A Decade of Business Process Management Conferences: Personal Reflections on a Developing Discipline

Wil M.P. van der Aalst

Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven, The Netherlands.
www.vdaalst.com

Abstract. The Business Process Management (BPM) conference series celebrates its tenth anniversary. This is a nice opportunity to reflect on a decade of BPM research. This paper describes the history of the conference series, enumerates twenty typical BPM use cases, and identifies six key BPM concerns: process modeling languages, process enactment infrastructures, process model analysis, process mining, process flexibility, and process reuse. Although BPM matured as a research discipline, there are still various important open problems. Moreover, despite the broad interest in BPM, the adoption of state-of-the-art results by software vendors, consultants, and end-users leaves much to be desired. Hence, the BPM discipline should not shy away from the key challenges and set clear targets for the next decade.

1 History of the BPM Conference Series

The first International Conference on Business Process Management (BPM 2003) took place in Eindhoven in the last week of June 2003, ten years ago. Therefore, I would like to take the opportunity to describe the origin of the conference series and look back on a decade of BPM research.

The direct trigger to organize the first BPM conference was generated by Grzegorz Rozenberg in his capacity as chair of the Petri Nets Steering Committee. He invited us to organize the 24th International Conference on Application and Theory of Petri Nets (Petri Nets 2003) in Eindhoven. Moreover, he stimulated me to organize a co-located event. This is how the idea for "BPM 2003" was born. The subtitle of BPM 2003 – "On the Application of Formal Methods to Process-Aware Information Systems" – illustrates the relation with Petri Nets conference. Together with Arthur ter Hofstede and Mathias Weske, I served as PC chair of BPM 2003 [11]. I was also PC chair of Petri Nets 2003 (together with Eike Best) [5]. Kees van Hee, Hajo Reijers, Eric Verbeek, and many others from the Technische Universiteit Eindhoven (TU/e) were involved in the organization of both conferences. BPM 2003 was remarkably successful considering it was organized for the first time: 77 papers were submitted of which 25 papers were accepted. Moreover, various BPM vendors and consultants participated. Carl Adam Petri gave a keynote, received a prestigious Royal medal ("Commandeur in de Orde van de Nederlandse Leeuw"), and it was interesting to see him talking with BPM vendors about workflows. Afterwards, we decided to continue organizing BPM conferences given the growing interest in the topic and enthusiasm of the BPM 2003 participants. We decided to disconnect the BPM conference from the Petri Nets conference (to clearly show that BPM is not linked to a specific formalism). Apparently, these were wise decisions; in subsequent years the BPM conference series evolved into one of the premier information systems conferences. The following list shows the ten BPM conferences organized thus far:

– BPM 2003 (Van der Aalst, Ter Hofstede, Weske, Reijers, Van Hee, et al.), Eindhoven, The Netherlands [11],
– BPM 2004 (Weske, Desel, Pernici, et al.), Potsdam, Germany [19],
– BPM 2005 (Godart, Perrin, Van der Aalst, Benatallah, Casati, Curbera, et al.), Nancy, France [4],
– BPM 2006 (Dustdar, Fiadeiro, Sheth, Rosenberg, et al.), Vienna, Austria [23],
– BPM 2007 (Rosemann, Dumas, Alonso, Dadam, Ter Hofstede, et al.), Brisbane, Australia [14],
– BPM 2008 (Pernici, Casati, Dumas, Reichert, et al.), Milan, Italy [22],
– BPM 2009 (Reichert, Dadam, Reijers, Eder, Dayal, et al.), Ulm, Germany [18],
– BPM 2010 (Zur Muehlen, Hull, Mendling, Tai, et al.), Hoboken, USA [31],
– BPM 2011 (Toumani, Rinderle-Ma, Wolf, Hacid, Schneider, et al.), Clermont-Ferrand, France [36], and
– BPM 2012 (Dumas, Kindler, Gal, Barros, et al.), Tallinn, Estonia [15].

Since 2005 the conference features co-located workshops. The main proceedings are published in Springer's Lecture Notes in Computer Science (LNCS) and the workshop proceedings are published in Springer's Lecture Notes in Business Information Processing (LNBIP). From 2003 to 2009, selected papers were invited for special issues of Data & Knowledge Engineering (DKE). Since 2010, each year the best papers are invited for a special issue of Information Systems (IS).

Although BPM 2003 was the first real BPM conference, there were some informal predecessor workshops. Together with Giorgio De Michelis and Skip Ellis, I organized the "Workflow Management: Net-based Concepts, Models, Techniques and Tools" (WFM'98) workshop [12]. This workshop was co-located with Petri Nets 1998 in Lisbon, Portugal. Together with Jörg Desel and Roland Kaschek, I also organized the "Software Architectures for Business Process Management" (SABPM'99) workshop [6]. This workshop was one of the pre-conference workshops of CAiSE 1999 in Heidelberg, Germany. Based on these events, I started to work with Jörg Desel and Andreas Oberweis on an edited book during my sabbatical at the University of Karlsruhe. In 2000, the book "Business Process Management: Models, Techniques, and Empirical Studies" [7] appeared. This LNCS volume can be seen as a direct predecessor of BPM 2003 given the topic and people involved. Therefore, I will include it in my later analysis.

2 Pre-BPM Era

Business Process Management (BPM) has various roots in both computer science and management science. Therefore, it is difficult to pinpoint the starting point of BPM.

However, it is obvious that BPM existed long before the term became popular. Therefore, I reflect on the origins of BPM by summarizing major developments before the conference in 2003.

Since the industrial revolution, productivity has been increasing because of technical innovations, improvements in the organization of work, and the use of information technology. Adam Smith (1723-1790) showed the advantages of the division of labor. Frederick Taylor (1856-1915) introduced the initial principles of scientific management. Henry Ford (1863-1947) introduced the production line for the mass production of "black T-Fords". It is easy to see that these ideas are used in today's BPM systems.

Around 1950 computers and digital communication infrastructures started to influence business processes. This resulted in dramatic changes in the organization of work and enabled new ways of doing business. Today, innovations in computing and communication are still the main drivers behind change in business processes. So, business processes have become more complex, heavily rely on information systems, and may span multiple organizations. *Therefore, process modeling has become of the utmost importance.* Process models assist in managing complexity by providing insights and by documenting procedures. Information systems need to be configured and driven by precise instructions. Cross-organizational processes can only function properly if there is a common agreement on the required interactions. As a result, process models are widely used in todays organizations.

In the last century many process modeling techniques have been proposed. In fact, the well-known Turing machine described by Alan Turing (1912-1954) can be viewed as a process model. It was instrumental in showing that many questions in computer science are undecidable. Moreover, it added a data component (the tape) to earlier transition systems. Petri nets play an even more prominent role in BPM as they are graphical and able to model concurrency. In fact, most of the contemporary BPM notations and systems use a token-based semantics adopted from Petri nets. Petri nets were proposed by Carl Adam Petri (1926-2010) in 1962. This was the first formalism able to model concurrency. Concurrency is very important as in business processes many things happen in parallel. Many cases may be handled at the same time and even within a case there may be various activities enabled or running concurrently. Therefore, a BPM system should support concurrency natively.

Since the seventies there has been consensus on the modeling of data (cf. the Relational Model by Codd [17] and the Entity-Relationship Model by Chen [16]). Conversely, process modeling is best characterized by the term "divergence". There is little consensus on the fundamental concepts. Despite the availability of established formal languages (e.g., Petri nets and process calculi) industry has been pushing adhoc/domain-specific languages. As a result there is a plethora of systems and languages available today (BPMN, BPEL, UML, EPCs, etc.).

Figure 1 sketches the emergence of BPM systems and their role in the overall information system architecture. Initially, information systems were developed from scratch, i.e., everything had to be programmed, even storing and retrieving data. Soon people realized that many information systems had similar requirements with respect to data management. Therefore, this generic functionality was subcontracted to a database system. Later, generic functionality related to user interaction (forms, buttons, graphs, etc.)

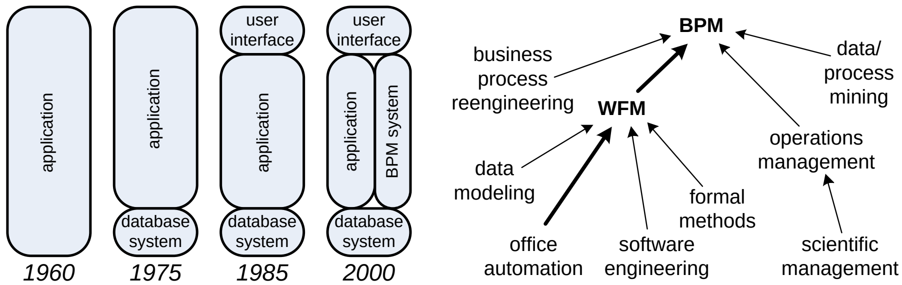

Fig. 1. Historic view on information systems development illustrating that BPM systems can be used to push processes out of the application (left, adapted from [1])) and an overview of some disciplines that contributed to the development of the BPM discipline (right).

was subcontracted to tools that can automatically generate user interfaces. This trend continued in different areas. BPM systems can be seen in this context: a BPM system takes care of process-related aspects. Therefore, the application can focus on supporting individual/specific tasks. In the mid-1990s many *Workflow Management* (WFM) systems became available. These systems focused on automating workflows with little support for analysis, flexibility, and management. BPM systems provide much broader support, e.g., by supporting simulation, business process intelligence, case management, etc. However, compared to the database market, the BPM market is much more diverse and there is no consensus on notations and minimal capabilities. This is not a surprise as process management is much more complex than data management.

A good starting point from a scientific perspective is the early work on *office information systems*. In the seventies, people like Skip Ellis, Anatol Holt, and Michael Zisman already worked on so-called office information systems, which were driven by explicit process models [2, 24, 26, 25, 29, 30, 34, 40, 42, 41]. Ellis et al. [24, 26, 25] developed office automation prototypes such as Officetalk-Zero, Officetalk-D and Officetalk-P at Xerox PARC in the late 1970s. These systems used Information Control Nets (ICN), a variant of Petri nets, to model processes. Office metaphors such as inbox, outbox and forms were used to interact with users. The prototype office automation system SCOOP (System for Computerizing of Office Processes) developed by Michael Zisman also used Petri nets to represent business processes [40, 42, 41]. It is interesting to see that pioneers in office information systems already used Petri-net variants to model office procedures. During the seventies and eighties there was great optimism about the applicability of office information systems. Unfortunately, few applications succeeded. As a result of these experiences, both the application of this technology and research almost stopped for a decade. Consequently, hardly any advances were made in the eighties. In the nineties, there was a clear revival of the ideas already present in the early office automation prototypes [8]. This is illustrated by the many commercial WFM systems developed in this period.

In the mid-nineties there was the expectation that WFM systems would get a role comparable to Database Management (DBM) systems. Most information systems subcontract their data management to DBM systems and there are just a few widely used products. However, despite the availability of BPM/WFM systems, process management is not subcontracted to such systems at a scale comparable to DBM systems. The application of "pure" BPM/WFM systems is still limited to specific industries such as banking and insurance. However, BPM/WFM technology is often hidden inside other systems. For example, ERP systems like SAP and Oracle provide workflow engines. Many other platforms include workflow-like functionality. For example, integration and application infrastructure software such as IBM's WebSphere provides extensive process support. In hindsight, it is easy to see why process management cannot be subcontracted to a standard BPM/WFM system at a scale comparable to DBM systems. As illustrated by the varying support of workflow patterns [9, 37], process management is much more complex than data management. *BPM is multifaceted, complex, and difficult to demarcate.* Given the variety in requirements and close connection to business concerns, it is often impossible to use generic BPM/WFM solutions. Therefore, BPM functionality is often embedded in other systems and BPM techniques are frequently used in a context with conventional information systems.

The first BPM conference in 2003 marked the transition from WFM to BPM [10]. Since then the BPM discipline matured. Today, the relevance of BPM is acknowledged by practitioners (users, managers, analysts, consultants, and software developers) and academics. This is illustrated by the availability of BPM systems, conferences, and books such as [3, 8, 13, 21, 28, 32–35, 39].

3 BPM Use Cases

One of the goals of this paper is to reflect on 10 years of BPM research by analyzing the proceedings of past BPM conferences (BPM 2003 - BPM 2011) and the edited book [7] that can be viewed as a predecessor of the first BPM conference (see Section 1). In total 289 papers were analyzed by tagging each paper with the *use cases* and *key concerns* described in the remainder.

Before conducting this analysis, I identified 20 use cases as shown in Figure 2. For example, use case *design model* (DesM) refers to the creation of a process model from scratch and use case *discover model from event data* (DiscM) refers to the automated generation of a process model using process mining techniques. Models constructed through use case DesM are descriptive (D), normative (N), and/or executable (E). This is denoted by the "$D|N|E$" tag in Figure 2. A model discovered through process mining (DiscM) is typically not normative as it is based on observed behavior (cf. "$D|E$" tag). Models can also be obtained through *selection* (SelM), *merging* (MerM), or *composition* (CompM). Most papers were tagged with one or two use cases. The tagging was based on the most important use case(s) the paper aims to support. For example, the paper "Graph Matching Algorithms for Business Process Model Similarity Search" [20] presented at BPM 2009 was tagged with the use case *select model from collection* (SelM) since the paper presents an approach to rank process models in a repository based on some initial model.

Use cases *design configurable model* (DesCM), *merge models into configurable model* (MerCM), and *configure configurable model* (ConCM) deal with configurable

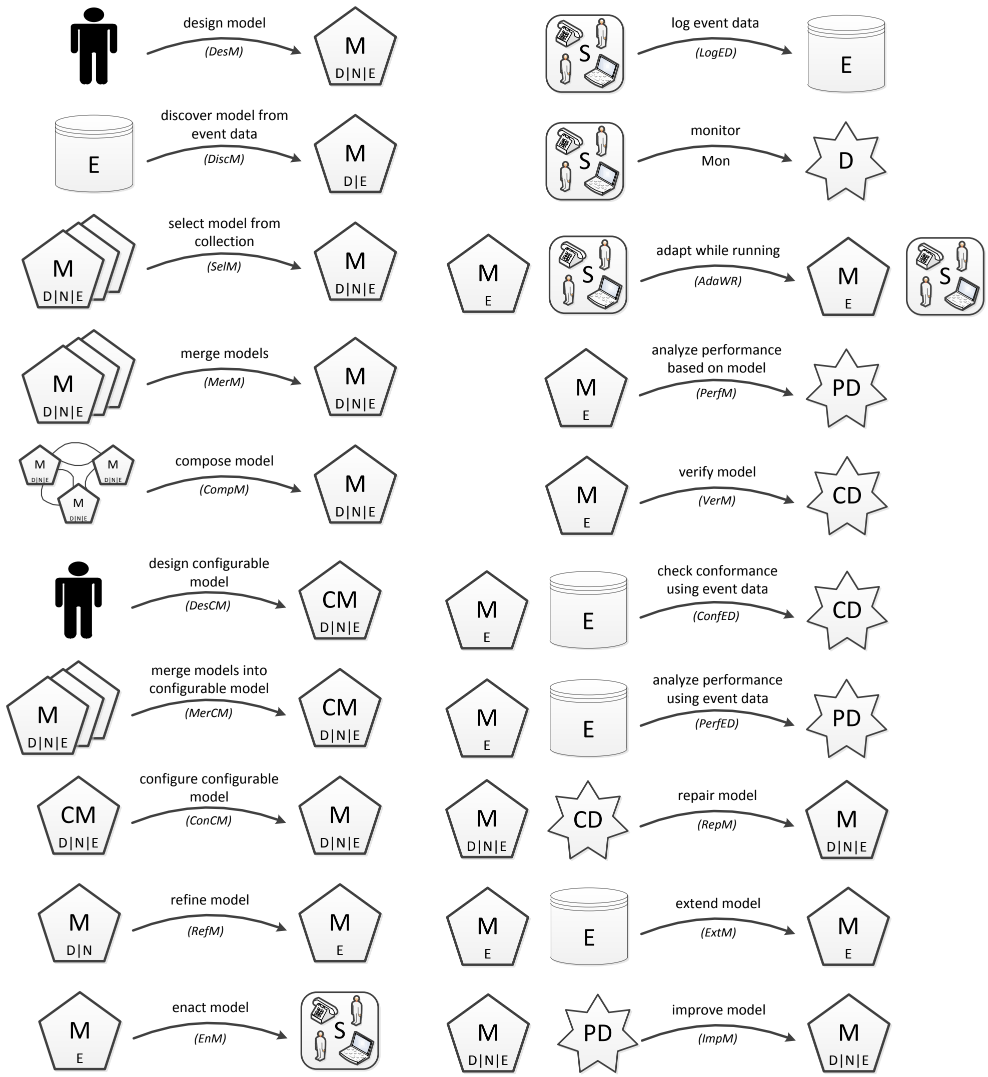

Fig. 2. Twenty BPM use cases: *M* = model, *E* = event data, *CM* = configurable model, *S* = information system, *D* = diagnostics, *CD* = conformance-related diagnostics, *PD* = performance-related diagnostics. Models can be tagged as descriptive (*D*), normative (*N*), or executable (*E*).

models. There models – also referred to as reference models – correspond to families of concrete process models (i.e., variants of the same process). Use case *refine model* (RefM) refers to the transformation of a descriptive or normative model into an executable model. Executable models can be enacted by applying use case *enact model* (EnM).

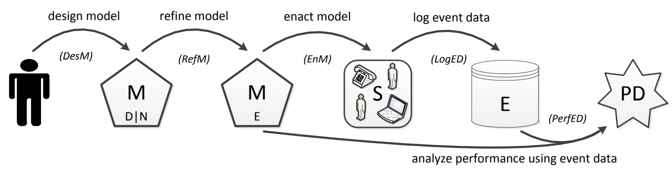

Fig. 3. Composite use case obtained by chaining five use cases: DesM, RefM, EnM, LogED, and PerfED.

Use cases can be chained together as shown in Figure 3. The last use case in the chain is *analyze performance using event data* (PerfED) which requires event data and an executable model. PerfED use may be used to uncover bottlenecks observed in reality. Use case *check conformance using event data* (ConfED) requires similar input but focuses on deviations rather than performance. Use case *analyze performance based on model* (PerfM) focuses on performance without using event data, e.g., model-based simulation used to analyze flow times, utilization, and bottlenecks. Use case *monitor* (Mon) refers runtime analysis without using any model (e.g., flow time analysis without looking "inside the process").

Use case *adapt while running* (AdaWR) refers to changing the system and model at runtime to provide flexibility, e.g., modifying a model and subsequently migrating process instances. The paper "Instantaneous Soundness Checking of Industrial Business Process Models" [27] also presented at BPM 2009 is a typical example of a paper tagged with use case *verify model* (VerM). In this paper 735 industrial business process models are checked for soundness (absence of deadlock and lack of synchronization) using three different approaches.

Use case *repair model* (RepM) changes a model based on conformance-related diagnostics, e.g., deviations are used to correct the model. Use case *improve model* (ImpM) is similar but focuses on performance-related diagnostics. For example, bottleneck analysis is used to redesign the process. Use case *extend model* (ExtM) refers to the process mining scenario where a model is enriched using information extracted from the event log [3].

Papers where typically tagged with one dominant use case, but sometimes more tags were used. In total, 367 tags were assigned (on average 1.18 use cases per paper). By simply counting the number of tags per use case and year, the relative frequency of each use case per year can be established. For example, for BPM 2009 four papers were tagged with use case *discover model from event data* (DiscM). The total number of tags assigned to the 23 BPM 2009 papers is 30. Hence, the relative frequency is $4/30 = 0.133$. Table 1 shows all relative frequencies including the one just mentioned. The table also shows the average relative frequency of each use case over all years. These averages are shown graphically in Figure 4.

Figure 4 shows that use cases *design model* (DesM) and *enact model* (EnM) are most frequent. This is not very surprising as these use cases are less specific than most other use cases. The third most frequent use case – *verify model* (VerM) – is more surprising (relative frequency of 0.144). An example paper having such a tag is [27]

Table 1. Relative importance of use cases in [7], [11], [19], [4], [23], [14], [22], [18], [31], [36], and [15]. Each of the 289 papers was tagged with, on average, 1.18 use cases (typically 1 or 2 use cases per paper). The table shows the relative frequency of each use case per year. The last row shows the average over 10 years. All rows add up to 1.

year	design model DesM	discover model from event data DiscM	select model from collection SelM	merge models MerM	compose model CompM	design configurable model DesCM	merge models into configurable model MerCM	configure configurable model ConCM	refine model RefM	enact model EnM	log event data LogED	monitor Mon	adapt while running AdaWR	analyze performance based on model PerfM	verify model VerM	check conformance using event data ConfED	analyze performance using event data PerfED	repair model RepM	extend model ExtM	improve model ImpM
2000	0.406	0.000	0.000	0.000	0.031	0.000	0.000	0.000	0.000	0.188	0.000	0.000	0.063	0.125	0.188	0.000	0.000	0.000	0.000	0.000
2003	0.306	0.028	0.056	0.000	0.056	0.000	0.000	0.028	0.000	0.222	0.028	0.000	0.139	0.000	0.111	0.000	0.000	0.000	0.028	0.000
2004	0.348	0.130	0.000	0.000	0.043	0.000	0.000	0.000	0.000	0.217	0.000	0.000	0.043	0.087	0.087	0.000	0.000	0.000	0.000	0.043
2005	0.216	0.039	0.039	0.000	0.098	0.000	0.000	0.000	0.000	0.294	0.000	0.000	0.059	0.078	0.137	0.000	0.000	0.000	0.000	0.039
2006	0.094	0.038	0.057	0.019	0.132	0.000	0.000	0.000	0.019	0.245	0.019	0.000	0.075	0.019	0.226	0.019	0.019	0.000	0.000	0.019
2007	0.231	0.128	0.026	0.026	0.051	0.026	0.026	0.077	0.077	0.077	0.026	0.026	0.026	0.051	0.103	0.000	0.026	0.000	0.000	0.000
2008	0.227	0.045	0.023	0.000	0.045	0.000	0.023	0.000	0.023	0.182	0.045	0.000	0.023	0.091	0.136	0.000	0.045	0.023	0.068	0.000
2009	0.167	0.133	0.067	0.000	0.033	0.000	0.033	0.000	0.133	0.133	0.033	0.000	0.000	0.033	0.167	0.033	0.033	0.000	0.000	0.000
2010	0.167	0.133	0.100	0.000	0.000	0.033	0.000	0.033	0.033	0.300	0.000	0.000	0.000	0.000	0.100	0.067	0.000	0.000	0.033	0.000
2011	0.061	0.091	0.121	0.030	0.000	0.000	0.061	0.000	0.000	0.121	0.000	0.061	0.000	0.000	0.182	0.152	0.030	0.061	0.030	0.000
average	0.222	0.077	0.049	0.007	0.049	0.006	0.014	0.014	0.029	0.198	0.015	0.009	0.043	0.048	0.144	0.027	0.015	0.008	0.016	0.010

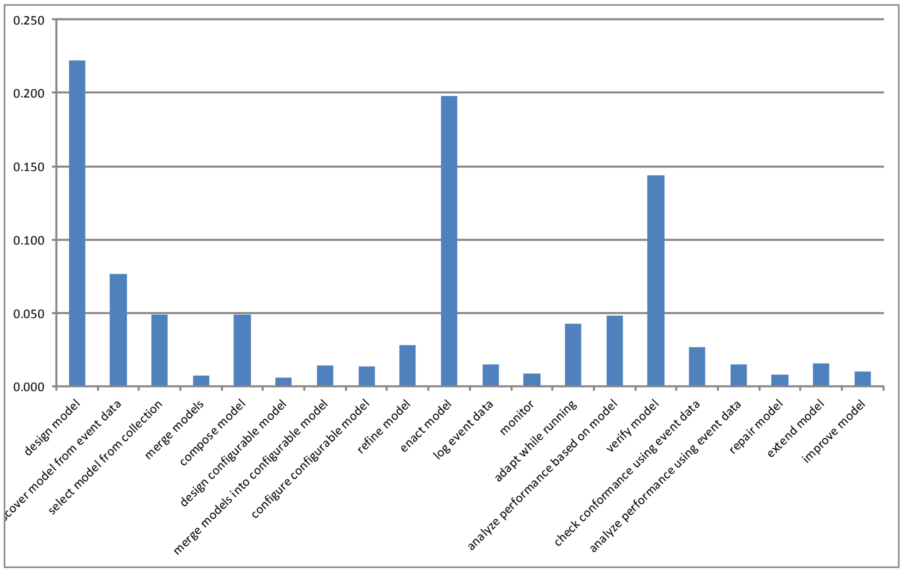

Fig. 4. Average relative importance of use cases (taken from Table 1).

which was mentioned before. Over the last decade there has been considerable progress in this area and this is reflected by various papers presented at BPM. In this context it is remarkable that the use cases *monitor* (Mon) and *analyze performance using event data* (PerfED) have a much lower relative frequency (respectively 0.009 and 0.015). Given the practical needs of BPM one would expect more papers presenting techniques to diagnose and improve the performance of business processes.

Figure 5 shows changes of relative frequencies over time. The graph shows a slight increase in process-mining related topics. However, no clear trends are visible due to the many use cases and small numbers. Therefore, all BPM papers were also analyzed based on the six key concerns presented next.

4 BPM Key Concerns

To provide a trend analysis at a coarser level of granularity, I also identified six *key concerns* before analyzing the 289 papers:

Process modeling languages The first concern is about the process modeling language to be used. Many papers propose a new notation or evaluate existing ones. There are often competing requirements, e.g., the language should be very expressive but simple at the same time [9]. Languages intended for automated process execution (e.g., BPEL) may be very different from languages mainly used for discussion and documentation (e.g., EPCs). Other languages may be tailored towards verification (e.g., WF-nets) or process mining (e.g., C-nets or hidden Markov chains).

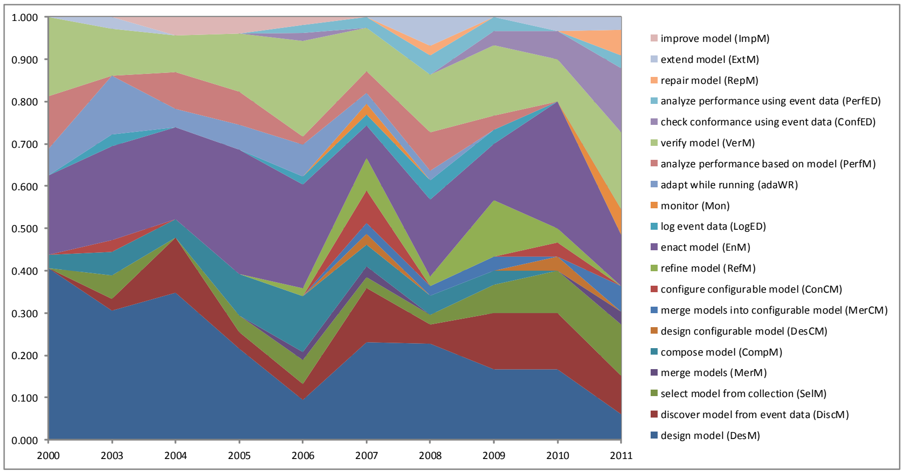

Fig. 5. Development of the relative importance of each use case plotted over time (derived from Table 1).

Process enactment infrastructures The second key concern is the creation of an infrastructure to execute, support, and monitor processes. Examples of topics dealing with this concern are the development of workflow engines, service-oriented computing, interoperability, cloud computing, enterprise application integration, work distribution systems, etc.

Process model analysis The third concern refers to the analysis of processes based on models without using event data. Papers addressing this concern are about topics such as soundness verification, simulation, model checking, queueing networks, controllability, etc.

Process mining The fourth concern refers to all analysis techniques that are driven by event data. For example, process discovery techniques that construct a model based on observed traces. Process mining is not limited to discovery and also includes conformance checking and extension [3]. Conformance checking can be used to check if reality, as recorded in an event log, conforms to the model and vice versa. Extension adds a new perspective to the process model by cross-correlating it with an event log.

Process flexibility The fifth concern addresses the problem that existing WFM/BPM systems tend to be inflexible. Process flexibility can be seen as the ability to deal with both foreseen and unforeseen changes, by varying or adapting those parts of the business process that are affected by them, while retaining the essential format of those parts that are not impacted by the variations [38]. Papers on case handling, adaptive workflows, late-binding, declarative languages, etc. all aim at adding flexibility.

Process reuse The sixth concern addresses the problem that (parts of) processes are often "reinvented" rather than reused. The challenge is to avoid duplicate modeling

Table 2. Relative importance of the six key concerns over the years. All rows add up to 1. The last row shows the average relative frequency over all years.

year	process modeling languages	process enactment infrastructures	process model analysis	process mining	process flexibility	process reuse
2000	0.355	0.161	0.290	0.000	0.161	0.032
2003	0.325	0.200	0.250	0.050	0.075	0.100
2004	0.286	0.238	0.238	0.143	0.048	0.048
2005	0.288	0.231	0.212	0.058	0.096	0.115
2006	0.154	0.308	0.288	0.096	0.077	0.077
2007	0.387	0.097	0.194	0.194	0.065	0.065
2008	0.324	0.108	0.297	0.135	0.081	0.054
2009	0.148	0.111	0.370	0.222	0.037	0.111
2010	0.240	0.240	0.200	0.160	0.000	0.160
2011	0.143	0.171	0.200	0.314	0.000	0.171
average	0.265	0.187	0.254	0.137	0.064	0.093

and implementation efforts. Configurable process models, reference models, process repositories, similarity search, etc. are typical approaches to promote reuse.

As for the use cases, the papers in [7], [11], [19], [4], [23], [14], [22], [18], [31], [36], and [15] were tagged with one, or sometimes more, key concerns. A total of 342 tags were assigned to the 289 papers (1.18 tag per paper on average). The tags were used to determine the relative frequencies listed in Table 2. For example, for BPM 2010 I tagged four papers with key concern *process reuse*. The total number of tags for BPM 2010 is 25. Hence, the relative frequency is $4/25 = 0.16$. The bottom row gives the average relative frequency of each concern over all 10 years. Both trends and averages are depicted graphically in Figure 6. As expected, the first three concerns are most frequent. The fourth and sixth concern (process mining and process reuse) are gaining importance, whereas the relative frequency of the process flexibility concern seems to decrease over time.

5 Reflections

Before discussing some insights obtained when tagging the 289 BPM papers, I first reflect on the analysis method used.

The tagging of a paper with use cases and key concerns is highly subjective. It is unlikely that another expert would come to the exact same tags for each paper. For example, to tag a paper one needs to decide what the key contribution of the paper is. Many papers are rather broad and difficult to classify. For example, papers on topics such as “Social BPM”, “BPM maturity”, and “BPM Security” cannot be tagged easily, because these topics seem orthogonal to the uses cases and key concerns. This explains

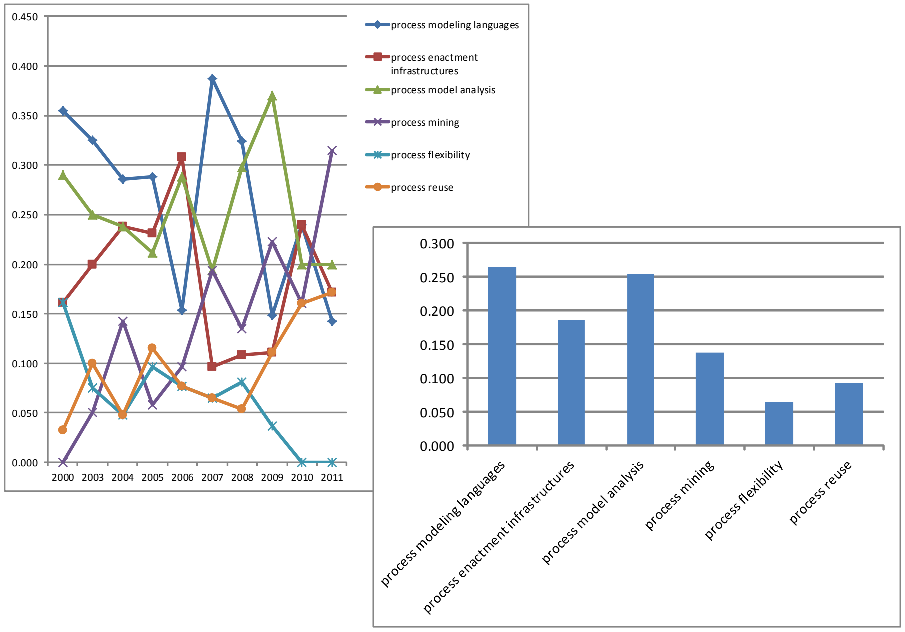

Fig. 6. Visualization of the results from Table 2 for each of the six key concerns: changes of relative importance over time (left) and average relative frequency (right).

why broad use cases like *design model* (DecM) and *enact model* (EnM) score relatively high.

The key concerns were identified before tagging the papers. In hindsight there seem to be at least three potentially missing concerns: *process integration*, *patterns*, and *collaboration*. Many papers are concerned with web services and other technologies (e.g., SaaS, PaaS, clouds, and grids) to integrate processes. These are now tagged as *process enactment infrastructures* (second concern). In the BPM proceedings there are various papers proposing new patterns collections or evaluating existing languages using the existing workflow patterns [9, 37]. These are now tagged as *process modeling languages* (first concern). Another recurring concern seems to collaboration, e.g., collaborative modeling or system development.

Given a process, different *perspectives* can be considered: the *control-flow* perspective ("What activities need to be executed and how are they ordered?"), the *organizational* perspective ("What are the organizational roles, which activities can be executed by a particular resource, and how is work distributed?"), the *case/data* perspective ("Which characteristics of a case influence a particular decision?"), and the *time* perspective ("What are the bottlenecks in my process?"), etc. The use cases and key concerns are neutral/orthogonal with respect to these perspectives. Although most papers focus on the control-flow perspective, there are several papers that focus on the organizational perspective, e.g., papers dealing with optimal resource allocations or role-based access control. It would have been useful to add additional tags to papers based on the perspectives considered.

Next, I reflect on the papers themselves. Comparing papers published in the early BPM proceedings with papers published in more recent BPM proceedings clearly shows that the BPM discipline progressed at a remarkable speed. The understanding of process modeling languages improved and analysis techniques have become much more powerful. Despite the good quality of most papers, some weaknesses can be noted when reflecting on the set of 289 BPM papers.

- Many papers introduce a new modeling language. The need for such new languages is often unclear, and, in many cases, the proposed language is never used again. A related problem is that many papers spend more time on presenting the context of the problem rather than the actual analysis and solution. For example, there are papers proposing a new verification technique for a language introduced in the same paper. Consequently, the results cannot be used or compared easily.
- Many papers cannot be linked to one of the 20 use cases of Section 3 in a straightforward manner. Authors seem to focus on originality rather than relevance and show little concern for real-life use cases. One could argue that such papers propose solutions for rather exotic or even artificial problems.
- Many papers describe implementation efforts; however, frequently the software is not available for the reader. Moreover, regrettably, many of the research prototypes seem to "disappear" after publication. As a result, research efforts get lost.
- Many papers include case studies, e.g., to test a new technique or system, which is good. Unfortunately, most case studies seem rather artificial. Often the core contribution of the paper is not really evaluated or the case study is deliberately kept vague.

6 Outlook

The BPM discipline clearly matured over the last decade. Nevertheless, there are many exciting open problems and BPM will remain highly relevant. Whereas it used to be in vogue to present a technique which just exists on paper, more recent BPM papers tend to describe an implementation of the ideas (proof-of-concept). Moreover, the importance of empirical evaluation seems to increase. Consider for example a new fictive verification technique X. It used to be acceptable to just present the idea and some proof of X's correctness. Now it is expected that X is implemented and some experimental results need to be provided in the paper. In the future authors will be required to include an empirical evaluation using large collections of process models and really compare results using benchmark examples. This is good development. Researchers should focus on the hard BPM problems rather than trying to come up with original (fake or artificial) problems. Hence, it would be good to organize competitions centering around challenging BPM problems that need to be solved urgently.

The "Big Data" wave provides new prospects for BPM research. Organizations are recording large amounts of event data and start understanding the potential value of such data. This is great opportunity to promote "evidence-based BPM", e.g., research ideas can be empirically evaluated using real data. This will increase the credibility of BPM research and help convincing practitioners to adopt new ideas.

This paper lists 20 use cases. As mentioned, for some papers it is unclear to see which use case the authors are trying to support. It would be interesting to require BPM authors to tag their paper with the use cases addressed by their work. This could serve as a reality check for the authors and help to structure the field (e.g., to find reviewers and related work). Such ideas would of course require further development of the use cases presented in this paper (including requests for input, open discussions, and consensus building).

Although BPM research is extremely relevant and its results are useable by many organizations, there are no strong industrial counterparts willing to invest in foundational BPM research. Established consultancy and software firms tend to be rather conservative and end-user organizations are not aware of advances in BPM research. Topics are hyped, but the actual realization of ideas often leaves much to be desired. As a result "BPM is everywhere and nowhere". Compare this to other domains (e.g., the high-tech industry) where there are fewer players, but these are able to invest in R&D and cannot survive by selling only buzzwords. Given these circumstances, it is important to develop *shared* open-source software as a means to influence practitioners. Unfortunately, many prototypes are developed from scratch and "fade onto oblivion" when the corresponding research project ends. Therefore, there is a need to maintain larger open-source software platforms shared between different groups. However, this is not easy. For example, open-source tools like ProM and YAWL have more than 400,000 lines of code. It is a challenge to sustain such efforts over a longer period of time. The BPM community would benefit from shared development efforts using a limited number of large open-source software platforms (instead of developing many throw-away prototypes), but this requires a change in research culture.

Given all of these challenges, I'm looking forward to a new decade of exhilarating and simulating BPM research!

References

1. W.M.P. van der Aalst. The Application of Petri Nets to Workflow Management. *The Journal of Circuits, Systems and Computers*, 8(1):21–66, 1998.
2. W.M.P. van der Aalst. Business Process Management Demystified: A Tutorial on Models, Systems and Standards for Workflow Management. In J. Desel, W. Reisig, and G. Rozenberg, editors, *Lectures on Concurrency and Petri Nets*, volume 3098 of *Lecture Notes in Computer Science*, pages 1–65. Springer-Verlag, Berlin, 2004.
3. W.M.P. van der Aalst. *Process Mining: Discovery, Conformance and Enhancement of Business Processes*. Springer-Verlag, Berlin, 2011.
4. W.M.P. van der Aalst, B. Benatallah, F. Casati, and F. Curbera, editors. *International Conference on Business Process Management (BPM 2005)*, volume 3649 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2005.
5. W.M.P. van der Aalst and E. Best, editors. *24th International Conference on Applications and Theory of Petri Nets (ICATPN 2003)*, volume 2679 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2003.

6. W.M.P. van der Aalst, J. Desel, and R. Kaschek, editors. *Software Architectures for Business Process Management (SABPM'99)*, Heidelberg, Germany, June 1999. Forschungsbericht Nr. 390, University of Karlsruhe, Institut AIFB, Karlsruhe, Germany.
7. W.M.P. van der Aalst, J. Desel, and A. Oberweis, editors. *Business Process Management: Models, Techniques, and Empirical Studies*, volume 1806 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2000.
8. W.M.P. van der Aalst and K.M. van Hee. *Workflow Management: Models, Methods, and Systems*. MIT press, Cambridge, MA, 2004.
9. W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Workflow Patterns. *Distributed and Parallel Databases*, 14(1):5–51, 2003.
10. W.M.P. van der Aalst, A.H.M. ter Hofstede, and M. Weske. Business Process Management: A Survey. In W.M.P. van der Aalst, A.H.M. ter Hofstede, and M. Weske, editors, *International Conference on Business Process Management (BPM 2003)*, volume 2678 of *Lecture Notes in Computer Science*, pages 1–12. Springer-Verlag, Berlin, 2003.
11. W.M.P. van der Aalst, A.H.M. ter Hofstede, and M. Weske, editors. *International Conference on Business Process Management (BPM 2003)*, volume 2678 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2003.
12. W.M.P. van der Aalst, G. De Michelis, and C.A. Ellis, editors. *Proceedings of Workflow Management: Net-based Concepts, Models, Techniques and Tools (WFM'98)*, Lisbon, Portugal, June 1998. UNINOVA, Lisbon.
13. W.M.P. van der Aalst and C. Stahl. *Modeling Business Processes: A Petri Net Oriented Approach*. MIT press, Cambridge, MA, 2011.
14. G. Alonso, P. Dadam, and M. Rosemann, editors. *International Conference on Business Process Management (BPM 2007)*, volume 4714 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2007.
15. A. Barros, A. Gal, and E. Kindler, editors. *International Conference on Business Process Management (BPM 2012)*, Lecture Notes in Computer Science. Springer-Verlag, Berlin, 2012.
16. P.P. Chen. The Entity-Relationship Model: Towards a unified view of Data. *ACM Transactions on Database Systems*, 1:9–36, Jan 1976.
17. E.F. Codd. A Relational Model for Large Shared Data Banks. *Communications of the ACM*, 13(6):377–387, June 1970.
18. U. Dayal, J. Eder, J. Koehler, and H. Reijers, editors. *International Conference on Business Process Management (BPM 2009)*, volume 5701 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2009.
19. J. Desel, B. Pernici, and M. Weske, editors. *International Conference on Business Process Management (BPM 2004)*, volume 3080 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2004.
20. R. Dijkman, M. Dumas, and L. Garcia-Banuelos. Graph Matching Algorithms for Business Process Model Similarity Search. In U. Dayal, J. Eder, J. Koehler, and H. Reijers, editors, *Business Process Management (BPM 2009)*, volume 5701 of *Lecture Notes in Computer Science*, pages 48–63. Springer-Verlag, Berlin, 2009.
21. M. Dumas, W.M.P. van der Aalst, and A.H.M. ter Hofstede. *Process-Aware Information Systems: Bridging People and Software through Process Technology*. Wiley & Sons, 2005.
22. M. Dumas, M. Reichert, and M.C. Shan, editors. *International Conference on Business Process Management (BPM 2008)*, volume 5240 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2008.
23. S. Dustdar, J.L. Fiadeiro, and A. Sheth, editors. *International Conference on Business Process Management (BPM 2006)*, volume 4102 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2006.

24. C.A. Ellis. Information Control Nets: A Mathematical Model of Office Information Flow. In *Proceedings of the Conference on Simulation, Measurement and Modeling of Computer Systems*, pages 225–240, Boulder, Colorado, 1979. ACM Press.
25. C.A. Ellis and G. Nutt. Workflow: The Process Spectrum. In A. Sheth, editor, *Proceedings of the NSF Workshop on Workflow and Process Automation in Information Systems*, pages 140–145, Athens, Georgia, May 1996.
26. C.A. Ellis and G.J. Nutt. Office Information Systems and Computer Science. *ACM Computing Surveys*, 12(1):27–60, 1980.
27. D. Fahland, C. Favre, B. Jobstmann, J. Koehler, N. Lohmann, H. Voelzer, and K. Wolf. Instantaneous Soundness Checking of Industrial Business Process Models. In U. Dayal, J. Eder, J. Koehler, and H. Reijers, editors, *Business Process Management (BPM 2009)*, volume 5701 of *Lecture Notes in Computer Science*, pages 278–293. Springer-Verlag, Berlin, 2009.
28. A.H.M. ter Hofstede, W.M.P. van der Aalst, M. Adams, and N. Russell. *Modern Business Process Automation: YAWL and its Support Environment*. Springer-Verlag, Berlin, 2010.
29. A. W. Holt. Coordination Technology and Petri Nets. In G. Rozenberg, editor, *Advances in Petri Nets 1985*, volume 222 of *Lecture Notes in Computer Science*, pages 278–296. Springer-Verlag, Berlin, 1985.
30. A.W. Holt, H. Saint, R. Shapiro, and S. Warshall. Final Report on the Information Systems Theory Project. Technical Report RADC-TR-68-305, Griffiss Air Force Base, New York, 1968.
31. R. Hull, J. Mendling, and S. Tai, editors. *International Conference on Business Process Management (BPM 2010)*, volume 6336 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2010.
32. S. Jablonski and C. Bussler. *Workflow Management: Modeling Concepts, Architecture, and Implementation*. International Thomson Computer Press, London, UK, 1996.
33. F. Leymann and D. Roller. *Production Workflow: Concepts and Techniques*. Prentice-Hall PTR, Upper Saddle River, New Jersey, USA, 1999.
34. M. zur Muehlen. *Workflow-based Process Controlling: Foundation, Design and Application of workflow-driven Process Information Systems*. Logos, Berlin, 2004.
35. M. Reichert and B. Weber. *Enabling Flexibility in Process-Aware Information Systems: Challenges, Methods, Technologies*. Springer-Verlag, Berlin, 2012.
36. S. Rinderle, F. Toumani, and K. Wolf, editors. *International Conference on Business Process Management (BPM 2011)*, volume 6896 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2011.
37. N. Russell, W.M.P.van der Aalst, A.H.M. ter Hofstede, and D. Edmond. Workflow Resource Patterns: Identification, Representation and Tool Support. In *Proceedings of the 17th Conference on Advanced Information Systems Engineering (CAiSE'05)*, volume 3520 of *Lecture Notes in Computer Science*, pages 216–232. Springer-Verlag, Berlin, 2005.
38. H. Schonenberg, R. Mans, N. Russell, N. Mulyar, and W.M.P. van der Aalst. Process Flexibility: A Survey of Contemporary Approaches. In *Advances in Enterprise Engineering I*, volume 10 of *Lecture Notes in Business Information Processing*, pages 16–30. Springer-Verlag, Berlin, 2008.
39. M. Weske. *Business Process Management: Concepts, Languages, Architectures*. Springer-Verlag, Berlin, 2007.
40. M.D. Zisman. *Representation, Specification and Automation of Office Procedures*. PhD thesis, University of Pennsylvania, Warton School of Business, 1977.
41. M.D. Zisman. Office Automation: Revolution or Evolution. *Sloan Management Review*, 13(3):1–16, 1978.
42. M.D. Zisman. Use of Production Systems for Modeling Asynchronous Concurrent Processes. *Pattern-Directed Inference Systems*, pages 53–68, 1978.