another.

Definition 4.12 (Reassignment metrics). Let L be a log. For $p_1, p_2 \in P$:

$$-p_1 \nabla_L^{\text{reassign}} p_2 = (\sum_{c \in L} |p_1 \nabla_c^{\text{reassign}} p_2|) / (\sum_{c \in L} (|c| - 1))$$

$$-p_1 \bar{\nabla}_L^{\text{reassign}} p_2 = (\sum_{c \in L \wedge p_1 \nabla_c^{\text{reassign}} p_2} 1) / |L|$$

$p_1 \nabla_L^{\text{reassign}} p_2$ is obtained by dividing the total number of reassignments from p_1 to p_2 in the event log by the maximum number of reassignments in the log. For example, if there are 10 events in a log and *John* has reassigned an activity to *Mike* once, $p_1 \nabla_L^{\text{reassign}} p_2$ equals 1/9. $p_1 \bar{\nabla}_L^{\text{reassign}} p_2$ ignores multiple reassignment within one instance.

In this section, we formalized the metrics introduced in Section 3.2. It is important to note that each of the metrics is derived from some log L and the result can be represented in terms of a weighted graph (P, R, W) , where P is the set of performers, R is the set of relations, and W is a function indicating the weight of each relation (see Section 3.1). For example, the basic hand-over of work metric \triangleright_L defines $R = \{(p_1, p_2) \in P \times P \mid p_1 \triangleright_L p_2 \neq 0\}$ and $W(p_1, p_2) = p_1 \triangleright_L p_2$. For the Hamming distance $R = \{(p_1, p_2) \in P \times P \mid p_1 \Delta_L^{HD} p_2 \neq 1\}$ and $W(p_1, p_2) = 1 - (p_1 \Delta_L^{HD} p_2)$. For the Pearson's correlation coefficient $R = \{(p_1, p_2) \in P \times P \mid p_1 \Delta_L^{PC} p_2 \geq \alpha\}$ (where α is some threshold value between -1 and 1) and $W(p_1, p_2) = (1 + (p_1 \Delta_L^{PC} p_2)) / 2$. In other words, given an event log L each metric results in a sociogram that can be analyzed using existing SNA tools.

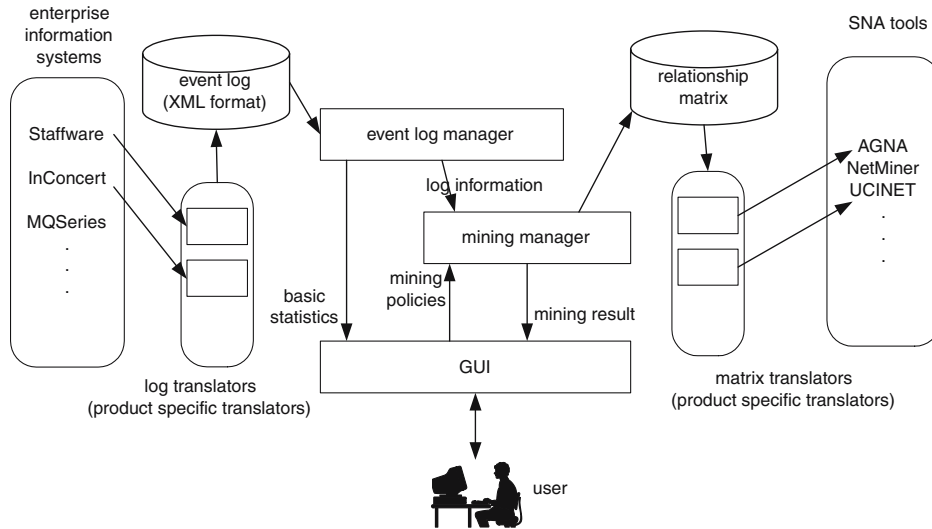


Figure 3. The architecture of MiSoN.

5. MiSoN

This section introduces our tool *MiSoN* (*Mining Social Networks*). MiSoN has been developed to discover relationships between individuals from a range of enterprise information systems including workflow management systems such as Staffware, InConcert, and MQSeries, ERP systems, and CRM systems. Based on the event logs extracted from these systems MiSoN constructs sociograms that can be used as a starting point for SNA. The derived relationships can be exported in a matrix format and used by most SNA tools. With such tools, we can apply several techniques to analyze social networks, e.g., find interaction patterns, evaluate the role of an individual in an organization, etc.

MiSoN has been developed using Java including XML-based libraries such as JAXB and JDOM, and provides an easy-to-use graphical user interface. Figure 3 shows the architecture of MiSoN. The mining starts from a tool-independent XML format which includes information about processes, cases, activities, event times, and performers. From enterprise information systems recording event logs, we can export to this XML format.

Figure 4 shows the XML schema describing this format. It is an extension of the DTD suggested in (Aalst et al., 2003). The schema has the *WorkflowLog* element as a root element. It has *Data*, *Source*, and *Process* elements. The *Source* element contains the information about software or system that was used to record the log (e.g., Staffware). The *Process* element represents the process where the process log belongs. Note that there may be multiple *Process* elements in a log. Each *Process* element may hold multiple *Process Instance* elements that correspond to cases. The *AuditTrailEntry* element represents a log line, i.e., a single event. It contains *WorkflowModelElement*, *EventType*, *Timestamp*, and *Originator* elements. For SNA, the *WorkflowModelElement*, *EventType*, and *Originator* elements are most important. The *WorkflowModelElement* refers to the activity (or subprocess) the event corresponds to. The *EventType* specifies the type of the event, e.g., *schedule* (i.e., a task becomes enabled for a specific instance), *assign* (i.e., a task instance is assigned to a user), *start* (the beginning of a task instance), *complete* (the

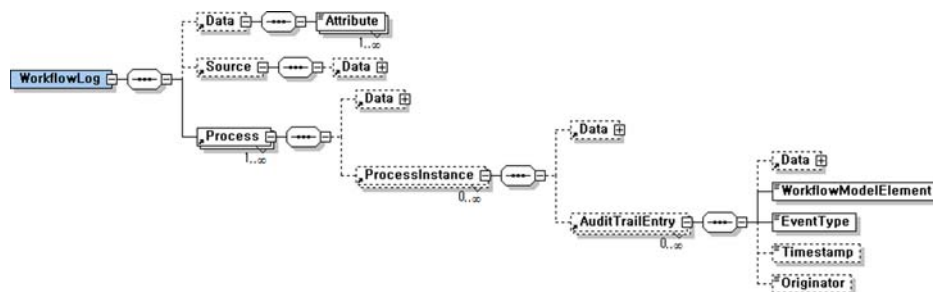


Figure 4. MiSoN workflow mining format (XML schema).

completion of a task instance), and *reassign* (as discussed in Section 4.4). In total, we identify 12 events. Last but not least the *Originator* element refers to the performer. To make the format more expressive, we define the *Data* element and other elements have it as a sub tags. If users want to specify more information than the basic elements, they can record the additional information using the *Data* element. Such information can be used for other types of process mining such as performance analysis, process knowledge extraction, etc.

After reading an event log that conforms to the XML schema, MiSoN provides functionalities for displaying user statistics and event log statistics. Using the metrics defined in Section 4, MiSoN constructs relationships between individuals. When calculating the relationships, the user can select suitable metrics and set relevant options. The result can be displayed using a matrix representation and a graph representation, but it can also be exported to SNA tools. Exported data contains the number of performers, names of performers, and a relationship matrix.

To illustrate MiSoN we have used an event log as generated with Staffware, which was converted to the XML format. For this log, we only consider the “released by” event type to make sociograms. This event corresponds to the *complete* event type in our XML format. We have tested MiSoN with several metrics mentioned in the previous section. Figure 5

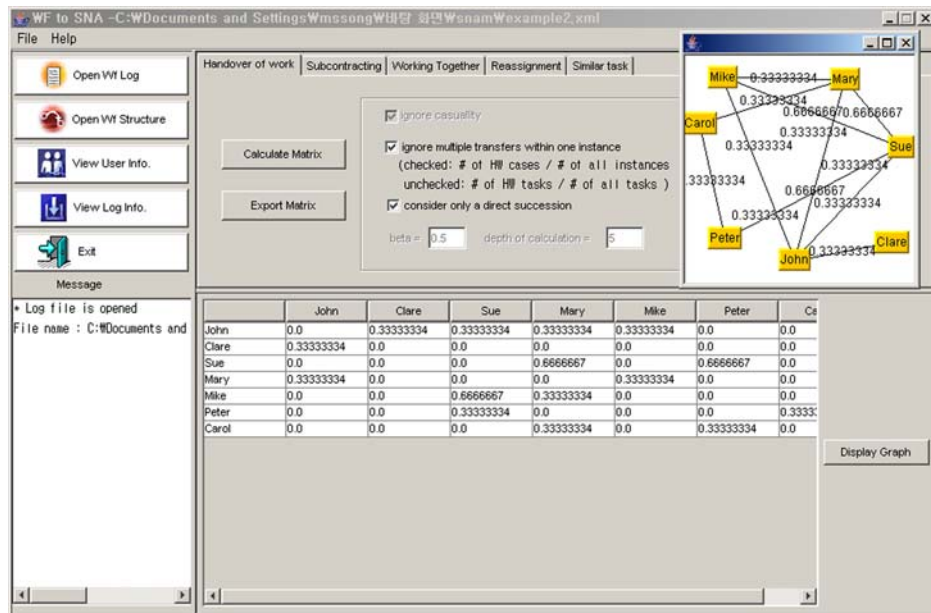


Figure 5. MiSoN screenshot showing a sociogram based on a Staffware log.

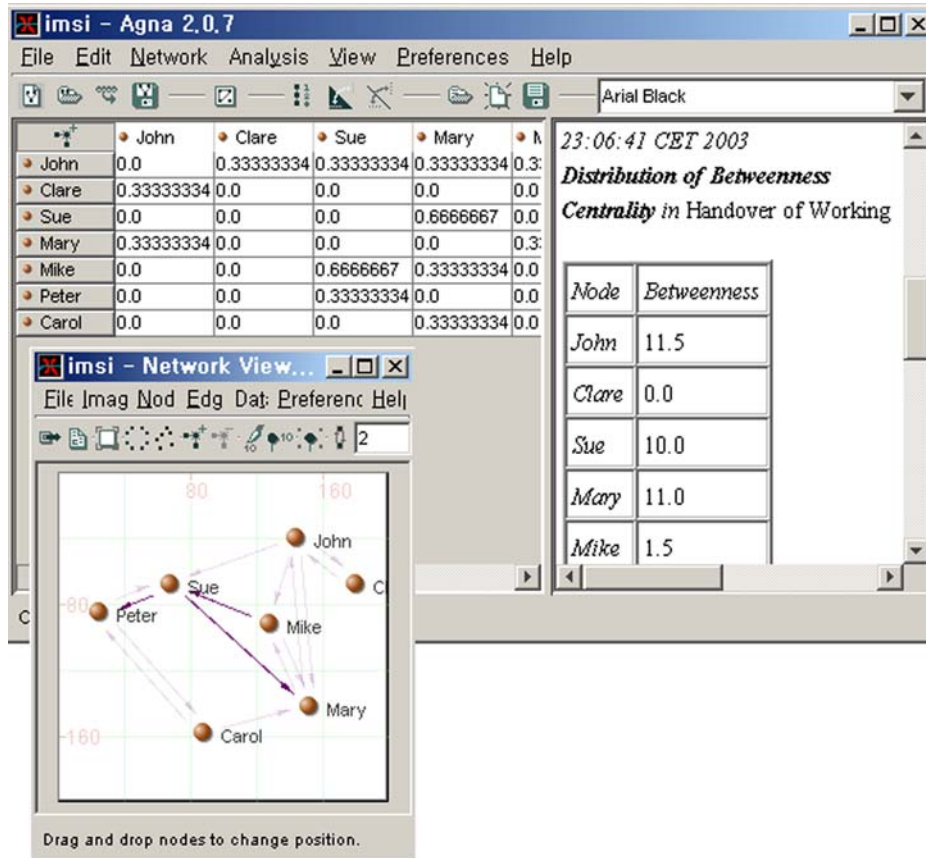


Figure 6. Screenshot of AGNA when analyzing the input from MiSoN.

shows a screenshot of MiSoN when displaying the mining result of handover of work metrics.

MiSoN can export the mining result using the AGNA-translator (but also other tools like UCINET and NetMiner). AGNA (cf. <http://www.geocities.com/imbenta/agna/>) is an SNA tool that allows for a wide variety of sociometric analysis techniques. For example, AGNA supports various notions of centrality including the Bavelas–Leavitt index described in Section 3.1. John and Sue have the highest Bavelas–Leavitt index (the value is 4.2), while Clare has the smallest value (2.8). Figure 6 shows the analysis using the tool AGNA. It also shows the network structure of result.

MiSoN can also export the mining result to other SNA tools like UCINET (cf. <http://www.analytictech.com>) and NetMiner (cf. <http://www.netminer.com>). In fact, in the case study described in the next section we will mainly use NetMiner to analyze the social network.

6. Case study

6.1. CONTEXT

To demonstrate how our metrics can be applied to real workflow logs and what kinds of analysis can be performed, we employed real workflow log data and carried out a case study. The case study we describe here involved one of the twelve provincial offices of the Dutch national public works department, employing about 1000 civil servants. For reasons of confidentiality, we cannot disclose the name of this specific office.

The office's primary responsibility is the construction and maintenance of the road and water infrastructure within its provincial borders. For this purpose, it subcontracts various parties such as road construction companies, cleaning companies, and environmental agencies. Also, the provincial office purchases services and products to support its construction and maintenance activities on the one hand (e.g., mechanical tools, fuel, and traffic signs) and its administrative activities on the other (e.g., office supplies).

The process we dealt with concerns the handling of invoices, as received by the provincial office in question. In general, the handling of an invoice involves several validation steps and, if the invoice is approved, it is completed by payment. On a yearly basis, the provincial office processes some 20,000 invoices from its various subcontractors and suppliers.

The provincial office has implemented its own proprietary workflow management system to support the processing of invoices. This system records transaction information between activities. We extracted a process log and analyzed it. Since the extracted data are also stored in a relational database, we first developed a translator which converts the process log in the database to an XML file using the format described in the previous section.

The process consists of 17 real activities, aside from logistic steps and splits. The log data contains 4,988 cases. The number of total log lines (i.e., events) is 33,603 and 43 employees participated in the process execution. The log holds no information about reassignments. Hence, we cannot apply the reassignment metrics presented in Section 4.4. However, all other metrics we discussed in Section 4 have been applied in this case study.

6.2. METRICS APPLICATION

We applied our metrics to the log data and derived several social networks. Moreover, by applying several SNA techniques, we tried to find the characteristics of the social network.

Figure 7 shows a social network which was derived by applying the handover of work metrics. The network represents how cases are transferred among performers. As indicated in Section 4, there are three refinements possible for the handover of work metrics. To generate this network, we take




Figure 7. Social network based on the handover of work metric \triangleright_L .

into account direct succession and multiple transfers in a case, but we ignore the real process structure, i.e., we use the metric \triangleright_L introduced in Definition 4.3. The network has 43 nodes and 406 links. The density of the network is 0.225 and it has no isolated nodes.

Table IV. Performers having high values for (1) betweenness, (2) in-closeness, (3) out-closeness, and (4) power when analyzing the social network shown in Figure 7

Ranking	Name	Betweenness	Name	In-closeness	Name	Out-closeness	Name	Power
1	user1	0.152	user1	0.792	user23	0.678	user4	4.102
2	user4	0.141	user4	0.792	user1	0.667	user1	2.424
3	user23	0.085	user16	0.75	user4	0.656	user30	1.964
4	user5	0.079	user23	0.689	user5	0.635	user17	1.957
5	user16	0.065	user2	0.667	user13	0.625	user7	1.774
6	user13	0.057	user15	0.618	user18	0.616	user8	1.394
7	user18	0.052	user5	0.609	user2	0.606	user2	1.347
8	user2	0.049	user7	0.592	user16	0.58	user23	1.098
9	user7	0.04	user13	0.568	user7	0.572	user16	1.058
10	user31	0.029	user18	0.568	user17	0.556	user18	0.581

In order to find people who are located in the center of the network, we calculate several centrality values such as betweenness, in and out closeness, and power (Bonacich, 1987) of each node. Normally, the nodes which are the most central have a powerful position in the network. Table IV shows the top 10 ranked performers among the people involved based on (1) *betweenness* (i.e., the extent to which a node lies between all other pair of nodes on their geodesic paths), (2) *in-closeness* (i.e., the inverse of the sum of distances from all the other nodes to a given node, which is then normalized by multiplying it by the number of nodes minus 1), (3) *out-closeness* (i.e., the normalized inverse of the sum of distances from a node to all the other nodes), and (4) *power* (i.e., Bonacich's metric based on the principle that nodes connected to

Table V. Summary of arc weights for various values of β

Beta	Sum	Average	Standard deviation	Minimum value	Maximum value
0.1	1.000025	0.000541	0.003269	0	0.086734
0.3	1.000091	0.000541	0.002895	0	0.074274
0.5	1.000001	0.000541	0.002631	0	0.065751
0.7	1.000011	0.000541	0.002522	0	0.063232
0.9	0.999979	0.000541	0.002586	0	0.067214




Figure 8. Social network based on subcontracting metric.

powerful nodes are also powerful (Bonacich, 1987)). In this table, we find that *user 1* and *user 4* have larger values than others in most measurements.⁵

When generating a social network related to the handover of metrics, we can also consider indirect succession using a “causality fall factor” β . By applying various values of β , we generate several social networks. Despite the value of β , the derived networks have the same structure except the weight of arcs. Table V shows the sum, average, standard deviation, minimum value, and maximum value of the arc weights based on different values of β . If we use a small β , the value of arcs between performers who have the relationship

Table VI. A list of people having a high degree of in-/out-closeness based on the sub-contracting network shown in Figure 8

Ranking	Name	In-closeness	Name	Out-closeness
1	user4	0.262	user4	0.262
2	user1	0.214	user1	0.214
3	user16	0.214	user7	0.167
4	user18	0.19	user13	0.143
5	user5	0.167	user5	0.167
6	user7	0.167	user16	0.214
7	user13	0.143	user18	0.19
8	user19	0.143	user14	0.095
9	user10	0.119	user23	0.119
10	user17	0.119	user27	0.119

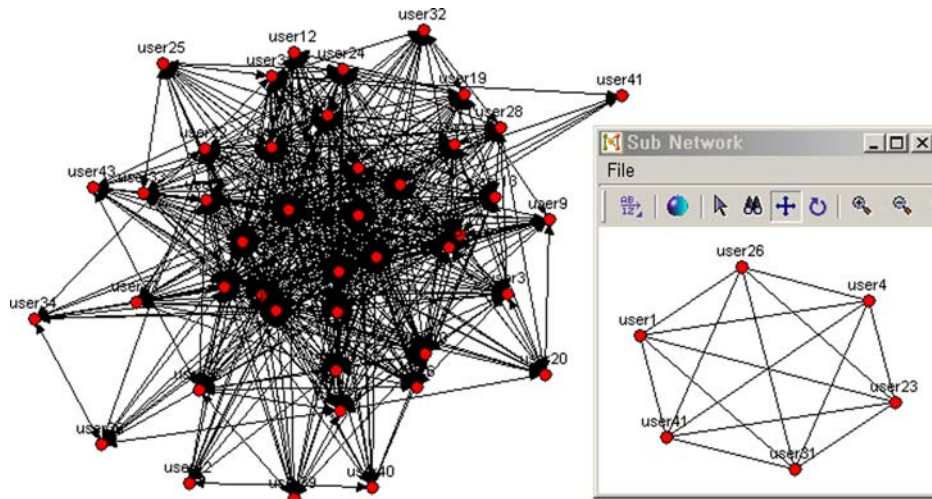


Figure 9. Social network based on the working together metric (left) and the ego network of *user41* (right).

of direct succession is larger than between others. However, if we use a large value of β , these differences decrease.

To find subcontracting relationships between people, we apply in-between metrics. Figure 8 shows the resulting social network. The network has 43 nodes and 146 links. The density of the network is 0.081 and eight nodes are isolated from the network. In this network, the direction of arcs is important. The start node of an arc represents a contractor, while the end node of an arc represents a subcontractor. Table VI shows the ten people of highest in-degree and out-degree of centrality (based on the in-closeness and out-closeness calculated by NetMiner).

Figure 9 shows the social network derived by applying the working together metrics and the ego network (Mitchell, 1969) corresponding to *user41*. In the ego network, the nodes represent the people working together with *user41* according to this metric. Note that *user41* works together with *user1*, *user4*, *user23*, *user26*, and *user31*. The average size of ego network of the generated network is 24.698 and the standard deviation of this value is 9.709.

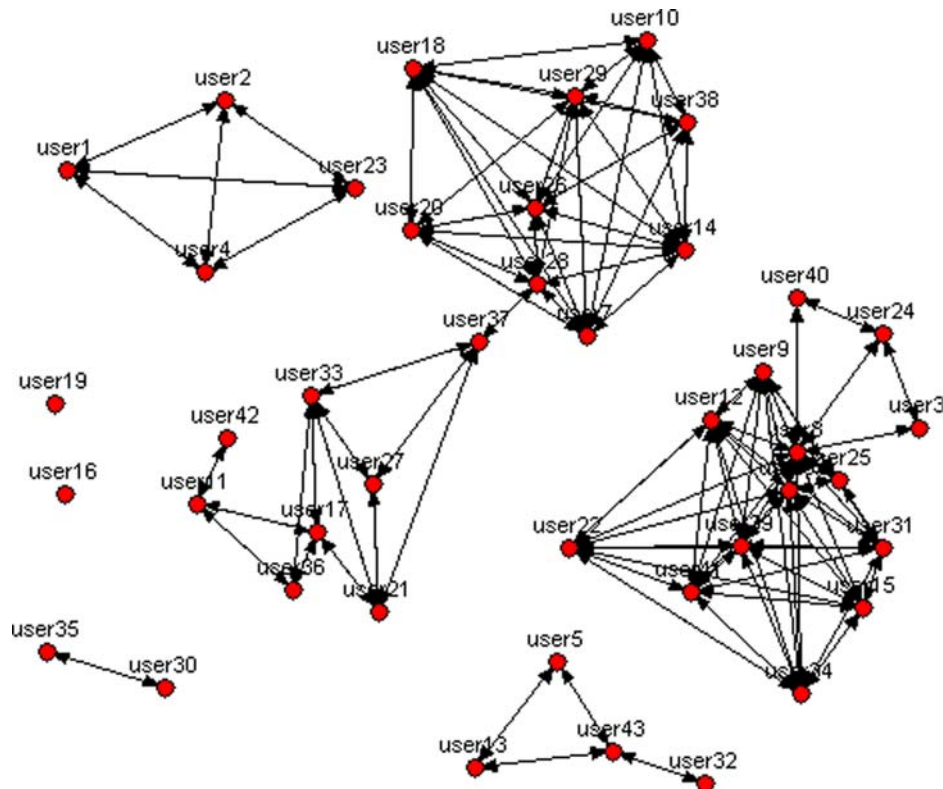


Figure 10. Social network derived from Pearson's correlation coefficient (threshold value 0.75).

This means that the social network suggests that an employee on average works with 24 people.

Applying the metrics based on joint activities, we calculate the distance between people. Figure 10 shows the social network which is derived by applying Pearson's correlation coefficient. From the performer by task matrix, we first apply $\log_{10}(x+1)$, then calculate the distances between people. We get 5 clusters and two isolated nodes. The nodes in the same cluster play the same role. In this case, the bridge node can be interpreted as a person who has multiple roles. In the network, *user8*, *user28*, *user37*, and *user43* have multiple roles.

Finally, we explore how cases are transferred among groups. To calculate case transfers among groups, we combine the handover of work metrics with a role model. In this case study, we use the results of correspondence analysis (Clausen, 1998) as a role model of performers. (Of course, we can also use the results of the metrics based on joint activities.) Correspondence analysis is frequently used in biological science to analyze ecological systems based on species scores for specific locations (Gauch, 1982). In this paper, we apply correspondence analysis to find relationships between activities and performers. We first make a performer by activity matrix from the workflow logs. Then, by applying correspondence analysis to the matrix, we derive the relationship between activities, between performers, and between activities and performers. Figure 11 shows the graphical result of applying correspondence analysis. In the figure, boxes represent activities and circles rep-

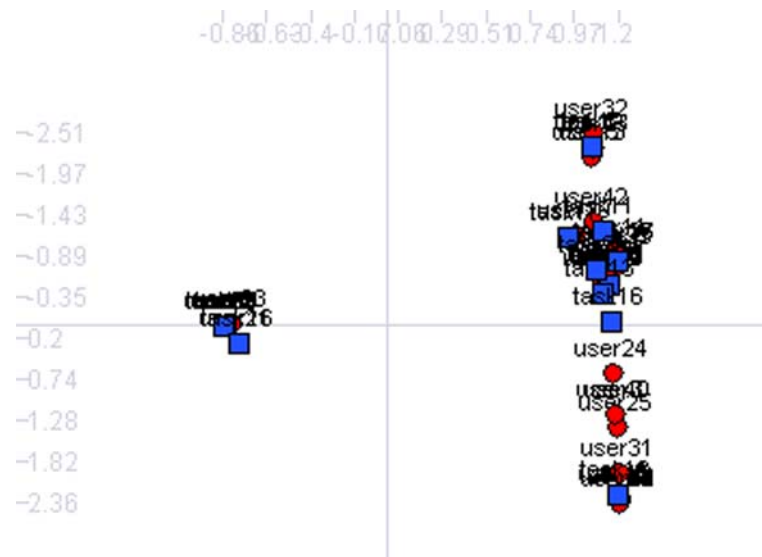


Figure 11. Graphical result of correspondence analysis.

Table VII. The result of correspondence analysis: users are clustered into five groups

Group	Performers	Activities
group1	user1, user2, user4, user16, user23, user30, user35	task2, task3, task15, task21, task22
group2	user3, user24, user25, user40	
group3	user5, user13, user32, user43	task8, task19
group4	user6, user8, user9, user12, user15, user22, user31, user39, user41	task18
group5	user7, user10, user11, user14, user17, user18, user19, user20, user21, user26, user27, user28, user29, user33, user34, user36, user37, user38, user42	task5, task7, task11, task13, task16, task17, task20

resents performers. Closely positioned nodes indicate a strong correspondence from a work handover perspective between the respective users and/or tasks. (Although the distance between user nodes and task nodes should not be interpreted as an absolute measure.) From this figure, performers and activities are classified into five groups. Table VII shows the results. In the remainder we will use these five groups as a role model.

Figure 12(a) shows the social network of handover of work metrics considering the role model given in Table VII. By putting the nodes in the same group closely, we have reconstructed the original network. And by summing

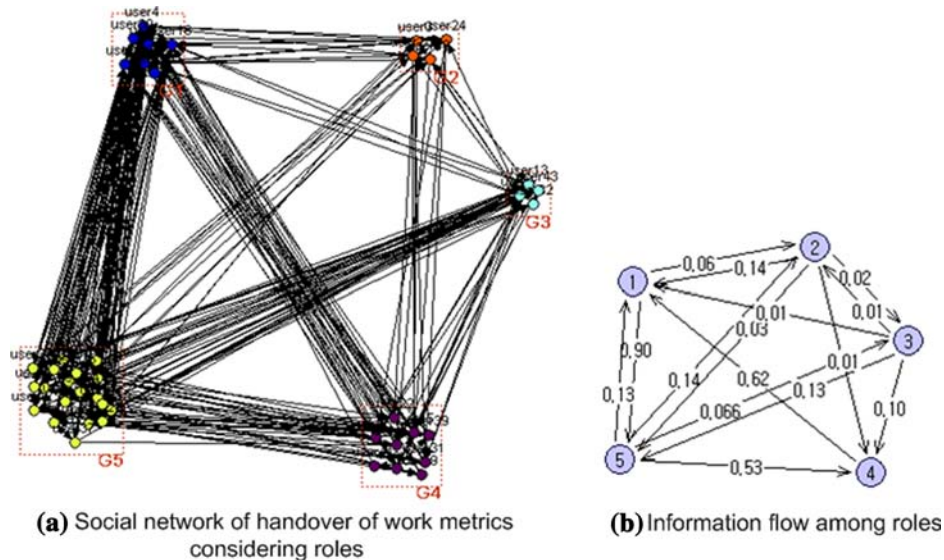


Figure 12. Social network based on handover of work metric using the five groups shown in Table VII.

Table VIII. Information flow between groups

From	To					Sum
	group1	group2	group3	group4	group5	
group1	1.330	0.058	0.002	0.002	0.895	2.287
group2	0.143	0.014	0.020	0.005	0.028	0.211
group3	0.014	0.005	0.002	0.104	0.030	0.154
group4	0.620	0.000	0.002	0.005	0.004	0.630
group5	0.132	0.135	0.134	0.526	0.617	1.545
Sum	2.239	0.212	0.160	0.642	1.574	4.827

up the weight of arcs between groups we derive the aggregated network shown in Figure 12(b).

Table VIII shows the information flow of the network according to the role model. It is also derived by summing up the weight of arcs between groups. For example, the value from *group1* to *group2* is calculated by adding up the weights of the arcs from nodes in *group1* to nodes in *group2*. Based on Table VIII we can make some observations. First, the highest value (1.330) is in the cell from *group1* to *group1*. It means that the handover of work within *group1* happened most frequently. Second the values from *group1* to *group5* (0.895), from *group4* to *group1* (0.620), and from *group5* to *group4* (0.529) have the highest values. It represents that more handover of work happened between these groups.

The goal of this subsection is not to provide a comprehensive overview of all the diagrams we developed or to provide very specific information about the studied process or organization in question. Its main purpose is to illustrate the various types of analysis possible. In the next section, we will reflect on the relevance of the different types of analysis for the organization in question.

6.3. ORGANIZATIONAL RELEVANCE

Prior to our analysis, the involved management did not express any specific needs or questions about the invoice handling process. And yet, they indicated that the handling of invoices is in the center of their attention. There are two main reasons for this. First of all, it is the single most distributed process under the responsibility of the public works department. For example, if invoices are related to some particular public works project, its project leader must personally certify that delivery has taken place before payment may happen. Project leaders, however, may reside at any location within the provincial borders (in contrast to the performers working at the administrative head office). The distributed nature of the process increases the probability of hand-over errors and work getting lost.

The second reason for the attention for the invoice process is the recent Dutch law about penalty interests. Parties that send invoices to public organizations and receive their rightful payment after more than 30 days are entitled to a compensation proportional to the invoice amount. Due to the current interest rate, this compensation exceeds commercial rates. Sluggish settlement of invoices directly affects the public works department's financial position.

Both issues have contributed to the decision to introduce workflow technology to support the invoice handling process, as this is expected to increase quality and decrease process lead time. Management expressed a general interest in results from SNA to learn about process execution, behavior of the involved parties, and potential opportunities for improvement.

After we applied our metrics to the log data and derived the social networks as shown, we presented the managers of the three most involved departments our analysis results in a joint session. Roughly speaking, the three departments are, respectively, responsible for the administrative, contractual, and financial aspects of the invoice handling. The goal of this meeting from a research perspective was threefold:

1. To validate our understanding of the process.
2. To generate feedback on our analyses.
3. To identify further analysis opportunities.

To determine whether we properly understood the process, we discussed the process model of the invoice handling process and the involved parties for each of the various steps. This led to no surprising new insights. We will reflect on the other two aims of the meeting in more detail.

Feedback on analyses. After we explained the various SNA notions we presented the results from our analyses as presented for a large part in Section 6.2. We started with discussing the top 5 and bottom 5 of the lists of performers as ordered on their scores on betweenness, in and out closeness, and power in, respectively, the social networks of handover, subcontracting, and working together metrics. Note that we used the lists that included the real names of the actors to facilitate meaningful feedback.

From the responses, we learned that, typically, performers with high scores (e.g. *user1* and *user4* in Table IV) work for the administrative department in supportive functions. This confirms a general insight that highly connected people often are assistants. Because the administrative department is responsible for both the preparation and completion of the handling of each invoice, its staff is involved in the handling of each case, giving them strong ties with other performers. The managers indicated, however, that not all of the people in these positions were present in the top of the lists, indicating that having a supportive function is not sufficient in itself to become highly connected.

Performers with low scores could be categorized as follows. First of all, project leaders were highly represented in the bottom of the lists (e.g. *user9*). As stated before, they play an isolated role in the handling of invoices, being solely responsible for certifying that goods have been delivered (and only if an invoice is related to a project at all). Other performers with limited formal verification responsibilities were identified as well (e.g. *user22*). The second category of relatively unconnected performers could be traced back to auxiliary logins (e.g. *user30*), used by system administrators and management to deal with exceptional circumstances. An example of an exceptional situation is an invoice that is being withdrawn while its processing has already started. The isolated “participation” of this category of users is therefore not very surprising. It did, however, make the managers conscious of the visibility of this type of irregular interference. One manager remarked: “So, auditors can derive this type of information too”. The third category turned out to be more surprising, as it involved senior positions in the contractual and financial departments (e.g. *user41*). At least nominally, they are expected to be actively involved in the process. Their low position could indicate that a large amount of work being executed with workflow technology is delegated to their juniors. Also, one of these performers would retire in a couple of weeks.

After the discussion of the lists of performers, we presented the social network indicating the distance between people (see Figure 10). The relations between users were readily recognized by the involved managers. For example, the subgraph of *user1*, *user2*, *user4* and *user23* concerned the group of highly-connected assistants at the administrative office we encountered earlier. Then, we took a closer look at the two isolated nodes. One of them – rather characteristically – turned out to be the system administrator (*user19*). The isolated position of the other node, *user16*, led to some excitement. At first, the isolation of this performer was not understood, as she was considered to perform an explicit role in the contractual handling of invoices. Then it occurred to one of the managers that the involved person was included in another cluster under a *different user name* as well (*user32*). The existence of such a situation was a complete surprise to the managers and considered highly undesirable for compliancy reasons.

The final result we obtained feedback on, involved the correspondence analysis, such as presented in Table VII and Figure 11. Managers readily recognized *group1*, *group3*, and *group5*. At the same time, they indicated that they did not differentiate themselves between most of the performers in *group2* and *group4*. This is in line with the observation that some performers from these groups are closely positioned to each other in Figure 10. For example, the positions of *user25* from *group2* and *user31* from *group4* nearly coincide. And yet, the strong correspondence between *group4* and *task18* indicates that a degree of performer specialization has taken place with re-

spect to this specific invoice check that had gone unnoticed with the concerned managers.

Further analyses. Aside from the various surprising aspects of the invoice handling process, the managers of the provincial office were most intrigued by the subcontracting analysis. After some discussion, they expressed their suspicion about three parts of the process where a similar yet undesirable “back-and-forth” behavior may take place. Specifically, they meant that a performer (the contractor) routes a work package to another performer (the subcontractor), who subsequently routes it back to the contractor or one of the contractor’s close colleagues, because the subcontractor feels the invoice is received in error. This, for example, takes place when an invoice related to a project is sent for verification to the wrong project leader. Each occurrence of this pattern is highly undesirable, as it slows down the processing of the invoice without making any progress. From an organizational perspective, it is just as unwelcome when the work package is routed back to the original contractor as to a colleague with a similar organizational role.

Our initial analysis did not cover this more general kind of subcontracting pattern, because it focused on the identity of the original contractor only (see Definition 4.6). To investigate the expressed suspicions we analyzed the mining log in various ways, using other than SNA techniques as well. Therefore, in the context of this paper, we will be brief about this additional analysis. It turned out that in the handling of over 17% of all invoices, at least once an undesired subcontracting takes place at either of the three identified places in the process. The exact distribution is shown in Figure 13. As can be seen, there are cases where 10 or more erroneous routings take place.




Figure 13. The distribution of undesired generalized subcontracting within the handling of invoices.

As a result of the additional analysis we carried out and discussed again, management of the provincial office re-enforced the existing procedure that staff, when in doubt, should contact the intended next performer by phone first. Especially in cases where hand-overs take place between performers at the head and regional offices, management felt that people acted too timidly with respect to this procedure.

7. Discussion

In this section, we demonstrated how our metrics can be applied to a real workflow log of a Dutch organization. Based on the metrics defined in Section 4, we derived various sociograms, some of which have been shown in this paper. Using the sociograms we applied SNA techniques such as betweenness, closeness, power, ego network, etc. We also showed the possibility of applying other analysis techniques such as correspondence analysis to compare users based on their “profile”.

Next we discussed the organizational relevance of our analyses. As we indicated, many of our findings corresponded with existing insights of the involved management, supporting the correctness of our analyses. At various points, our analyses came as a surprise. These particularly concerned senior performers who did not seem very connected, the clear visibility of the actions of irregular performers, and a degree of unnoticed performer specialization that had taken place. In addition, we found it interesting to observe how our analysis results triggered the management to identify and define additional questions. This, in our eyes, strongly supports the relevance and viability of process mining in an organizational context, even though our additional analyses extended beyond SNA.

As discussed in Section 2.4, ethical and legal issues play an important role in the practical application of process mining in general and SNA analysis in particular. One concern we certainly felt is that the validation and discussion of our analysis results required us to disclose the identity of the involved performers. Our tools can anonymize the sociograms without any problems and we can also aggregate them at a group level. However, in our discussions with the involved management, it was helpful to be able to refer to specific people. We informed the management that it is illegal to perform actions towards individuals based on the presented results. Because of the clear value of this type of analysis, the managers expressed their intent to ask for the consent of their employees for the use of future analyses. Note that the re-enforced policy that resulted from our additional analysis was neither based on information obtained on individual performers, nor did it affect any individual more than others.

8. Related work

Related work can be divided in two categories: process mining and SNA. However, before discussing the directly related work in detail, we would like to position this work in the broader Computer Supported Cooperative Work (CSCW) domain. The work presented in this paper may contribute to the ongoing discussions in the CSCW domain on “technology push” (i.e., “What is possible?”) versus “social pull” (i.e., “What is desirable from a sociological point of view?”). There exist many definitions of the term CSCW: some emphasize the support of work processes while other emphasize the fact that people work in groups (Ellis et al., 1991; Ellis, 2000). Within the CSCW domain there has been a constant struggle between more technological-oriented views and more sociological-oriented views. A nice illustration is the so-called “Winograd- Suchman debate” in the early nineties (Winograd and Flores, 1986; Suchman, 1994; Winograd, 1994; Malone, 1995). Winograd and Flores advocated the use of a system called the “coordinator”, a system based on Speech Act theory (i.e., the language/action perspective) in-between e-mail and workflow technology (Winograd and Flores, 1986; Winograd, 1994). People like Suchman and others argued that such systems are undesirable as they “carry an agenda of discipline and control over an organization’s members” (Suchman, 1994). Clearly, process mining and discovering social networks based on event logs add another dimension to this discussion. We would like to emphasize that the goal of process mining is not to control people. However, it can be used to monitor and analyze the behavior of people and organizations. Clearly, such technology triggers ethical questions as discussed in Section 2.4. Although it is important to balance between both views, our approach can be seen as a “technology push” approach triggered by the event logs present in contemporary information systems.

8.1. RELATED WORK ON PROCESS MINING

The idea of process mining is not new (Agrawal et al., 1998; Cook and Wolf, 1998; Aalst et al., 2003) but has been mainly aiming at the control-flow perspective. The idea of applying process mining in the context of workflow management was first introduced in Agrawal et al. (1998). This work is based on workflow graphs, which are inspired by workflow products such as IBM MQSeries Workflow (formerly known as Flowmark). Cook and Wolf have investigated similar issues in the context of software engineering processes. In Cook and Wolf (1998) they describe three methods for process discovery: one using neural networks, one using a purely algorithmic approach, and one Markovian approach. Schimm (2000) has developed a mining tool suitable for discovering hierarchically structured workflow processes. Herbst and Karagiannis also address the issue of process mining in the context of workflow

management using an inductive approach (Herbst, 2000, 2001). They use stochastic task graphs as an intermediate representation and generate a workflow model described in the ADONIS modeling language. Most of the approaches have problems dealing with parallelism and noise. Our work in Aalst and Dongen (2002) and Aalst et al. (2004) is characterized by the focus on workflow processes with concurrent behavior (rather than adding *ad-hoc* mechanisms to capture parallelism). In Weijters and Aalst (2003) a heuristic approach using rather simple metrics is used to construct so-called “dependency/frequency tables” and “dependency/frequency graphs”. These are then used to tackle the problem of noise. The approaches described in Aalst and Dongen (2002), Aalst et al. (2004) and Weijters and Aalst (2003) are based on the α algorithm.

Process mining in a broader sense can be seen as a tool in the context of Business Process Intelligence (BPI). In Grigori et al. (2001) and Sayal et al. (2002), a BPI toolset on top of HP’s Process Manager is described. The BPI toolset includes a so-called “BPI Process Mining Engine”. However, this engine does not provide any techniques as discussed before. Instead it uses generic mining tools such as SAS Enterprise Miner for the generation of decision trees relating attributes of cases to information about execution paths (e.g., duration). In order to do workflow mining it is convenient to have a so-called “process data warehouse” to store audit trails. Such a data warehouse simplifies and speeds up the queries needed to derive causal relations. In Mühlen and Rosemann (2000), Zur Muehlen describes the PISA tool which can be used to extract performance metrics from workflow logs. Similar diagnostics are provided by the ARIS Process Performance Manager (PPM) (IDS Scheer, 2002). The latter tool is commercially available and integrates some of our ideas. A customized version of PPM is the Staffware Process Monitor (SPM) (Staffware, 2002) which is tailored towards mining Staffware logs. Note that none of the latter tools is extracting models, i.e., the results do not include control-flow, organizational or social network related diagnostics. The focus is exclusively on performance metrics.

For more information on process mining we refer to a special issue of Computers in Industry on process mining (Aalst and Weijters, 2004) and the survey paper (Aalst et al., 2003). Note that although quite some work has been done on process mining from event logs none of the approaches known to the authors have incorporated the social dimension as discussed in this paper.

8.2. RELATED WORK ON SNA

Since the early work of Moreno (1934), sociometry, and SNA in particular, have been active research domains. There is a vast amount of textbooks, research papers, and tools available in this domain (Moreno, 1934; Bavelas, 1948; Feldman, 1977, 1979, 1987; Burt and Minor, 1983; Bernard et al., 1990;

Scott, 1992; Wasserman and Faust, 1994; Nemati and Barko, 2003). There have been many studies analyzing organizational activity based on insights from social network analysis. However, some of these studies typically have an adhoc character and sociograms are typically constructed based on questionnaires rather than using a structured and automated approach as described in this paper. More structured approaches are often based on the analysis of e-mail interaction and additional electronic sources. Several studies have generated sociograms from email logs in organization (Ogata et al., 2001; Nardi et al., 2002; Farnham et al., 2004a, b; Fisher and Dourish, 2004) to analyze the communication structure. Such studies have resulted in the identification of relevant, recurrent aspects of interaction in organizational contexts (Begole et al., 2002; Fisher and Dourish, 2004). However, these studies are unable to relate the derived social networks to a particular workflow process, as the analyzed data does not reveal to what activity or case it applies.

Most tools in the SNA domain take sociograms as input. MiSoN is one of the few tools that generate sociograms as output. The only comparable tools are tools to analyze e-mail traffic, cf. BuddyGraph (<http://www.buddygraph.com>), MetaSight (<http://www.metasight.co.uk/>), and ContactMap (Nardi et al., 2002). However, these tools monitor unstructured messages and cannot distinguish between different activities (e.g., work-related interaction versus social interaction).

As indicated in the introduction, this paper extends the results presented in Aalst and Song (2004). Unlike Aalst and Song (2004), this paper provides concrete metrics, a more elaborate description of MiSoN, and a case study illustrating the applicability of the approach.

9. Conclusions

This paper presents an approach, concrete metrics, and a tool to extract information from event logs and construct a sociogram which can be used to analyze interpersonal relationships in an organization. Today many information systems are “process aware” and log events in some structured way. As indicated in the introduction, workflow management systems register the start and completion of activities, ERP systems log all transactions (e.g., users filling out forms), call center and CRM systems log interactions with customers, etc. These examples have in common that there is some kind of event log. Unfortunately, the information in these logs is rarely used to derive information about the process, the organization, and the social network. In this paper, we focus on the latter aspect and present an approach to discover sociograms. These sociograms are based on the observed behavior and may use events like the transfer of work or delegation from one individual to another. MiSoN can interface with commercial systems such as Staffware

and standard SNA tools like AGNA, UCINET and NetMiner, thus allowing for the application of the ideas presented in this paper.

The approach presented in this paper has a number of obvious limitations. First of all, we can only monitor the events that are actually logged. This implies that some interactions may not be visible. Moreover, people may find ways to work around the system. Second, the system may enforce certain interaction patterns. If workers are completely controlled by the system, the discovered sociograms reflect the system rather than the organization. Fortunately, most systems offer a lot of flexibility when it comes to the selection and ordering of work-items. Even workflow management systems allow for a pull mechanism where workers select work-items from a shared pool in any order. Moreover, other types of process-aware information systems (e.g., ERP, CRM, PDM systems) tend to allow for even more flexibility. In addition to the two limitations mentioned, it is important to note that the sociogram and its analysis are just a starting point for a deeper investigation.

Besides an approach and tool to derive sociograms from event logs, the paper also presents a case study conducted within a Dutch national public works department. The case study shows that the event logs in real organizations allow for social network analysis. Moreover, in this particular case the analysis results provide relevant, surprising organizational information. The established results and resulting discussions have formed the basis for additional process mining to deal with managerial concerns, resulting in the re-enforcement of organizational policies. In the future, we plan to repeat our analysis within the public works department and apply our approach in many other organizations as well. It would be interesting to compare the results we obtain on the basis of event logs to results of the analysis of other communication means usage e.g., e-mail. This would provide an even richer view on organizational interaction and process improvement opportunities.

We also investigate extensions of the approach using filtering techniques and more advanced forms of clustering. For example, we now abstract from the results of activities. If activities or cases can be classified as successful or unsuccessful, important or unimportant, standard or special, etc., this information could be used when building sociograms.

Recently, MiSoN has been integrated in the ProM framework.⁶ The ProM framework allows for various types of process mining, i.e., given a log it is possible to not only derive sociograms but also process models. The ProM framework also provides an LTL checker that can check properties expressed in Linear Temporal Logic (LTL) (Manna and Pnueli, 1991). This allows for all kinds of questions, e.g., checking the 4-eyes principle (two tasks for the same case need to be executed by different people to avoid fraud). This LTL checker can be used to answer more detailed questions based in insights generated from the SNA analysis. In the context of the ProM framework also a prototype of an e-mail analysis tool has been developed. Based on a user's

Inbox located on some Exchange server, the prototype can translate the e-mails to the XML format described in this paper. However, since e-mails may refer to different processes and there are no explicit pointers to tasks and cases, and heuristics and/or conventions need to be used. Therefore, we only consider this as a means to provide more context to the SNA analysis based on true event logs.

Acknowledgements

Minseok Song visited the Department of Technology Management at Eindhoven University of Technology with funding by the BK21 program. He would like to thank the Ministry of Education of Korea for its financial support through the BK21 program. The authors would also like to thank Ton Weijters, Boudewijn van Dongen, Ana Karla Alves de Medeiros, Andriy Anikolov, Laura Maruster, Eric Verbeek, Anne Rozinat, Christian Günther, Monique Jansen-Vullers, Michael Rosemann, and Peter van den Brand for their on-going work on process mining techniques and tools at Eindhoven University of Technology.

Notes

1. Note that MiSoN has been embedded in the ProM framework and can be downloaded from <http://www.processmining.org>.
2. Clearly the number of events in Table I is too small to establish these assumptions accurately. However, for the sake of argument we assume that the things that did not happen will never happen, cf. Section 2.3.
3. In the context of the ProM framework (<http://www.processmining.org>), we offer a tool that can assign randomly generated names to workers before starting the analysis.
4. These are the event types used by ProM and MinSoN, for more information see <http://www.processmining.org>.
5. Note that the real user names are changed into anonymous identifiers like *user1*. Although during our analysis and interaction with the organization real user names were used, we abstract from the real user names in this paper to ensure privacy and confidentiality.
6. See <http://www.processmining.org> for more information and to download the tools used in this paper.

References

- Agrawal, R., D. Gunopulos and F. Leymann (1998): Mining Process Models from Workflow Logs. *Sixth International Conference on Extending Database Technology*, pp. 469–483.
- Bavelas, A. (1948): A Mathematical Model for Group Structures. *Human Organization*, vol. 7, pp. 16–30.

- Begole, J., J. Tang, R. Smith and N. Yankelovich (2002): Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups. In C. Neuwirth and T. Rodden (eds.): *Proceedings of the 2002 ACM conference on Computer Supported Cooperative Work*. New York, NY, USA: ACM Press, pp. 334–343.
- Bernard, H., P. Killworth, C. McCarty, G. Shelley and S. Robinson (1990): Comparing Four Different Methods for Measuring Personal Social Networks. *Social Networks*, vol. 12, pp. 179–216.
- Bonacich, P. (1987): Power and Centrality: A family of Measures. *American Journal of Sociology*, vol. 92, pp. 1170–1182.
- Burt, R. and M. Minor (1983): *Applied Network Analysis: A Methodological Introduction*. Newbury Park CA: Sage <http://www.cbppweb.nl/index.htm>.
- Clausen, S.E. (1998): *Applied Correspondence Analysis: An Introduction*, Sage Publications.
- CBP n.d., College Bescherming Persoonsgegevens (CBP; Dutch Data Protection Authority).
- Cook, J. and A. Wolf (1998): Discovering Models of Software Processes from Event-Based Data. *ACM Transactions on Software Engineering and Methodology*, vol. 7, no. 3, pp. 215–249.
- Culotta, A., R. Bekkerman and A. McCallum (2004): Extracting Social Networks and Contact Information from Email and the Web. *Proceedings of the First Conference on Email and Anti-Spam (CEAS)*.
- Ellis, C. (2000): *An Evaluation Framework for Collaborative Systems*. Boulder, USA: University of Colorado, Department of Computer Science, Technical Report, CU-CS-901-00.
- Ellis, C., S. Gibbs and G. Rein (1991): Groupware: Some Issues and Experiences. *Communications of the ACM*, vol. 34, no. 1, pp. 38–58.
- Farnham, S., S. Kelly, W. Portnoy and J. Schwartz (2004a): Wallop: Designing Social Software for Co-Located Social Networks. *Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04)*. CA: IEEE Computer Society Press, Los Alamitos.
- Farnham, S., W. Portnoy and A. Turski (2004b): Using Email Mailing Lists to Approximate and Explore Corporate Social Networks, In D. McDonald, S. Farnham and D. Fisher (eds.): *Proceedings of the CSCW'04 Workshop on Social Networks*.
- Feldman, M. (1987): Electronic Mail and Weak Ties in Organizations. *Office: Technology and People*, vol. 3, pp. 83–101.
- Fischer, L. (ed.) (2001): *Workflow Handbook 2001, Workflow Management Coalition*. Lighthouse Point, FL: Future Strategies.
- Fisher, D. and P. Dourish (2004): Social and Temporal Structures in Everyday Collaboration. In E. Dykstra-Erickson and M. Tscheligi (eds.): *Proceedings of the 2004 Conference on Human Factors in Computing Systems (CHI2004)*. New York, NY, USA: ACM Press, pp. 551–558.
- Freeman, L. (1977): A Set of Measures of Centrality Based on Betweenness. *Sociometry*, vol. 40, pp. 35–41.
- Freeman, L. (1979): Centrality in Social Networks: Conceptual Clarification. *Social Networks*, vol. 1, pp. 215–239.
- Gauch, H.G. (1982): *Data Analysis in Community and Landscape Ecology*, Cambridge University Press.
- Grigori, D., F. Casati, U. Dayal and M. Shan (2001): Improving Business Process Quality through Exception Understanding, Prediction, and Prevention. In P. Apers, P. Atzeni, S. Ceri, S. Paraboschi, K. Ramamohanarao and R. Snodgrass (eds.): *Proceedings of 27th International Conference on Very Large Data Bases (VLDB'01)*, Morgan Kaufmann, pp. 159–168.

- Herbst, J. (2000): A Machine Learning Approach to Workflow Management. *Proceedings 11th European Conference on Machine Learning, Vol. 1810 of Lecture Notes in Computer Science*. Berlin: Springer-Verlag, pp. 183–194.
- Herbst, J. (2001): *Ein induktiver Ansatz zur Akquisition und Adaption von Workflow-Modellen*, PhD thesis, Universität Ulm.
- Hulsman, B. and P. Ippel (1994): *Personeelsinformatiesystemen: De Wet Persoonsregistraties toegepast*. The Hague: Registratiekamer.
- IDS Scheer (2002): ARIS Process Performance Manager (ARIS PPM): Measure, Analyze and Optimize Your Business Process Performance (whitepaper). IDS Scheer, Saarbruecken, Gemany, <http://www.ids-scheer.com>.
- Jablonski, S. and C. Bussler (1996): *Workflow Management: Modeling Concepts, Architecture, and Implementation*. London, UK: International Thomson Computer Press.
- Leymann, F. and D. Roller (1999): *Production Workflow: Concepts and Techniques*. Upper Saddle River, New Jersey, USA: Prentice-Hall PTR.
- Malone, T. (1995): Commentary on Suchman article and Winograd response. *Computer Supported Cooperative Work*, vol. 3, no. 1, pp. 37–38.
- Manna, Z. and A. Pnueli (1991): *The Temporal Logic of Reactive and Concurrent Systems: Specification*. New York: Springer-Verlag.
- Mitchell, J. (1969): The Concept and Use of Social Networks. In J. Mitchell (ed.): *Social Networks in Urban Situations*. Manchester: Manchester University Press, pp. 1–50.
- Moreno, J. (1934): *Who Shall Survive?*. Washington, DC: Nervous and Mental Disease Publishing Company.
- Mühlen, M. and M. Rosemann (2000): Workflow-based Process Monitoring and Controlling – Technical and Organizational Issues. In R. Sprague (ed.): *Proceedings of the 33rd Hawaii International Conference on System Science (HICSS-33)*. Los Alamitos, CA: IEEE Computer Society Press, pp. 1–10.
- Nardi, B., S. Whittaker, E. Isaacs, M. Creech, J. Johnson and J. Hainsworth (2002): Integrating Communication and Information Through ContactMap. *Communications of the ACM*, vol. 45, no. 2, pp. 89–95.
- Nemati, H. and C. Barko (2003): *Organizational Data Mining: Leveraging Enterprise Data Resources for Optimal Performance*. Hershey, PA, USA: Idea Group Publishing.
- Ogata, H., Y. Yano, N. Furugori and Q. Jin (2001): Computer Supported Social Networking For Augmenting Cooperation. *Computer Supported Cooperative Work*, vol. 10, no. 2, pp. 189–209.
- Sauerwein, L. and J. Linnemann (2001): *Guidelines for Personal Data Processors: Personal Data Protection Act*. The Hague: Ministry of Justice.
- Sayal, M., F. Casati, U. Dayal and M. Shan (2002): Business Process Cockpit. *Proceedings of 28th International Conference on Very Large Data Bases (VLDB'02)*. Morgan Kaufmann, pp. 880–883.
- Schimm, G. (2000): Generic Linear Business Process Modeling. In S. Liddle, H. Mayr and B. Thalheim (eds.): *Proceedings of the ER 2000 Workshop on Conceptual Approaches for E-Business and The World Wide Web and Conceptual Modeling Vol. 1921 of Lecture Notes in Computer Science*. Berlin: Springer-Verlag, pp. 31–39.
- Scott, J. (1992): *Social Network Analysis*. Newbury Park, CA: Sage.
- Smith, M. (1999): Invisible Crowds in Cyberspace: Measuring and Mapping the Social Structure of Usenet. In M. Smith and P. Kollock (eds.): *Communities in Cyberspace: Perspectives on New Forms of Social Organization*, Routledge Press.
- Staffware (2002): Staffware Process Monitor (SPM), <http://www.staffware.com>.
- Suchman, L. (1994): Do Categories Have Politics? The Language /Action Perspective Reconsidered. *Computer Supported Cooperative Work*, vol. 2, no. 3, pp. 177–190.

- van der Aalst, W. and B. Dongen (2002): Discovering Workflow Performance Models from Timed Logs. In Y. Han, S. Tai and D. Wikarski (eds.): *International Conference on Engineering and Deployment of Cooperative Information Systems (EDCIS 2002)*, Vol. 2480 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag, pp. 45–63.
- van der Aalst, W. and K. Hee (2002): *Workflow Management: Models, Methods, and Systems*. Cambridge, MA: MIT press.
- van der Aalst, W. and M. Song (2004): Mining Social Networks: Uncovering Interaction Patterns in Business Processes. In J. Desel, B. Pernici and M. Weske (eds.): *International Conference on Business Process Management (BPM 2004)*, Vol. 3080 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag, pp. 244–260.
- van der Aalst, W. van der and A. Weijters (eds.) (2004): *Process Mining*. Amsterdam: Elsevier Science Publishers, Special Issue of *Computers in Industry*, Vol. 53, no. 3.
- van der Aalst, W., B. van Dongen, J. Herbst, L. Maruster, G. Schimm and A. Weijters (2003): Workflow Mining: A Survey of Issues and Approaches. *Data and Knowledge Engineering*, vol. 47, no. 2, pp. 237–267.
- van der Aalst, W., A. Weijters and L. Maruster (2004): Workflow Mining: Discovering Process Models from Event Logs. *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 9, pp. 1128–1142.
- Wasserman, S. and K. Faust (1994): *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.
- Weijters, A. and W. Aalst (2003): Rediscovering Workflow Models from Event-Based Data using Little Thumb. *Integrated Computer-Aided Engineering*, vol. 10, no. 2, pp. 151–162.
- Winograd, T. (1994): Categories, Disciplines, and Social Coordination. *Computer Supported Cooperative Work*, vol. 2, no. 3, pp. 191–197.
- Winograd, T. and F. Flores (1986): *Understanding Computers and Cognition: A New Foundation for Design*. Norwood: Ablex.