

Multi-party Business Process Resilience By-Design: A Data-Centric Perspective

Pierluigi Plebani¹(✉), Andrea Marrella², Massimo Mecella²,
Marouan Mizmizi¹, and Barbara Pernici¹

¹ Politecnico di Milano - DEIB, Piazza Leonardo da Vinci 32, 20133 Milan, Italy
{pierluigi.plebani,marouan.mizmizi,barbara.pernici}@polimi.it

² Sapienza Università di Roma - DIAG, via Ariosto 25, 00185 Roma, Italy
{marrella,mecella}@diag.uniroma1.it

Abstract. Nowadays every business organization operates in ecosystems and cooperation is mandatory. If, on the one side, this increases the opportunities for the involved organizations, on the other side, every actor is a potential source of failures with impacts on the entire ecosystem. For this reason, *resilience* is a feature that multi-party business processes today must enforce. As resilience concerns the ability to cope with unplanned situations, managing the critical issues is usually a run-time task.

The aim of this work is to emphasize *awareness on resilience* in multi-party business processes also at design-time, when a proper analysis of involved data allows the process designer to identify (possible) failures, their impact, and thus improving the process model. Using a data-centric collaboration-oriented language for processes, i.e., OMG CMMN – Case Management Model and Notation, as modeling notation, our approach allows the designer to model a flexible business process that, at run-time, results easier to manage in case of failures.

Keywords: Process resilience · Artifact-centric modeling · Levels of resilience · CMMN - Case Management Model and Notation

1 Introduction

The adoption of service oriented architectures and workflow automation (a.k.a. orchestration), while enabling and making easier the integration among heterogeneous systems, has also reduced the difficulties in digitizing the communications among different organizations. As a result, digital business ecosystems have been proposed as a paradigm for enabling the cooperation among these organizations [20]; they can be conceptualized in terms of *multi-party business processes*: every actor performs some internal tasks (private view) and communicates with the other actors if some information is needed to perform the internal tasks or if some results need to be notified to make the others able to perform their own tasks (external view, also referred to as choreography). Although

this communication is a great opportunity for organizations, the resulting interdependencies are also difficult to manage, especially when some failures occur: a party could stop working for internal reasons and all the parties which depend on the information that the failing one is responsible for, might fail as well, creating a domino-effect.

A proper design of *resilient business processes* becomes fundamental. Generally speaking, resilience concerns the ability for a system to cope with unplanned situations in order to keep carrying out its mission [6]. In particular, making a multi-party business process resilient means to help the organization to cope with the complexity of the processes and to avoid, limit or mitigate possible failures that might affect the technological infrastructure as well as the involved organizational structure [4].

Usually, satisfying resilience requirements is considered as a mainly run-time issue, as it is related to the ability to cope with *unplanned situations*. In the literature [22], several approaches have been proposed to keep business processes running even when some unplanned exceptions occur, by enacting countermeasures. If we focus on *what to do* in case of failure, this approach seems to be the only possibility. However, if we focus on *what is affected* when a failure occurs, some improvements can be done also at design-time.

The aim of this work is to propose a *design-time and data-centric approach* for improving the resilience of multi-party business processes. Data are considered as “first class citizens” of our approach, as their unavailability might determine the failure of the processes. Depending on the data characteristics and the impacts of their possible unavailability, we propose a way to classify process models in terms of resilience by defining a set of *levels of resilience*. To achieve this goal, instead of focusing on the process activities (control flow), thus modeling the process using an activity-centric notation like OMG BPMN – Business Process Model and Notation, we adopt OMG CMMN – Case Management Model Notation [21] as a basic notation, and we enhance it to better cope with the data life-cycle definition in a process.

The rest of the paper is organized as follows: Sect. 2 introduces a motivating case study – to be used all along the paper – in which resilience aspects are considered. Section 3 defines the concept of resilience in multi-party business processes and proposes an approach to specify different levels of resilience. Section 4 defines the modeling approach based on CMMN able to support the definition of business processes according to the proposed levels of resilience. Section 5 illustrates the relevant literature related to resilient processes, and Sect. 6 presents a critical discussion about our approach, threat to validity and possible extensions.

2 Motivating Example

Smart devices have been adopted by several organizations to increase the effectiveness of business processes. For instance, in the logistics domain, smart devices provide real-time monitoring of goods transportation in terms of their position

or state (e.g., temperature, humidity). Although the advantages of the adoption of smart devices are clear, there are also some side-effects in terms of system reliability. In fact, smart devices are prone to failure due to their limitations in terms of computational power and energy autonomy. Moreover, in some cases they are operating in extreme conditions (e.g., meteorological stations on top of mountains), thus they might stop working without any previous notice.

Implications of the use of sensors in processes are illustrated through the example shown in Fig. 1, illustrating a real case study involving the *ShopAnalyser* company and *Shop Inc.*, one of its clients.




Fig. 1. Running example overview.

The ShopAnalyser company offers products and services to physical shops/commercial centers willing to monitor and analyze the behavior of their customers while they are walking inside their premises. To this aim, ShopAnalyser sells innovative sensors able to capture the probe packets periodically sent by cellphones and to localize and track the position of cellphones. In this way, assuming that a cellphone belongs to exactly one customer, the sensor is able to track the behavior of the customer inside the area and, correlating the MAC addresses, it realizes when the same customer periodically visit the shop. The analytics required to understand the customers' behaviors are offered by ShopAnalyser as a service to all the shops which buy its sensors. More specifically, ShopAnalyser produces one report every week to the shops, and they use these reports as a basis for defining or improving their marketing strategies.

Shop Inc. decides to acquire sensors and the analytics service from ShopAnalyser. The owner of Shop Inc., through its maintenance personnel, is responsible for the installation and physical maintenance of the sensors: ShopAnalyser delivers the sensors to Shop Inc., which installs them in the shop and configures them to send collected data to the data center of ShopAnalyser. Some status leds are embedded in the sensors to make the owners of the shops aware about possible malfunctioning: problems in the behavior of the sensors, when the probe packets sent by the cellphones are not collected correctly (in this case the Shop informs the ShopAnalyser, which will enact some repair action, such as sending substituting sensors), or to signal connection problems (i.e., the sensor are working, however the data cannot be sent to ShopAnalyser). ShopAnalyser is responsible for the data analysis, which produces a weekly report, and for the identification of the sensors malfunctioning which cannot be detected directly by the shops,

i.e., data captured by sensors and sent to the data center which are unrealistic (e.g., one hundred cellphones identified in the same tiny shop at the same time).

Although some actions are in place to cope with the malfunctioning of sensors, in the case study the focus is mainly on signaling possible failures: e.g., if a sensor stops working then a replacement is provided; if the network connection is interrupted, then the ISP – Internet Service Provider – is called to resume the connection. Actually, these occurring failures could have a more significant impact as they affect the data availability. In fact, during the down time, an amount of sensor data is not collected so it is not represented in the data set used for the analysis. As a consequence, the report used for marketing purposes might become not realistic.

To model multi-party business processes, as the one of the case study, activity-centric modeling languages such as BPMN are usually adopted. Even if this type of languages results more intuitive for the process designers, this approach has some limitations wrt specifying process resilience. As an example, the order of activities during exception handling is loosely specified: when addressing process resilience, the designer should specify recovery activities, and the order in which they are performed is usually decided at run-time based on considerations about the status of the process. Other approaches, as declarative modeling, rely on an open-world assumption, thus leaving room for supporting situations that cannot be planned at design-time [9]. In this work we adopt an artifact-based language, i.e., CMMN – Case Management Model and Notation [21], which aims to become the de-facto standard for artifact-based modeling. However, as discussed in the next section, also this language has some limitations when defining the data aspects, thus next sections will also propose some possible extensions.

Figure 2 presents the CMMN model of the ShopAnalyser case study¹. The outer box “Shop Improvement” represents the *case plan model*, i.e., the complete behavior of the process. Inside the case plan model, there are three *stages*: “Sensor data acquisition”, “Data analysis”, and “Marketing analysis”. Stages can be informally defined as a group of *tasks* (drawn as rounded boxes) organized according to an implicit or explicit control flow. Stages could also be decorated with *entry* and/or *exit conditions* represented, respectively, by an empty or filled diamond that specifies Boolean expressions predicating on data managed by the tasks in the stage or some events to occur. When these conditions become true, the stage opens (in case of entry condition) or terminates (in case of exit condition). Entry and exit conditions can be applied to tasks, stages, and case plan models. As an example, the “Reading values” task starts only when the sensors have been installed. The “Data analysis” stage opens every week and terminates when a new report is produced by the “Data mining” task. Finally, once the conversion rate² obtained by executing all the activities is considered suffi-

¹ The following discussion of the CMMN diagram does not intend to be exhaustive. For a complete description of the standard, the reader should refer to the official specification document [21].

² In marketing, the conversion rate measures the ratio between visitors and effectively paying customers. Therefore it is a measure of the goal of Shop Inc.

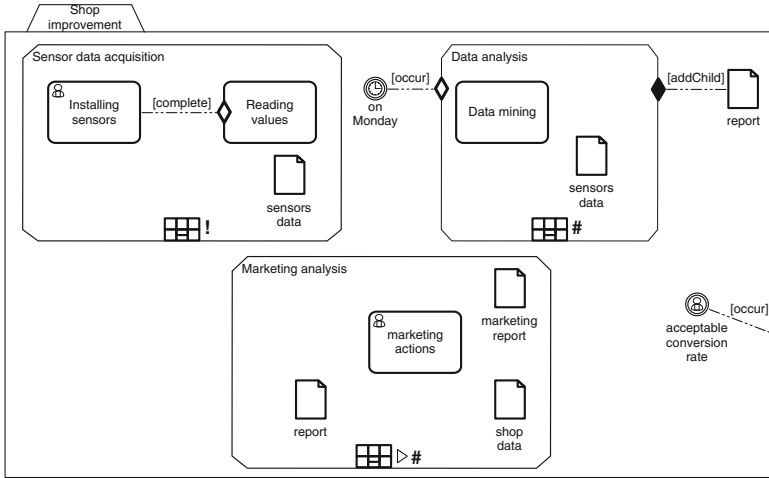


Fig. 2. CMMN diagram of the case study process

cient, then the business process concludes. Finally, *case plan items* (i.e., “Sensors data”, “Report”, and “Shop data”) are included in the stages which use them. It is worth noting that, according to the reported diagram, since the moment in which the sensors have been installed, the sensor reading task keeps running till the time in which the expected conversion rate has been achieved. At the same time, the marketing analysis is not coordinated with the other activities as it is performed by analysing the reports produced by the ShopAnalyser.

3 Multi-party Business Process Resilience

During the process enactment unplanned situations might occur. Depending on the nature of the raised issues, the magnitude of their impact varies and one or more activities may be involved. At the same time, different countermeasures can be taken to mitigate these negative effects. As an example, as for many reasons the sensors might not be able to communicate with ShopAnalyser, an alternative source of information about the number of clients in the shops might be considered, to be able to equally infer customers’ behaviors in the reports. Alternative ways to collect such information may include the ability of counting the number of persons entering the shop, which might be available from other unrelated applications, such as video surveillance. In this way, ShopAnalyser will not have gaps in the analysis, but only lower quality data. Other ways to improve the final reports may include algorithms to fill in the gaps of sensor information, based for instance on sales prediction algorithms applied when sensor data have not been collected.

Similarly to what is usually done in emergency management [17,27], where a *preparedness phase* aims to improve the systems by learning from the previous emergencies, we propose an approach which helps the process designers in

improving their process models by considering the previous experiences in failures generated by data unavailability. In particular, we propose an approach to categorize resilience characteristics, then to define resiliency levels, and to model the resilience improvement aspects from a modeling perspective.




Fig. 3. Problem setting

3.1 Data Perspective on Resilience

As previously introduced, our approach analyzes the multi-party business process resilience from a *data perspective*: data dependencies among the involved parties and relationships between process activities and data are taken into account to identify the *sources of possible failures*, and how the process can be better modeled to make it resilient with respect to these failures.

To this aim, in order to set the boundaries of our problem, we define a multi-party business process in terms of (see Fig. 3):

- *Parties*: actors involved in the process. Each of them participates in the business process to achieve a personal goal. All the parties are interested in making the process up and running without problems, as their personal goals also depends on the resilience of the whole process. As an example, Shop Inc. wants to make the marketing strategy more effective by increasing the conversion rate. On the other side, ShopAnalyser wants to sell a good service to its customers. Although the concept of *role* that is related to process participants is included in the CMMN standard, no graphical notation able to explicitly include parties is defined as of today. In this paper, we do not address this issue of lack of graphical constructs for parties, thus we do not propose any extension concerning the modeling of parties.
- *Tasks*: a task is a unit of work performed by a party, which consumes data as input and produces data as output. The data produced by a task must be required by at least another party. In multi-party business processes, we are more interested in the dependencies among the parties, rather than to internal executions of processes by each party, thus we are not including tasks which are internal to a single party.

- *Data*: units of storage used by the data producer to store/write data and by the data consumer to read such data. Producers and consumers are parties performing tasks. Data can also be used to verify the entry and exit conditions, thus to realize when a stage or task starts or terminates.

Resilience of this type of processes depends on both the reliability of the tasks and the lack of data availability. The *reliability of the task* concerns the possibility that one or more tasks cannot be executed: i.e., the required infrastructure to perform the job is not available, also including the human resources for which the unavailability of data can block the execution of manual tasks. On the other side, *lack of data availability* is a situation in which the data consumed by a task are not available. This situation can occur for different reasons. Firstly, it may be directly connected to the task reliability, as all the tasks by definition produce data and these data are relevant for at least one of the participating parties, and problems on tasks may have also the side effect to make data unavailable. Moreover, there are situations in which tasks are properly working, but the returned data, although available, do not have a sufficient quality level to enable processing, thus they can be considered unavailable. Completeness, timeliness, and accuracy are some of quality parameters through which we can define the acceptable level of data quality for considering the data available [5]. For this reason, the definition of the data could be coupled with the definition of quality levels that are considered acceptable for a task that is using such data.

3.2 Levels of Resilience

Having bounded our space of analysis and identified the possible sources of failure, we aim to classify multi-party business processes in terms of their degree of resilience. We define levels of resilience on the basis of the ability of the multi-party process to adjust the possible unexpected failures. As it will be discussed in Sect. 5, other proposals in the literature have been put forward to define resilience for processes, e.g., [28]. However, here we do not focus on the structure of the process or its components and instances, but we aim to classify the way resilience can be considered and obtained, in terms of preparedness to unexpected events which might be caused or have impact on data availability. In particular, the following four levels of designed resilience have been identified:

- **Level 0 – None.** At this level business processes are designed without taking into account the data unavailability that might cause failures during the execution. As a consequence, also countermeasures to be adopted in case of critical situations are not defined. The designed process only reflects the wishful scenario where it is assumed that all the parties correctly execute their tasks and all the data are transferred among them as expected. Although a process design of this type can be useful to define the agreement between the parties, no support is given to the resilience.
- **Level 1 - Failure-awareness.** A first step for improving the process design is to make the process aware that there are possible sources of failure, so there

will be the need to make it resilient. In this work, we consider failures caused by data unavailability, which might impact on one or more tasks of the same party that is producing such data, or tasks performed by other parties. For this reason, failure-aware business processes are designed to have a clear map of which are the relevant data subject to failures, as well as the impact of these failures. The analysis of potential failures depends on several factors: amount of data, how the data are collected, how the data are stored. As an example, data stored on a local server have a probability of failure that is lower than data stored on a smart device connected to a wireless network. Similarly, if data created by one party and used by several parties becomes unavailable, the impact of this failure will be greater than the one produced by data created and consumed by the same party.

- **Level 2 – Identifying alternatives for data and goals.** For processes classified in this level, the model of the process makes an initial attempt to overcome possible failures, whose nature and impact have been defined with the previous level. In more detail, there are two aspects to be taken into account:
 - *Alternative Data:* based on the information about the source of failures and the potential impact of these failures, the designer can decide to include in the process model the alternative data. In this way, starting from the data having more probability of failures and greater impact, the designer has to specify if there are alternative data sources and how to reach them. A more precise model requires an analysis of the gap between the quality of the data in the original data source with respect to the quality of the data in the alternative data source. For instance, in case the sensors installed in Shop Inc. stops working, the process model indicates as an alternative source other services, e.g. installed door counter and/or Google Popular Times or even historical data stored in a different, but accessible, place. The issue of quality of data has been extensively addressed in traditional information systems, e.g., [5], but the quality of big data (which includes sensor-generated data) is still to be precisely defined [10].
 - *Alternative Goal:* as the process resilience implies to mitigate the effect of a failure, a possible mitigation include revising the initial expectations of the process to achieve a given goal. The designer defines, for each party, a new goal that represents a status that can terminate the execution of the process in an acceptable way. If the initial goal corresponds to the optimal goal, the alternative goal could be considered as a best-effort goal. As an example, ShopAnalyser realizing that the data coming from the sensors contain errors, instead of releasing full reports with all details, it can decide to release for a reduced-price an incomplete report.

It is worth noting that the business process models at this level do not prescribe any specific actions to cope with the failures at run-time. For this reason, a model at this level only supports who is in charge of executing the process, to select, in case of failures, new data sources as well as to decide to

consider satisfactory the result of the execution even if the initial goal is not possible to be fulfilled, accepting a weaker goal.

- **Level 3 – Defining alternative actions.** At this level, processes have been designed by considering also actions to be taken in case of failures. Design-time mechanisms are conceived to be able to (semi)-automatically move the process to an acceptable state when unexpected or unplanned failures occur. Based on the information about the alternatives (both data and goal), the designer can embed in the business process how these alternatives could be effectively managed. New tasks can be added to the process to express the activities to be performed in order to improve the quality of the data alternatives to a quality level equivalent to the original service. Taking as example the problems of missing data, the previous level suggests to include the door counter and the Google Popular Times in the list of possible alternatives. At this level, the process designer should specify if the alternative data should be considered as they are produced, or if additional actions must be taken, e.g., to combine both services into a reliable assessment of the indoor occupancy for Shop Inc.

With these levels of resilience, we aim at supporting the process designer in understanding if the resilience is modeled, and if there is room to improve the process model by specifying possible alternative solutions. As an example, once the designer understands that the modeled processes are at level 0, the first step should be to start considering the evolution of the data in the process.

4 Modeling Resilience

In this section we discuss, for each level previously introduced, which is the practical impact of using CMMN as modeling language. In this way, we are able to highlight which are the current possibly missing constructs and their semantics. Moreover, we propose an extension of CMMN able to improve the specification of which data are used and in which way, in order to better analyze the possible failures and the impacts. Concerning the extensions proposed hereafter, at this stage, we do not intend to be complete and formal. Our attempt is to verify that CMMN has the potentiality for being used to model resilient business processes. A precise definition of the new constructs will be considered in future work.

Level 0 - None. CMMN standard is sufficient to express the basic scenario where resilience is not considered at all. The model of the business process for the ShopAnalyser case study, shown in Fig. 2, belongs to this level.

Level 1 - Failure Awareness. One of the main shortcomings of CMMN is the poor semantics about data. In the current version, data are defined in terms of *CaseFileItems* with no restrictions about the format and the nature of the represented data. On the one side, this allows maximum flexibility in modeling various scenarios. On the other side, no information about the link between tasks and data is provided, unless data are attached to the entry and exit conditions as predicates in the boolean expressions.




Fig. 4. Level 1 (failure awareness) compliant process model.

To overcome this limitation, we propose to extend CMMN allowing the connections between tasks and CaseFileItems also be annotated with the actions performed on the data: e.g., *create*, *read*, *update*, *delete*. It is also possible to link the data to the events that are defined in terms of these data (i.e., to predicate on). The use of this extension in the case study is shown in Fig. 4. The new elements in the model allow the designer to identify the data that might have impact in case of their unavailability, e.g., the lack of sensors' data will have more impact than the lack of the shops' data, as the former can cause a domino effect affecting all the tasks in the process.

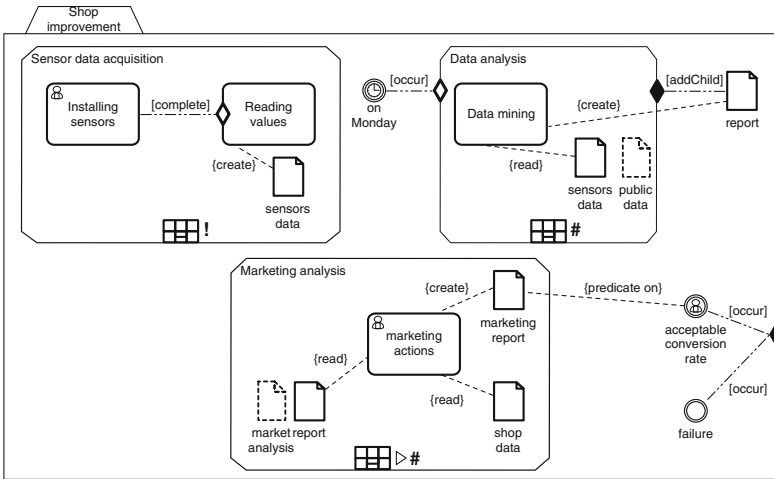


Fig. 5. Level 2 compliant process model.

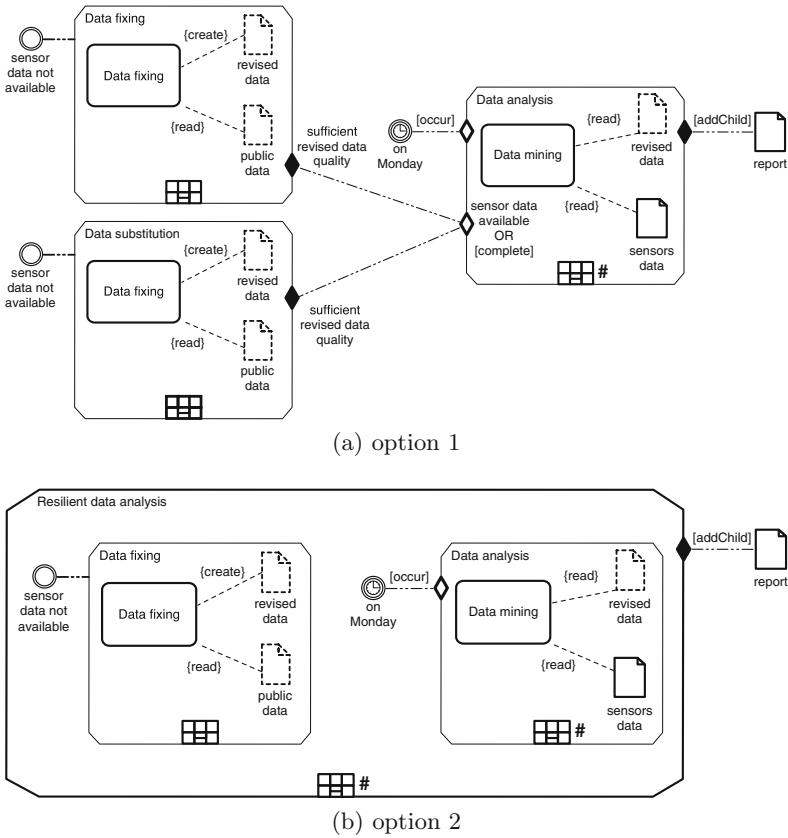


Fig. 6. Level 3 compliant process models.

Level 2 - Identifying Alternatives for Data and Goals. To cope with *alternative data*, we propose to add a new icon with a shape identical to a CaseFileItem, but with a dashed border strictly attached to the original data source. Conversely, the definition of the *alternative goals* does not require any extension to CMMN, as the usage of events that define the existence of a failure can be combined with the expression defining the alternative goal.

In the example in Fig. 5, two alternative sources are defined: public data as alternative for the sensor data and public market analysis to be used instead of the report produced by the data analysis task.

Level 3 - Defining Alternative Actions. Figure 6 shows two possible process models which exploit the CMMN extension proposed above to increase process resilience. For this level, we do not need to add further constructs to CMMN. In the first case, reported on top of the figure, the designer is assuming that in case of failure in acquiring the sensor data, the data analysis task cannot be executed until either “Data fixing” or “Data substitution” has terminated. In particular, exploiting the existence of alternative data sources, the data substitution simply

replaces the data source. This task can be considered concluded only if the quality of the data now provided is considered sufficient for the data analysis. On the other side, the data fixing implements data quality algorithms to improve the data quality as required by the data analysis. It has to be noted that, according to this model, the data analysis potentially might never start.

The process designer could also propose a different approach, shown in the lower part of the figure, where the data fixing and data analysis are included in the same stage. In this case, data analysis and data fixing work in parallel trying to achieve a common goal, i.e., the report delivery.

5 Related Work

Research on *resilient systems* encompasses several disciplines, such as psychology [29], ecology [11], sociology [3] and engineering [14]. In information systems, *resilience engineering* has its roots in the study of safety-critical systems [14], i.e., systems aimed to ensure that organizations operating in turbulent and interconnected settings achieve high levels of safety despite a multitude of emerging risks, complex tasks, and constantly increasing pressures. A system is considered as resilient if its capabilities can be adapted to new organizational requirements and changes that have not been explicitly incorporated into the existing system's design [19]. In the BPM field, cf. [19,23], this means that respective business processes are able to automatically adapt themselves to such changes. Over the last years, change management in BPM has been mainly tackled through the notions of *process flexibility* [22] and *risk-aware BPM* [25,26].

On the one hand, research on process flexibility has focused on four major flexibility needs, namely (i) *variability* [12,13], (ii) *looseness* [2,16], (iii) *adaptation* [18,24], and (iv) *evolution* [7,8]. The ability to deal with changes makes process flexibility approaches a required but not sufficient mean for the building of resilient BPM systems. In fact, there exists a (seemingly insignificant but) relevant gap between the concepts of flexibility and resilience: (i) process flexibility is aimed at producing “reactive” approaches that reduce failures from the outset or deal with them at run-time if any “known” disturbance arises; (ii) process resilience requires “proactive” techniques accepting and managing change “on-the-fly” rather than anticipating it, in order to allow a system to address new emerging and unforeseeable changes with the potential to cascade. On the other hand, while relatively close to the concept of risk-aware BPM, which evaluates operational risks on the basis of historical threat probabilities (with a focus on the “cause” of disturbances and events), resilient BPM shifts attention on the “realized risks” and its consequences, to improve risk prevention and mitigation, and therefore aim at complementing conventional risk-aware approaches.

Surprisingly, the fact is that there exists only a limited number of research works investigating resilience of BPM systems [4,30,31], and they are all at conceptual level. For example, the work of Antunes and Mourao [4] derives a set of fundamental requirements aimed at supporting resilient BPM. The approach of Zahoransky et al. [31] investigates the use of process mining [1] to create probability distributions on time behavior of business processes. Such distributions

can be used as indicators to monitor the level of resilience at run-time and indicate possible countermeasures if the level drops. Finally, the work [30] provides a support framework and a set of measures based on the analysis of previous process executions to realize and evaluate resilience in the BPM context.

If compared with the aforementioned works, our research aims at providing concrete indicators to measure the resilience of a multi-party business process by focusing on the data exchanged between the activities composing the process, an aspect neglected in the existing approaches to process resilience. We believe that such indicators can provide a reliable mean for evaluating in advance the impacts of potential disturbances and improving decision making at run-time.

6 Discussion

The levels of resilience presented in this paper, and the practical guidelines on how to achieve them during the design of processes, namely by precisely modeling in CMMN, are a concrete methodological tool to support process designers to be aware of how resilient are the processes they are working on. At design-time, it is important to be aware of failures, and to identify data and goal alternatives, in order to be able to design alternative actions. On the one side, flexible approaches cope with exceptional situations during run-time, but only a deep awareness during design-time can make really the process *resilient-by-design*.

Clearly our work should be extended and validated in many aspects. However we consider it as an important starting point in deeply investigating how to make better resilient processes. On the one side, a precise formalization of the modeling constructs to be used in order to achieve each level, and patterns to be used, is crucial in order to make the overall approach effective. On the other side, a validation is needed, in which to compare, by adopting empirical approaches [15], processes at different levels and the real resilience they achieve during enactment. Measuring resilience of multi-party business processes is not an easy task, and no measurable indicators exist nowadays in this context. Our aim is to be able to correlate our levels with a qualitative notion of “a process is more resilient of another one”, and this is only possible through a large collection of case studies (models and execution traces) on top of which to perform quantitative correlation analysis. To this aim, the levels of resilience introduced in this paper go in the direction of providing a reference framework which represents an important input to the research and practitioners’ community. In fact, adopting and extending a well known standard, i.e., CMMN, gives the opportunity to develop approaches able to provide this quantitative analyses.

7 Concluding Remarks

In this paper we have discussed the concept of multi-party resilient processes, and we have presented a possible way of classifying them on the basis of four levels, based on how data and goals are taken into account when considering possible ways to cope with changes. The originality of the proposed approach is

in considering resilience at design-time, during process modeling, and not mainly as a run-time issues, when exceptions and anomalous events should be faced during enactment. We have shown a practical way to achieve the levels during modeling, by using and extending the newly introduced standard CMMN for artefact-centric processes. After discussing relevant work, we have provided a discussion about the limitations and possible extensions of our work, which is a promising initial step towards defining effective resilient processes.

Acknowledgments. This work is partly supported by the Italian projects Social Museum e Smart Tourism (CTN01.00034.23154), NEPTIS (PON03PE.00214.3), RoMA - Resilience of Metropolitan Areas (SCN.00064), ITS2020 (CTN01.00176.166195), and by the Sapienza project “Data-aware Adaptation of Knowledge-intensive Processes in Cyber-Physical Domains through Action-based Languages”.

References

1. van der Aalst, W.M.P.: Process Mining: Data Science in Action. Springer, Heidelberg (2016)
2. van der Aalst, W.M.P., Pesic, M., Schonenberg, H.: Declarative workflows: Balancing between flexibility and support. *Comput. Sci. - R&D* **23**(2), 99–113 (2009)
3. Adger, W.N.: Social and ecological resilience: are they related? *Prog. Hum. Geogr.* **24**(3), 347–364 (2000)
4. Antunes, P., Mourõ, H.: Resilient business process management: framework and services. *Expert Syst. Appl.* **38**(2), 1241–1254 (2011)
5. Batini, C., Scannapieco, M.: Data and Information Quality - Dimensions, Principles and Techniques. Springer, Heidelberg (2016)
6. Caralli, R.A., Allen, J.H., White, D.W.: CERT Resilience Management Model: A Maturity Model for Managing Operational Resilience. Addison-Wesley, Boston (2010)
7. Casati, F., Ceri, S., Pernici, B., Pozzi, G.: Workflow evolution. *Data Knowl. Eng.* **24**(3), 211–238 (1998)
8. Dadam, P., Rinderle, S.: Workflow evolution. In: Liu, L., Özsu, M.T. (eds.) *Encyclopedia of Database Systems*, pp. 3540–3544. Springer, Heidelberg (2009)
9. Fahland, D., Lübke, D., Mendling, J., Reijers, H.A., Weber, B., Weidlich, M., Zugal, S.: Declarative versus imperative process modeling languages: the issue of understandability. In: Halpin, T., Krogstie, J., Nurcan, S., Proper, E., Schmidt, R., Soffer, P., Ukor, R. (eds.) *BPMDS/EMMSAD-2009*. LNBIP, vol. 29, pp. 353–366. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-01862-6_29](https://doi.org/10.1007/978-3-642-01862-6_29)
10. Firmani, D., Mecella, M., Scannapieco, M., Batini, C.: On the meaningfulness of “Big Data Quality”. *Data Sci. Eng.* **1**(1), 6–20 (2016)
11. Gunderson, L.H.: Ecological resilience-in theory and application. *Ann. Rev. Ecol. Syst.* **31**, 425–439 (2000)
12. Hallerbach, A., Bauer, T., Reichert, M.: Capturing variability in business process models: the Provop approach. *J. Softw. Maint. Evol. Res. Pract.* **22**(6–7), 519–546 (2009)
13. vom Hallerbach, A., Bauer, T., Reichert, M.: Configuration and management of process variants. In: Brocke, J., Rosemann, M. (eds.) *Handbook on Business Process Management*. International Handbooks on Information Systems, vol. 1, pp. 237–255. Springer, Heidelberg (2010)

14. Hollnagel, E., Woods, D.D., Leveson, N.: Resilience Engineering: Concepts and Precepts. Ashgate Publishing, Ltd., Farnham (2007)
15. Johannesson, P., Perjons, E.: An Introduction to Design Science. Springer, Heidelberg (2014)
16. Marrella, A., Lespérance, Y.: Synthesizing a library of process templates through partial-order planning algorithms. In: Nurcan, S., Proper, H.A., Soffer, P., Krogstie, J., Schmidt, R., Halpin, T., Bider, I. (eds.) BPMDS/EMMSAD-2013. LNBIP, vol. 147, pp. 277–291. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-38484-4_20](https://doi.org/10.1007/978-3-642-38484-4_20)
17. Marrella, A., Mecella, M., Russo, A.: Collaboration on-the-field: suggestions and beyond. In: 8th International Conference on Information Systems for Crisis Response and Management (ISCRAM) (2011)
18. Marrella, A., Mecella, M., Sardiña, S.: Intelligent process adaptation in the SmartPM system. ACM TIST **8**(2), 25 (2017)
19. Müller, G., Koslowski, T.G., Accorsi, R.: Resilience - a new research field in business information systems? In: Abramowicz, W. (ed.) BIS 2013. LNBIP, vol. 160, pp. 3–14. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-41687-3_2](https://doi.org/10.1007/978-3-642-41687-3_2)
20. Nachira, F., Nicolai, A., Dini, P., Le Louarn, M., Rivera Leon, L. (eds.): Digital Business Ecosystems. European Commission DG Information Society and Media, Office for Official Publications of the European Communities, Luxembourg (2007). ISBN 92-79-01817-5
21. OMG: Case Management Model and Notation, Version 1.0, May 2014. <http://www.omg.org/spec/CMMN/1.0>
22. Reichert, M., Weber, B.: Enabling Flexibility in Process-Aware Information Systems - Challenges, Methods, Technologies. Springer, Heidelberg (2012)
23. Rosemann, M., Recker, J.: Context-aware process design exploring the extrinsic drivers for process flexibility. In: 7th International Workshop on Business Process Modeling, Development, and Support (BPMDS) (2006)
24. Sadiq, S.W., Orlowska, M.E.: On capturing exceptions in workflow process models. In: Abramowicz, W., Orlowska, M.E. (eds.) BIS 2000, pp. 3–19. Springer, London (2000)
25. Suriadi, S., Weiß, B., Winkelmann, A., et al.: Current research in risk-aware business process management: overview, comparison, and gap analysis. Commun. Assoc. Inf. Syst. **34**(1), 933–984 (2014)
26. Tjoa, S., Jakoubi, S., Goluch, G., Kitzler, G., Goluch, S., Quirchmayr, G.: A formal approach enabling risk-aware business process modeling and simulation. IEEE Trans. Serv. Comput. **4**(2), 153–166 (2011)
27. Van De Walle, B., Turoff, M., Hiltz, S.R.: Information Systems for Emergency Management. M.E. Sharpe, Armonk (2009)
28. de Vrieze, P., Xu, L.: Resilience analysis of collaborative process management systems. In: Afsarmanesh, H., Camarinha-Matos, L.M., Lucas Soares, A. (eds.) PRO-VE 2016. IAICT, vol. 480, pp. 124–133. Springer, Cham (2016). doi:[10.1007/978-3-319-45390-3_11](https://doi.org/10.1007/978-3-319-45390-3_11)
29. Yates, T.M., Masten, A.S.: Fostering the Future: Resilience Theory and the Practice of Positive Psychology. Wiley, Hoboken (2004)
30. Zahoransky, R.M., Brenig, C., Koslowski, T.: Towards a process-centered resilience framework. In: 10th International Conference on Availability, Reliability and Security (ARES), pp. 266–273. IEEE (2015)
31. Zahoransky, R.M., Koslowski, T., Accorsi, R.: Toward resilience assessment in business process architectures. In: Bondavalli, A., Ceccarelli, A., Ortmeier, F. (eds.) SAFECOMP 2014. LNCS, vol. 8696, pp. 360–370. Springer, Cham (2014). doi:[10.1007/978-3-319-10557-4_39](https://doi.org/10.1007/978-3-319-10557-4_39)