



EdgeMiner: Distributed Process Mining at the Data Sources

Julia Andersen
Kiel University
Germany
jro@cs.uni-kiel.de

Patrick Rathje
Kiel University
Germany
pra@cs.uni-kiel.de

Christian Imenkamp
University of Bayreuth
Germany
christian.imenkamp@uni-bayreuth.de

Agnes Koschmider
University of Bayreuth
Germany
agnes.koschmider@uni-bayreuth.de

Olaf Landsiedel
Kiel University
Germany
ol@cs.uni-kiel.de

Abstract

Process mining is moving beyond mining traditional event logs and nowadays includes, for example, data sourced from sensors in the Internet of Things (IoT). The volume and velocity of data generated by such sensors make it increasingly challenging to efficiently process the data by traditional process discovery algorithms, which operate on a centralized event log. This paper presents EdgeMiner, an algorithm for distributed process mining operating directly on sensor nodes on a stream of real-time event data. In contrast to centralized algorithms, EdgeMiner tracks each event and its predecessor and successor events directly on the sensor node where the event is sensed and recorded. As EdgeMiner aggregates direct successions on the individual nodes, the raw data does not need to be stored centrally, thus improving scalability and privacy. We analytically and experimentally show the correctness of EdgeMiner. Our evaluation results show that EdgeMiner determines predecessors for each event efficiently, reducing the communication overhead by up to 96% compared to querying all nodes. Further, we show that the number of queried nodes stabilizes after relatively few events, and batching predecessor queries in groups reduces the average queried nodes per event to less than 2.5%.

CCS Concepts

• **Computer systems organization** → **Sensor networks**; • **Information systems** → **Data stream mining**; • **Security and privacy** → *Privacy-preserving protocols*.

Keywords

Distributed Process Mining, Event Data Stream, In-Network Processing, Scalability, Internet of Things

ACM Reference Format:

Julia Andersen, Patrick Rathje, Christian Imenkamp, Agnes Koschmider, and Olaf Landsiedel. 2025. EdgeMiner: Distributed Process Mining at the Data Sources. In *The 40th ACM/SIGAPP Symposium on Applied Computing (SAC '25), March 31-April 4, 2025, Catania, Italy*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3672608.3707873>



This work is licensed under a Creative Commons Attribution 4.0 International License. *SAC '25, March 31-April 4, 2025, Catania, Italy*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0629-5/25/03
<https://doi.org/10.1145/3672608.3707873>




Figure 1: While traditional process mining (left) collects all events at a central entity, EdgeMiner (right) processes them directly at the source, and only exchanges aggregates (partial footprint matrixes), increasing scalability and privacy.

1 Introduction

Mining the never-ending data stream of events sensed by the Internet of Things (IoT) provides deep insights into, for example, processes in smart manufacturing, smart homes, and healthcare and is known as process mining [2, 10, 14, 17]. However, already today, IoT devices generate significantly more data than we can store or transfer to the cloud [12, 28, 37]. This high data velocity is a key challenge for process mining algorithms, which traditionally operate on static event logs, i.e., assume that data is readily available for mining at a central location with vast storage and computing capabilities [29]. Parsing such event logs, established algorithms commonly construct so-called footprint matrices (FMs), in which they store the order of events, so-called directly follows relations [15, 31, 36]. Next to the scalability challenge, collecting events at a central location causes privacy concerns [11, 19, 24].

In this paper, we introduce EdgeMiner, an algorithm for distributed process mining: it composes FMs, the key building block of process mining, in a distributed manner within a network of sensor nodes. For each event that a node in EdgeMiner detects, it locally stores the event as well as the predecessor and successor of this event and updates a local FM accordingly, see Fig. 1. Thus, over time, each node constructs a partial FM that represents the view of this node on the direct successions of this process. To derive the complete FM of the process, one of possibly multiple (central) entities queries the nodes for their partial FMs and combines these into a complete matrix. As a result, EdgeMiner enables (a) scalable mining, as a node only interacts with its predecessors after each event and, when queried, only exchanges aggregates, i.e., partial

FMs, with the requesting entity and (b) privacy-preserving process mining, as nodes locally merely store their own as well as predecessor and successor events, and a requesting entity only receives aggregate FMs instead of raw event streams. Further, EdgeMiner guarantees an FM identical to the FM when using a centralized approach, as we prove in this paper.

Analyzing real-world datasets, our evaluation shows that activities statistically only have a few predecessor activities, which rarely change throughout the lifetime of a process. For example, even in the large Hospital Log [32], comprising more than 600 distinct activities, the average number of predecessors per activity remains modest, resulting in an average number of less than 7. As a result, for each event, a node can very efficiently determine the predecessor, reducing the communication overhead by up to 96% compared to the naive baseline of querying all nodes in the network. Further, by batching predecessor queries into groups of, for example, 40 events, we show that EdgeMiner can reduce the average number of queried nodes per event to less than 2.5% of all nodes. As a final result, we show that for most datasets, an intermediate FM already complies with the final FM by more than 90% after less than 5% of the events of the dataset.

In summary, the contributions of this paper are as follows.

- We introduce EdgeMiner, a distributed process mining algorithm enabling distributed process discovery and conformance checking by operating directly where the events originate.
- We ensure that the nodes only share aggregated event data with a requesting entity.
- We analytically and experimentally show the correctness of the merge of EdgeMiner.
- We demonstrate that EdgeMiner's interim results quickly converge to the footprint matrix of the whole event log.
- We show that EdgeMiner reduces the message count for the predecessor determination by up to 96% compared to the naive approach and, for example, by 97.5% when batching queries with a batch size of 40.
- We release EdgeMiner's implementation as open-source¹.

The remainder of the paper is structured as follows: Section 2 introduces the necessary background. Next, Section 3 presents related work on distributed process mining. Afterwards, we introduce the design of EdgeMiner in Section 4. Section 5 holds our evaluation of EdgeMiner. We conclude the paper in Section 6.

2 Preliminaries and Background

In this section, we give a brief introduction to process mining and provide definitions of events, traces, event logs, direct successions, and footprint matrices. In process mining, we discover a process model from an event log (process discovery) or check that a process follows a given model (conformance checking) [25, 27, 29, 30]. A common way to represent a process model is a Petri net [22, 23].

Events and Event Logs: Process mining algorithms operate on event logs. Each event in the log typically contains at least three attributes: an activity, a timestamp, and a case ID. The case ID is used to distinguish different process instances. A trace is a sequence of events ordered by timestamp corresponding to a particular case.

¹Available at: <https://github.com/ds-kiel/EdgeMiner>

Direct Successions: In the following, let L be an event log over the activity set T . We define direct successions and therewith definable relationships between two activities, following the definition by Gatta et al. [11].

Definition 2.1 (Direct Succession). Activity b directly succeeds a in event log L , i.e. $a >_L b$, if and only if there exists a trace $\sigma = (t_0, \dots, t_{n-1}) \in L$ and $j \in \{0, \dots, n-2\}$ such that $t_j = a \wedge t_{j+1} = b$.

Definition 2.2 (Causality, No Direct Succession). With the help of Definition 2.1 we now define causality (1) and no direct succession (2) between two activities.

$$a \rightarrow_L b \quad \Leftrightarrow \quad a >_L b \wedge b \not>_L a \quad (1)$$

$$a \#_L b \quad \Leftrightarrow \quad a \not>_L b \wedge b \not>_L a \quad (2)$$

Footprint Matrix: A footprint matrix (FM) stores information regarding the direct succession relationships between activities within the event log. For event log L , the corresponding footprint matrix FM has dimension $|T_L| \times |T_L|$, where each activity in L is mapped to a specific row and column in the matrix, i.e., we can enumerate the activities such that each activity number i is an element of $\{0, \dots, |T_L| - 1\}$. An entry denotes how often the activity corresponding to the row was followed by the activity of the column. This way, $FM_{ij} \in \mathbb{N}_0$ denotes the entry of the footprint matrix that describes the relationship between activities i and j .

$$FM_{ij} > 0 \quad \Leftrightarrow \quad i >_L j$$

$$FM_{ij} = 0 \quad \Leftrightarrow \quad i \not>_L j$$

For our design, we define partial footprint matrices containing only one specific column: For $i, j \in \{0, \dots, |T_L| - 1\}$, we define FM^i as the footprint matrix that only contains column i and FM_j^i as the entry of column i in row j .

Alpha Miner: To illustrate process mining and its use of FMs, we recapitulate the main steps of the seminal process discovery algorithm Alpha Miner (Algorithm 1). Later, we discuss how newer discovery algorithms extend the alpha miner's concept and its FMs.

First, the algorithm denotes all activities occurring in the event log L as T_L . Each activity of this set corresponds to a transition in the final output model computed by the algorithm. Then, the algorithm determines two sets, T_I and T_O , which contain the start activities of the traces in L and their end activities, respectively. Subsequently, the algorithm determines all direct successions in L , which it stores in an FM. This FM is the basis for further calculations in the algorithm and the resulting Petri net. Next, the algorithm generates a set X_L , which comprises pairs of activity sets. Each element of the first set is directly followed by every element of the second set, although not vice versa. Moreover, activities within each set do not directly succeed one another. The Alpha Miner minimizes the set of pairs, followed by the derivation of Petri net places from the minimized pairs in the next step. Finally, the algorithm adds arcs between these places and transitions, forming the resulting Petri net.

To form the set X_L , we need to know the \rightarrow_L and $\#_L$ relationships between activities. These relations solely depend on the direct successions between the activities. An FM stores exactly those direct successions. For the Alpha Miner, the specific number of direct successions is unimportant, just whether a succession exists.

Algorithm 1 Alpha Miner

-
- 1: $T_L = \{t \in T \mid \exists \sigma \in L : t \in \sigma\}$
 - 2: $T_I = \{first(\sigma) \in T \mid \forall \sigma \in L\}$
 - 3: $T_O = \{last(\sigma) \in T \mid \forall \sigma \in L\}$
 - 4: $X_L = \{(A, B) \mid A \subseteq T_L \wedge B \subseteq T_L$
 $\quad \wedge \forall a \in A \forall b \in B : a \rightarrow_L b$
 $\quad \wedge \forall a_1, a_2 \in A : a_1 \#_L a_2 \wedge \forall b_1, b_2 \in B : b_1 \#_L b_2\}$
 - 5: $Y_L = \{(A, B) \in X_L \mid \forall (A', B') \in X_L : A \subseteq A' \wedge B \subseteq B'$
 $\quad \Rightarrow (A, B) = (A', B')\}$
 - 6: $P_L = \{p_{(A,B)} \mid (A, B) \in Y_L\} \cup \{i_L, o_L\}$
 - 7: $F_L = \{(a, p_{(A,B)}) \mid (A, B) \in Y_L \wedge a \in A\}$
 $\quad \cup \{(p_{(A,B)}, b) \mid (A, B) \in Y_L \wedge b \in B\}$
 $\quad \cup \{(i_L, t) \mid t \in T_I\} \cup \{(t, o_L) \mid t \in T_O\}$
 - 8: $\alpha(L) = (P_L, T_L, F_L)$
-

Heuristics Miner and Inductive Miner: Newer process discovery algorithms extend the concept of the Alpha Miner while keeping the FM as a key building block. For example, the Heuristics Miner starts with constructing a dependency graph [36]. To build the dependency graph, it calculates a so-called dependency measure between each two activities, depending on the number of direct successions between the activities. As an FM stores all direct successions, the miner uses the FM for this calculation. The Inductive Miner also relies on direct successions stored in an FM. The core concept is to construct a directly-follows graph, which visually represents the FM along with start and end activity sets. From this graph, a process tree is formed by repeatedly partitioning the nodes of the directly-follows graph, capturing the structure of the process step by step.

3 Related Work

Today's approaches to distributed process discovery focus primarily on multiple compute nodes to speed up the processing of ever-growing, centralized event logs: Van der Aalst [28] presents one of the first methods for distributed process discovery. He introduces a technique for computing partial and overlapping process models as Petri nets from partial event logs and defines a way to merge them to a Petri net of the complete event log. In contrast to EdgeMiner, he creates models rather than FMs in a distributed manner. Others [9] leverage Map-Reduce [7] to parallelize the Alpha Miner, distributing the computational workload across multiple compute nodes to handle the ever-growing size of event logs. However, their approach still relies on an event log that is initially stored at a central location. Later works [8] bring Map-Reduce to the Flexible Heuristics Miner [35]. Gatta et al. [11] extend the primary motivation of the previous work (i.e., performance and scalability) to include privacy concerns in the healthcare domain but still focus on one event log per hospital. In contrast, EdgeMiner emphasizes scalability, i.e., processing data directly at the source, and privacy preservation by only sending aggregate information, i.e., partial FMs, to the requesting node.

In the context of conformance checking, the Single-Entry Single-Exit (SESE) breaks down large process models and event logs into smaller sub-processes to analyze them independently [21]. Similarly, passage-based decomposition [26] and Projected Conformance Checking (PCC) [16] split logs into sets of activities that share direct successor-predecessor relationships: By solving each set independently and merging the results, it distributes, for example, conformance checking across multiple nodes. Although these approaches efficiently handle large and complex event logs, they remain centralized. One recent work presents conformance checking in a distributed manner also based on footprint matrices, yet – in contrast to EdgeMiner – their work does not cover the underlying mechanism to distribute the calculation of the matrices [1].

Streaming process mining [3] takes a different approach to distributed process mining. It focuses mainly on the *real-time* analysis of streams of events (i.e., event streams). Generally, those algorithms maintain a computational state according to previous events [4, 5, 13, 34] and employ approximations such as lossy counting [5]. While those methods are also motivated by scalability and performance, they are not designed to compute FMs at the sensor nodes.

In contrast, EdgeMiner distributes the FM computation as each sensor node computes a partial FM, enabling data aggregation at an earlier stage. As a result of processing events directly at their sources, EdgeMiner increases scalability, reduces privacy risks, and enables both process discovery and online conformance checking.

4 Design

In this section, we begin by discussing our assumptions and setting. Next, we give a design overview of EdgeMiner, detail its design, introduce optimizations, prove correctness, and discuss how EdgeMiner enables the distribution of well-known process mining algorithms.

4.1 Assumptions and Setting

4.1.1 Assumptions on Events and Event Logs. In adherence to common practice, we assume events with the attributes case ID, activity name, and timestamp, whereas the case ID is unique for each case. We assume timestamps within the same case to be distinct, ensuring unique FMs. In our evaluation in Section 5, we show that our assumptions are both realistic and applicable in numerous real-world scenarios.

4.1.2 EdgeMiner Setting. EdgeMiner targets typical IoT application settings: A network of IoT nodes that are, for example, integrated into machines in a smart factory or smart health setting. These sensors can communicate with each other, either directly or via the Internet. Further, we assume that membership management allows nodes to gain knowledge of other nodes and how to reach them, as is common in IoT applications. Finally, we assume that their clocks are well synchronized, using, for example, established clock-synchronization algorithms such as NTP [20]. This is essential, as accurate timestamps ensure proper event ordering. Without loss of generality, we assume each node detects exactly one type of activity to ease simplicity and comprehensibility. Note that by introducing virtual nodes, we can easily extend the problem scenario to allow each node to be responsible for multiple activity types.

4.2 EdgeMiner Algorithm

EdgeMiner consists of two phases: Phase 1 operates continuously upon each event detected, while Phase 2 executes upon a footprint matrix request. In Phase 1, nodes communicate to determine predecessors when recording an event and record their directly-follows relations in their local FMs. As a result, every node continuously updates its local, partial FM. Other than Phase 1, Phase 2 runs on demand whenever a central entity requests the FM. Then, all nodes send their partial FMs to the requesting entity, which constructs the FM.

4.2.1 Phase 1 – Event Ordering & Partial Footprint Matrices. For simplicity, we first introduce a naive version of EdgeMiner, which is not communication efficient. Later, we introduce optimizations that ensure efficiency. When an event e on node i is triggered, node i executes the following steps (see Fig. 2). First, node i stores event e locally, including its case ID, the activity name, and timestamp. Then, it sends all other nodes a message containing e to determine the predecessor event.

When receiving the request from i , a node checks its storage for events in the same case that have no successor yet and whose timestamps are smaller than the timestamp of e . If such an event exists, the node responds with the potential predecessor event data. If there are multiple responses, node i picks the event with the latest timestamp that is still smaller than the own one, as predecessor. We denote the node with the predecessor event as node j . Node i stores the newly determined predecessor locally, updates its local footprint matrix FM^i by increasing the count in the respective cell, and informs j . Subsequently, node j saves e as the successor event of its own event. Events that do not have successors for a long duration are flagged as end activities.

Suppose a node has to choose its predecessor from several potential predecessor events. In that case, this indicates that events are happening in quick succession and (a) communication is delayed or (b) multiple nodes are identifying predecessors in parallel. This may lead to temporarily incorrectly set predecessors. EdgeMiner automatically rectifies such inconsistencies during the algorithm's run as soon as the node that belongs to the incorrectly set predecessor event receives a request from the node of its correct successor: The node sets the correct successor and informs the incorrect successor. Then, the incorrect one starts the predecessor search again. A node then corrects the counters in its local FM. The rectification process converges since each trace has a single start event and because the nodes compare the timestamps of set successor events with the timestamps of incoming predecessor-search-requests.

Formally, each node $k \in \{0, \dots, n-1\}$ only keeps track of its predecessor activities in its FM. Therefore, each local FM has dimension $n \times 1$, with n the number of activities. The value at index $\ell \in \{0, \dots, n-1\}$ describes how often k directly follows ℓ in all of the cases, formally $\ell >_L k$. It is necessary to increment FM_ℓ^k by 1 each time k detects $\ell >_L k$ and subtract 1 if a direct succession gets fixed. If a node does not find a predecessor, it determines that it recorded the start of a trace and logs the event as a start event. In EdgeMiner, a node searches for its predecessor rather than its successor since it can be sure that the preceding event already occurred and, thus, the node can search for it immediately.

4.2.2 Phase 2 – Requesting a Footprint Matrix. This phase begins when an outside entity requests the FM. During this phase, the outside entity requests the local FMs of all nodes and start and end events within any case; see Fig. 3. Subsequently, each node responds with its local FM. The requesting entity receives all responses, and it assembles the footprint matrix $FM^{all} \in \mathbb{N}_0^{n \times n}$ by concatenating the columns FM^i , $i \in \{0, \dots, n-1\}$ of all n nodes, ordered by activity ID. Next, this FM serves as a process model for conformance checking, which compares the next events, traces, or sets of traces. Further, process discovery algorithms such as the Alpha Miner [31], the Heuristics Miner [36], or the Inductive Miner [15] can use this assembled FM for process discovery. Please note that the outside entity only receives aggregate information, i.e., the partial FMs. Raw event data is only processed on the sensor nodes, ensuring scalability and privacy.

4.3 Optimizations

Next, we introduce three optimizations: Most-Frequent-Predecessors (MFPs), batching, and a sliding window to enhance efficiency.

4.3.1 Most-Frequent-Predecessors (MFPs). In Phase 1 of EdgeMiner, a node – upon detecting an event – sends requests to all other nodes to determine the event's predecessor. This results in $n-1$ requests where n denotes the number of distinct nodes, i.e., activities in the event log. To reduce communication overhead, EdgeMiner contacts nodes in order of their likelihood of being the predecessor, i.e., it begins with the node that in the past was the most frequent predecessor according to its local FM. We show in a dataset analysis in Section 5 that activities of real-world processes statistically only have a few predecessor activities, which further only rarely change throughout the lifetime of a process. For example, even for datasets with a substantial number of activities, such as the Hospital Log [32] comprising more than 600 distinct activities, the average count of predecessors per activity remains modest, averaging fewer than 7 predecessors. As a result, for each event, a node can very efficiently determine the predecessor, reducing the communication overhead by up to 96% compared to querying all nodes in the network and thereby increasing scalability. Additionally, our findings demonstrate that even datasets with numerous activities, such as the Hospital Log, stabilize in the number of requested nodes after only a relatively small number of events.

4.3.2 Batching. To further minimize communication overhead in EdgeMiner, we batch querying of predecessor events. Rather than sending individual requests for each event, nodes accumulate several events before dispatching them in a single message. The nodes continue to use MFP Requesting. As soon as the node finds the predecessor of an event, it removes the event from the batch. The node ceases to query additional nodes once the batch is empty. For example, by batching predecessor queries in batches of 40 events, we show that EdgeMiner can reduce the average number of queried nodes per event to less than 2.5% of all nodes.

4.3.3 Sliding Window. To reduce resource demands, EdgeMiner utilizes a sliding window: Once the local event log reaches a given memory size or information age, the node discards it.




Figure 2: Phase 1 – Event Ordering and Partial Footprint Matrix Construction in EdgeMiner: Without a central entity, nodes determine the order of events collaboratively using message passing. In this example, event e_1 is a start event. Therefore, after the node detects the event, it queries the other nodes for the predecessor event. In our example, it does not get any positive responses and, thus, denotes it detected a start event. Upon sensing event e_2 , the detecting node queries for the predecessor event and receives a response, listing event e_1 as the predecessor. It stores this information in its local FM.




Figure 3: Phase 2 – Requesting a Footprint Matrix: A central entity requests partial FMs and start/end activity flags from all nodes. Upon receiving the data, it concatenates the matrices, forms start and end activity sets, and computes the footprint matrix to derive the final process model, e.g., using the Alpha Miner [31].

4.4 Correctness

After introducing the algorithmic design of EdgeMiner, we show its correctness. We claim that merging the partial FMs in EdgeMiner outputs the same FM as when calculating the FM based on a full event log. Here, we discuss the intuition of the proof. Appendix A holds a formal version.

By definition of the FM, there is one column and row per activity. An entry of the FM describes how often the activity of the row was followed by the activity of the column. A column of an FM, therefore, contains all predecessor relations of a specific activity. In EdgeMiner, the calculation of these columns is distributed to the nodes, and they are updated whenever the predecessor of an activity is found. Every predecessor relation of a node, and therefore activity, is found by the design of EdgeMiner, and no direct successions are added to the FMs of multiple nodes but precisely one. Thus, locally created columns in concatenated form ordered by activity ID must correspond precisely to the FM obtained when the entire log is considered and all direct successions are entered into one large matrix.

In our experimental evaluation in Section 5, we show (a) that both EdgeMiner and established central algorithms result in identical FMs and thereby process models and (b) that in four out of five datasets an intermediate FM already complies with the FM of the whole event log by more than 90% after less than 5% of the events of the datasets.

4.5 Distributed Mining with EdgeMiner

By distributing the computation of FMs, we can efficiently implement key process mining algorithms, such as, for example, the Alpha Miner, Heuristics Miner, and Inductive Miner, while also supporting online conformance checking in a scalable and privacy-preserving manner [1]. Below, we outline how EdgeMiner distributes each of these algorithms.

4.5.1 Distributed Alpha Miner. The Alpha Miner relies on identifying direct successions in event logs (cf. Section 2). Hence, it builds on causal relations between activities. We leverage EdgeMiner to distribute the computation of the FM from which the Alpha Miner constructs the process model. The Alpha Miner expects an FM consisting only of ones and zeros, i.e., if there is a directly-follows relation or not. As EdgeMiner counts the number of successions, we set FM entries greater than zero to one.

4.5.2 Distributed Heuristics Miner. In contrast to the Alpha Miner, the Heuristics Miner not only considers whether a direct succession is present but also considers the frequency and dependency measure of activity relationships. With a request, EdgeMiner merges all partial FMs to a global FM at the requesting entity, which includes the frequencies of the direct successions. From these, it calculates the dependency measures of the activity relationships. Next, the Heuristics Miner constructs a process model based on these frequencies and dependency measures.

4.5.3 Distributed Inductive Miner. The Inductive Miner first constructs a directly-follows graph and then recursively splits it into subprocesses to derive a process tree as a process model. To obtain the directly-follows graph, the Inductive Miner needs the direct successions present in the event log as well as the set of start activities and the set of end activities. EdgeMiner offers direct successions in the form of an FM along with the two sets. Just as for the Alpha Miner, the frequencies of the direct successions are not relevant, so all entries greater than zero are set to one.

4.5.4 Distributed Footprint-Based Online Conformance Checking. Online conformance checking continuously monitors process executions and compares them against a reference model in real time. Using EdgeMiner, each node generates and maintains a local, partial FM that captures the current state of the process at that node. As new events are detected, the local FMs are updated, and each

node can use its FM to perform conformance checking against a predefined process model [1].

5 Evaluation

In this section, we evaluate EdgeMiner in terms of correctness, accuracy, and performance. We first discuss implementation, baselines, and datasets in Section 5.1. Next, Section 5.2 analyzes these datasets. In Section 5.3, we experimentally evaluate the correctness of EdgeMiner. In Section 5.4, we examine how MFP Requesting reduces the number of queried nodes for Phase 1 of the algorithm compared to naive querying. Section 5.5 estimates FM sizes during an algorithm run, and Section 5.6 reflects on the evaluation results.

5.1 Evaluation Setup

5.1.1 Implementation. We implement EdgeMiner in Python and use a network of Docker containers to simulate the distributed nature of EdgeMiner.

5.1.2 Baselines. In our experiments, we compare EdgeMiner with and without MFP Requesting to two baselines: (1) "Query All" represents the upper bound where all nodes are queried, while (2) a theoretical oracle-like approach serves as the lower bound. The latter requires no requests if the predecessor corresponds to the node itself and only one request to identify the predecessor otherwise. Further, we compare EdgeMiner to traditional process discovery algorithms [15, 31, 36] to validate that the resulting FMs are identical.

5.1.3 Datasets. We conduct experiments on five real-world datasets (event logs): Hospital Log [32], Sepsis Cases [18] and an IoT Event Log for Process Mining in Smart Factories² [17]. These process-mining event logs naturally support geographical distribution over multiple nodes. To further substantiate our evaluation, we also include the BPI Challenge 2017 [33] and the Road Traffic Fine Management Process [6] datasets. In our setting, each activity in a dataset represents a node in an IoT network.

5.2 Dataset Analysis

We begin our experimental evaluation by analyzing real-world event logs; see Table 1. Our analysis shows that even in large datasets like the Hospital Log with over 600 distinct activities, the average number of preceding activities per activity stays below 7, showcasing the potential of MFP Requesting to reduce the average number of queried nodes per event. Further, our analysis reveals that event logs like the Smart Factories dataset contain numerous self-loops resulting from different lifecycle transitions and states within the recorded events. This characteristic benefits EdgeMiner, as self-loops do not require communication.

5.3 Experimental Correctness & Fitness

In addition to the analytical correctness proof in Section 4.4, we experimentally validate that the resulting FMs of EdgeMiner are identical to the ones computed by today's centralized algorithms³. We test both on randomly generated artificial datasets and common

²We use the `MainProcess.xes` file of the non-cleaned version.

³We use implementations available in `pm4py` (<https://pm4py.fit.fraunhofer.de>).




Figure 4: Average Number of Nodes Queried with and without MFP Requesting (including standard error): MFP Requesting and knowledge of the start events reduce communication demands by a factor of 7.5 to 30 depending on the dataset.

real-world event logs; see Section 5.1.3. For all datasets, the resulting FMs are identical at the end of the log; see Fig. 5.

Next, we examine how the merged FMs converge over time: Fig. 5 demonstrates that nearly all datasets reach 90% fitness within less than 5% of the trace. Here, the Smart Factories dataset is an exception, reaching 90% fitness after processing approximately 41% of the events, as it is quite small, see Table 1. The results indicate that querying interim results can be highly informative, potentially eliminating the need to process all events before requesting the first FM. Moreover, our analysis suggests that reducing the sliding window for discarding old data has minimal impact on the resulting FM.

5.4 Number of Queried Nodes

To assess EdgeMiner's performance, we analyze its two phases separately. For the first experiments, we assume a batch size of one.





5.4.1 Comparison Across Datasets. Fig. 4 and Table 1 illustrate the average number of nodes queried per event to find the predecessor for each dataset. Query All serves as the naive baseline for messaging every node in the network. With MFP Requesting, nodes require significantly fewer requests to identify their predecessor event. We achieve an additional message reduction when the process has a designated start event, making predecessor requests for the respective node redundant. Overall, the reduction increases with the number of activities, influenced by factors such as an expanding average case length, diminishing significance of start events, varying numbers of self-loops relative to total events, and the presence of just a few highly frequent predecessors.

Fig. 6 shows Cumulative Distribution Functions (CDF) for all datasets. The Hospital Log has the steepest increase, with over 90% of events needing queries to no more than 5% of the nodes. Conversely, the Traffic Fine dataset has very few instances where the predecessor can be identified after querying just 5% of the nodes. This deviation stems from the limited number of 11 activities in the Traffic Fine dataset. Hence, within the 5% threshold, only self-loops are considered.

5.4.2 Stabilization over Time. We now delve deeper into the specifics of the Hospital Log, as it is by far the largest one in terms of activities. Fig. 8 illustrates how the number of queried nodes stabilizes in the Hospital Log after fewer than 3,000 events. There are still fluctuations, for example, at around 5,000 events. This event log is a

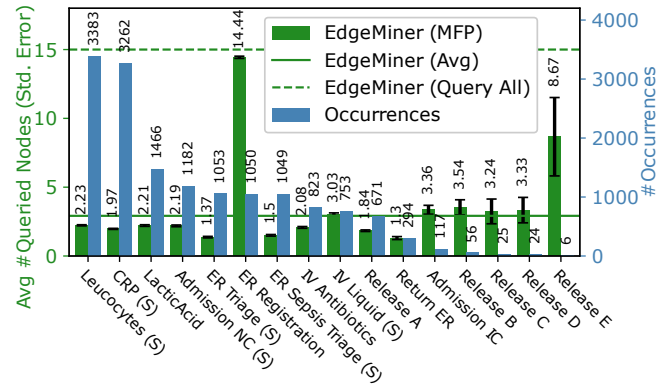
Table 1: Dataset Properties and EdgeMiner Performance (standard deviation in parentheses if not stated otherwise).

Event Log	Events	Activities (Start)	Cases	Length	Self-Loops	Predecessors	MFP Ratio	Reduction
Traffic Fine	561,470	11 (1)	150,370	3.73	4,306	6.36 (2.71)	0.80 (0.17)	64.94%
Sepsis	15,214	16 (6)	1,050	14.49	1,034	7.19 (2.67)	0.53 (0.20)	81.30%
Smart Factories	8,607	21 (1)	271	31.76	5,854	4.62 (3.33)	0.69 (0.07)	92.15%
BPIC 2017	1,202,267	26 (1)	31,509	38.16	448,205	6.85 (5.30)	0.79 (0.18)	93.40%
Hospital	150,291	624 (29)	1,143	131.49	26,542	6.78 (12.25)	0.69 (0.28)	96.13%

**Figure 5: Fitness Over Time: Intermediate FMs quickly converge to the centrally computed FM. The BPIC 2017 dataset, for example, already has a fitness of over 90% after 200 events.****Figure 6: Cumulative Distribution Functions: A steeper curve indicates a higher proportion of events that queried few nodes to identify their predecessors, reflecting greater efficiency.****Figure 7: Relationship between Batch Size and Average Number of Queried Nodes per Event (normalized): Batching effectively reduces the number of queries.****Figure 8: Number of Nodes Queried Over Time in the Hospital Log (moving average; green line: window size 50, step 1; orange line: window size 200, step 100): The moving average stabilizes after around 3,000 events. The reference line (Query All) indicates the maximum number of requested nodes per event. The plots show the first 10,000 of over 150,000 events.**

good example of stabilization due to its large activity set. Datasets with more activities take longer to stabilize, as each activity must occur multiple times for the MFPs to be identified.

5.4.3 Activity-wise Analysis. Next, we evaluate the number of queries per activity in the Sepsis Cases log; see Fig. 9. In contrast to other event logs, the Sepsis Cases do not have one designated start event but six different ones. However, about 95% of cases start with "ER Registration", which explains why the corresponding bar clearly stands out above the others. Although "Release E" shows an increased query count, it is not a start activity. The increased query count results from the limited occurrence (only six times), indicating non-converged predecessor stabilization. The situation is similar for the four activities to its left. They occur infrequently, which also explains the higher standard error.

**Figure 9: Average Number of Queried Nodes in the Sepsis Dataset: Displayed is the average number of queried nodes per activity (left y-axis) alongside the number of occurrences of each activity (right y-axis). The x-axis displays the different activities of the event log and highlights the start activities (S). We plot the number of requested nodes when querying all nodes and the average number of queried nodes over all events using MFP Requesting. The diagram shows that activities with the highest average queries occur the least (excluding start activities).**

5.4.4 Batching. Fig. 7 illustrates the reduction in the average number of queried nodes per event as the batch size increases. For example, at a batch size of 40, all datasets achieve an average number of nodes queried per event of less than 2.5% of the nodes in the network. For the Traffic Fine dataset, 2.5% corresponds to 0.25 nodes, whereas for the Hospital Log, it is 15.6.

Table 2: EdgeMiner’s Efficiency and Privacy Across Datasets Against Centralized PM (CS/CE = Central Storage/Entity).

		Centr.	Naive EdgeMiner	EdgeMiner with MFP & Batching (10)
# Msg. per Event	Traffic Fine	1	21	1.76
	Sepsis	1	31	2.32
	Smart Factories	1	41	2.30
	BPIC 2017	1	51	1.76
	Hospital	1	1247	26.46
Privacy Enablers	Storage	None	No CS	No CS
	Msg. Scope	To CE	In-Network	In-Network

5.4.5 Communication Cost and Privacy. Table 2 compares the centralized approach with the naive EdgeMiner algorithm and EdgeMiner using MFP Requesting and batching regarding communication overhead and privacy. The centralized approach requires fewer messages per event but at the cost of reduced privacy. For the naive EdgeMiner, the number of messages per event equals $2 \cdot (n - 1) + 1$ with n the number of activities/nodes. All nodes except the node itself are contacted, all other nodes reply, and the chosen node gets a confirmation message. For EdgeMiner with MFP Requesting and batching, we compute $2 \cdot k + 1$ with k the average number of queried nodes per event (see Fig. 4). In the centralized approach, event data is sent to a central entity right after the event, which requires only one message per event. However, this approach relies on a central storage system where all events are collected, presenting privacy and scalability risks. Although the number of messages per event for EdgeMiner with MFP Requesting and batching is higher than in the centralized approach, it is a considerable improvement compared to the naive EdgeMiner. We use a batch size of 10 events. As we show in Section 5.4.4, batching with larger batch sizes reduces the number of queried nodes and, therefore, the number of messages per event even further. The trade-off presented here is clear: while the EdgeMiner variants, particularly EdgeMiner with MFP Requesting and batching, cause higher communication overhead, they offer substantial scalability and privacy benefits by eliminating central storage and limiting non-aggregated communication to within the network.

5.5 Footprint Matrix Sizes

In the second phase of EdgeMiner, nodes transmit their data to a central entity, generating a fixed number of messages. These messages include the node’s local FM along with two additional Booleans indicating whether a node recorded the start or end activity of a trace. In EdgeMiner, each node only keeps track of its predecessors in its FM, resulting in transmitting FMs across the network of a maximal size of $n \times 1$, where n represents the total number of activities. Even for datasets like the Hospital Log, which has an activity count of 624, each node transmits only a list consisting of at most 624 elements. Compared to existing distributed process discovery algorithms [9, 11], which often require sending

up to quadratic FMs to a central entity, our approach is characterized by a compact message size. Moreover, storing such matrix sizes on IoT nodes is feasible even with constrained resources. For long-running processes, EdgeMiner can discard event data related to own, predecessor, and successor events once a node’s partial FM has been collected.

5.6 Evaluation Discussion

EdgeMiner leverages data locality to enhance its efficiency. Unlike existing distributed process discovery algorithms [9, 11], which typically collect the event data completely before processing it and creating FMs or process models, EdgeMiner adopts a proactive strategy by processing data at runtime and at its source. When the FM is requested, EdgeMiner has already started the process discovery algorithm by collecting partial FMs at each node thereby speeding up mining. Consequently, our strategy provides a significant time advantage, especially when handling large datasets with hundreds of activities and thousands of events. When looking at conformance checking, EdgeMiner enables online conformance checking by distributing the computation of FMs and making the partial FMs available at the data sources.

6 Conclusion

In this paper, we introduce EdgeMiner, a novel approach tailored to enable distributed process discovery and conformance checking in IoT networks. By distributing computations to IoT and edge devices, EdgeMiner addresses scalability and privacy concerns inherent in centralized approaches. We discuss how EdgeMiner’s design principles support common process mining algorithms that construct footprint matrices or record direct successions, such as the Alpha, Heuristics, and Inductive Miner [15, 31, 36]. Moreover, integrating Most-Frequent-Predecessor Requesting reduces communication overhead, ensuring efficient event handling within the network. The batching of predecessor queries further minimizes communication overhead.

Our empirical evaluations showcase EdgeMiner’s stark reduction of messages in real-world scenarios compared to a naive predecessor determination approach. EdgeMiner achieves a reduction in communication for determining predecessor events by up to 96% compared to querying all network nodes after each event. By batching predecessor queries in groups of, for example, 40 events, experiments show an average reduction in the number of queried nodes per event to less than 2.5% of all nodes. We analytically and experimentally demonstrate the correctness of the algorithm and show how the intermediate results of EdgeMiner quickly converge to the FM of the full event log. We highlight how EdgeMiner provides a significant time advantage over algorithms lacking pre-processing capabilities while still producing the same FMs as the centralized algorithms. By avoiding the storage of full event traces and sensitive personal data in central locations, EdgeMiner enables data privacy through aggregation that eliminates personal identification.

Acknowledgments

This work received funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), grant 496119880.

References

- [1] Julia Andersen, Patrick Rathje, and Olaf Landsiedel. 2024. Check my Flow: Distributed Conformance Checking at the Source. In *International Conference on Business Process Management: Business Processes Meet the Internet-of-Things*. Springer.
- [2] Yannic Bertrand, Jochen De Weerd, and Estefania Serral. 2022. A Bridging Model for Process Mining and IoT. In *Process Mining Workshops*, Jorge Munoz-Gama and Xixi Lu (Eds.). Springer International Publishing, Cham.
- [3] Andrea Burattin. 2022. *Streaming Process Mining*. Springer International Publishing, Cham, 349–372. https://doi.org/10.1007/978-3-031-08848-3_11
- [4] Andrea Burattin, Alessandro Sperduti, and Wil M. P. van der Aalst. 2012. Heuristics Miners for Streaming Event Data. *CoRR* abs/1212.6383 (2012). arXiv:1212.6383 <http://arxiv.org/abs/1212.6383>
- [5] Andrea Burattin, Alessandro Sperduti, and Wil M. P. van der Aalst. 2014. Control-flow discovery from event streams. In *2014 IEEE Congress on Evolutionary Computation (CEC)*. 2420–2427. <https://doi.org/10.1109/CEC.2014.6900341>
- [6] Massimiliano de Leoni and Felix Mannhardt. 2015. Road Traffic Fine Management Process.
- [7] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: Simplified Data Processing on Large Clusters. *Commun. ACM* 51, 1 (2008).
- [8] Joerg Evermann. 2016. Scalable Process Discovery Using Map-Reduce. *IEEE Transactions on Services Computing* 9, 3 (2016).
- [9] J. Evermann and G. Assadipour. 2014. Big Data Meets Process Mining: Implementing the Alpha Algorithm with Map-Reduce. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing*. Association for Computing Machinery.
- [10] Carlos Fernandez-Llatas et al. 2015. Process Mining Methodology for Health Process Tracking Using Real-Time Indoor Location Systems. *Sensors* 15, 12 (2015).
- [11] R. Gatta et al. 2019. On the Feasibility of Distributed Process Mining in Healthcare. In *Computational Science – ICCS 2019*. Springer.
- [12] Mouzhi Ge, Hind Bangui, and Barbara Buhnova. 2018. Big Data for Internet of Things: A Survey. *Future Generation Computer Systems* 87 (2018).
- [13] Marwan Hassani, Sergio Siccha, Florian Richter, and Thomas Seidl. 2015. Efficient Process Discovery From Event Streams Using Sequential Pattern Mining. In *2015 IEEE Symposium Series on Computational Intelligence*. 1366–1373. <https://doi.org/10.1109/SSCI.2015.195>
- [14] Dominik Janssen, Felix Mannhardt, Agnes Koschmider, and Sebastiaan J van Zelst. 2020. Process model discovery from sensor event data. In *International Conference on Process Mining*. Springer.
- [15] Sander JJ Leemans, Dirk Fahland, and Wil MP Van Der Aalst. 2013. Discovering block-structured process models from event logs—a constructive approach. In *Application and Theory of Petri Nets and Concurrency: 34th International Conference, PETRI NETS 2013, Milan, Italy, June 24–28, 2013. Proceedings 34*. Springer, 311–329.
- [16] Sander JJ Leemans, Dirk Fahland, and Wil MP Van der Aalst. 2018. Scalable process discovery and conformance checking. *Software & Systems Modeling* 17 (2018), 599–631.
- [17] Lukas Malburg, Joscha Grüger, and Ralph Bergmann. 2022. Dataset: An IoT-Enriched Event Log for Process Mining in Smart Factories.
- [18] Felix Mannhardt. 2016. Sepsis Cases - Event Log.
- [19] Felix Mannhardt et al. 2019. Privacy-Preserving Process Mining: Differential Privacy for Event Logs. *Business & Information Systems Engineering* 61 (2019).
- [20] David L Mills. 1985. *Network time protocol (NTP)*. Technical Report.
- [21] Jorge Munoz-Gama, Josep Carmona, and Wil M.P. van der Aalst. 2014. Single-Entry Single-Exit decomposed conformance checking. *Information Systems* 46 (2014), 102–122. <https://doi.org/10.1016/j.is.2014.04.003>
- [22] T. Murata. 1989. Petri nets: Properties, analysis and applications. *Proc. IEEE* 77, 4 (1989), 541–580. <https://doi.org/10.1109/5.24143>
- [23] James L. Peterson. 1977. Petri Nets. *ACM Comput. Surv.* 9, 3 (Sept. 1977), 223–252. <https://doi.org/10.1145/356698.356702>
- [24] Anastasiia Pika et al. 2020. Privacy-Preserving Process Mining in Healthcare. *International journal of environmental research and public health* 17, 5 (2020).
- [25] Wil M.P. van der Aalst. 2011. *Process mining : discovery, conformance and enhancement of business processes*. Springer, Germany. <https://doi.org/10.1007/978-3-642-19345-3>
- [26] Wil MP van der Aalst and HMW Verbeek. 2014. Process Discovery and Conformance Checking Using Passages. *Fundamenta Informaticae* 131, 1 (2014).
- [27] Wil M. P. van der Aalst. 2012. Process Mining: Overview and Opportunities. *ACM Transactions on Management Information Systems* 3, 2 (2012).
- [28] Wil M. P. van der Aalst. 2013. Decomposing Petri nets for process mining: A generic approach. *Distributed and Parallel Databases* 31, 4 (2013).
- [29] Wil M. P. van der Aalst. 2016. *Process Mining: Data Science in Action* (2 ed.). Springer.
- [30] Wil M. P. van der Aalst et al. 2012. Process Mining Manifesto. In *Business Process Management Workshops: BPM 2011 International Workshops, Clermont-Ferrand, France, Revised Selected Papers, Part I 9*. Springer.
- [31] Wil M. P. van der Aalst, T. Weijters, and L. Maruster. 2004. Workflow mining: Discovering process models from event logs. *IEEE Transactions on Knowledge and Data Engineering* 16, 9 (2004).
- [32] Boudewijn van Dongen. 2011. Real-life event logs - Hospital log.
- [33] Boudewijn van Dongen. 2017. BPI Challenge 2017.
- [34] Sebastiaan J. van Zelst and et. al. 2019. Online conformance checking: relating event streams to process models using prefix-alignments. *Int. J. Data Sci. Anal.* 8, 3 (2019), 269–284. <https://doi.org/10.1007/s41060-017-0078-6>
- [35] A Weijters and Joel Tiago S Ribeiro. 2011. Flexible Heuristics Miner (FHM). In *Symposium on computational intelligence and data mining (CIDM)*. IEEE.
- [36] A. Weijters, Wil M. P. van der Aalst, and A. K. Alves de Medeiros. 2006. *Process Mining with the Heuristics Miner-algorithm*. Eindhoven University of Technology.
- [37] Arkady B. Zaslavsky, Charith Perera, and Dimitrios Georgakopoulos. 2013. Sensing as a Service and Big Data. *CoRR* abs/1301.0159 (2013). arXiv:1301.0159

A Appendix: Formal Correctness Proof

We introduce a formal correctness proof. For this, we show that merging partial FM's in EdgeMiner outputs the same FM as calculating the FM based on a full event log.

We begin by introducing notations before defining our claim. Let FM^{all} denote the FM computed from an entire event log L and $n := |T_L|$, where T_L denotes the set of all activities in L . Further, let FM^i for $i \in \{0, \dots, n-1\}$ be the partial FM computed by node i , and L_i the partial event log that node i has knowledge of. Since i only keeps track of its predecessor events, its footprint matrix FM^i contains only one column. In the following correctness claim, we only consider the footprint matrices and omit start and end activities. The design of EdgeMiner ensures that a node considers an activity only a start or end activity if it corresponds to the start or end event of an entire trace. Further, we assume the footprint matrices only consist of ones and zeroes. The correctness for footprint matrices with entries in \mathbb{N}_0 follows from the fact that EdgeMiner detects all direct successions and adds a direct succession to exactly one partial footprint matrix.

CLAIM 1. *Merging the partial footprint matrices FM^0, \dots, FM^{n-1} by concatenating them ordered by activity ID results in the same FM as the FM^{all} , formally $(FM^0 \ \dots \ FM^{n-1}) = FM^{all}$.*

PROOF. To prove the claim, we show that all direct successions represented in the partial FM's can be found in FM^{all} and vice versa. Thus proving the claim of equivalence between the merged partial FM's and the FM with unrestricted knowledge.

By definition, $FM^{all} \in \{0, 1\}^{n \times n}$. As already discussed, for FM^i , $i \in \{0, \dots, n-1\}$, we assume without loss of generality that $FM^i \in \{0, 1\}^{n \times 1}$. It directly follows $(FM^0 \ \dots \ FM^{n-1}) \in \{0, 1\}^{n \times n}$.

Let $i >_L j$, $i, j \in \{0, \dots, n-1\}$ be a direct succession encoded in the merged partial footprint matrix $(FM^0 \ \dots \ FM^{n-1})$. Due to the concatenation of the partial matrices, it holds $FM^j_i = 1$. FM^j contains all direct successions between its activity and its preceding activities. This means there is a trace in the event log L in which i is followed by j . By definition, FM^{all} captures all direct successions in L , meaning it must hold $FM^{all}_{ij} = 1$.

It remains to show that the opposite direction holds. Let $k >_L \ell$, $k, \ell \in \{0, \dots, n-1\}$ be a direct succession represented in FM^{all} , formally $FM^{all}_{k\ell} = 1$. Since each node keeps track of all its predecessor activities in its local FM, we know that $FM^k_\ell = 1$. Merging all partial FM's only means concatenating them ordered by activity ID. Therefore, it follows $(FM^0 \ \dots \ FM^{n-1})_{k\ell} = 1$. Thus proving the claim. \square