Adaptive Online Optimization of Business Processes

Maksym Avramenko[1]

[1]University of Tartu, Narva mnt 18, 51009 Tartu, Estonia

Abstract

Operational business process optimization aims to improve the performance of ongoing cases in near real time, e.g., by reallocating resources. While simulation can be useful to support such decisions, most simulation approaches focus on long-term analysis, where stochastic effects average out over time and the precise starting state of the simulation does not have a major effect on the simulation forecasts. To apply simulation in short-term operational settings, it is crucial to model the process in its exact current state and to quantify the uncertainty of the simulation forecasts. This doctoral project adopts a three-phased approach to address these challenges. First, we will develop a method to start a simulation from the current state of a process, discovered from an event log. Second, we will design accuracy metrics and uncertainty modeling techniques to evaluate the reliability of short-term simulation forecasts. Finally, we will explore how short-term simulations can support the optimization of ongoing cases of a process. In this way, the doctoral thesis will contribute to closing the gap between real-time process monitoring and data-driven operational decision support in business process management.

Keywords

Business Process Optimization, Business Process Simulation, Business Process Management, Process Mining

1. Introduction

Business process optimization is a discipline within Business Process Management (BPM), where analysts seek to improve performance metrics such as cycle time, resource utilization, and throughput [1, 2]. Typically, analysts concentrate on *long-term or tactical changes* by, for example, adding steps to a workflow, deploying automation, or reorganizing staffing, that require multiple weeks or months to implement and evaluate [3]. The benefits of these modifications usually appear only after the process reaches a *steady state*. This means that key performance indicators, such as average throughput, become stable. At this point, the number of active *process instances*, or so-called *cases*, also levels out. Each case represents a single execution of the process for a specific customer, transaction, or request.

Recent studies have highlighted the growing relevance of *operational, near-term decisions* in BPM [4, 5]. Simulation is a common approach used to support decision-making in such settings, as it allows analysts to evaluate the impact of alternative actions before they are implemented [6]. While simulation is traditionally applied for long-term planning, it can also be used to support short-term optimization. In contexts where disruptions arise suddenly, such as a surge in incoming requests or a shortfall in available resources, organizations may need to make immediate adjustments [3]. In these cases, *short-term simulation* enables analysts to forecast near-future outcomes by modeling the process exactly as it stands at a specific point in time. For example, in loan application process, it can estimate clearance times under alternative staffing or scheduling [7].

To assess the performance of business processes, simulation is often used as an analytical tool. Simulation models are typically evaluated using *event logs*, which record the execution history of a business process. These logs are composed of timestamped activity events, grouped by individual cases, capturing which activities were started or completed, by whom, and when. Simulation models that assume a *steady state* face two main limitations [7, 8]. First, they ignore cases that are already active when the simulation starts. Second, they fail to account for recent changes in the process, such as workload spikes or resource shifts. To address this, many conventional approaches use a method called *warm-up*. In a warm-up, the simulation starts from an empty state and runs until the system reaches a

ICPM Doctoral Consortium and Demo Track 2025, October 20-24, 2025, Montevideo, Uruguay

maksym.avramenko@ut.ee (M. Avramenko)
0009-0004-0725-7623 (M. Avramenko)

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

steady state. Once the process stabilizes, performance metrics are recorded, assuming that the initial imbalance no longer affects the results. However, this method does not fully reflect the real current state of the process. Although this assumption is suitable for long-term analysis, it frequently breaks down when the objective is to predict short-term performance, such as for a few days or weeks [9]. Ignoring partially completed cases and ongoing resource allocations may lead to under- or overestimation of short-term outcomes. Failure to capture these *ongoing states* can distort simulation results [10].

To address these limitations, we aim to develop a *short-term simulation and optimization framework* that initializes directly from the *real process state*, thereby incorporating partially executed tasks and current resource allocations. Our work is structured around three main research directions. First, we introduce a technique for *process state computation* that reconstructs the exact execution state of ongoing cases from *partial event logs*, meaning live, ongoing logs that contain only the activities executed so far for each case. This method enables a simulation engine to continue execution from an accurate, real-time state rather than from an artificial or empty initial condition. Second, we define *evaluation metrics* to measure the accuracy of short-term simulation results and introduce methods to model and quantify *uncertainty*, helping assess how reliable the predictions are. Third, we investigate how short-term simulation can be used to support *optimization actions*, such as real-time resource reallocation or task prioritization. By simulating potential interventions in advance, decision-makers can select those that are most likely to mitigate workload spikes or emerging bottlenecks.

2. Related Work

Developing a system for short-term operational optimization involves three core contributions: (1) computing the current state of a process and initializing short-term simulation from ongoing cases; (2) evaluating the accuracy and uncertainty of short-term simulation outcomes; and (3) leveraging simulation results to optimize the process at runtime through targeted interventions.

Simulations from the current process state was introduced in [3], by using data from WFMS or ERP systems to build simulation models and extract real-time states. This idea was extended using event graphs in [7, 8, 11]. The usefulness of short-term simulation has been shown in e.g., healthcare [5, 4] and service operations [12]. However, they assume that the current process state is either fully available (e.g., through a WFMS) or manually configured. An alternative is to approximate the starting state from historical data. For instance, Pourbafrani et al. [9] proposed estimating periodic ”average” states to initialize simulation runs. While this reduces initialization bias, it assumes process stability and fails to capture sudden deviations or process drift. In contrast, our approach derives the current state directly from ongoing cases, allowing for more realistic initialization in dynamic environments.

Short-term simulation is useful if its predictions are accurate and trustworthy. Most BPS tools, however, rely on aggregate metrics designed for long-term analysis, such as average cycle time or throughput. These metrics do not reflect short-term dynamics and are not sensitive to inaccuracies at the case level. Some prior work evaluates simulation logs using control-flow abstractions [13] or trace similarity measures [14], but are not for the time-sensitive nature of short-term predictions.

In parallel, research in predictive process monitoring has introduced tools for estimating uncertainty, such as prediction intervals and confidence bounds [15]. However, these are rarely integrated into simulation environments and are often limited to point predictions (e.g., next activity or remaining time). We build on these ideas by defining accuracy metrics for ongoing cases and incorporating uncertainty quantification to capture both aleatoric (stochastic) and epistemic (model-based) uncertainty, using techniques inspired by ensemble modeling and probabilistic forecasting [16, 17].

Recent work in prescriptive process monitoring explore simulation-supported decision-making. For example, De Moor et al. [18] introduce SimBank, a benchmarking environment that enables the evaluation of prescriptive interventions through process simulations. However, most approaches either assume fixed process conditions or depend heavily on predictive models instead of leveraging simulation engines. Our research addresses this gap by combining online simulation with operational optimization, allowing decision-makers to evaluate, compare, and select short-term interventions, such as resource reallocation, task prioritization, or rerouting - based directly on live execution data.

3. Research Plan

The doctoral research will proceed in three phases (Table 1): (1) reconstructing the current state of ongoing processes from event logs; (2) measuring the reliability and uncertainty of short-term simulation outcomes; and (3) using these insights to support short-term operational optimization. Each phase progressively advances from accurate state estimation to simulation-based decision support.

Phase	Research Question & Contribution	Status
Phase 1: Process State Computation and Short-Term Simulation Initialization	**RQ1:** How can we reconstruct the current execution state of a business process from an event log of ongoing cases to support short-term simulation? **Contribution 1:** A formalization of the notion of process state for ongoing business process simulations, a method to discover this state from an event log of ongoing cases, and an approach to initialize a simulation engine directly from the discovered state.	Approach completed with results being evaluated.
Phase 2: Accuracy Metrics and Uncertainty Modeling	**RQ2:** How can we measure the quality and reliability of short-term simulation outcomes starting from the current state? **Contribution 2:** A set of evaluation metrics tailored to assess the quality of short-term simulation outcomes, and a method to model the uncertainty of these outcomes.	Planned for next phase of work
Phase 3: Operational Simulation-Driven Optimization	**RQ3:** How can simulation be used to support short-term optimization of ongoing processes by detecting performance drops, evaluating potential interventions, and guiding operational decisions? **Contribution 3:** A framework for using short-term simulation to detect potential performance issues in ongoing processes, evaluate the impact of alternative interventions, and support operational decision-making through targeted short-term optimizations.	Planned.

Table 1
Research phases, questions, and progress

3.1. Phase 1: Process State Computation and Short-Term Simulation Initialization

The first research question (**RQ1**) asks: *How can we reconstruct the current execution state of a business process from an event log of ongoing cases to support short-term simulation?*

Contribution 1 addresses RQ1 by introducing a formalization of the state of an ongoing business process simulation, a method to discover this state from an event log of ongoing cases, and an approach to initialize a simulation directly from the discovered state [19]. This allows short-term simulations to start from the actual execution context, instead of an artificial or empty starting point. Figure 1 illustrates the steps of the approach to discover the state of a business process.

The state computation approach uses two inputs: (1) a workflow graph (BPMN-style), and (2) an event log with ongoing traces. It follows three steps that produce a detailed state for each ongoing case, including arrival rate, active flows, enabled and running activities and resources.

1. **Estimate control-flow state:** Uses an *n-gram index* to link the last *m* activities of each case to control-flow markings.
2. **Estimate enablement information:** A *concurrency oracle* identifies causal relationships and enablement times.
3. **Derive case arrival, activity start and resource information:** Extracts start timestamps and resources from the log, as well as last case arrival information completing the execution state.

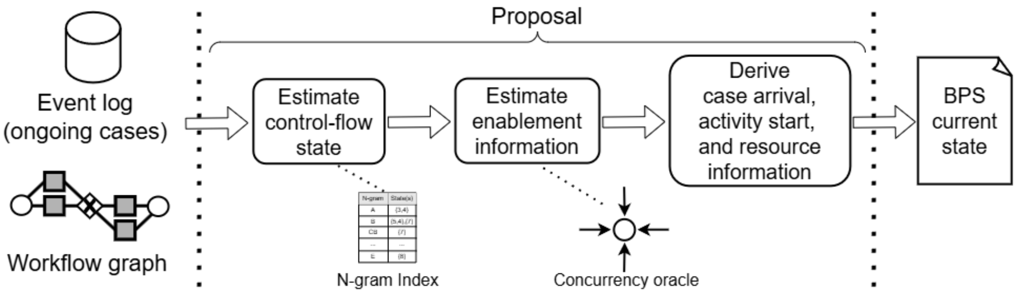

Figure 1: Overview of the proposed approach for reconstructing the process state.

The reconstructed state initializes the control-flow, enables activities at their actual enablement times, and queues ongoing activity completions with adjusted remaining times. Resource states are set directly from the log, ensuring that the simulation begins from the real execution state. We implemented this in Python and integrated it into the Prosimos simulation engine [20].

We evaluated the approach on both synthetic and real-world processes. For each process, we compared our approach ("ProcState") to a traditional "warm-up" simulation that reaches a similar workload level. We assessed accuracy with three metrics: ongoing case distance (OCD) which measures workload accuracy, n-gram distance (NGD) which measures control-flow sequence similarity and remaining cycle time distance (R-CTD) which compares remaining times of ongoing cases.

In *synthetic logs*, ProcState consistently achieved lower OCD and comparable or better NGD and R-CTD, especially in unstable/circadian workloads. In *real-life logs*, ProcState exactly matched the workload (OCD = 0) and achieved comparable control-flow and timing accuracy, though BPIC17 showed slightly less accurate timing due to unmodeled long delays.

These results confirm that initializing short-term simulation from the reconstructed state improves realism and reliability for operational decision support. Future work will explore more robust methods to estimate missing information and quantify the confidence of the reconstructed state. Extending support for advanced workflow constructs, such as inclusive gateways and event-based behaviors, is also an important direction for improvement.

3.2. Phase 2: Accuracy Metrics and Uncertainty Modeling

The second research question (**RQ2**) asks: *How can we measure the quality and reliability of short-term simulation outcomes starting from the current process state?*

Contribution 2 focuses on developing evaluation techniques that go beyond traditional performance metrics and incorporate a comprehensive view of simulation uncertainty. In long-term simulations, aggregated indicators such as average cycle time or throughput offer a stable performance picture. However, in short-term scenarios, even small inaccuracies in modeling ongoing cases or resource states can lead to substantial forecast errors. Our work in this phase includes two key directions:

- **Short-term accuracy metrics:** We will design specialized metrics to compare simulated and actual event logs over a limited prediction horizon. Examples include:
  - *Remaining events per case error*, measuring the deviation in expected case completion steps.
  - *Remaining cycle time error*, quantifying time prediction errors for ongoing cases.
  - Other measures focusing on case completion order consistency, resource assignment accuracy, utilization, etc.

  Short-term accuracy evaluation requires careful tuning of time horizons and other parameters. Metrics must be sensitive enough to detect relevant deviations, yet not overreact to noise.
- **Uncertainty modeling:** We will study methods to quantify and communicate both aleatoric and epistemic uncertainty in short-term simulation. While repeated simulation runs from the same reconstructed state can capture *aleatoric uncertainty* (i.e., inherent randomness), they do not

account for *epistemic uncertainty* - the uncertainty caused by model limitations or inaccuracies in the reconstructed process state.

Preliminary results from Phase 1 show that simulation performance is highly sensitive to such epistemic uncertainty. In particular, we observe rapid degradation of prediction quality when models have low accuracy. Therefore, in this thesis, we will explore methods to jointly model both types of uncertainty. To this end, we draw inspiration from recent work in the machine learning field that leverages ensemble-based techniques to capture and separate aleatoric and epistemic uncertainty [16, 17]. Our goal is to adapt these techniques to the context of business process simulation and evaluate their applicability for supporting reliable operational decision-making.

In summary, this phase aims to establish a foundation for trustworthy short-term simulation by developing both tailored accuracy metrics and structured uncertainty modeling methods. These tools will help evaluate how well simulation results align with reality and how much confidence can be placed in their predictions. This foundation is essential for the next phase, where simulation outcomes will be used to guide operational optimization decisions.

3.3. Phase 3: Simulation-Driven Optimization

The third research question (**RQ3**) asks: *How can simulation results be used to guide short-term optimization actions, such as resource reallocation or scheduling adjustments?*

Contribution 3 focuses on leveraging short-term simulation to support operational interventions, such as reassigning staff across resource pools or prioritizing certain cases, to mitigate workload spikes or delays. These interventions are simulated directly from the reconstructed state developed in Phase 1 and evaluated using the accuracy and uncertainty metrics defined in Phase 2.

The goal is to help decision-makers compare the effectiveness of intervention options and understand the uncertainty associated with their outcomes. A simulation might predict a significant performance gain, but if that forecast is highly uncertain, the intervention may not be reliable in practice. We will also assess the trade-off between predicted performance gains, uncertainty, and computational cost. Short-term optimization often requires simulating hundreds of candidate interventions, and each candidate may need to be simulated multiple times to estimate uncertainty. To address this challenge, we will study lightweight optimization strategies that identify interventions while minimizing the number of required simulation runs. Additionally, we will focus on specific categories of intervention types that occur frequently in practical settings (e.g., task reprioritization or dynamic resource reassignment).

Finally, we aim to evaluate the proposed framework using realistic case studies. Where possible, we will collaborate with industry partners to test the approach on historical execution data and obtain expert feedback. While validating on live processes is challenging, even semi-realistic or retrospective analyses can help refine the methods and assess their practical value.

4. Conclusion

This research focuses on short-term simulation and operational decision support in business process management. We propose a method to reconstruct the current state of ongoing processes from partial event logs. This allows simulation to start from real execution conditions instead of assuming a steady state. We also define metrics and model the uncertainty to assess forecast reliability and plan to use simulation to evaluate short-term optimization actions. By linking real-time monitoring with decision-making, this work supports better short-term process decisions.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] W. M. P. van der Aalst, Business process simulation survival guide, in: Handbook on Business Process Management 1, Introduction, Methods, and Information Systems, 2nd Ed, Springer, 2015, pp. 337–370.

[2] M. Dumas, M. L. Rosa, J. Mendling, H. A. Reijers, Fundamentals of Business Process Management, Second Edition, Springer, 2018.

[3] H. A. Reijers, W. M. P. van der Aalst, Short-term simulation: Bridging the gap between operational control and strategic decision making, in: IASTED International Conference on Modeling and Simulation (MS), 1999, pp. 417–423.

[4] L. Bedoya-Valencia, E. Kirac, Evaluating alternative resource allocation in an emergency department using discrete event simulation, Simul. 92 (2016) 1041–1051.

[5] F. Pegoraro, E. A. P. Santos, E. de Freitas Rocha Loures, G. da Silva Dias, L. M. dos Santos, R. O. Coelho, Short-term simulation in healthcare management with support of the process mining, in: WorldCIST Vol. 2, Springer, 2018, pp. 724–735.

[6] D. Chapela-Campa, M. Dumas, From process mining to augmented process execution, Software and Systems Modeling 22 (2023) 1977–1986.

[7] M. T. Wynn, M. Dumas, C. J. Fidge, A. H. M. ter Hofstede, W. M. P. van der Aalst, Business process simulation for operational decision support, in: BPM Workshops 2007, volume 4928 of *LNCS*, Springer, 2007, pp. 66–77.

[8] A. Rozinat, M. T. Wynn, W. M. P. van der Aalst, A. H. M. ter Hofstede, C. J. Fidge, Workflow simulation for operational decision support, Data Knowl. Eng. 68 (2009) 834–850.

[9] M. Pourbafrani, N. Lücking, M. Lucke, W. M. P. van der Aalst, Steady state estimation for business process simulations, in: BPM Forum 2023, volume 490 of *LNBIP*, Springer, 2023, pp. 178–195.

[10] D. Chapela-Campa, M. Dumas, Efficient online computation of business process state from trace prefixes via n-gram indexing, IEEE Transactions on Services Computing (2025) 1–15.

[11] Y. Liu, H. Zhang, C. Li, R. J. Jiao, Workflow simulation for operational decision support using event graph through process mining, Decis. Support Syst. 52 (2012) 685–697.

[12] I. Khodyrev, S. Popova, Discrete modeling and simulation of business processes using event logs, in: Proceedings of the International Conference on Computational Science (ICCS), volume 29 of *Procedia Computer Science*, Elsevier, 2014, pp. 322–331.

[13] S. J. J. Leemans, D. Fahland, W. M. P. van der Aalst, Discovering block-structured process models from event logs containing infrequent behaviour, in: N. Lohmann, M. Song, P. Wohed (Eds.), Business Process Management Workshops - BPM 2013 International Workshops, Revised Papers, volume 171 of *Lecture Notes in Business Information Processing*, Springer, 2013, pp. 66–78.

[14] D. Chapela-Campa, I. Benchekroun, O. Baron, M. Dumas, D. Krass, A. Senderovich, A framework for measuring the quality of business process simulation models, Inf. Syst. 127 (2025) 102447.

[15] I. Teinemaa, M. Dumas, M. L. Rosa, F. M. Maggi, Outcome-oriented predictive process monitoring: Review and benchmark, ACM Trans. Knowl. Discov. Data 13 (2019) 17:1–17:57.

[16] M. Chan, M. Molina, C. Metzler, Estimating epistemic and aleatoric uncertainty with a single model, in: A. Globersons, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. M. Tomczak, C. Zhang (Eds.), Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, 2024.

[17] L. M. C. Cabezas, V. S. Santos, T. Ramos, R. Izbicki, Epistemic uncertainty in conformal scores: A unified approach, CoRR abs/2502.06995 (2025).

[18] J. D. Moor, H. Weytjens, J. D. Smedt, J. D. Weerdt, Simbank: from simulation to solution in prescriptive process monitoring, CoRR abs/2506.14772 (2025).

[19] M. Avramenko, D. Chapela-Campa, F. Milani, M. Dumas, What's coming next? short-term simulation of business processes from current state, Proceedings of the 7th International Conference on Process Mining (ICPM 2025) - (2025). To be published.

[20] O. López-Pintado, M. Dumas, J. Berx, Discovery, simulation, and optimization of business processes with differentiated resources, Inf. Syst. 120 (2024) 102289.