



Modeling and Analyzing Release Trajectory based on the Process of Issue Tracking

Hani Abdeen, Houari Sahraoui

► To cite this version:

Hani Abdeen, Houari Sahraoui. Modeling and Analyzing Release Trajectory based on the Process of Issue Tracking. [Technical Report] 1374, DIRO, Université de Montréal. 2015, pp.11. hal-01132768

HAL Id: hal-01132768

<https://inria.hal.science/hal-01132768>

Submitted on 17 Mar 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Modeling and Analyzing Release Trajectory based on the Process of Issue Tracking

Hani Abdeen and Houari Sahraoui
DIRO, Université de Montréal, Montréal(QC), Canada
{abdeenha,sahraoui}@iro.umontreal.ca

Abstract—Software release development process, that we refer to as “release trajectory”, involves development activities that are usually sorted in different categories, such as incorporating new features, improving software, or fixing bugs, and associated to “issues”. Release trajectory management is a difficult and crucial task. Managers must be aware of every aspect of the development process for managing the software-related issues. Issue Tracking Systems (ITSs) play a central role in supporting the management of release trajectory. These systems, which support reporting and tracking issues of different kinds (such as “bug”, “feature”, “improvement”, etc.), record rich data about the software development process. Yet, recorded historical data in ITSs are still not well-modeled for supporting practical needs of release trajectory management.

In this paper, we describe a sequence analysis approach for modeling and analyzing releases’ trajectories, using the tracking process of reported issues. Release trajectory analysis is based on the categories of tracked issues and their temporal changing, and aims to address important questions regarding the co-habitation of unresolved issues, the transitions between different statuses in release trajectory, the recurrent patterns of release trajectories, and the properties of a release trajectory.

Index Terms—Software Evolution, Software Release, Issue Tracking Systems, Release Trajectory, Sequence-Analysis

I. INTRODUCTION

Software release trajectory is part of the software evolution process that leads from the software initial development, or from a released version, to a new release available to the software users [1], [2]. It includes the development activities such as adding new features, improving software, and fixing bugs. These activities change the software state [3]. For example, fixing a bug B changes the software state from “defected by B” to “cleaned of B”, and fixing all identified bugs changes the software state from “defected” to “clean”.

Release managers have the responsibility to control the release schedule and document the release state. Therefore, they need supporting tools to help them to understand the release trajectory and to improve their process planning and control skills. In this work, we are interested in techniques that can characterize the process of development activities that lead to a software release, and help software managers to better understand the evolution of software state through the software release. Unlike existing work on understanding software evolution and code changes (e.g., [4], [5], [6], [7], [8]), we seek an approach that describes the evolution of software state from release to release based on concrete diagnostics reported by the software managers. The approach should help software managers in understanding the software state in the beginning of a release trajectory (e.g., the software was suffering some bugs), and what transitions were achieved to reach the state which demarcates the underlined release.

Nowadays, Issue Tracking Systems (ITSs) play a central role in supporting software-development management [9]. ITSs are used by software developers and managers to report, track and manage issues of different kinds (such as “bug”, “new feature”, “improvement”, “technical task”, “test”, “documentation” etc.). Recorded issues in ITSs summarize the development activities that lead to software release, and the types of recorded issues represent the categories of development activities. Indeed, an opened, and accepted, issue of type “bug” means that the software suffers from the underlined bug. When this bug issue is closed and recorded as fixed, this means that the current release has sorted out a bug by performing specific development activity relevant to the bug fix. Hence, release notes are usually generated from ITSs recorded data, which include information about the fixed/implemented issues and their types [10].

Nevertheless, recorded historical data in ITSs are modeled and presented in a way that does not support adequately practical needs of release trajectory comprehension. Actually, well known ITSs, such as JIRA, provide helpful tools for filtering and sorting issues according to different criteria. For instance, users can list all resolved issues within a period of time and sort them according to their closing dates. Despite these facilities, it is not trivial to retrieve the evolution process of issues tracking in ITSs. For instance, managers may be interested in addressing the following questions: what type of issues was closed first, “bug”, “feature” or “improvement”? and what type was closed next? when and for how long time “improvement” issues were unresolved (i.e., open)? what type(s) of unresolved issues co-existed together in a given period? when and how this state was changed?

For example, Figure 1 shows the list of resolved (and fixed) issues for the project Apache-Solr [11] in a period of 11 days, between 1-Jan-2015 and 11-Jan-2015. This list includes only 30 issues that are sorted from first to last resolved issue. Users can easily observe that the improvement issue SOLR-6483 (in the top left line) was resolved before the bug issue SOLR-6906 (in the second line), as well as before the improvement issue SOLR-1723. However, it is not trivial to observe that, according to this list, improvement issues were open since 15-Jan-2010 (the creation date of issue SOLR-1723) and the life-cycle of improvement queries span without interruptions until 11-Jan-2015 (when the issue SOLR-6324, in the bottom right line, was closed). Indeed, during all this period, there was at least one improvement issue unresolved yet. In the same vein, it is not trivial to observe that bug fixing queries span the period from 22-Oct-2014 (the creation date of issue SOLR-6643) to 11-Jan-2015 (when the issue SOLR-6946 was

T	Key	Status	Resolution	Created	Resolved
	SOLR-6483	RESOLVED	Fixed	04/Sep/14	01/Jan/15
	SOLR-6906	RESOLVED	Fixed	02/Jan/15	02/Jan/15
	SOLR-6905	RESOLVED	Fixed	02/Jan/15	02/Jan/15
	SOLR-6797	RESOLVED	Fixed	26/Nov/14	03/Jan/15
	SOLR-2811	RESOLVED	Implemented	04/Oct/11	03/Jan/15
	SOLR-6735	RESOLVED	Fixed	13/Nov/14	04/Jan/15
	SOLR-6127	RESOLVED	Fixed	02/Jun/14	04/Jan/15
	SOLR-6907	RESOLVED	Fixed	03/Jan/15	04/Jan/15
	SOLR-6908	RESOLVED	Fixed	04/Jan/15	04/Jan/15
	SOLR-6666	RESOLVED	Fixed	29/Oct/14	05/Jan/15
	SOLR-6896	RESOLVED	Fixed	29/Dec/14	06/Jan/15
	SOLR-6830	RESOLVED	Fixed	08/Dec/14	06/Jan/15
	SOLR-6766	RESOLVED	Fixed	19/Nov/14	06/Jan/15
	SOLR-6874	RESOLVED	Fixed	21/Dec/14	06/Jan/15
	SOLR-6761	RESOLVED	Fixed	19/Nov/14	07/Jan/15
	SOLR-6918	RESOLVED	Fixed	06/Jan/15	07/Jan/15

T	Key	Status	Resolution	Created	Resolved
	SOLR-6925	CLOSED	Fixed	07/Jan/15	08/Jan/15
	SOLR-6643	RESOLVED	Fixed	22/Oct/14	08/Jan/15
	SOLR-6839	RESOLVED	Fixed	11/Dec/14	08/Jan/15
	SOLR-6945	RESOLVED	Fixed	09/Jan/15	10/Jan/15
	SOLR-2035	RESOLVED	Fixed	10/Aug/10	10/Jan/15
	SOLR-1723	RESOLVED	Fixed	15/Jan/10	10/Jan/15
	SOLR-4512	RESOLVED	Fixed	28/Feb/13	10/Jan/15
	SOLR-6932	RESOLVED	Fixed	08/Jan/15	10/Jan/15
	SOLR-6793	RESOLVED	Fixed	25/Nov/14	10/Jan/15
	SOLR-6950	RESOLVED	Fixed	10/Jan/15	10/Jan/15
	SOLR-6824	RESOLVED	Fixed	06/Dec/14	11/Jan/15
	SOLR-4580	RESOLVED	Fixed	14/Mar/13	11/Jan/15
	SOLR-6946	RESOLVED	Fixed	09/Jan/15	11/Jan/15
	SOLR-6324	RESOLVED	Fixed	05/Aug/14	11/Jan/15

Fig. 1. List of resolved (fixed) issues (of types “bug”, “improvement”, or “new feature”) in Apache-Solr project between 1-Jan-2015 and 11-Jan-2015, as recorded in JIRA (<http://issues.apache.org/jira/issues>). Issues are sorted from first resolved issue (top-left) to latest resolved issue (bottom-right).

closed), also without interruptions. Additionally, queries for new features span the period from 14-Mar-2013 to 11-Jan-2015 (the creation and closing dates of the issue SOLR-4580).

In this paper, we propose a sequence analysis approach for modeling and analyzing the releases’ trajectories using the tracking process of reported issues in ITSs. It employs sequence-analysis techniques for analyzing releases’ trajectories, generating a new kind of reports summarizing the trajectories and mining recurrent patterns. Our approach aims to address important questions regarding the co-habitation of unresolved issues, the transitions between different software statuses in release trajectory, the recurrent patterns of release trajectories, and the properties of a release trajectory.

To evaluate our approach, we performed a descriptive study of releases’ trajectories using a fairly large dataset of releases, 84 releases, that we extracted from the historical evolution of three large software projects. Our study shows that our approach contribute to the understanding of the development processes of releases providing new kinds of reports that can assist managers in better documenting releases, in auditing the behavior of releases development (that is implicitly the behavior of involved software developers), and in identifying recurrent patterns of development paths of releases.

This paper is organized as follows. Sections II and III present respectively the background and the approach for modeling and analyzing releases’ trajectories. The evaluation of our approach is presented in Section IV. Finally, Section V present the relevant existing work before concluding in Section VI.

II. BACKGROUND

Past studies of software evolution (such as [8], [12], [13], [14], [15]) widely employed time-series analysis to study historical software change events. In time-series analysis, continuous data represented as successive data points are required, and techniques such as ARIMA (autoregressive integrated

moving average) are used to fit models to time-series data to allow a deeper understanding of data evolution [8], [15]. To apply time-series analysis, data points must be uniformly spaced, and hence a suitable time unit (period length) must be specified for obtaining sufficiently stationary time series.

An alternative to time-series analysis is the sequence analysis (also known as trajectory analysis), which is the statistical study of succession of states or events that are usually ordered temporally¹ [16]. Unlike time-series analysis, sequence analysis is based on analysis of categorical data, rather than continuous data. It has been widely used in social science to analyze and understand biographical trajectories, trajectories of cohabitation, housing, career, etc. Sequence analysis aims at identifying sequences of states in the developmental process, where each individual is characterized by his unique trajectory. A trajectory is defined as a string of states with specific order and each state has a specific duration. Similarities between trajectories can then be measured based on shared sub-sequences. Sequence analysis can address questions about the types of developmental processes, the structural properties of a process, and the factors influencing different types of development paths.

In this paper, we are interested in learning how a software state evolves over time through a software release based upon software tracked issues and the temporal changing of their status (from open/unresolved to closed/resolved). More specifically, we are interested in: 1) understanding whether development processes of software releases follow similar trajectories; 2) identifying patterns in releases’ trajectories; 3) and identifying frequent and consistent transitions between distinct states in releases’ trajectories.

¹A sequence may also reflect other types of order, such as spatial order, preference order, hierarchical order, etc.

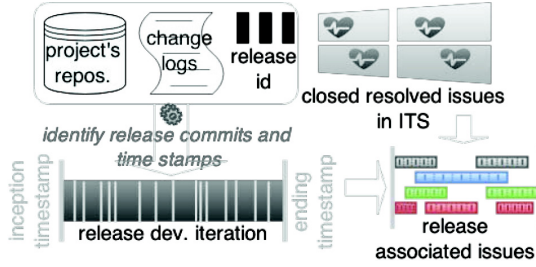


Fig. 2. Overview of required data by our approach.

III. RELEASE TRAJECTORY ANALYSIS APPROACH

In this section, we describe our sequence-analysis approach for modeling and analyzing release trajectory based on the evolution process of software state. Our approach relies on the software diagnostics as reported by the software managers and developers in the form of issues, and on the tracking process of reported issues in the issue tracking system.

By definition the trajectory of a given software release, R , is the evolution process of the software state starting from the first moment registered in the release development process, the *inception timestamp* of R , and ending by the moment that demarcates the development iteration leading to the targeted release, the *ending timestamp* of R . According to this definition, a release is recognized by its identifier (id) and its inception and ending timestamps, and characterized by the software states (as well as the transitions between states) that occurred between these timestamps. Hence, a release is concerned only by the resolved issues (i.e., the performed changes) between those timestamps [10].

In the remainder of this section, we describe the required data before detailing our approach for modeling and analyzing software release trajectory.

A. Required Data

Figure 2 provides an overview of required data to apply our approach. First of all, for analyzing the trajectory of a given software release, the inception and ending timestamps of the underlined release must be provided. By definition, these timestamps are respectively the timestamps of the start (earliest) and end (latest) commit revisions for the underlined release. Therefore, we need to access the change log of the studied system and retrieve the commits on both the trunk and branches. Unfortunately, identifying the earliest and latest commits for releases usually requires human intervention [17]. For certain software projects these commits can be identified easily thanks to some standard tags in the commit log, that are generated automatically by release preparation tools. For instance, when Maven release plugin is used to release a project, this automatically inserts two commits in the project change log with specific log messages: “[maven-release-plugin] prepare release *RELEASE_ID*” (this is the latest commit for the release *RELEASE_ID*), “[maven-release-plugin] prepare for next development iteration” (this is the earliest commit for the next release). If commit logs do not provide such specific tags, we need to map commits for releases using the approach of Shobe *et al.* [17]. Thus, we need to determine the names of releases and their representative branches in the project’s repository, as described in [17].

To retrieve the release trajectory, our approach relies on the reported issues in the software ITS that are resolved, closed, and associated to the considered release. Here, we consider that an issue is “resolved” when it was stated explicitly as treated and resolved –i.e., the issue resolution parameter is explicitly set to “fixed”, “implemented”, or other similar terms; but not to “invalid”, “not a problem”, “not complete”, or other similar terms. For some closed issues, we may find that developers have manually associated them to the concerned release, but unfortunately this is not always the case. Therefore, we select all issues that were closed and resolved during the development iteration of the considered release, that is determined by the inception and ending timestamps of the release.

After identifying the resolved issues, issues can be filtered according to their types as “included” or “excluded”, and included issues can be sorted as “recurrent (relevant)” or “non-recurrent (not relevant)”. In this work, all standard issue types² are included, e.g., “Bug”, “Dependency Upgrade”, “Documentation”, “new Feature”, “Improvement”, “Task”, “Test”, “Wish” etc., whilst the “sub-task” issue type is excluded.

Note that the list of standard issue types may vary from one project to another, and it could be very long. However, only a small subset of the aforementioned types can be considered as recurrently reported in different software projects and reported in release notes. These are: “Bug”, “new Feature”, “Improvement”, and “Task”. The first three types have software development story nature (e.g., Bug issue: “Unhandled exception thrown from within ...”; Improvement issue: “Reduce object allocations”; new Feature issue: “Add a policy interface for targeting ...”), whilst Task issues have technical nature (e.g., Task issue: “Upgrade to latest JPA 2.0 TCK”). Recurrent issue types are processed by our approach separately and modeled explicitly in the releases’ trajectory, whilst non-recurrent types, such “Documentation” and “Wish”, are aggregated all together in one virtual type.

Furthermore, in this work, we assume that sub-task issues should be excluded as they are related to other issues having standard types. More precisely, as the name “sub-task type” indicates, sub-task issues are useful for splitting up a parent issue having a standard type into a number of smaller tasks that can be assigned and tracked separately³. Hence, a resolved sub-task issue means that only a partial development activity has been achieved regarding the parent issue. We assume that the software state does not change until resolving the parent issue itself.

Once the aforementioned data about releases and issues are collected, our approach for modeling releases’ trajectories can be applied. Additionally, when commit logs are available and associated to (tagged by) the recorded issues, our approach can use commits to provide further insights on a release trajectory from commits-activity perspective.

²In this paper we mainly use the terminology of the Atlassian JIRA when we refer to issues types, categories, parameters.

³Refer to JIRA documentation: <https://confluence.atlassian.com/display/JIRA/Creating+a+Sub-Task>

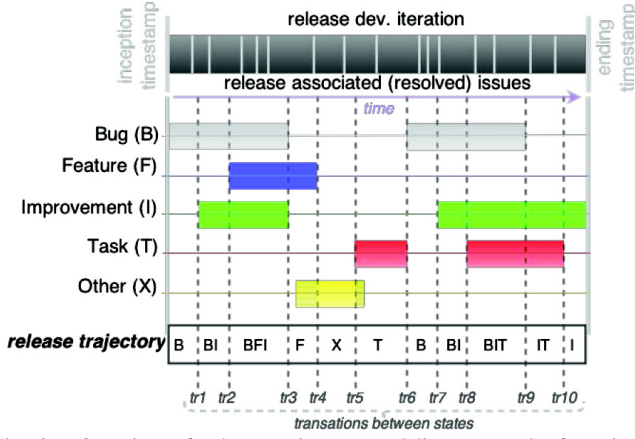


Fig. 3. Overview of release trajectory modeling approach: from issues tracking processes to a trajectory of states.

B. Release Trajectory Modelling

Figure 3 depicts our approach for mapping the closed resolved issues associated to a given release to a trajectory (sequence) of distinct states representing the release's trajectory. In this representation, each state has its specific duration, and the duration of the release's trajectory is the period between its inception date and ending date. The modeling algorithm is summarized as follows:

- **Identify atomic states and transitions** w.r.t. each considered issue type: For each issue type, p , issues whose lifecycles overlap, regarding their creation and closing dates, are grouped together and aggregated to form one atomic state characterized by p . That state is starting at an *opening timestamp* ($open_t$), the earliest creation date of issues forming the state, and ending at a *closing timestamp* ($close_t$), the latest closing date of issues forming the state. A state can indeed be formed of one issue if it does not overlap with other issues of its type. In such a case, the state's $open_t$ and $close_t$ are respectively the creation date and closing date of that issue. These timestamps $open_t$ and $close_t$ represents two atomic transitions, that are respectively moving into (i.e., opening) and moving out (i.e., closing) the underlined state. The output of this step is a map between considered issue types and their lists of opening and closing transitions. In Figure 3, the Bug issues which are associated to the analyzed release form two atomic Bug states. The first one goes from the inception timestamp to $tr3$ and the second starts at $tr6$ and ends at $tr9$. These states indicate periods when the undertaken release is diagnosed as suffering from bugs, and the period between transitions $tr3$ and $tr6$ shows that no bug issue associated to the release was open. Hence, from the perspective of the analyzed release, this period indicates a healthy/clean state, in which the release developers did not diagnose bugs that may concern the undertaken release.
- **From atomic states and transitions to global ones** w.r.t. all considered issue types: In this step, the lists of atomic opening and closing transitions for all considered issue types are merged into one list, and sorted from the old one to the new one. Then, the algorithm processes the transitions so that atomic states which occur between two successive transitions from one global, aggregated, state delimited by

those transitions. The output is then a list of ordered global states along with the transitions between those states (that is the union of all atomic transitions) forming the effective trajectory of the underlined release. By definition, a global state is said *atomic* if it concerns only of issue type, and is said *complex* if it concerns different types of issues. Consequently, the complexity of a state can be defined by the number of issue types that are involved in the state. In Figure 3, the analyzed release is stated as suffering from bugs from its inception timestamp. At the transition $tr1$, the release is still diagnosed as suffering from bugs, and also certain improvements are required. Hence, $tr1$ is a transition from an atomic state, B, to a (more) complex one, $BI = B + I$. At $tr2$, another symptom is identified when the release requires additional features, whilst certain bugs fixing and improvements are still required. The release keeps this last state for a relatively long time until $tr3$ when both identified bugs and required improvements are sorted out, but the release still requires additional features. We observe that by the end of the release development iteration only improvements were required.

- **Refine the release trajectory with regard to non-relevant issue types:** As mentioned earlier, issue types are sorted into two categories relevant/recurrent and non-relevant/non-recurrent. Issues which belong to non-relevant types are then then processed all together assigned to one virtual type, annotated by X in Figure 3. Moreover, when an atomic state of the non-relevant type X overlaps with other states that are of relevant types, then the global state in the overlapping period is formed only by the states of relevant types. Hence, atomic states of the non-relevant type can impose global states only when other states do not occur. This aims at reducing noises of non-relevant issues on the release trajectory, as well as reducing the complexity of visualization techniques that can be used to analyze release trajectories. Figure 3 shows that there is one atomic state of the non-relevant type X, and this state overlaps with an atomic state of type F, between $tr3$ and $tr4$, and another atomic state of type T, between $tr5$ and $tr6$. However, the global states in those periods of overlapping are formed only by the relevant issue types F and T. The period between $tr4$ and $tr5$ is the only one with the global state of type X as no other relevant states occurred during that period.
- **Normalize the duration of releases' trajectories:** A final, optional, step is to normalize the duration of releases' trajectories so that the duration of a state within a release trajectory is proportional to the release duration. Since releases can have completely different durations, this normalization is important when investigating recurrent patterns of developmental paths. In this paper, we normalize trajectories before analyzing them.

A release trajectory can also be summarized as a sequence of ordered distinct states where the states' duration is ignored. This representation, namely *distinct-successive-state* representation (DSS for short), is useful to focus on the number of states and transitions that occur in a release trajectory. Figure 4 shows the DSS representation of the release's trajectory represented in Figure 3, where the trajectory length reflects only the number of transitions in the release's trajectory.

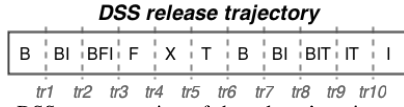


Fig. 4. The DSS representation of the release's trajectory in Figure 3.

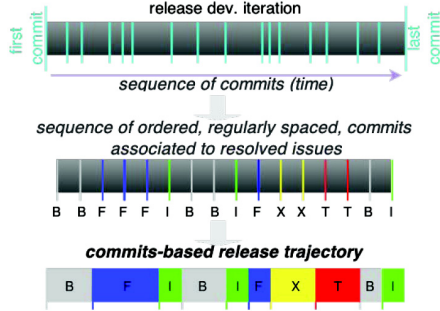


Fig. 5. Example of commits-based release trajectory.

C. Commits-based Release Trajectory

Up to now, the representation of a release's trajectory is fully based on reported statuses of the release's associated issues in the issue tracking system. This representation focuses on the temporal evolution of software states regardless of chronology of detailed changes that lead to changing the software states. These detailed changes are recorded in historical code repositories, such as SVN and Git, in the project's change log which consists of a list of ordered commits. Assuming that each commit in the project's log includes only changes for one specific task (i.e., the project's commits do not include tangled changes [18], [19]), and the project's recorded issues are tagged explicitly in the commits' logs, then both commits and types of recorded and tagged issues can provide further insight on releases' trajectory. Consequently, a new kind of releases' trajectories can be constructed, namely *commits-based* (or *activity-based*) trajectories.

Figure 5 shows an example of commits-based release trajectory, where the release's development iteration is represented by a sequence of 17 commits, ordered chronologically. All of these commits are explicitly tagged in their logs by issues associated to the underlined release. Commits which cannot be connected to the release's issues are not considered. In the first step, commits are annotated by the representative types of their associated issues. Then, they are regularly spaced regardless of the elapsed time between them. In other words, the unit of duration of commits-based releases' trajectories is "commit", not time. For instance, the duration of the release's trajectory in Figure 5 is 17 commits. The first two commits are annotated by the issue type Bug (B), the next three commits are annotated by the issue type new Feature (F), and so on and so forth until the last commit, annotated by the issue type Improvement (I). Finally, successive commits which have the same annotation type, such as the first two commits (B), are merged to form one state, and the duration of the state is measured by the number of merged commits. This leads to the commits-based release's trajectory which is represented similarly to the normal (issues-based) release's trajectory. The main differences in commits-based releases trajectories is that 1) the trajectory duration is measured by the number of involved commits, and 2) the trajectory's states are all *atomic* as states cannot overlap.

Similarly to normal releases' trajectories, the commits-based trajectories can also be normalized to have the same duration, so that the duration of a state within a commits-based trajectory is proportional to the number of commits involved in the trajectory. Furthermore, commits-based trajectories can also have the DSS representation where the duration of states is ignored, and then the trajectory's length reflects the number of transitions between the trajectory's states.

D. Summarizing Releases' Trajectories

Once releases' trajectories are constructed, sequence-analysis techniques can be applied for summarizing and analyzing the trajectories. In this section, we describe the main techniques and properties that we believe relevant to the comprehension of software releases' trajectories.

First of all, each release trajectory is characterized by three basic properties: 1) distinct states involved in it (the states in the DSS format), 2) the number of transitions between those states, and 3) the durations of those states.

Sequence-analysis provide a wide range of techniques for synthesizing global information for a family of trajectories (e.g., the trajectories of releases of a considered software project), and computing overall descriptive statistics. In our context, we mainly focus on techniques that allow software managers to understand the overall picture of releases' trajectories and capturing valuable pieces of information that can assist them while planning next releases. More specifically, we focus on techniques for estimating the probability to switch at a given position from one state s_i to another one s_j (i.e., transition rates between states), and for summarizing the trajectories of software releases in one picture, said the trajectory made of the most frequent state at each position (i.e., the trajectory of modal states).

- *Transition rates:* The transition rate between two distinct states s_i and s_j is the probability to switch at a given position from s_i to s_j , $p(s_j | s_i)$. This probability, as implemented in TraMineR [20], is based on the occurrences of this transition in the analyzed trajectories as well as on the occurrences s_i not followed by s_j . It is important to note that the transition rates between two states (from s_i to s_j , and from s_j to s_i) are usually not equal.
- *The trajectory of modal states:* An interesting summary over the releases' trajectories is the trajectory made of the most frequent state at each position p , with regard to all analyzed releases of a given software project. In the trajectory of modal states, each position is characterized by one state (i.e., the most frequent state at that position) and the frequency of that state at that position. This trajectory provides an overall picture of the analyzed releases' trajectories, and shows at which positions the trajectories of analyzed releases behave similarly (the frequency of modal states at those positions is high), and at which positions they behave differently (where the frequency of modal states is relatively low).

TABLE I
INFORMATION ABOUT COLLECTED DATA FROM STUDY SOFTWARE PROJECTS

Project	Releases	#Releases	Period	Avg #days per release	#Issues -resolved-	#Commits	%Issues tracked in Commits	%Commits tagged by Issues
Solr	1.1 – > 4.10.3	32	Jan-06 to Dec-14	101 days	3619	9296	3036/3619 = 84%	4841/9296 = 52%
OpenJPA	0.9.6 – > 2.3.0	17	Fev-06 to Nov-13	161 days	1553	7219	1348/1553 = 87%	4042/7219 = 56%
Struts2	2.1.2 – > 2.3.21	35	Oct-07 to Dec-14	72 days	1317	4325	916/1317 = 70%	1263/4325 = 29%
		84	$\approx 7\text{years}$		6489	20840		

E. Mining Patterns of Releases' Trajectories

Although the modal –most frequent– trajectory may provide a useful summary over a family of releases' trajectories, it is still of limited interest since it captures only one general pattern of the family of the analyzed releases. A more useful approach consists in identifying sets of distinct (sub) families of trajectories and identifying recurrent patterns of releases' trajectories. For this purpose, clustering methods can be employed to identify homogeneous clusters of releases' trajectories. To this end, we need techniques that can assess similarities/distances between releases' trajectories.

Several sequence similarity (dissimilarity) measures have been proposed in the literature. Those measures can be classified into two categories: 1) based on the count of matching states and sub-sequences between the two compared trajectories, and 2) based on the minimal cost of transforming one of the compared trajectories into the other. One of the most used dissimilarity measure for identifying groups of life-course trajectories with similar patterns is the Optimal Matching (OM) metric, that belongs to the second category. For computing the OM between two trajectories, the minimal cost for transforming one into the other should be assessed. This cost depends on the allowed transforming operations and their individual costs. The OM distance metric allows two transforming operations: 1) the substitution of one state by another one, and 2) the insertion or deletion of a state. Hence, the OM between two trajectories Tra_1 and Tra_2 is defined as the minimal cost (in terms of insertions, deletions and substitutions) for transforming Tra_1 into Tra_2 or Tra_2 into Tra_1 .

Usually, the cost of insertion or deletion operation of any state is set as a constant, independently of the state and its position in the trajectory. As for the costs of substitution operations are depending on the concerned states, and hence they are usually organized in a square symmetrical matrix of dimension A , namely the substitution cost matrix (SCM), where A is the size of the alphabet of possible distinct states in the analyzed trajectories. In our context, for issues-based releases' trajectories, $A = 16$. SCM can be specified manually, which can be effort-prone and error-prone activity. Hence, it is recommended to generate the SCM automatically using statistical descriptive information about the relations between the distinct states in the alphabet. In this paper, the substitution costs for two states s_i and s_j are determined from the estimated transition rates as $2 - p(s_i | s_j) - p(s_j | s_i)$, the cost is high when changes between s_i and s_j are rarely observed and lower when they are frequently observed.

IV. EVALUATION

Using sequence-analysis of releases' trajectories, one can address important questions, such as: “do releases' trajectories of a given software project follow the same development standard or methodology?”, “what are the frequent typologies of releases' trajectories that can represent standards?”, “when and why some releases follow turbulent chaotic trajectories or persist in one fixed state?” and finally, “what are the relations between the topological and longitudinal properties of releases' trajectories and other measurable properties of releases?”

To address these questions empirically, a large set of release trajectories must be analyzed. Here, the contribution of sequence-analysis techniques is to assist in performing exploratory tasks that consist of summarizing the dataset, extracting relevant information to the question of interest, and organizing the trajectories into representative homogeneous groups that capture patterns of trajectories. Then, the resulting summaries and patterns can be used in classical statistical analysis techniques, such as in regression-like models.

In this section, we evaluate the utility of our approach via a descriptive study of releases' trajectories using a fairly large dataset of releases, extracted from the historical evolution of three large software projects. Our study aims to put in evidence the utility of our approach for understanding the developmental paths of releases and extracting new kinds of reports that can address some of the aforementioned questions.

A. Projects and Data Collection

We extracted the data from three large Apache projects, namely Solr [11], OpenJPA [21] and Struts [22]. The issues of analyzed projects are collected from JIRA (<http://issues.apache.org/jira/>), and the change logs for Solr and OpenJPA are collected from their SVN repositories, while the change logs for Struts2 are collected from its Git repository. In total, we collected issue and commit data associated to 84 releases of those projects, covering an evolution period of around 7 years for each project. Table I shows for each project the number of considered releases, the average of releases' duration in days, and the number of issues and commits associated to those releases. As it can be seen, a large percentage of the considered issues are explicitly referenced in the commits' logs of the analyzed projects. However, the percentage of commits that are explicitly tagged with issues is relatively small. For instance, only 52% of the Solr analyzed commits are tagged with issues that can explicitly explain the nature of development activity associated to them. In this study, only these commits can be used in our approach for modeling the commits-based releases' trajectories.

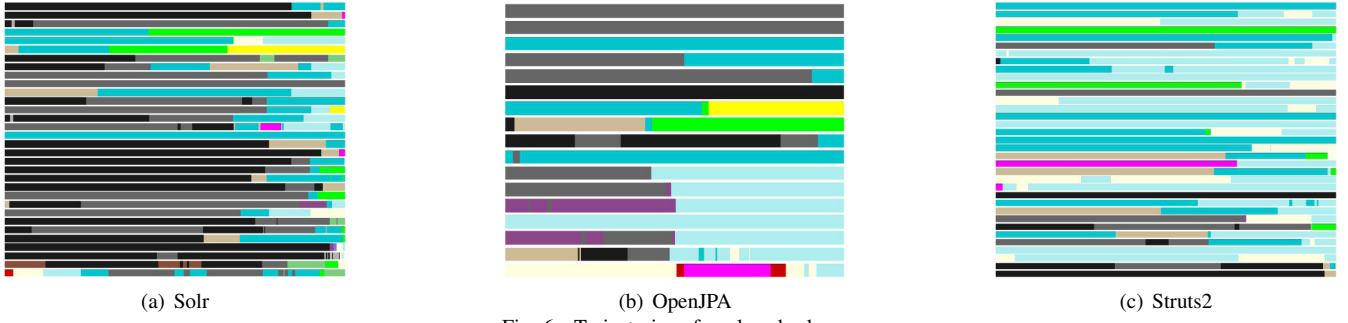


Fig. 6. Trajectories of analyzed releases

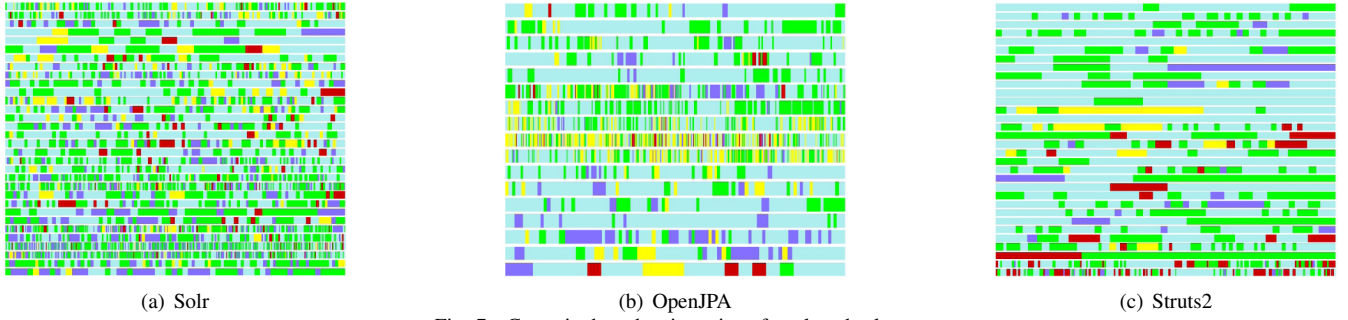


Fig. 7. Commits-based trajectories of analyzed releases

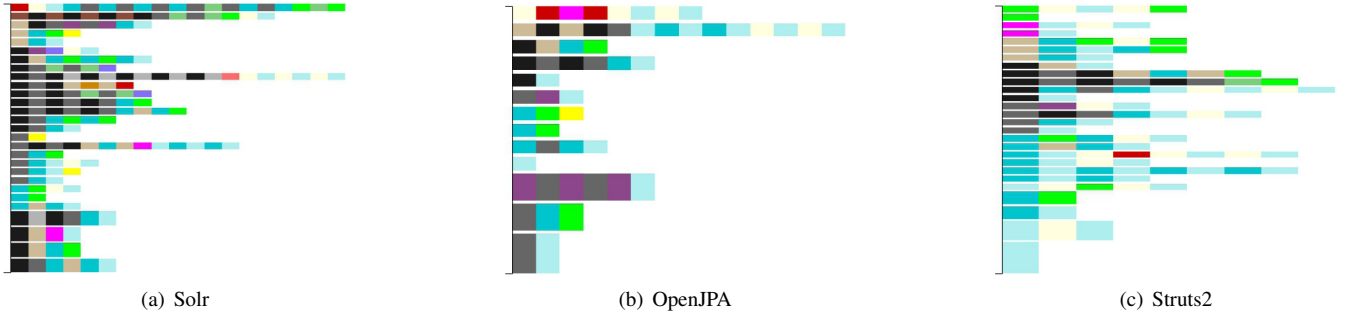


Fig. 8. Cumulative frequency of releases' trajectories in Distinct Successive State (DSS) format

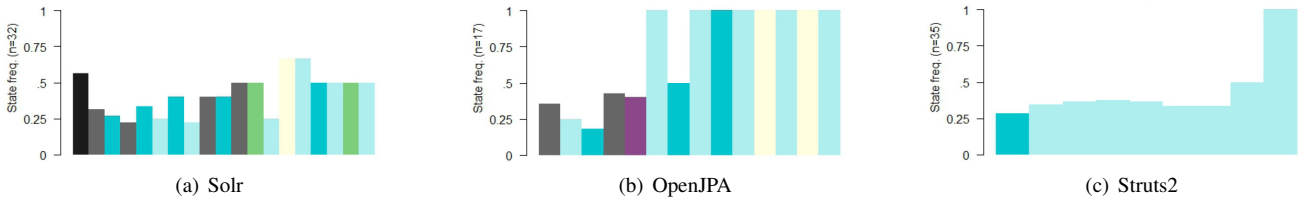


Fig. 9. Modal -most frequent- states' trajectory based on DSS format of releases trajectories

B. Overview of Releases' Trajectories

Figure 6 shows the trajectories of the analyzed releases. In this figure, we easily observe that there are different shapes of releases' trajectories, particularly when comparing the releases of the three analyzed projects. For instance, at a first glance on releases of the Solr project Figure 6(a), we observe that for almost all releases the complex states BIFT (black color state) and BIF (gray color state) are omnipresent along the trajectories. This reflects that the Solr software is frequently in complex states whereas issues of all considered types (bug, B, improvement, I, new feature, F, and technical task, T) stay

cohabiting for a long time during the development process, and developers work on these different issues in a parallel fashion. More precisely, the commits-based releases' trajectories for Solr in Figure 7(a) shows that the development activities shift very frequently from one activity type to another one. Here, we merely observe in some places a continuity in the nature of performed commits. Indeed, only in some releases, we can observe that successive commits have focused on bug fixing (turquoise color state) or on performing some improvement (green color state).

Fortunately, the releases of the Struts2 project have different shapes than the ones of Solr, as shown in Figure 6(c) for the

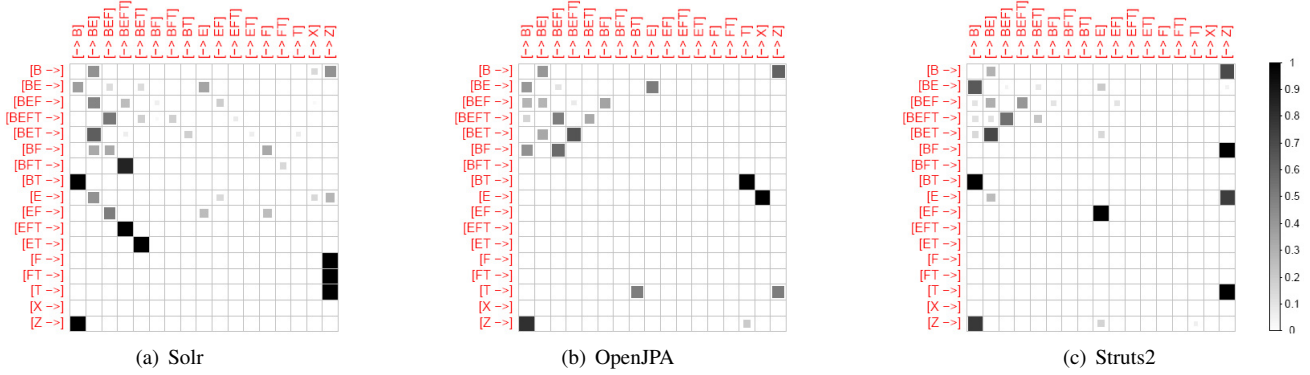


Fig. 10. Transition rates (probabilities) between distinct states in releases trajectories

issues and Figure 7(c) for the commits. This reflects that the Solr and Struts2 projects adopt different strategies of development. In Struts2, we observe that almost all releases are about fixing bugs and performing improvements, except for a small number of releases, especially the first two releases (at the bottom of figures). The trajectories of these two releases are of similar nature of the Solr releases. In fact, these releases, 2.1.2 and 2.1.3, are among the earliest releases of the Struts Action Framework 2.0. They represent an important transition in the life cycle of Struts as they demarcate the period when Struts and WebWork projects have decided to join in one project, that is Struts2. In the Struts2 project, we also observe that in some releases, technical task and/or new feature inquiries cohabit with bug fixing ones (see fuchsia and gray color states). However, we can observe in the commits-based trajectories that development activities associated to these inquiries are performed in a relatively continuous fashion.

As for the OpenJPA project, we observe in Figure 6(b) and Figure 7(b) that the releases' trajectories can be classified into two categories. In the first category, the releases' trajectories are similar to those of the Solr project, whereas issues of all types cohabit together for long periods, and the development activities shift very frequently from one development nature to another. These releases occur mainly after the mid life of OpenJPA, precisely, when OpenJPA has passed from 1.x.x to 2.x.x. After this transition, a relatively large number of new features have been incorporated in OpenJPA (mainly in 2.1.0) and a large number of improvements have been performed⁴. In the second category, which involves the earliest releases and 1.x.x releases, the releases' trajectories are more similar to those of the Struts2 project. Still, the turbulence of development activity, commits-based, states is more visible in OpenJPA than in Struts2.

Regardless of the states' duration in the trajectories, Figure 8 shows the distinct successive state (DSS) format of releases' trajectories. We easily observe that the projects Solr and Struts include both long and short trajectories, whilst almost all trajectories of OpenJPA are short. We also observe that the releases of Struts2 project are different from those of Solr and OpenJPA as they (almost all) begin with relatively simple states, that are only requesting bug fixing and/or improvements. Moreover, a distinguishable property

of Struts2 releases' trajectories is that many releases have periods where no issue is opened, indicated by the letter Z for Zen (or clean) state, and thereafter bug or improvement issues appear. These states may reflect periods of relaxation, or reflection, in releases' development. The DSS trajectories, can be summarized by the modal –most frequent– states trajectories as in Figure 9. Here, the Solr modal trajectory in Figure 9(a) shows that more than 50% of Solr releases begin with the complex state which is associated to all types of considered issues, the state BIFT. However, this state is shared only in the beginning of releases' trajectories, and most releases end with states related to bug issues. We also observe that at each position, at least 25% of the Solr releases share one, modal, state. The same observations can be reported to the modal trajectory of OpenJPA, as shown in Figure 9(b). Unlike both Solr and OpenJPA, Figure 9(c) shows that the modal state in almost all positions in the Struts2 trajectory is related to bug issues only. Still that, the frequency of this modal state is not high in almost all the positions.

C. Transition Rates

An important information that helps software managers in better understanding the evolution of software states through a new release is around the frequency and consistency of transitions between different states. That is the transition rate, or probability, between software states. Figure 10 shows in form of matrix the probabilities of transitions between different states in the releases' trajectories of the three analyzed projects. The first important information that we observe is that the transition probability between the same two states can considerably differ from one project to another one. This reflects that the transition probability between two states cannot be generalized with regard to the nature of states. It rather depends on the evolution path of the analyzed project and its development methodology.

In the Solr project, Figure 10(a), we observe several decisive transition probabilities showing that the involved states are strongly connected in the evolution process of software states. For instance, we observe that when Solr state is BT (i.e., requesting bug fixing and performing technical task) the probability is then very strong to sort out the technical tasks first (the software will move to the state B). However, if new features are requested in addition to BT (i.e., the Solr state is BFT), then the probability is very strong that Solr will move the most complex state requesting software improvement/enhancement

⁴Refer to <http://openjpa.apache.org/builds/2.0.0/apache-openjpa-2.0.0/RELEASE-NOTES.html>

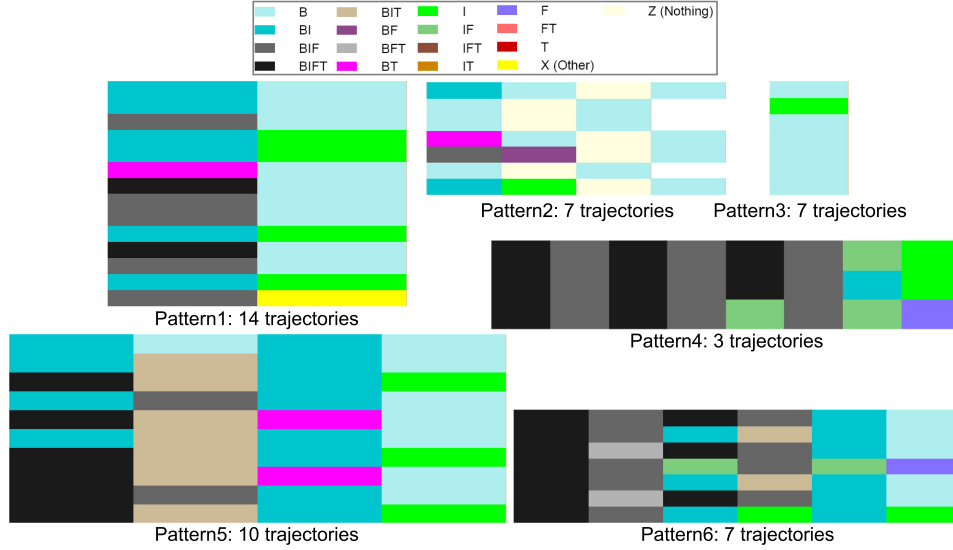


Fig. 11. Identified patterns of releases' trajectories –based on DSS format and the transition rates between distinct states with regard to all analyzed trajectories–

(move to the BIFT state). The transition to the BIFT state also seems very probable if Solr is in the state IFT. Hence, we deduce that when the states F and T cohabit together with another state, either B or I, then Solr will move to the most complex state BIFT, whereas all issues of all kind cohabit together. Finally, we observe that when the state of Solr is only around new feature issues (F) or technical tasks issues (T), then it is very probable to sort out these issues before opening new ones, and Solr will move to a clean state (Z). However, if Solr is in a clean state (Z), then it is very probable that the first issues to be opened are bug issues, and Solr will move the state B.

In the same vein, Figure 10(b) shows that when the release state is concerned about bug and technical issues (BT) the probability is then strong to close bug issues first (move to T state). However, this is not the case in Struts2 as Figure 10(b) shows, where the probability is high to close technical issues before bug issues. We also observe that when OpenJPA is concerned about only improvement/enhancement issues, then it is most probable to close these issues and open directly other, non-recurrent, types of issues (move to X state).

D. Patterns of Releases Trajectories

In this section, we attempt to identify recurrent paths, or sub-paths, of software state evolution that can help in categorizing releases' trajectories in a limited number of groups. Such groups reflect recurrent patterns of releases' trajectories.

To identify patterns of releases' trajectories, we performed an automatic analysis of the 84 trajectories of the three projects, using clustering techniques. The aim is to identify families of homogeneous releases' trajectories that can represent patterns in the universe of software releases. In our analysis, we use the Optimal Matching (OM) distance metric, that we presented in Section III-E. Recall that OM between two trajectories relies on the cost of transformations that should be applied to transform one trajectory into the other. To estimate this cost, we used the transition probabilities between distinct states based on all analyzed releases' trajectories. In this way, trajectories which are grouped (clustered) together

are not necessarily supposed to be *identical* or to share *identical* relatively long subsequences of states. Indeed, this criterion is relaxed so that two *distinct* states can be considered very similar if the transition probability between them is high, in the analyzed population of releases' trajectories.

Figure 11 summarizes the results of our analysis. Six groups of homogeneous releases' trajectories were identified, involving in total 48 releases. These groups can be placed in order of importance according to the number of grouped trajectories in each one. The larger a group of trajectories is, the more important/recurrent is the pattern.

The most recurrent pattern in our dataset is *Pattern1* which involves 14 releases trajectories. This pattern has the following main characteristics. The trajectories consist of only one transition (two states), and they all end by atomic states related to bug (B) or improvement (I) issues. They start by more complex states sharing in almost all cases the aforementioned atomic states B and I, such as the states BI, BIF and BIFT. This pattern reflects that these states are strongly connected, at least in our sample. A simpler case of *Pattern1* is *Pattern3* which involves 7 trajectories, that are all consisting of only one state. This pattern represents releases which are about bug fixing or performing improvements (but not both), as in the ending state of *Pattern1*.

The second recurrent pattern is *Pattern5* which involves 10 releases trajectories. All these trajectories consist of 3 transitions (4 distinct states) and, as in *Pattern1*, they all end by atomic states related to bug (B) or improvement (I) issues. This pattern is different than the other ones mainly with regard to its second state, where in almost all involved trajectories, the second state is around bug and improvement issues that are mixed with technical issues (BIT), most frequently, or with new feature issues (BIF). The same observation can be reported to the first state in the pattern, but here either both types of issues, new feature issues and technical issues, are cohabiting together with bug and improvement issues (the most complex state BIFT), or non of them is involved, and then the state is BI.

The main differences in `Pattern4` and `Pattern6` is that these patterns involve releases that all start by complex states. This is generally true during the first two states, when the software state is related to at least three types of issues, that are bug, improvement and new feature issues. These releases actually play the role of “issues collectors”, as they inherit a complex mixture of issues that were opened before their inception dates and take in charge the cleaning of inherited state. In `Pattern4`, the aforementioned three types of issues cohabit together after several transitions, in which the development path is mainly about sorting out technical issues and then opening new technical issues, to be sorted out next.

The final pattern that we identified is `Pattern2` which involves 7 releases trajectories. In this pattern, all trajectories consist of 3 transitions, and almost all of them are only around bug fixing and performing improvements. The main distinguishable property of this pattern is that all involved trajectories traverse clean/Zen states (no opened issues) before facing new identified bugs, that have to be resolved during the development iterations of concerned releases.

V. RELATED WORK

The most related work to our proposal is around analyzing and documenting software releases. In this direction, Hindle *et al.* [3] analyze the behavior of projects around the time of release. They use four classes to categorize software revisions in the source control system: source code, testing, build and documentation. A revision belong to one of the aforementioned categories based on the files associated to the revision. Then, they use these classes to identify behaviors that occur around release time (i.e., before or after releasing a new version), that they call as release patterns. In a recent work, Moreno *et al.* [10] propose an automated approach for generating detailed release notes, ARENA. ARENA relies on the code changes occurred in the commits performed between two releases. It summarizes those changes and link them to other information extracted from commits’ notes and from the description of issues associated to the previous release. Our present work is complementary to the work in [10] as it can enrich the release notes with new kind of reports about the tracking process of issues. From another perspective, Khomh *et al.* [23] investigate empirically whether rapid short release cycles improve the quality of the released software, comparing the crash rates and the proportion of post-release bugs of the rapidly released versions with those having a traditional (long) release cycle. The study shows that the number of post-release bugs remain comparable with both types of release cycle, but with shorter release cycles bugs are fixed faster.

Another body of related work is around analyzing software evolution based on change logs for understanding the evolution process and identifying common evolution phases. In this research direction, Xing and Stroulia [6] present an approach for understanding evolution phases using structural differencing algorithm to analyze changing class models over time and build the system’s evolution profile. The resulting sequence of changes is then analyzed to gain insight about the system’s evolution. Barry *et al.* [7] propose a method to identify volatility-based patterns of software evolution. Volatility is approximated by computing the amplitude, dispersion and pe-

riodicity of software changes at regular interval in the software history. Each period is defined by a volatility class, and sequence analysis is applied to reveal similar pattern in time. Kemerer *et al.* [8] use also sequence analysis for mining patterns of software evolution from change logs. In their approach, change events are categorized into 30 categories that are based on (1) three subjective classifications for changes (maintenance tasks: corrections, adaptations, and enhancements); (2) six basic categories for changed modules (data handling, logic, structure, computation, user interface, and module interface); (3) and three types of modifications that can be performed in “enhancement” events (add, change or delete). To categorize events, only the text/comments describing the change logs are used. In the same vein, Hindle *et al.* [4] propose an approach to classify large commits into 5 categories of maintenance tasks: corrective, adaptive, perfective, feature addition, and non-functional improvement. However, Kothari *et al.* [5] argue that it is not possible to categorize software changes in a confident and automated fashion using the aforementioned specific categories of maintenance tasks. They propose an automated approach that examines the temporal evolution of source code to identify *Change Clusters*. These change clusters classify code change activities as either a software maintenance or a new development. To identify change clusters, their approach first identifies a subset of canonical changes that best represent the modification activities in a given time period. Then, using clustering techniques, it classifies all other changes in the underlined time period as being similar to one of the identified canonical clusters. The authors use change clusters to identify trends in the development process.

In a different direction, D’Ambros and Lanza [24] propose a visualization to analyze and characterize the evolution of software entities, at different granularity levels. Besides, there is a family of research papers that focus on analyzing software change logs for identifying commits that contain tangled changes [18] or peripheral modifications [25], providing further insights on the nature of commits [26], [27], or for identifying the commit window of a release [17].

VI. CONCLUSION

We present a sequence-analysis approach to model and analyze the evolution of software states through a release, namely the release trajectory. The states are identified based on the recorded issues in the ITS, and the release’s trajectory is based on the tracking process of these issues. We consider the recorded issues as concrete diagnostics of the software state, and we analyze the evolution of their statuses during the life-cycle of releases. We also adapted our approach to model the tracking of software issues in the software change logs as development-based, commits-based, releases’ trajectory. Based on our model, we used sequence analysis techniques to summarize release trajectories and address important questions regarding the co-habitation of unresolved issues and the transition probability between different software states, and for mining recurrent patterns of release trajectories. The study opens perspectives for further detailed investigations of development paths of releases and possible relations between the properties of release trajectories and other properties that can be associated to software releases, such as the cost and productivity.

REFERENCES

- [1] A. van der Hoek, R. S. Hall, D. Heimbigner, and A. L. Wolf, "Software release management," in *Proceedings of the 6th European SOFTWARE ENGINEERING Conference Held Jointly with the 5th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. ESEC '97/FSE-5. New York, NY, USA: Springer-Verlag New York, Inc., 1997, pp. 159–175.
- [2] G. Blokdijk and I. Menken, *Release Management Best Practice Handbook: Building, Running and Managing Effective Software Release Management and Support - Ready to Use Supporting Documents Bringing ITIL Theory into Practice*. London, UK, UK: Emereo Pty Ltd, 2008.
- [3] A. Hindle, M. Godfrey, and R. Holt, "Release pattern discovery via partitioning: Methodology and case study," in *Mining Software Repositories, 2007. ICSE Workshops MSR '07. Fourth International Workshop on*, 2007, pp. 19–26.
- [4] A. Hindle, D. M. German, M. W. Godfrey, and R. C. Holt, "Automatic classification of large changes into maintenance categories," in *Int. Conf. on Prog. Comp.*, 2009.
- [5] J. Kothari, T. Denton, A. Shokoufandeh, S. Mancoridis, and A. Hassan, "Studying the evolution of software systems using change clusters," in *Program Comprehension, 2006. ICPC 2006. 14th IEEE International Conference on*, 2006, pp. 46–55.
- [6] Z. Xing and E. Stroulia, "Understanding phases and styles of object-oriented systems' evolution," in *Soft. Maint., 2004. Proc.. 20th IEEE Int. Conf. on*, Sept 2004, pp. 242–251.
- [7] E. Barry, C. Kemerer, and S. Slaughter, "On the uniformity of software evolution patterns," in *Intl Conf. on Softw. Eng.*, 2003, pp. 106–113.
- [8] C. F. Kemerer and S. Slaughter, "An empirical approach to studying software evolution," *IEEE Trans. Softw. Eng.*, vol. 25, no. 4, pp. 493–509, Jul. 1999.
- [9] O. Baysal, R. Holmes, and M. W. Godfrey, "Situational awareness: Personalizing issue tracking systems," in *Proceedings of the 2013 International Conference on Software Engineering*, ser. ICSE '13. Piscataway, NJ, USA: IEEE Press, 2013, pp. 1185–1188.
- [10] L. Moreno, G. Bavota, M. Di Penta, R. Oliveto, A. Marcus, and G. Canfora, "Automatic generation of release notes," in *Proc. of the 22nd ACM SIGSOFT Int. Symp. on Foundations of Software Engineering*. ACM, 2014, pp. 484–495.
- [11] Apache Solr: "is an open source enterprise search server based on the lucene java search library". [Online]. Available: <http://lucene.apache.org/solr/>
- [12] I. Herraiz, "A statistical examination of the evolution and properties of libre software," in *Software Maintenance, 2009. ICSM 2009. IEEE International Conference on*, Sept 2009, pp. 439–442.
- [13] A. Hindle, M. Godfrey, and R. Holt, "Mining recurrent activities: Fourier analysis of change events," in *Software Engineering - Companion Volume, 2009. ICSE-Companion 2009. 31st International Conference on*, May 2009, pp. 295–298.
- [14] Z. Lin, "Understanding and simulating software evolution," in *Proceedings of the 2013 International Conference on Software Engineering*, ser. ICSE '13. Piscataway, NJ, USA: IEEE Press, 2013, pp. 1411–1414. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2486788.2487020>
- [15] M. Goulão, N. Fonte, M. Wermelinger, and F. B. e Abreu, "Software evolution prediction using seasonal time analysis: A comparative study," in *16th European Conference on Software Maintenance and Reengineering, CSMR 2012, Szeged, Hungary, March 27-30, 2012*, T. Mens, A. Cleve, and R. Ferenc, Eds. IEEE Computer Society, 2012, pp. 213–222.
- [16] J.-A. Gauthier, , and P. Blanchard, P. Blanchard, F. Bühlmann, and J.-A. Gauthier, Eds. Amsterdam: Springer, 2014.
- [17] J. F. Shobe, M. Y. Karim, M. B. Zanjani, and H. Kagdi, "On mapping releases to commits in open source systems," in *Proc. of the 22Nd Int. Conf. on Prog. Comp.*, ser. ICPC 2014. ACM, 2014, pp. 68–71.
- [18] H. Kirinuki, Y. Higo, K. Hotta, and S. Kusumoto, "Hey! are you committing tangled changes?" in *Proc. of the 22Nd Int. Conf. on Prog. Comp.*, ser. ICPC 2014. New York, NY, USA: ACM, 2014, pp. 262–265.
- [19] K. Herzig and A. Zeller, "The impact of tangled code changes," in *Proceedings of the 10th Working Conference on Mining Software Repositories*, ser. MSR '13. IEEE Press, 2013, pp. 121–130. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2487085.2487113>
- [20] A. Gabadinho, G. Ritschard, N. S. Müller, and M. Studer, "Analyzing and visualizing state sequences in R with TraMineR," *Journal of Statistical Software*, vol. 40, no. 4, pp. 1–37, 2011. [Online]. Available: <http://www.jstatsoft.org/v40/i04/>
- [21] Apache OpenJPA: "is a java persistence project". [Online]. Available: <http://openjpa.apache.org/>
- [22] Apache Struts 2: "is a web framework for creating web-based java applications". [Online]. Available: <http://struts.apache.org/>
- [23] F. Khomh, T. Dhaliwal, Y. Zou, and B. Adams, "Do faster releases improve software quality?: An empirical case study of mozilla firefox," in *Proceedings of the 9th IEEE Working Conference on Mining Software Repositories*. IEEE Press, 2012, pp. 179–188.
- [24] M. D'Ambros and M. Lanza, "Software bugs and evolution: a visual approach to uncover their relationship," in *Software Maintenance and Reengineering, 2006. CSMR 2006. Proceedings of the 10th European Conference on*, March 2006, pp. 10 pp.–238.
- [25] N. Kusunoki, K. Hotta, Y. Higo, and S. Kusumoto, "How much do code repositories include peripheral modifications?" in *20th Asia-Pacific Software Engineering Conference*, 2013, pp. 19–24.
- [26] L. Hattori and M. Lanza, "On the nature of commits," in *23rd IEEE/ACM Int. Conf. on Automated Software Engineering*, 2008, pp. 63–71.
- [27] N. Dragan, M. L. Collard, M. Hammad, and J. I. Maletic, "Using stereotypes to help characterize commits," in *IEEE 27th International Conference on Software Maintenance*, 2011, pp. 520–523.