

Hybrid Process Technologies in the Financial Sector

Søren Debois², Thomas Hildebrandt², Morten Marquard¹, and Tijs Slaats^{1,2} *

¹ Exformatics A/S, Dag Hammarskjölds Allé 13, DK-2100 Copenhagen Ø, Denmark
{mmq, ts}@exformatics.com, <http://www.exformatics.com>

² IT University of Copenhagen, Rued Langgards Vej 7, DK-2300 Copenhagen S, Denmark
{debois, hilde, tslaats}@itu.dk, <http://www.itu.dk>

Abstract. Danish mortgage credit institutes deal with highly variable and knowledge-intensive processes. At the same time these processes are required to be strictly conformant to current regulations and laws. In addition different divisions of the business are interested in different views on the same process: whereas the IT department implementing the processes would like a complete view that shows the underlying business rules and supports all variants, the end users are only interested in a local view that (1) shows only the aspects of the process that they are responsible for and (2) only shows the variants of the process that are relevant to them. This paper reports on a project we undertook with such a credit institute where we investigated and addressed these issues by providing a hybrid solution, allowing processes to be modelled using our constraint-based modelling tools, but also supporting flow-based views of both the entire process and specific variants.

Keywords: Hybrid Process Technologies, Declarative Modelling, DCR Graphs, Financial Processes, Process Flexibility, Knowledge Work

1 Introduction

This paper describes an investigation by Danish Adaptive Case Management vendor Exformatics A/S into the usability of its declarative workflow modelling and execution engine for application within the financial sector. This investigation was carried out in collaboration with the Process Modelling Group of the IT University of Copenhagen and Danish mortgage credit institution BRFKredit.

The investigation indicated that:

- BRFKredit uses process modelling primarily as an internal communication tool.
- Different stakeholders have radically different use for the resulting process models.
- Different stakeholders prefer somewhat different process notations.
- Many stakeholders are content with representing processes as “representative traces”.

* Authors listed alphabetically. This work is supported in part by the Computational Artifacts project (VELUX 33295, 2014-2017), by the Danish Agency for Science, Technology and Innovation, by the CFIR innovation network, by Exformatics A/S, and by the IT University of Copenhagen.

- Moreover, such representative traces should contain only activities “relevant” to the business case at hand.

In order to accommodate these diverse requirements of process models, Exformatics has extended their declarative modelling tools to derive from the model representative traces and other relevant flow-based visualisations. Through this extension the tools now support a hybrid modelling approach where processes are modelled declaratively based on their underlying business rules, but the behaviour supported by the model can be visualised in more familiar flow-based notations, both in full and as representative traces. The new hybrid features and the broader perspectives of the technology were well received by BRFKredit, but a more thorough evaluation with more users and case studies is of course needed to make any firm conclusions on their use in practice.

2 Background

Exformatics A/S is a Danish vendor of Adaptive Case Management (ACM) solutions with a customer base of approximately 40 organisations, both Danish and international and both private and public. The core product of Exformatics is a case handling system aimed at knowledge workers such as lawyers working with intellectual property protection or real estate management; engineers and project managers designing and constructing large power plants; marketing employees planning campaigns for broadcasting; as well as case workers in various areas such as HR, political hearings, and workforce political issues.

Exformatics strongly believes in the need for formal workflow modelling notations as a means to

- Strengthen communication of requirements from client organisations.
- Strengthen communication within client organisations post deployment.
- Expedite development.
- Enable run-time system updates.

However, for its current clientele of knowledge workers, Exformatics has found flow-based modelling notations insufficiently flexible. The cases handled by Exformatics’ clients are at the same time uniform yet never actually equal—two different marketing campaigns are at the end of the day not radically different, however, the devil is in the detail, and the detail invariably differs a great deal. Hence, Exformatics has adopted and co-developed a declarative workflow modelling notation, DCR Graphs [5], originally conceived at the Process Modelling Group of the IT University of Copenhagen, headed by Thomas Hildebrandt.

Exformatics has invested considerable resources into bringing DCR Graphs from merely an academic vision to a practical, marketable product, in the form of a Process Execution Engine and Workflow Modelling Toolkit [7]. The former has been deployed in Exformatics’ solutions, most notably with the inclusion of a complete model of the workflow of the Danish foundation Dreyers Fond [3].

However, so far, the use of formal workflow modelling has, from the perspective of the customer, been an implementation detail; a technical trick Exformatics uses to speed

up its implementation. It is the vision of Exformatics that (suitably authorised) expert end-users should be able to update or even create models themselves, with the Exformatics Process Engine automatically reflecting such updates. With that vision in mind, Exformatics has been looking for a mature process-intensive knowledge worker-heavy company, preferably in the financial sector, with which to experiment with the possible future directions of the Exformatics' Workflow Modelling Toolkit.

Such a company was found in Danish mortgage credit institution BRFkredit in late 2014. A formal collaboration was established between Exformatics A/S, BRFkredit, IT University of Copenhagen and a fourth partner, within the purview of and financially supported by the Copenhagen Fintech Innovation and Research Network (CFIR).

BRFkredit is a Danish mortgage credit institution lending against collateral on owner-occupied homes, commercial properties and subsidised housing. On the Danish housing market mortgage loans are not typically taken out from a bank, but instead from specialised credit institutions who issue mortgage bonds that pool together many borrowers and investors, thereby spreading out the associated risk. BRFkredit is one such mortgage credit institution.

Within BRFkredit, loans for residential purposes account for the majority of all lending; corporate lending is loans for office and business properties, private rental and cooperative housing. BRFkredit finances the current lending by issuing covered bonds and mortgage bonds continuously, i.e. as tap issues.

BRFkredit A/S is owned by Jyske Bank A/S through the holding company BRFHolding A/S. BRFkredit has a market share of the total Danish mortgage market of approximately 8% Jyske Bank and BRFkredit is Denmark's fourth largest financial institution. BRFkredit has lending of around DKK 213 billion (approx. EUR 28 billion) distributed on around 120.000 mortgage loans, managed by just over 750 full time employees³.

3 Objectives

Exformatics chose to adopt a declarative notation because of a strong belief in their client's need for flexibility: knowledge worker end-users are the experts and should only be inhibited in their possible actions if required by laws or business rules. As has been commonly claimed in the academic literature [8], Exformatics believes that declarative notations are better suited at describing such laws and business rules than imperative or flow-based notations.

However, declarative notations have been shown to be more challenging to understand for end-users than more common flow-based notations such as BPMN [15]. Hence:

- The primary objective of the investigation for Exformatics was to gain an understanding of how the DCR Graph process modelling can be made more accessible to expert end-users.

³ BRFkredit in Brief. Accessed August 2015 at <http://www.brf.dk/Investors/About-BRFkredit/Additional-information/BRFkredit-in-brief>

- One secondary objective was to gain a better understanding of the motivation for and role of manual process modelling in financial institutions, and the applicability of DCR Graphs to the same.
- Another secondary objective was to carry out a practical (yet anecdotal) test of the hypothesis that simulation, potentially collaborative simulation, can be a strong tool for expert end-users' work with process models. For the latter, Exformatics places a high priority on support for simulation in its tool-chain [7]. Believing that the ability to simulate and "play through" a process will help users better understand the ramifications of a particular declarative model.

4 Related Work

The direction taken in this project relates closely to the recently initiated work on hybrid business process modelling notations and technologies [10], which aims to combine the strengths of the flow- and constraint-based process modelling paradigms. A currently common approach in this field is to provide hybrid modelling notations that combine both flow- and constraint-based elements [1, 11–13]. Our own approach on the other hand uses the two paradigms separately: a constraint-based notation is used to model the process, whereas a flow-based notation is used to give more insight into the behaviour supported by the models. This relates closely to recent work [2, 9] on mapping from the declarative Declare [8] notation to Petri nets, however contrary to the work presented here we are not aware of these techniques being used in commercial tools or applied to industrial case studies.

5 Investigation

The investigation itself took the form of a sequence of full- and half-day meetings in early spring 2015, primarily at Exformatics, occasionally at ITU or BRFkredit. During the meetings we discussed the needs of BRFkredit for process modeling and the required extensions to the DCR Graphs tools to meet those needs. In the present case study, we report only on the conclusion of these discussions.

The objective of process modeling within BRFkredit is this: process models are constructed for the purpose of communication within the company. The constructed models are then used by several different stakeholders:

1. The IT department uses process descriptions as partial requirements specifications for IT systems supporting new and updated financial products.
2. Caseworkers use process models as road-maps for their daily work.
3. Management (at multiple levels) uses process models as abstracted views of "what goes on in the company".

It is interesting to note at this point that for BRFkredit process modeling has enough value as a communication tool alone to merit allocating resources to their construction and maintenance.

However, BRFkredit reported that their use of such models faces two major challenges:

1. Different stakeholders use different process model notations.
2. The level of abstraction appropriate for a model depends on the stakeholder consuming it.

We treat each of these in turn.

Different stakeholders use different notations. Attempts to introduce a small subset of BPMN as a standard modelling notation used everywhere in BRFkredit has not been successful: Most departments, including IT, find that notation unhelpful. It is important to know that this is not because those stakeholders are adverse to process modelling; on the contrary, some departments have on their own initiative produced extensive and comprehensive models of their processes for internal consumption.

These models appear to have a number of commonalities across departments.⁴

1. The models are trace models, i.e., each model describes a single “happy-path” run of the process in question. The models typically include very little or no possibility of choosing between or reordering activities.
2. The models heavily emphasize roles, whether occupied by humans or IT systems. Diagrams are invariably some form of swim-lane diagrams, typically produced in Powerpoint.

The emphasis in point 1 above on single traces over branching models is interesting, and in line with recent research on understandability of process models [4, 14]. It appears that for the majority of stakeholders understandability far outweighs precision or generality when it comes to usefulness of models as tools for communication. For discussion and communication, a representation of a single “happy path”—the most common variant of a particular process—is far more helpful than a branching model like BPMN. The greater precision afforded by the branching model sours discussion by bringing in irrelevant detail.

Where detail is required—when process models are used as requirements specification for IT, or where process model are used as road-maps for caseworkers—BRFkredit simply produces more single-trace models. Eg., at BRFkredit, the various processes for granting various kinds of mortgages has been enumerated to more than a thousand. BRFkredit explicitly mentions that this approach is very burdensome in the face of changes in internal processes, eg., in response to changes in business requirements or the regulatory climate. It seems very likely that the large number of processes can profitably be represented as minor variations on a very small number of core processes.

It is worth noting that in a large company such as BRFkredit, difficulties on agreeing on notation go beyond mere process notation. E.g., a seemingly simple and precise notion such as “client” has different meaning between departments: For the Sales department, it is a person who might be interested in obtaining a mortgage; for the Loan Monitoring department, it is a person who has an active loan with BRFkredit and so forth.

To the outside observer, the combination of differences in (ad hoc) model notations and terminology has the unfortunate consequence that one may be presented with two

⁴ For confidentiality reasons, we cannot include actual examples of the various custom models in this document.

diagrams, e.g. one depicting a process as seen by IT and one as seen by sales, and utterly fail to realise that the two diagrams depict aspects of the same process.

The level of abstraction appropriate for a model depends on the stakeholder consuming it. This is most easily explained by example, for instance:

1. Caseworkers need process descriptions that show only the process variant they are currently working on, and need not know any detail of underlying IT processes.
2. Managers sometimes need process descriptions that are precise about interactions between departments, but do not contain details about what goes on inside departments.
3. Managers at other times need process descriptions that describe only the part of processes pertaining to particular (financial) products or product lines.
4. IT needs process descriptions that contain every possible process variant and full detail on interaction with IT systems, but does not care at all about human interactions (eg., prospective client meetings).

6 Proposed technology

Reconciling differences in terminology in a large organization is a problem beyond the scope of process modeling; accordingly, we focus on notation.

Regarding differences in notations. We contend that the difference in process notations employed in BRFKredit does not arise from any inherent difference in preference in modeling notations; rather, it arises from the absence of a single notation suitable for all stakeholders. BPMN is apparently not it; not for want of flexibility, but simply because of its plethora of symbols and, to the layman, less than completely intuitive semantics.

Recall the above observation that most of the ad-hoc diagrams we have been presented with in reality presents only a single trace with precise distinction between roles. That notation, then, is apparently the appropriate notation for the majority of stakeholders. As such, we envision a mechanism for presenting process models in terms of a very small number of representative traces. This idea sits well with the idea of declarative or constraint-based process modeling: a declarative model is a concise representation of a typically large number of admissible traces, with semantics that allow us to compute efficiently whether a trace is admissible. If the processes of BRFKredit were represented as a single or, more realistically, small number of general processes, one could extract from these models representative traces “representing” the process in internal communications.

This idea begets the question: Which traces? Among all the possible traces admitted by our hypothesised general model, how do we find an appropriate set of representatives?

We contend that in a given process model, we can name activities the execution of which are in fact the objective of the process. In the case of BRFKredit, the objective of,

say, an instance of an mortgage application process is the eventual evaluation of that application. Variants of that process arise as different opportunities present themselves for reaching that goal. For instance, one variant is the one in which the value of the property securing the mortgage can be appraised statistically, without an actual inspection. A different variant arises when the property is in a insufficiently uniform neighbourhood, whence the constraints of the process model forbid the statistical appraisal.

In summary, we believe that—at least in this instance—the objectives of the process can be defined as the execution of a particular activity (“assess loan application”); and that the variants of the process can be identified by which activities are executed towards that objective (“statistical appraisal” or “on-site appraisal”).

This yields a method for identifying relevant traces: domain experts, which must be consulted anyway when constructing the model, help figuring out which activities characterise the objective of the process, and the key activities identifying (collections of) process variants.

Regarding differences in appropriateness of abstractions. For the single-trace model representatives suggested above, this is simply a matter of suitably projecting the trace in question, i.e., leaving out activities that are unwarranted at the desired level of abstraction. Following the examples from above:

Example 1: Caseworkers need process descriptions that show only the process variant they are currently working on, and need not know any detail of underlying IT processes.

Proposed solution: Given a complete trace, we simply remove all activities not directly involving the caseworker.

Example 2: Managers sometimes need process descriptions that are precise about interactions between departments, but do not contain details about what goes on inside departments. **Proposed solution:** Again, given a complete trace, we simply remove all activities not adjacent to an activity of a different department.

Example 3: Managers at other times need process descriptions that describe only the part of processes pertaining to particular (financial) products or product lines. **Proposed solution:** This relates to how the trace is generated; looking for a trace that concludes in, say, “assess additional loan application” partially fulfils this requirement.

Example 4: IT needs process descriptions that contain every possible process variant and full detail on interaction with IT systems, but does not care at all about human interactions (eg., prospective client meetings). **Proposed solution:** In this case, where we do need branching structure, the picture is less clear. For IT, process descriptions often play a role as part of a requirements specification, and thus the process model must describe all and only the behaviour of the desired system: However, we may assume a minimal sophistication with formal models and so use the general DCR model as the appropriate model. For DCR graphs, the possibilities for semantically well-founded projection is well-studied, and so getting rid of “human interactions” in the above example amounts to employing one of the known sound projection methods [6].

7 Proof-of-concept implementation

To present the above ideas to BRFkredit staff in practice, Exformatics has extended its existing workflow modeling tool with a proof-of-concept analysis tool for (a) presenting projected traces and (b) generating minimal traces executing or not executing given activities. We give a brief walk-through of this tool.

The tool presupposes a single DCR Graph-based process model encompassing a family of processes, e.g., a particular class of mortgage loan application processes. For confidentiality reasons, we cannot report on an actual such model here; for clarity, we have constructed a fictional process heavily inspired by the actual processes at BRFkredit, but simplified slightly. This DCR Graph model is presented in Figure 1 and can be found as a public graph (BPM 2015 BRF Example) online at <http://dcrgraphs.net>.

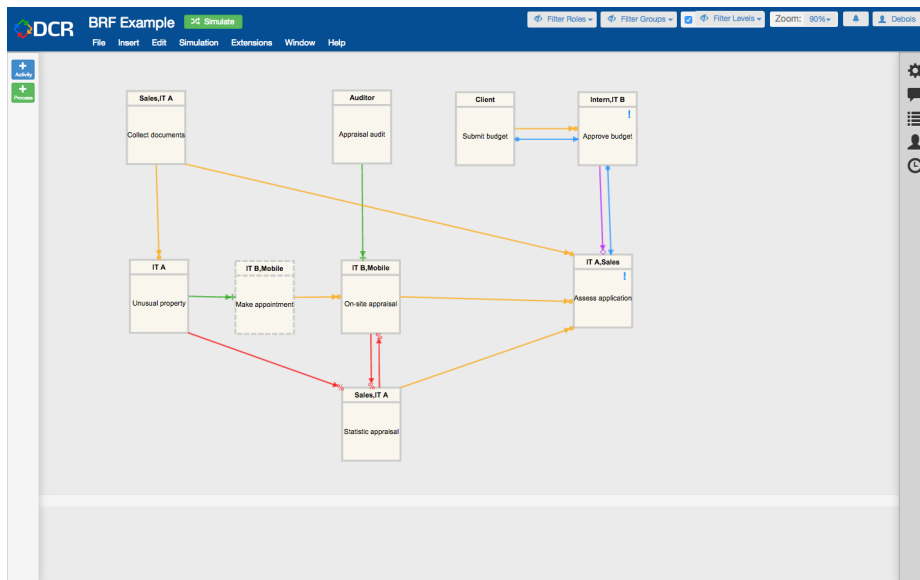


Fig. 1. DCR Graph model of a (fictitious) mortgage loan application process.

Boxes indicate activities, arrows constraints. Each box is labeled (in the middle) with the name of the activity and (in the top) with the roles participating in the activity. The objective of the process is the execution of the activity “Assess application”. The model makes this requirement explicit with the “!” on that box, signifying that the activity is required for the workflow to be considered complete.

The arrows with a bullet in the end (coloured yellow in the tool) indicate conditions, so we cannot “Assess application” before we have executed amongst others “Collect documents”. Arrows with a bullet in the beginning (coloured blue in the tool) indicate required responses, so whenever “Submit budget” is executed—i.e., whenever the

prospective client updates his budget—“Approve budget” must be executed again. The graphical notation for conditions and responses is consistent with the notation for precedence and response constraints used in DECLARE [8].

The arrows with % and + in the end (and coloured red and green in the tool) indicate respectively exclusion and inclusion; arrows with diamonds in the end (and coloured purple in the tool) represent “milestones”; however, it will not be necessary to understand these in detail to read the remainder of this report, so we omit further details and refer instead the reader to e.g. [3] for a brief overview.

This model is only potentially suitable for the IT stakeholders as a requirements specification for implementation purposes. Accordingly, using Exformatics Workflow Editing Tool’s plug-in infrastructure, we have constructed a plug-in that gives different perspectives on this model to be used by the other stake holders (Case Worker and Management). One such perspective is the full state-space of the DCR Graph model in a flow-graph notation; this representation includes the full detail of the model, including branching structure (decision points) and can thus be helpful for implementors, but it is generally way too detailed. An example is in Figure 2 showing the complete but rather overwhelming picture.

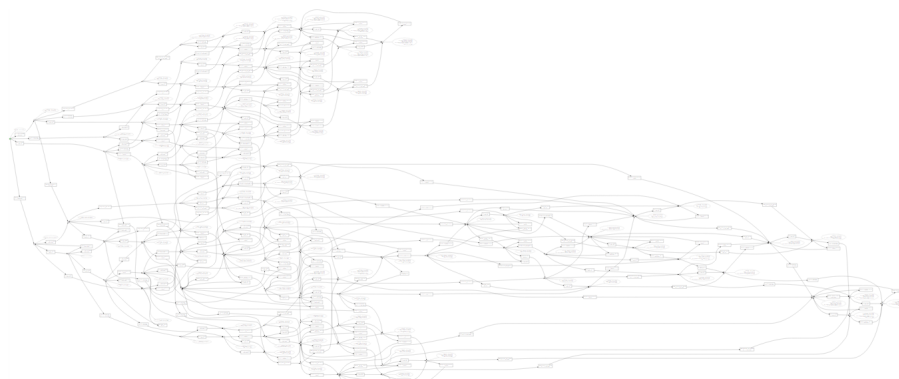


Fig. 2. Flow-graph representation of the model of Figure 1

For the stakeholders interested in representative traces, the prototype tool has a panel for specifying such, illustrated in Figure 3 below. One specifies under “Scope” the activity that should be the starting point of the trace; the ending point; optionally activities that must or cannot occur along the way. Under “Perspective”, one may indicate an optional projection onto specific roles or groups. The tool will then find the shortest trace satisfying the given constraints by simply searching through the transition graph. For instance, suppose we are interested in variants of the loan application process in which the property in question does need an on-site appraisal. Figures 4 and 5 below shows the input constraints and the resulting trace. Starting from the same constraints, we may restrict our attention to the mobile consultant, by clearing “Role” boxes in Fig-

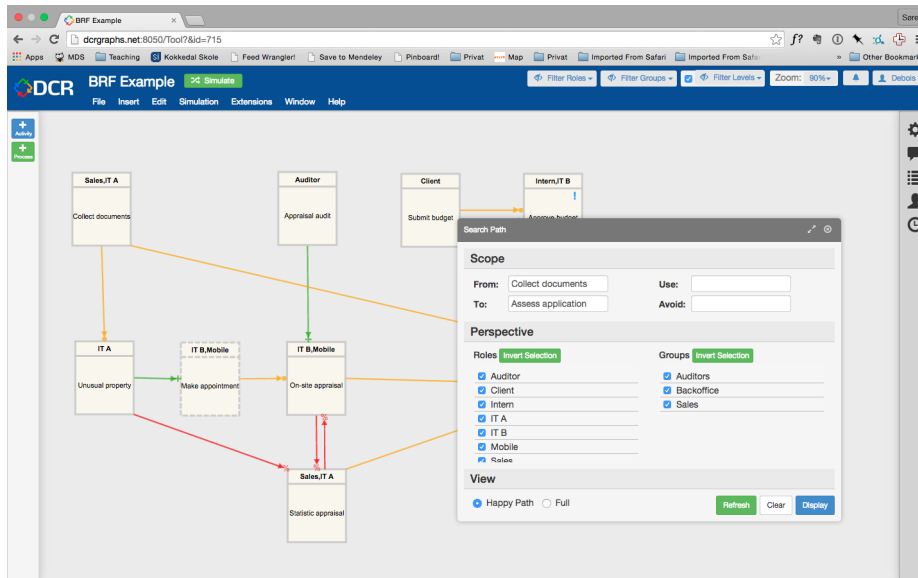


Fig. 3. Pop-up panel for specifying single-trace parameters.

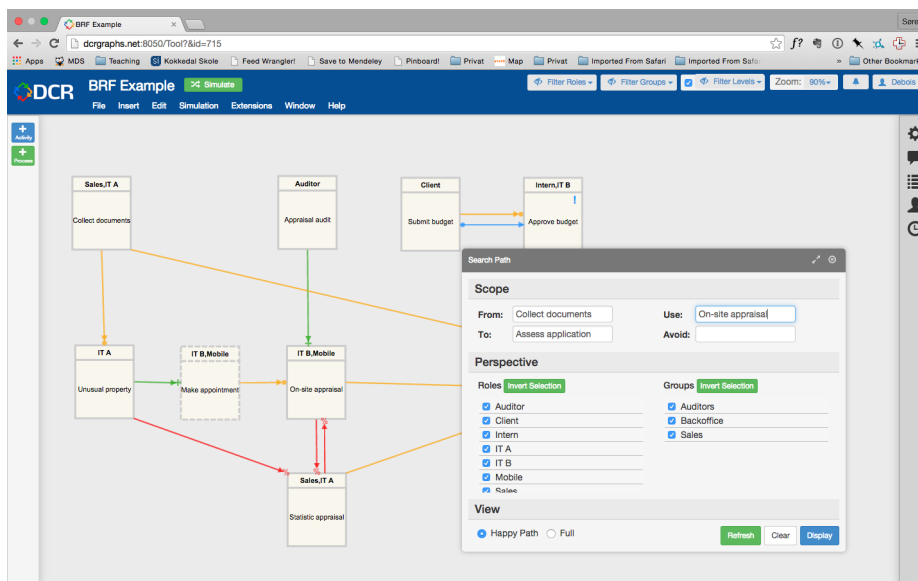


Fig. 4. Specification of process variant requiring on-site appraisal.

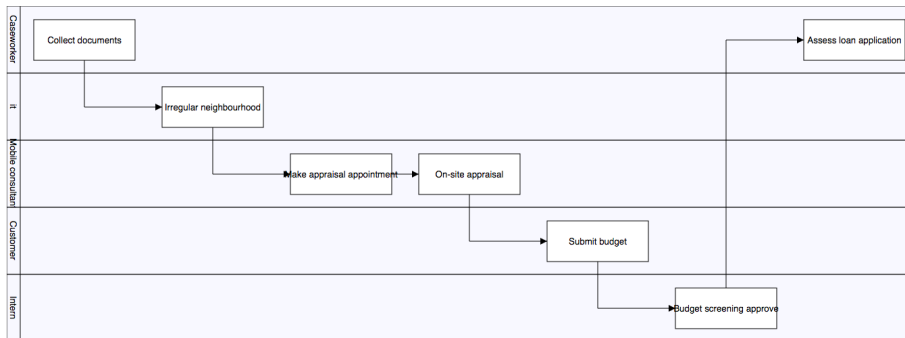


Fig. 5. Trace resulting from the parameters entered in Figure 4.

ure 4. Keeping only, say, the “Mobile Consultant” and the “Caseworker”, we obtain the trace in Figure 6.

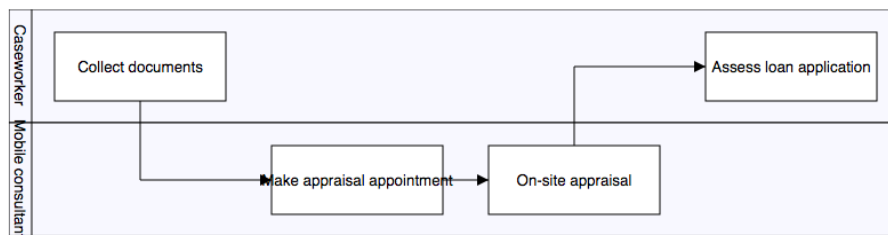


Fig. 6. Trace resulting from the parameters entered in Figure 4 plus projection to “Mobile Consultant” and “Caseworker”.

8 Outcome & General lessons learned

Exformatics A/S is currently extending the plug-in architecture for its Workflow Editing Tool to encompass the proof-of-concept technology reported here as an APP, as well as evolving the proof-of-concept to an important feature of its current offering.

BRFkredit is currently undertaking an in-house evaluation of the prototype, in which actual internal processes are being modelled in Exformatics proof-of-concept implementation.

This case demonstrates anecdotally a clear need for different visualisation of processes for different stakeholders. Also, the proof-of-concept implementation of the semi-automatic generation of “representative traces” was well-received by BRFkredit. Both of these points are likely independent of the concrete case and notation used, although we believe that the possibility of producing the projections depends on the

use of a formal declarative model such as DCR Graphs. The solution is available at <http://www.dcrgraphs.net>.

Acknowledgements: We are grateful to the reviewers for constructive comments.

References

1. Giuseppe De Giacomo, Marlon Dumas, Fabrizio Maria Maggi, and Marco Montali. Declarative process modeling in BPMN. In *Proceedings of the 27th International Conference on Advanced Information Systems Engineering (CAiSE 2015)*, 2015.
2. Johannes De Smedt, Seppe KLM Vanden Broucke, Jochen De Weerd, and Jan Vanthienen. A full R/I-net construct lexicon for declare constraints. *Available at SSRN 2572869*, 2015.
3. Søren Debois, Thomas Hildebrandt, Morten Marquard, and Tijs Slaats. A case for declarative process modelling: Agile development of a grant application system. In *International Workshop on Adaptive Case Management and other non-workflow approaches to BPM*, 2014.
4. Cornelia Haisjackl, Stefan Zugal, Pnina Soffer, Irit Hadar, Manfred Reichert, Jakob Pinggera, and Barbara Weber. Making Sense of Declarative Process Models: Common Strategies and Typical Pitfalls. In Selmin Nurcan, Henderik A. Proper, Pnina Soffer, John Krogstie, Rainer Schmidt, Terry Halpin, and Ilia Bider, editors, *BMMDS/EMMSAD*, volume 147 of *Lecture Notes in Business Information Processing*, pages 2–17, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
5. Thomas Hildebrandt and Raghava Rao Mukkamala. Declarative event-based workflow as distributed dynamic condition response graphs. In *Post-proceedings of PLACES 2010*, 2010.
6. Thomas Hildebrandt, Raghava Rao Mukkamala, and Tijs Slaats. Safe distribution of declarative processes. In *Proceedings of the 9th international conference on Software engineering and formal methods, SEFM'11*, pages 237–252, Berlin, Heidelberg, 2011. Springer-Verlag.
7. Morten Marquard, Muhammad Shahzad, and Tijs Slaats. Web-based modelling and collaborative simulation of declarative processes. In *Proceedings of 13th International Conference on Business Process Management (BPM 2015)*, 2015.
8. Maja Pesic, Helen Schonenberg, and Wil M. P. van der Aalst. DECLARE: Full Support for Loosely-Structured Processes. In *Proceedings of the 11th IEEE International Enterprise Distributed Object Computing Conference*. IEEE Computer Society, 2007.
9. Johannes Prescher, Claudio Di Ciccio, and Jan Mendling. From declarative processes to imperative models. In *Proceedings of the 4th International Symposium on Data-driven Process Discovery and Analysis (SIMPDA 2014), Milan, Italy, November 19-21, 2014.*, pages 162–173, 2014.
10. Hajo A. Reijers, Tijss Slaats, and Christian Stahl. Declarative modeling – An academic dream or the future for BPM? In Florian Daniel, Jianmin Wang, and Barbara Weber, editors, *Proceedings of 11th International Conference on Business Process Management (BPM 2013)*, volume 8094 of *Lecture Notes in Computer Science*, pages 307–322. Springer, 2013.
11. Shazia Sadiq, Wasim Sadiq, and Maria Orłowska. Pockets of flexibility in workflow specification. In Hideko S.Kunii, Sushil Jajodia, and Arne Sølvberg, editors, *Conceptual Modeling — ER 2001*, volume 2224 of *Lecture Notes in Computer Science*, pages 513–526. Springer Berlin Heidelberg, 2001.
12. W.M.P. van der Aalst, M. Adams, A.H.M. ter Hofstede, M. Pesic, and H. Schonenberg. Flexibility as a service. In *Database Systems for Advanced Applications*, volume 5667 of *Lecture Notes in Computer Science*, pages 319–333. Springer, 2009.
13. Michael Westergaard and Tijs Slaats. Mixing paradigms for more comprehensible models. In *Business Process Management - 11th International Conference, BPM 2013, Beijing, China, August 26-30, 2013. Proceedings*, pages 283–290, 2013.

14. Stefan Zugal, Jakob Pinggera, Barbara Weber, Jan Mendling, and Hajo A. Reijers. Assessing the Impact of Hierarchy on Model - A Cognitive Perspective. In *EESMod*, 2011.
15. Stefan Zugal, Pnina Soffer, Cornelia Haisjackl, Jakob Pinggera, Manfred Reichert, and Barbara Weber. Investigating expressiveness and understandability of hierarchy in declarative business process models. *Software & Systems Modeling*, June 2014.