



Concurrency and Objects Matter! Disentangling the Fabric of Real Operational Processes to Create Digital Twins

Wil M. P. van der Aalst^{1,2}(✉)

¹ Process and Data Science (Informatik 9), RWTH Aachen University,
Aachen, Germany

wvdaalst@pads.rwth-aachen.de

² Fraunhofer-Institut für Angewandte Informationstechnik (FIT),
Sankt Augustin, Germany

Abstract. Process mining dramatically changed the way we look at process models and operational processes. Even seemingly simple processes like Purchase-to-Pay (P2P) and Order-to-Cash (O2C) are often amazingly complex, and traditional hand-made process models fail to capture the true fabric of such processes. Many processes are inherently concurrent and involve interaction between different objects (customers, suppliers, orders, items, shipments, payments, machines, workers, etc.). Process mining uses event data to construct process models that can be used to diagnose performance and compliance problems. If such models reflect reality well, they can be used for forward-looking forms of process mining, including predictive analytics, evidence-based automation, and what-if simulation. The ultimate goal is to create a “digital twin of an organization” that can be used to explore different improvement actions. This paper provides a high-level overview of the different process mining tasks followed by a more detailed discussion on concurrency and object-centricity in process mining.

Keywords: Process mining · Event data · Concurrency · Digital twins

1 Towards a Digital Twin of an Organization

The desire to adequately describe operational processes has been around since the 1890-ties when the field of *scientific management* emerged. Scientific management is also known as Taylorism, named after its pioneer Frederick Winslow Taylor (1856–1915) who tried to systematically improve economic efficiency, especially labor productivity. Taylor systematically observed how people work and can be seen as the “first process miner” using pen and paper (see Fig. 1). In 1950 computers started to influence business processes. However, the systematic use of data about operational processes is much more recent [1].

OPERATION (What-Is-Now) Execution

OPERATION	DATE	TIME	...
DEPARTMENT Construction	11.11.11	11.11.11	...
MEN Mku Flakoty	11.11.11	11.11.11	...
MATERIALS Sand, requiring no pick	11.11.11	11.11.11	...
IMPLEMENTERS The 25th of Cambridge road London	11.11.11	11.11.11	...
CONDITIONS Dry, sunny, for construction	11.11.11	11.11.11	...
on average 1 hour 10 min of sand at 200 ft from road	11.11.11	11.11.11	...

MATERIAL	Vertical	Horizontal	ACTUAL CONTINUOUS WORK WITH NO ALLOWANCE FOR REST OR OTHER STOPS										ALLOWING FOR REST AND OTHER NECESSARY STOPS				
			Time to throw	Time to shovel	Time waiting with back	Time of walk-operation	Total time of complete	Volume of shovel-full	Weight of shovel-full	No. of shovel-fulls per minute	No. cubic yards per hour	No. shovel-fulls per minute	No. cubic yards per hour	No. shovel-fulls per minute	No. cubic yards per hour		
Sand or loam	4	5	0.073	0.011	0.011	0.011	0.104	0.16	16	9.6	3.3	7,200	10	7.4	2.6	7,100	
	6	5	0.073	0.011	0.011	0.011	0.110	0.14	14	8.4	2.7	7,200	10	6.6	2.9	5,500	
	8	5	0.073	0.011	0.011	0.011	0.119	0.13	13	7.6	2.3	5,100	10	6.0	1.4	3,940	
	10	5	0.073	0.011	0.011	0.011	0.126	0.11	11	7.6	1.9	5,100	10	6.0	1.1	3,500	
	12	5	0.073	0.011	0.011	0.011	0.134	0.10	10	7.6	1.7	5,100	10	6.0	1.0	3,400	
	14	5	0.073	0.011	0.011	0.011	0.142	0.09	9	7.6	1.5	4,500	10	5.9	0.9	3,300	
	16	5	0.073	0.011	0.011	0.011	0.150	0.08	8	7.6	1.3	3,900	10	5.7	0.8	3,200	
	18	5	0.073	0.011	0.011	0.011	0.158	0.07	7	7.6	1.1	3,300	10	5.5	0.7	2,900	
	20	5	0.073	0.011	0.011	0.011	0.166	0.06	6	7.6	0.9	2,700	10	5.3	0.6	2,600	
	30	5	0.073	0.011	0.011	0.011	0.200	0.05	5	7.6	0.6	2,100	10	5.0	0.4	2,100	
Loam, gravelly	4	5	0.092	0.011	0.011	0.011	0.133	0.14	14	10.0	3.7	8,100	10	6.7	1.8	6,800	
	6	5	0.092	0.011	0.011	0.011	0.148	0.10	10	10.0	2.7	6,000	10	6.0	1.2	5,100	
	8	5	0.092	0.011	0.011	0.011	0.163	0.08	8	10.0	1.9	5,100	10	5.9	1.0	4,200	
	10	5	0.092	0.011	0.011	0.011	0.178	0.07	7	10.0	1.5	4,500	10	5.7	0.9	4,100	
	12	5	0.092	0.011	0.011	0.011	0.193	0.06	6	10.0	1.1	3,900	10	5.5	0.8	3,800	
	14	5	0.092	0.011	0.011	0.011	0.208	0.05	5	10.0	0.9	3,300	10	5.3	0.7	3,400	
	16	5	0.092	0.011	0.011	0.011	0.223	0.04	4	10.0	0.7	2,700	10	5.1	0.6	3,100	
	18	5	0.092	0.011	0.011	0.011	0.238	0.03	3	10.0	0.5	2,100	10	4.9	0.5	2,600	
	20	5	0.092	0.011	0.011	0.011	0.253	0.02	2	10.0	0.3	1,500	10	4.7	0.4	2,100	
	30	5	0.092	0.011	0.011	0.011	0.320	0.01	1	10.0	0.1	500	10	4.5	0.2	1,100	
Gravel, medium	4	5	0.084	0.011	0.011	0.011	0.115	0.12	12.0	8.7	2.8	8,510	10	6.7	1.8	6,800	
	6	5	0.084	0.011	0.011	0.011	0.127	0.10	10.0	14.2	7.9	1.8	6,720	10	6.5	1.7	5,710
	8	5	0.084	0.011	0.011	0.011	0.140	0.08	8.0	11.4	7.1	1.3	4,850	10	5.5	1.0	3,740
	10	5	0.084	0.011	0.011	0.011	0.152	0.07	7.0	10.0	6.0	1.0	3,750	10	5.0	0.9	3,070
	12	5	0.084	0.011	0.011	0.011	0.164	0.06	6.0	9.0	5.0	0.9	3,300	10	4.9	0.8	2,700
	14	5	0.084	0.011	0.011	0.011	0.176	0.05	5.0	8.0	4.0	0.8	3,000	10	4.8	0.7	2,600
	16	5	0.084	0.011	0.011	0.011	0.188	0.04	4.0	7.0	3.0	0.7	2,700	10	4.6	0.6	2,400
	18	5	0.084	0.011	0.011	0.011	0.200	0.03	3.0	6.0	2.0	0.6	2,400	10	4.5	0.5	2,200
	20	5	0.084	0.011	0.011	0.011	0.212	0.02	2.0	5.0	1.0	0.5	2,100	10	4.4	0.4	2,110
	30	5	0.084	0.011	0.011	0.011	0.264	0.01	1.0	4.0	0.2	0.3	1,100	10	4.2	0.3	1,100

Data sheet used in time study to analyze performance of workers and to compare different approaches.

Detailed analysis of shoveling earth data

Fig. 1. Analyzing event data to improve operational processes is not that new. This is illustrated by some of the tables in [33]. Frederick Winslow Taylor can be seen as the first “process miner” using manually collected event data.

The desire to build computer models that mimic organizations and processes is also not that new. Since the 1960-ties so-called *discrete event simulation* tools have been available with SIMULA [11] as one of the first influential examples. In discrete event simulation it is common to estimate parameters and distributions based on observed data (e.g., service times and arrival rates). However, one still needs to model the process by hand. The first comprehensive approaches to automatically learn complete simulation models from event data became available around 2008 [30,31]. Based on event logs, it is possible to learn a control-flow model (transition system, Petri net, of BPMN model) that is enriched with information about resources, data, and time using replay or alignment techniques [30,31].




Fig. 2. Process mining provides a concrete approach to create a digital twin of an organization and its operational processes. A key element is the creation of a model based on event data that is able to mimic reality as well as possible. Such a model needs to be able to capture concurrency and interacting objects (customers, workers, products, orders, payments, shipments, etc.).

The notion of a *digital twin* is part of the Industry 4.0 development facilitated through advanced data analytics (machine learning, process mining, etc.) and the Internet of Things (IoT) connectivity [15,20]. The notion can be described

as an effortless integration of the “real reality” and a “modeled reality” in both directions. The “modeled reality” is based on the “real reality”, but may also influence the “real reality”. This is one of the key concepts in the *Internet of Production* (IoP) developed at RWTH Aachen University [6]. In IoP, *process mining* plays a key role. Gartner coined the term *digital twin of an organization* to indicate that the desire to create a digital twin is not limited to specific Industry 4.0 applications [19]. The goal is to create a virtual representation of an organization and its operational processes (including assets such as architectures, infrastructures, roles, responsibilities, products, etc.) to assess the impact of change in a controlled environment. Note that this is only a vision that is still far away from reality. However, it illustrates the role that models will need to play in the future.




Fig. 3. The difference between a digital model, a digital shadow, and a digital twin.

Figure 3 illustrates the difference between (a) a digital model, (b) a digital shadow, and (c) a digital twin. Building a discrete event simulation model (using e.g. Arena, AnyLogic, CPN Tools, FlexSim, Vensim, or Simul8) in a classical way corresponds to the *digital model* notion in Fig. 3(a). The dashed lines show that the model is created by hand. There is no automated connection between reality and the digital model. Moreover, insights generated by the simulation model do not automatically lead to concrete actions. The *digital shadow* notion in Fig. 3(b) uses a model, driven by data automatically extracted from reality. If such a connection is automated, it is often possible and desirable to update the model continuously. If reality changes, also the model changes. However, insights and diagnostics still need to be translated into actions manually. The *digital twin* notion in Fig. 3(c) shows that there is an automated and real-time connection between reality and the model(s) in both directions. As a result, the digital twin directly influences reality, possibly without human intervention.

It is important to note that many of these ideas have been realized in the context of process mining, albeit with a focus on individual processes in isolation [1]. Most process mining techniques aim to create a digital shadow, as indicated in Fig. 3(b). This ranges from control-flow discovery (from mining directly follows

graphs [2] to scalable inductive mining [23]) to automatically creating simulation models (e.g., [30,31]). However, under the umbrella term of “operational support” [1], process mining also aims to impact the process automatically in real-time. An early example is the work presented in [32], where workflow technology is connected to process mining. In [32] YAWL is used as a workflow management system, ProM as a process mining system, and CPN Tools as the simulation engine. ProM is used to learn a faithful simulation model from the event data of YAWL and/or the models in YAWL. At any point in time, the current state of the YAWL workflow system can be loaded into the simulation model and simulated using CPN Tools. This concept is termed *short-term simulation* because rather than focusing on the steady-state behavior, the focus is on transient behaviors and answering what-if questions. Commercial process mining tools increasingly support what we call “action-oriented process mining”. This means that diagnostics are turned into actions. The recent release of the Celonis EMS (Execution Management System), which embeds a low-code workflow management system, illustrates this trend.

The above shows that the idea of creating a digital twin was already realized in the field of process mining long before the term became “in vogue”. However, existing approaches typically focus on well-defined processes that are considered in isolation. We are still far away from creating a realistic “digital twin of an organization”. In this paper, we focus on two of the many challenges to create such digital twins:

- *Concurrency*. Organizations are like distributed systems or social systems. The different parts operate autonomously but need to synchronize at selected points in time. Although most organizations and systems are highly concurrent, the dominant paradigm is still the highly sequential Turing machine model created in 1936 which does not allow for concurrency. The von Neumann architecture defined in 1945 is based on the Turing machine and also views computation as a sequential process. Moreover, automata, transition systems, Markov chains, and many other representations of behavior do not support concurrency. If concurrency is supported, it is often added as an afterthought. Representations that start from concurrency, like Petri nets, are still the exception. Consider for example a Petri net without places and just transitions. Even people familiar with Petri nets have difficulties to accept that such a Petri net allows for any behavior (and that Petri nets are much more declarative than commonly assumed). Although organizations are highly concurrent, event logs are viewed as sequential (i.e., events are assumed to be totally ordered). This complicates the creation of a digital twin from event data.
- *Object-centricity*. Most of the modeling notations used (e.g., BPMN, Workflow Nets, UML activity diagrams, etc.) assume a single case notion. However, events may involve a variety of objects. Consider for example batching where in one event many objects are affected or an assembly step where a collection of objects is transformed into a new composite object. When drawing for example a BPMN model one needs to pick one case notion (the process

instance). In many applications this is not so easy. Consider for example the hiring process of new employees. Is the vacancy the case or the application? One can also consider the classical example of ordering books from Amazon. One order may include multiple books, a shipment may contain books of different orders, and an order may involved multiple shipments. Possible case notions are order, book, and shipment. It is impossible to create a digital twin of an organization without being able to represent the different objects and their interactions.

For example, imagine a car factory producing hundreds of cars per day with each car assembled from thousands of components. Process models that do not allow for concurrency and object-centricity are clearly unable to describe such a factory as a digital twin.

The remainder of this paper is organized as follows. Section 2 present a short high-level introduction to process mining. Section 3 discusses event logs and the importance of concurrency and object-centricity. Section 4 concludes this short paper.

2 Process Mining: A Top-Down View

In recent years, we could witness an uptake in process mining. There used to be a gap between *process science* (i.e., tools and techniques to improve operational processes) and *data science* (i.e., tools and techniques to extract value from data). Mainstream machine learning and data mining techniques do not consider operational processes. Business Process Management (BPM) and Operations Research (OR) tend to start from models rather than data. Process mining bridges this gap [1].

Currently, there are over 35 commercial process mining vendors (ABBYY Timeline, ARIS Process Mining, BusinessOptix, Celonis Process Mining, Disco/Fluxicon, Everflow, Lana, Mavim, MPM, Minit, PAFnow, QPR, etc.) and process mining is applied in most of the larger organizations in countries such as Germany and The Netherlands. Example application domains include: finance (Rabobank, Hypovereinsbank, etc.), telecom (Deutsche Telekom, Vodafone, etc.), logistics (Vanderlande, etc.), production (BMW, Siemens, Fiat, Bosch, etc.), food (Edeka, etc.), fashion (Zalando, etc.), energy (E-on, etc.), transport (Uber, DB, Lufthansa, etc.), healthcare (AstraZenica, Medtronic, etc.), consulting (Deloitte, EY, KPMG, etc.), and IT systems (Dell, IBM, ServiceNow, etc.).

Figure 4 shows a high-level overview of process mining. Event data need to be extracted from information systems. Such data can be explored, filtered, and cleaned. Process discovery tools transform event data into process models (e.g., BPMN, Petri nets, and UML activity diagrams). There are simple approaches like creating so-called Directly-Follows-Graphs (DFGs) that do not discover concurrency thus having obvious problems [2]. The Alpha algorithm was the first to discover concurrent processes [7]. This approach provides some guarantees, but most processes do not satisfy the assumptions described in [7]. After

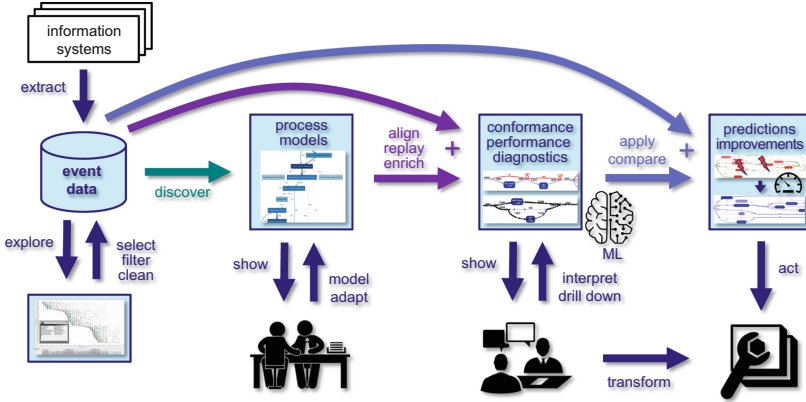


Fig. 4. Overview of the process mining pipeline.

the Alpha algorithm, dozens of more sophisticated algorithms were proposed [1, 9, 21–23, 34]. Using replay and alignment techniques it is possible to relate process models (hand-made or discovered) with event data. This can be used to discover differences between reality and model [1, 10, 29]. Moreover, the model can be extended with additional perspectives, e.g., organizational aspects, decisions, and temporal aspects. This way, detailed performance analyses are possible. Root-cause analysis can be performed for both conformance and performance problems. It is always possible to relate observations to the original event data. Such evidence-based diagnostics aid discussions about root-causes and possible improvements. The right-hand side of Fig. 4 refers to forward-looking techniques aimed at improving the processes. Process models extended with additional perspectives (organizational aspects, decisions, and temporal aspects) can be used to predict conformance and performance problems. As described in [1], predictions can be used to generate recommendations. Figure 4 shows that Machine Learning (ML) techniques can be used in this step. These may range from novel deep learning approaches (e.g., artificial recurrent neural networks like LSTM) to more traditional approaches like logistic regression and decision-tree learning.

It should be noted that process mining techniques are different from mainstream Machine Learning (ML) techniques. However, as Fig. 4 shows, process mining can be used to generate ML problems. The current trend is to make process mining techniques more action-oriented, e.g., automatically trigger a corrective workflow when a problem emerges.

The process mining manifesto [17] published a decade ago lists 11 challenges. Most of these challenges still exist and are still relevant. Figure 5 maps eight challenges onto the overview used before (Fig. 4). These are partly overlapping with the challenges listed in [17], e.g., basic tasks like data extraction and process discovery remain challenging. The reader interested in applications of process mining is recommended to read [26] with experience reports from Siemens, BMW, Uber, ABB, Bayer, and several other organizations.

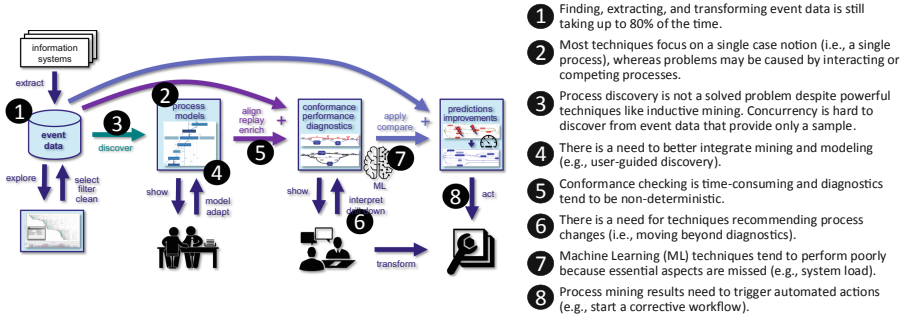


Fig. 5. Some of the challenges encountered when applying process mining.

The top-left corner and bottom-right corner show the interface between the real systems and organization on the one hand and process mining technology on the other hand. These correspond to the solid arrows in Fig. 3(c) used to explain the notion of a digital twin. Using state-of-the-art process mining tools it is possible to create a digital twin with limited scope (e.g., a single process). Process mining is probably the most concrete technology available to create digital twins. Most of the proposals are merely visions or application specific.

3 Process Mining: A Bottom-Up View

After providing a high-level view on process mining, we focus on concurrency and object-centricity. These are essential to create digital twins that properly reflect real organizations. To illustrate these concepts, we use Petri nets. However, it is good to note that the ideas are generic and not notation-specific.

3.1 Petri Nets

Figure 6 shows an accepting labeled Petri net eight places (p_1, p_2, \dots, p_8) and seven transitions (t_1, t_2, \dots, t_7) with initial marking $[p_1]$ and final marking $[p_8]$. We assume that the reader is familiar with the semantics of Petri nets [5, 13, 25, 27, 28]. However, to make the paper more self-contained, we informally explain the behavior of an accepting labeled Petri net. A transition t is *enabled* in a marking if each of its input places contains at least one token. An enabled transition t may *fire*, i.e., one token is removed from each of the input places $\bullet t$ and one token is produced for each of the output places $t \bullet$. This way the Petri net can move from one marking to the next. For example, in the marking shown in Fig. 6 (with a token in p_1) only t_1 is enabled. Firing t_1 means the consumption of one token and the production of three tokens. The resulting marking is $[p_2, p_3, p_4]$. In this marking, four transitions are enabled: There is a choice between t_5 or t_6 because both compete for the token in p_4 . The ordering of t_2 , t_4 , and t_5 or t_6 is not fixed. The transitions in Fig. 6 are labeled, e.g.,

executing t_6 correspond to taking an X-ray. Moreover, next to the initial marking indicated by the black token in place p_1 , there is also a final target marking with just a token in p_8 . We are interested in firing sequences leading from $[p_1]$ to $[p_8]$. Three examples are: $\sigma_1 = \langle t_1, t_2, t_4, t_5, t_7 \rangle$, $\sigma_2 = \langle t_1, t_2, t_4, t_6, t_7 \rangle$, and $\sigma_3 = \langle t_1, t_2, t_3, t_2, t_3, t_2, t_4, t_6, t_7 \rangle$. There are infinitely many firing sequences due to the loop. If we block the loop and do not fire transition t_3 , there are 12 possible firing sequences.

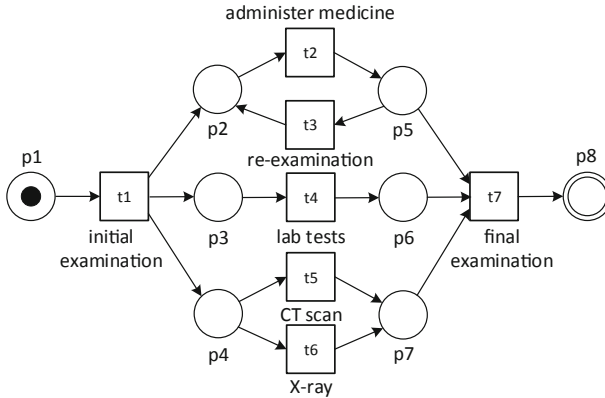


Fig. 6. An accepting labeled Petri net eight places (p_1, p_2, \dots, p_8) and seven transitions (t_1, t_2, \dots, t_7).

Figure 7 shows three example *runs* of the accepting labeled Petri net. Places in Fig. 7 correspond to tokens in Fig. 6, and transitions in Fig. 7 correspond to transition firings in Fig. 6. A run of a Petri net corresponds to a partial order. For example, r_1 in Fig. 7 does not impose an ordering on the three middle activities. The transition labels in Fig. 7 refer to the transitions in Fig. 6, e.g., t_{21} , t_{22} , and t_{23} in run r_3 refer to transition t_2 (administer medicine). For a formal definition of the runs of a Petri net, we again refer to standard literature [12, 14, 27]. Typically, the number of runs is much smaller than the number of firing sequences. For example, if we block the loop and do not fire transition t_3 , then there are only two runs (r_1 and r_2) whereas there where 12 possible firing sequences (e.g., σ_1 is one of the six firing sequences corresponding to run r_1). Run r_3 corresponds $7 * 6 = 42$ firing sequences.

The fact that run r_3 corresponds to 42 firing sequences illustrates the challenge of discovering concurrency. If we assume that t_3 is executed at most 5 times, then there are $2(1 + 1 + 1 + 1 + 1 + 1) = 12$ runs and $2(13 * 12 + 11 * 10 + 9 * 8 + 7 * 6 + 5 * 4 + 3 * 2) = 812$ firing sequences. Even when our event log has information about thousands of traces, it is extremely unlikely that one can witness all 812 variants (especially when not all variants have an equal probability). This illustrates that one cannot ignore concurrency, because it will lead to an explosion of possible interleavings of which just a fraction will be witnessed.

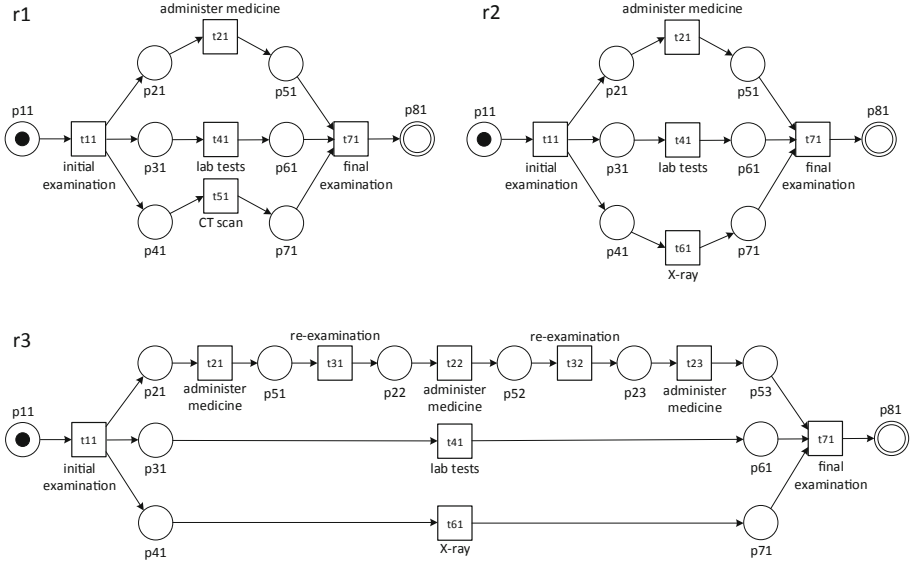


Fig. 7. Three example runs of the accepting Petri net: $r1$, $r2$, and $r3$. Run $r3$ corresponds to 42 firing sequences.

3.2 Object-Centric Partially-Ordered Event Logs

Next to the problem of concurrency, we also need to deal with events referring to collections of objects. This is analogous to moving from a classical Petri net to a Colored Petri Net (CPN) [5, 18]. In a CPN, tokens have values and can present different objects. In a classical Petri net, tokens are indistinguishable and transitions cannot consume or produce a variable number of tokens.

Techniques to discover Petri nets from event data assume precisely one case identifier per event [3, 4]. These case identifiers are used to correlate events, and the resulting discovered Petri net aims to describe the life-cycle of individual cases. In reality, there are often multiple intertwined case notions, and it is impossible to pick a single case notion to describe the whole process. For example, events may refer to mixtures of orders, items, packages, customers, and products. A package may refer to multiple items, multiple products, one order, and one customer. Therefore, we need to assume that each event refers to a collection of objects, each having a type (instead of a single case identifier). Such *object-centric event logs* are closer to data in real-life information systems (e.g., SAP, Salesforce, Oracle, etc.). From an object-centric event log, we want to discover an *object-centric Petri net* with places that correspond to object types and transitions that may consume and produce collections of objects of different types. Such object-centric Petri nets visualize the complex relationships among objects of different types.

In the remainder, we present object-centric event logs as defined in [3, 4]. Note that this is a simplified version of the later *OCEL standard* (see ocel-

standard.org) which also adds attributes to objects [16]. OCEL also provides JSON/XML serializations of object-centric event logs and intends to overcome the limitations of the XES standard [8]. Recall that is the official IEEE standard for storing and exchanging event data assuming a single case notion.

Definition 1 (Universes). *We define the following universes (based on [3, 4]):*

- \mathbb{U}_{ei} is the universe of event identifiers,
- \mathbb{U}_{act} is the universe of activity names (also used to label transitions in an accepting Petri net),
- \mathbb{U}_{time} is the universe of timestamps,
- \mathbb{U}_{ot} is the universe of object types (also called classes),
- \mathbb{U}_{oi} is the universe of object identifiers (also called entities),
- $type \in \mathbb{U}_{oi} \rightarrow \mathbb{U}_{ot}$ assigns precisely one type to each object identifier,
- $\mathbb{U}_{omap} = \{omap \in \mathbb{U}_{ot} \not\rightarrow \mathcal{P}(\mathbb{U}_{oi}) \mid \forall_{ot \in dom(omap)} \forall_{oi \in omap(ot)} type(oi) = ot\}$ is the universe of all object mappings indicating which object identifiers are included per type,¹
- \mathbb{U}_{att} is the universe of attribute names,
- \mathbb{U}_{val} is the universe of attribute values,
- $\mathbb{U}_{vmap} = \mathbb{U}_{att} \not\rightarrow \mathbb{U}_{val}$ is the universe of value assignments,² and
- $\mathbb{U}_{event} = \mathbb{U}_{ei} \times \mathbb{U}_{act} \times \mathbb{U}_{time} \times \mathbb{U}_{omap} \times \mathbb{U}_{vmap}$ is the universe of events.

An event $e = (ei, act, time, omap, vmap) \in \mathbb{U}_{event}$ is characterized by a unique event identifier ei , the corresponding activity act , the event's timestamp $time$, and two mappings $omap$ and $vmap$ for respectively object references and attribute values.

Definition 2 (Event Projection). *Given $e = (ei, act, time, omap, vmap) \in \mathbb{U}_{event}$, $\pi_{ei}(e) = ei$, $\pi_{act}(e) = act$, $\pi_{time}(e) = time$, $\pi_{omap}(e) = omap$, and $\pi_{vmap}(e) = vmap$.*

$\pi_{omap}(e) \in \mathbb{U}_{ot} \not\rightarrow \mathcal{P}(\mathbb{U}_{oi})$ maps a subset of object types onto sets of object identifiers for an event e . An *object-centric event log* is a collection of *partially ordered events*. Event identifiers are unique, i.e., two events cannot have the same event identifier.

Definition 3 (Object-Centric Event Log). *$L = (E, \preceq_E)$ is an event log with $E \subseteq \mathbb{U}_{event}$ and $\preceq_E \subseteq E \times E$ such that:*

- \preceq_E defines a partial order (reflexive, antisymmetric, and transitive),
- $\forall_{e_1, e_2 \in E} \pi_{ei}(e_1) = \pi_{ei}(e_2) \Rightarrow e_1 = e_2$, and
- $\forall_{e_1, e_2 \in E} e_1 \preceq_E e_2 \Rightarrow \pi_{time}(e_1) \leq \pi_{time}(e_2)$.

¹ $\mathcal{P}(\mathbb{U}_{oi})$ is the powerset of the universe of object identifiers, i.e., object types are mapped onto sets of object identifiers. $omap \in \mathbb{U}_{ot} \not\rightarrow \mathcal{P}(\mathbb{U}_{oi})$ is a partial function. If $ot \notin dom(omap)$, then we assume that $omap(ot) = \emptyset$.

² $\mathbb{U}_{att} \not\rightarrow \mathbb{U}_{val}$ is the set of all partial functions mapping a subset of attribute names onto the corresponding values.

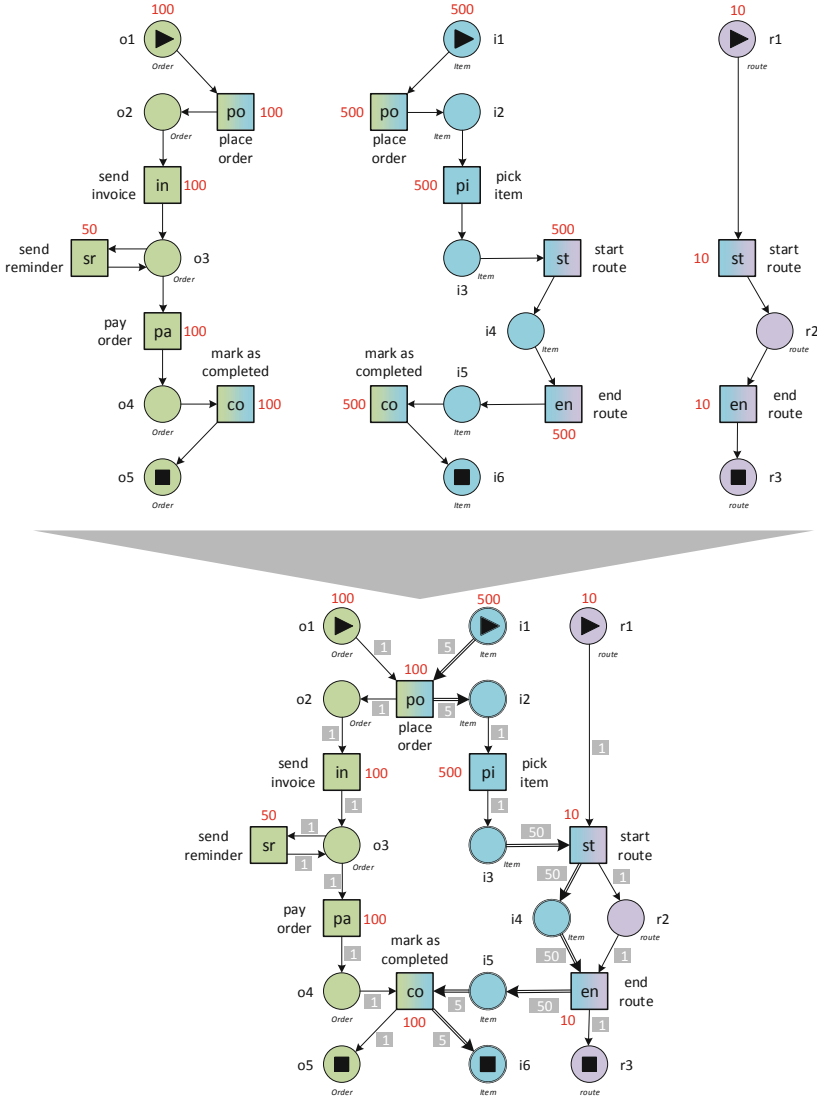


Fig. 8. An Object-Centric Petri Net (OCPN) can be learned by first learning a classical Petri net per object type and then merging the nets while correcting the multiplicities using variable arcs (a detailed derivation of the process model was presented in [4]).

Definition 3 allows for partially ordered event logs. Many process mining techniques require a total order, e.g., events are ordered based on timestamps and when two events have the same timestamp we assume some order. However, there are process discovery techniques that take into account causalities [3,24]. These can exploit such partial orders. There may be many reasons to use partially

ordered event logs: efficiency, imprecise timing information, uncertainty, and explicit partial order information (e.g., based on data flow analysis). As argued before, it is unreasonable to assume that all possible interleavings will indeed be present in the event log. Instead of a partial order one can also use the stricter notion of a weak order. This is particularly suitable when one has imprecise timestamps (e.g., events on the same day cannot be ordered).

3.3 Object-Centric Petri Nets

In this paper, we argued that concurrency and objects matter. To progress the field of process mining, we cannot assume that events are totally ordered and can be correlated using a single case notion. Hence, we need process mining techniques and process model representations handling concurrency and object-centricity as first-class citizens. In [4], we presented an approach to automatically learn a so-called *Object-Centric Petri Net* (OCPN) given an object-centric event log (e.g., in OCEL format [16]). A detailed explanation of the approach to discover OCPNs is beyond the scope of this short paper. Therefore, we only show the example depicted in Fig. 8.

Let $L = (E, \preceq_E)$ be an event log. The events in E refer to objects. Therefore, given a specific object o of type ot , it is possible to create a partial order of all events that refer to o . (E_o, \preceq_{E_o}) , with $E_o = \{e \in E \mid o \in \pi_{omap}(e)(ot)\}$ and $\preceq_{E_o} = \preceq_E \cap (E_o \times E_o)$, defines the corresponding partial order. Hence, we can group all partial orders of events corresponding to objects of a given type ot to get the required input for a standard process discovery algorithm. Note that the same event may appear in multiple partial orders. Next, we can learn a process model per object type. For simplicity, we assume that we discover a labeled accepting Petri net per object type satisfying the constraint that labels of visible transition are unique. There may be silent transitions (i.e., transitions that do not refer to an activity). However, there cannot be two transitions referring to the same activity.

The top part of Fig. 8 shows three labeled accepting Petri nets discovered for 100 orders, 500 items, and 10 routes. These three models happen to be sequential, but could have been concurrent. The initial and final markings are denoted by the places with the play and stop symbol. Next, the three labeled accepting Petri nets are merged into an OCPN. Since the visible transitions are unique, merging is trivial. However, the annotations need to be modified. In an OCPN there is a one-to-one correspondence between transition firings and events. A single event (i.e., transition occurrence) may involve a *variable number of objects* (e.g., one order may have any number of items). This is indicated by the double arcs in the lower part of Fig. 8. For example, on average one execution of *place order* corresponds to five items and one order. On average one execution of *start route* corresponds to 50 items and one route. For more details, we refer to [4].

The discovery Object-Centric Petri Nets (OCPNs) from object-centric event logs in OCEL format is still in its infancy. However, the topic is important because in most applications of process mining one faces the problem of one-to-many and many-to-many relations between different types of objects relevant

for an organization. Processes are intertwined and difficult to separate. Figure 8 shows that it is possible to create one, more holistic, process model that is showing the interactions between the different types of objects. Actually, the term “process model” may be misleading in the context of OCPNs that may represent collections of interacting processes.

4 Conclusion

To create a “digital twin of an organization” we need to disentangle the fabric of real operational processes. Process mining provides many of the ingredients to make such a step. In this paper, we provided a high-level overview of process mining and linked it to historical developments in the field of scientific management and simulation. As shown, there have been early examples of digital twins (or at least digital shadows) in the field of process mining. We mentioned, for example, the work combining the process mining framework ProM, the workflow management system YAWL, and CPN Tools as the simulation engine [32]. This enabled new forms of “short-term simulation” that can be used to see the effects of decisions given the current state and historic information.

However, we are far away from fully capturing the fabric of real operational processes in a single model. An important prerequisite is the proper handling of concurrency and entangled objects. One event may refer to many objects and organizations are highly concurrent. It is unrealistic to assume that one can witness all interleavings of highly concurrent processes. Therefore, we elaborated on Object-Centric Petri Nets (OCPNs) and OCEL as a format for exchanging object-centric event logs [16].

Future research needs to address the challenges described in this paper. Compared to the pen-and-paper analyses done by Frederick Winslow Taylor and colleagues more than a century ago, we have booked tremendous progress. The detailed event data available today provide unprecedented opportunities to create digital twins (provided we are able to concurrency and object-centricity properly).

Acknowledgments. The author thanks the Alexander von Humboldt (AvH) Stiftung for supporting his research.

References

1. van der Aalst, W.M.P.: Data science in action. In: Process Mining. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49851-4_1
2. van der Aalst, W.M.P.: A practitioner’s guide to process mining: limitations of the directly-follows graph. In: International Conference on Enterprise Information Systems (Centeris 2019). Procedia Computer Science, vol. 164, pp. 321–328. Elsevier (2019)
3. van der Aalst, W.M.P.: Object-centric process mining: dealing with divergence and convergence in event data. In: Ölveczky, P.C., Salaün, G. (eds.) SEFM 2019. LNCS, vol. 11724, pp. 3–25. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30446-1_1

4. van der Aalst, W.M.P., Berti, A.: Discovering object-centric Petri nets. *Fund. Inform.* **175**(1–4), 1–40 (2020)
5. van der Aalst, W.M.P., Stahl, C.: *Modeling Business Processes: A Petri Net Oriented Approach*. MIT Press, Cambridge (2011)
6. van der Aalst, W.M.P., Brockhoff, T., Ghahfarokhi, A.F., Pourbafrani, M., Uysal, M.S., van Zelst, S.J.: Removing operational friction using process mining: challenges provided by the Internet of Production (IoP). In: Hammoudi, S., Quix, C., Bernardino, J. (eds.) *DATA 2020. CCIS*, vol. 1446, pp. 1–31. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-83014-4_1
7. van der Aalst, W.M.P., Weijters, A.J.M.M., Maruster, L.: Workflow mining: discovering process models from event logs. *IEEE Trans. Knowl. Data Eng.* **16**(9), 1128–1142 (2004)
8. Acampora, G., Vitiello, A., Di Stefano, B., van der Aalst, W., Günther, C., Verbeek, E.: IEEE 1849: the XES standard - the second IEEE standard sponsored by IEEE Computational Intelligence Society. *IEEE Comput. Intell. Mag.* **12**(2), 4–8 (2017)
9. Augusto, A., Conforti, R., Marlon, M., La Rosa, M., Polyvyanyy, A.: Split miner: automated discovery of accurate and simple business process models from event logs. *Knowl. Inf. Syst.* **59**(2), 251–284 (2019)
10. Carmona, J., van Dongen, B., Solti, A., Weidlich, M.: *Conformance Checking: Relating Processes and Models*. Springer, Cham (2018). <https://doi.org/10.1007/978-3-319-99414-7>
11. Dahl, O.J., Nygaard, K.: SIMULA: an ALGOL based simulation language. *Commun. ACM* **1**, 671–678 (1966)
12. Desel, J.: Validation of process models by construction of process nets. In: van der Aalst, W., Desel, J., Oberweis, A. (eds.) *Business Process Management. LNCS*, vol. 1806, pp. 110–128. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45594-9_8
13. Desel, J., Esparza, J.: *Free Choice Petri Nets. Cambridge Tracts in Theoretical Computer Science*, vol. 40. Cambridge University Press, Cambridge (1995)
14. van Dongen, B.F., Desel, J., van der Aalst, W.M.P.: Aggregating causal runs into workflow nets. In: Jensen, K., van der Aalst, W.M., Ajmone Marsan, M., Franceschinis, G., Kleijn, J., Kristensen, L.M. (eds.) *Transactions on Petri Nets and Other Models of Concurrency VI. LNCS*, vol. 7400, pp. 334–363. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35179-2_14
15. Fuller, A., Fan, Z., Day, C., Barlow, C.: Digital twin: enabling technologies, challenges and open research. *IEEE Access* **8**, 108952–108971 (2020)
16. Ghahfarokhi, A.F., Park, G., Berti, A., van der Aalst, W.M.P.: *OCEL Standard* (2021). www.ocel-standard.org
17. van der Aalst, W., et al.: Process mining manifesto. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) *BPM 2011. LNBIP*, vol. 99, pp. 169–194. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28108-2_19
18. Jensen, K.: *Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use. Monographs in Theoretical Computer Science An EATCS Series*, Springer, Heidelberg (1997). <https://doi.org/10.1007/978-3-642-60794-3>
19. Kerremans, M., Kopcho, J.: Create a digital twin of your organization to optimize your digital transformation program. *Research Note G00379226* (2019). www.gartner.com
20. Kritzinger, W., Karner, M., Traar, G., Henjes, J., Sihm, W.: Digital twin in manufacturing: a categorical literature review and classification. *IFAC-PapersOnLine* **51**(11), 1016–1022 (2018). 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018

21. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs - a constructive approach. In: Colom, J.-M., Desel, J. (eds.) *PETRI NETS 2013*. LNCS, vol. 7927, pp. 311–329. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38697-8_17
22. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs containing infrequent behaviour. In: Lohmann, N., Song, M., Wohed, P. (eds.) *BPM 2013*. LNBIP, vol. 171, pp. 66–78. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06257-0_6
23. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Scalable process discovery and conformance checking. *Softw. Syst. Model.* **17**(2), 599–631 (2018)
24. Lu, X., Fahland, D., van der Aalst, W.M.P.: Conformance checking based on partially ordered event data. In: Fournier, F., Mendling, J. (eds.) *BPM 2014*. LNBIP, vol. 202, pp. 75–88. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-15895-2_7
25. Murata, T.: Petri nets: properties, analysis and applications. *Proc. IEEE* **77**(4), 541–580 (1989)
26. Reinkemeyer, L.: *Process Mining in Action: Principles, Use Cases and Outlook*. Springer, Berlin (2020). <https://doi.org/10.1007/978-3-030-40172-6>
27. Reisig, W.: *Petri Nets: Modeling Techniques, Analysis, Methods, Case Studies*. Springer, Berlin (2013). <https://doi.org/10.1007/978-3-642-33278-4>
28. Reisig, W., Rozenberg, G. (eds.): *Lectures on Petri Nets I: Basic Models*. Lecture Notes in Computer Science, vol. 1491. Springer, Berlin (1998). <https://doi.org/10.1007/3-540-65306-6>
29. Rozinat, A., van der Aalst, W.M.P.: Conformance checking of processes based on monitoring real behavior. *Inf. Syst.* **33**(1), 64–95 (2008)
30. Rozinat, A., Mans, R.S., Song, M., van der Aalst, W.M.P.: Discovering colored petri nets from event logs. *Int. J. Softw. Tools Technol. Transf.* **10**(1), 57–74 (2008)
31. Rozinat, A., Mans, R.S., Song, M., van der Aalst, W.M.P.: Discovering simulation models. *Inf. Syst.* **34**(3), 305–327 (2009)
32. Rozinat, A., Wynn, M., van der Aalst, W.M.P., ter Hofstede, A.H.M., Fidge, C.: Workflow simulation for operational decision support. *Data Knowl. Eng.* **68**(9), 834–850 (2009)
33. Taylor, F.W.: *The Principles of Scientific Management*. Harper and Brothers Publishers, New York (1919)
34. van Zelst, S.J., van Dongen, B.F., van der Aalst, W.M.P., Verbeek, H.M.W.: Discovering workflow nets using integer linear programming. *Computing* **100**(5), 529–556 (2018)