

# Interests, Difficulties, Sentiments, and Tool Usages of Concurrency Developers: A Large-Scale Study on Stack Overflow

Mehdi Bagherzadeh<sup>a,\*</sup>, Syed Ahmed<sup>a</sup>, Srilakshmi Sripathi<sup>a</sup>, Raffi  
Khatchadourian<sup>b</sup>

<sup>a</sup>Department of Computer Science and Engineering, Oakland University, MI, USA

<sup>b</sup>Department of Computer Science, City University of New York, Hunter College, NY, USA

---

## Abstract

**Context** Software developers are increasingly required to write concurrent code that is correct. However, they find correct concurrent programming challenging. To help these developers, it is necessary to understand concurrency topics they are interested in, their difficulty in finding answers for questions in these topics, their sentiment for these topics, and how they use concurrency tools and techniques to guarantee correctness. Interests, difficulties, sentiments, and tool usages of concurrency developers can affect their productivity.

**Objective** In this work, we conduct a large-scale study on the entirety of Stack Overflow to understand interests, difficulties, sentiments, and tool usages of concurrency developers.

**Method** To conduct this study, we take the following major steps. First, we develop a set of concurrency tags to extract concurrency questions and answers from Stack Overflow. Second, we group these questions and answers into concurrency topics, categories, and a topic hierarchy. Third, we analyze popularities, difficulties, and sentiments of these concurrency topics and their correlations. Fourth, we develop a set of race tool keywords to extract concurrency questions about data race tools and group these questions into race tool topics. We focus on data races because they are among the most prevalent concurrency bugs. Finally, we discuss the implications of our findings for the practice, research, and education of concurrent software development, investigate the relation of our findings with the findings of the previous work, and present a set of example questions that developers ask for each of our concurrency and tool topics as well as categories.

**Results** A few findings of our study are: ❶ questions that concurrency developers ask can be grouped into a hierarchy with 27 concurrency topics under 8

---

\*Corresponding author.

Email addresses: [mbagherzadeh@oakland.edu](mailto:mbagherzadeh@oakland.edu) (Mehdi Bagherzadeh),  
[sfahmed@oakland.edu](mailto:sfahmed@oakland.edu) (Syed Ahmed), [ssripathi@oakland.edu](mailto:ssripathi@oakland.edu) (Srilakshmi Sripathi),  
[raffi.khatchadourian@hunter.cuny.edu](mailto:raffi.khatchadourian@hunter.cuny.edu) (Raffi Khatchadourian)

major categories, ❷ *thread safety* is among the most popular concurrency topics and *client-server concurrency* is among the least popular, ❸ *irreproducible behavior* is among the most difficult topics and *memory consistency* is among the least difficult, ❹ *data scraping* is among the most positive concurrency topics and *irreproducible behavior* is among the most negative, ❺ *root cause identification* has the most number of questions for usage of data race tools and *alternative use* has the least. While some of our findings agree with those of previous work, others sharply contrast.

**Conclusion** The results of our study can not only help concurrency developers but also concurrency educators and researchers to better decide where to focus their efforts, by trading off one concurrency topic against another.

**Keywords:** Concurrency topics, concurrency topic hierarchy, concurrency topic difficulty, concurrency topic popularity, concurrency topic sentiment, data race tool usage, Stack Overflow.

---

## 1. Introduction

Software developers are increasingly required to write concurrent code to satisfy not only functional but also nonfunctional requirements of their software. For example, software with a graphical user interface must be concurrent to satisfy a functional requirement that it should be able to display more than one window at a time. Similarly, software with timing constraints must be concurrent to be able to satisfy a nonfunctional requirement that it should provide a better performance. Software is concurrent if its computations can potentially run at the same time and otherwise is sequential. In addition to satisfying its functional and nonfunctional requirements, the concurrent software must be correct. For example, with the initial value of  $x = 0$ , the software with the code  $x = x + 1$ , that is executed by two concurrent threads, should set the value of  $x$  to 2, which is the correct and desirable value. However, a concurrency bug, such as a data race bug, can set the value of  $x$  to the incorrect and undesirable value 1. This data race bug happens when both threads first read the initial value 0 of  $x$  and then increment and set it to the value 1 separately. A data race happens when two concurrent computations access the same variable of the program and at least one of these accesses writes to the variable. However, most developers think sequentially and find writing concurrent programs that are correct challenging [1, 2].

To help concurrency developers with writing concurrent code that is correct, it is necessary to understand concurrency topics they are interested in, their difficulty in finding answers for questions in these topics, their sentiment for these topics, and how they use concurrency tools and techniques to guarantee correctness. Interests, difficulties, sentiments, and tool usages of concurrency developers can affect their productivity [3, 4, 5, 6]. This understanding can not only help concurrency developers and their practice but also the research and education of concurrent software development by allowing members of these

communities to better decide when and where to focus their efforts [7, 8, 9, 10, 11, 12, 13, 14]. Without such understanding, practitioners may not prepare for similar difficulties, researchers may make incorrect assumptions about interests of practitioners, and educators may teach the wrong concurrency topics.

With more than 3 million developer participants and 38 million questions and answers, written in 2 billion words, Stack Overflow [15] has become a large and popular knowledge repository for developers to ask questions, receive answers, and learn about a broad range of topics. This makes Stack Overflow a great source to learn about interests, difficulties, sentiments, and tool usages of concurrency developers [16, 17, 7, 8, 11, 12, 14, 18].

In this work, we conduct a large-scale study on the entirety of Stack Overflow to understand interests, difficulties, sentiments, and tool usages of concurrency developers by answering the following research questions:

- **RQ1. Concurrency topics** What concurrency topics do developers ask questions about?
- **RQ2. Concurrency topic hierarchy** What categories do these concurrency topics belong to?
- **RQ3. Concurrency topic popularity** What topics are more popular among concurrency developers?
- **RQ4. Concurrency topic difficulty** What topics are more difficult to find answers to their questions?
- **RQ5. Concurrency topic sentiment** What topics are more sentimental among concurrency developers?
- **RQ6. Correlations of topic popularity, difficulty & sentiment** Are there correlations between topic popularity, difficulty, and sentiment?
- **RQ7. Concurrency tool topics** What topics do developers ask questions about when using data race tools and techniques?

For tool usage, we focus on the tools for finding and fixing data race bugs. We focus on data races because data race bugs are the not only one of the most prevalent bugs in concurrent software [1, 16] but also notoriously difficult to find and fix where some of these bugs can have extreme consequences. In a recent empirical study of concurrency bugs, data race bugs are about half of the concurrency bugs [19] in the open source software. Data races were the culprits in the Therac-25 disaster [20], the Northeastern electricity blackout of 2003 [21], and the mismatched NASDAQ Facebook share prices of 2012 [22]. Therefore, numerous tools and techniques [19] have been developed to help concurrency developers with data races.

To answer our research questions, we take the following major steps. First, we develop a set of concurrency tags to extract concurrency questions and answers from Stack Overflow. Second, we use topic modeling [23] and card sorting [24] techniques to group these questions and answers into concurrency topics. Third, we use card sorting [24] to group similar topics into concurrency categories and construct a topic hierarchy. Fourth, we use well-known metrics that are proposed and used by the previous work [11, 12, 8, 9, 11, 25, 7] to measure the popularities, difficulties, and sentiments of these topics and analyze their correlations. Fifth, we develop a set of race tool keywords to extract concurrency

questions about data race tools and techniques. Sixth, we use card sorting to group these questions into race tool topics. Finally, we discuss the implications of our findings for the practice, research, and education of concurrent software development. We also investigate the relation of our findings with the findings of the previous work and present a set of examples of questions that developers ask for each of our concurrency and tools topics and categories.

A few findings of our study are:

**Concurrency topics** (1) Concurrency questions that developers ask can be grouped into 27 concurrency topics ranging from *multithreading* to *parallel computing*, *mobile concurrency* to *web concurrency*, and *memory consistency* to *runtime speedup*. (2) Developers ask more about the *correctness* of their concurrent software than their *performance*.

**Concurrency topic hierarchy** (3) Concurrency questions can be grouped into a hierarchy with 8 major categories: *concurrency models*, *programming paradigms*, *correctness*, *debugging*, *basic concepts*, *persistence*, *performance*, and *graphical user interface (GUI)*.

**Concurrency topic popularity** (4) *Thread safety* is among the most popular concurrency topics and *client-server concurrency* is among the least popular.

**Concurrency topic difficulty** (5) *Irreproducible behavior* is among the most difficult concurrency topics and *memory consistency* is among the least difficult.

**Concurrency topic sentiment** (6) Most questions of all concurrency topics are neutral in sentiment. (7) *Data scraping* is among the most positive concurrency topics and *process life cycle management* is among the least positive. (8) *irreproducible behavior* is among the most negative concurrency topics and *process life cycle management* is among the least negative.

**Correlation of popularity, difficulty and sentiment** (9) There is a statistically significant negative correlation between the popularity, difficulty, and negative sentiment of concurrency topics.

**Race tool topics** (10) Concurrency questions about data race tools can be grouped into 4 topics: *root cause identification*, *general*, *configuration & execution*, and *alternative use*. (11) Developers ask the most race tool questions about *root cause identification* and the least about *alternative use*.

**Relation of our findings with previous work** (12) Our findings relate to several findings of previous work, including the works by Pinto et al. [18], Barua et al. [7], Rosen and Shihab [11], Lu et al. [1], Guzman et al. [26], Sinha et al. [27], Tourani et al. [28], Wang et al. [29], and Bagherzadeh and Khatchadourian [17]. While some of our findings agree with the findings of these previous works, others sharply contrast.

### 1.1. Extensions to Our Previous Work

This work is a significantly revised and extended version of our previous work [30] with the following major extensions:

1. Addition of the new research question **RQ5** and the sentiment analysis of 245,541 concurrency questions and answers for all 27 concurrency topics;
2. Addition of the new research question **RQ7** and the development of the race tool keywords, extraction and grouping of race tool questions into

topics, and presenting a set of example questions that developers ask about race tool topics;

3. Extension of our previous research question **RQ6** with the analysis of 15 new correlations between sentiments, popularities, and difficulties of concurrency topics;
4. Extension of our previous implications with the implications of our new findings for the practice, research, and education of concurrent software development;
5. Extension of our previous relations with the relations of our new findings for **RQ5–RQ7** with the findings of the previous work;
6. Extension of our previous data collection with the distribution of concurrency questions and answers over different programming languages;
7. Addition of 10 more findings, to our 12 previous findings, and 10 more relations with the findings of previous works, to our 16 previous relations;
8. Extension of our related work with previous works related to the sentiment analysis and usage of data race tools and techniques.

All the data used in this study are available at <https://goo.gl/uYCQPU>.

### 1.2. Outline

Section 2 presents our methodology for the collection and analysis of our data. Section 3 discusses the answers to our research questions and investigates the relation of our findings with the findings of the previous work. Section 4 shows the implications of our findings for the practice, research, and education of concurrent software development. Section 5 discusses threats to the validity of our study and their mitigation strategies. Section 6 discusses related works. Section 7 concludes the paper.

Throughout this paper, the difficulty of a topic is solely defined as the difficulty for finding answers to the questions in that topic.

## 2. Methodology

Figure 1 shows an overview of our methodology for the collection of concurrency questions and answers from Stack Overflow and their analysis.

### 2.1. Data Collection

**Step 1: Download Stack Overflow dataset** In this step of our methodology, we download the Stack Overflow dataset  $\mathcal{S}$  which is publicly available through Stack Exchange Data Dump [31]. The data set  $\mathcal{S}$  includes question and answer posts and their metadata. The metadata of a post includes its identifier, its type (question or answer), title, body, tags, creation date, view count, score, favorite count, and the identifier of the accepted answer for the post, if the post is a question. An answer to a question is accepted if the developer who posted the question marks it as accepted. The question can have at most 1 accepted answer. The question can have 1 to 5 tags that are used to specify what the question is about.

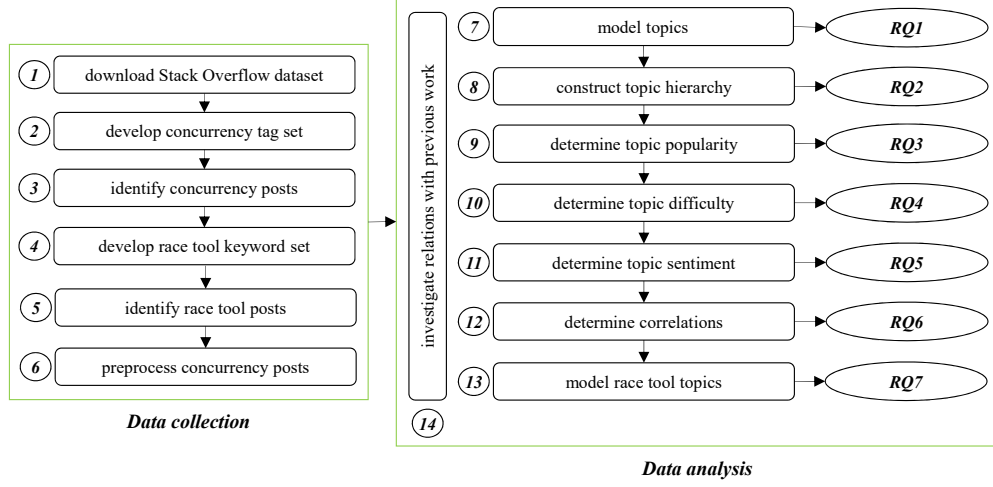


Figure 1: An overview of the methodology of our study.

Our dataset  $\mathcal{S}$  includes 38,485,046 questions and answers posted over a time span of over 9 years from August 2008 to December 2017 by 3,589,412 developer participants of Stack Overflow. Among these posts, 14,995,834 (39%) are questions and 23,489,212 (61%) are answers. Among these answers 8,034,235 (21%) are marked as accepted answers.

**Step 2: Develop concurrency tag set** In this step, we develop a set of concurrency tags  $\mathcal{T}$  to identify and extract concurrency questions and answers in Stack Overflow. First, we start with a tag set  $\mathcal{T}_{init}$  that includes our initial concurrency tags. Second, we extract questions  $\mathcal{P}$  from our dataset  $\mathcal{S}$  whose tags match a tag in  $\mathcal{T}_{init}$ . Third, we construct a set of candidate tags  $\mathcal{T}$  by extracting tags of the questions in  $\mathcal{P}$ . Finally, we refine  $\mathcal{T}$  by keeping tags that are significantly relevant to concurrency and exclude others. We use two heuristics  $\mu$  and  $\nu$ , that are also used by the previous work [11, 12, 17], to measure the significance and relevance of a tag  $t$  in  $\mathcal{T}$ .

$$\text{(Significance)} \quad \mu = \frac{\text{number of questions with tag } t \text{ in } \mathcal{P}}{\text{number of questions in } \mathcal{P}}$$

$$\text{(Relevance)} \quad \nu = \frac{\text{number of questions with tag } t \text{ in } \mathcal{P}}{\text{number of questions with tag } t \text{ in } \mathcal{S}}$$

A tag  $t$  is significantly relevant to concurrency if its  $\mu$  and  $\nu$  are higher than specific thresholds. Our experiments using a broad range of thresholds for  $\mu$  and  $\nu$  show that  $\mu = 0.1$  and  $\nu = 0.01$  allow for a significantly relevant set of concurrency tags. Our threshold values are in line with the values used by previous work [12, 7, 17]. Table 1 shows the concurrency tag set  $\mathcal{T}$  for select significance and relevance values  $\mu$  and  $\nu$  with our tagset  $\mathcal{T}$  in gray.

We use the initial tag set  $\mathcal{T}_{init} = \{multithreading\}$  to construct the tag set  $\mathcal{T}$  below using our tag development approach. To select the set of our initial tags in  $\mathcal{T}_{init}$ , we manually inspect the top 100 most used tags on Stack Overflow and select all the tags that are related to concurrency. Our set  $\mathcal{T}_{init}$  includes a tag *multithreading*, that is used by previous work [1, 18] as well.

$$\mathcal{T} = \left\{ \begin{array}{l} \textit{concurrency} \textit{ locking} \textit{ multiprocessing} \textit{ multithreading} \textit{ mutex} \textit{ parallel-} \\ \textit{processing} \textit{ pthreads} \textit{ python-multithreading} \textit{ synchronization} \textit{ task-} \\ \textit{parallel-library} \textit{ thread-safety} \textit{ threadpool} \end{array} \right\}$$

Our concurrency tag set  $\mathcal{T}$  with 12 tags is broad and includes generic tags, such as *concurrency* and *synchronization*, and specific tags, such as *thread safety* and *thread pool*

**Step 3: Extract concurrency posts** In this step, we extract Stack Overflow questions whose tag set contains a tag in  $\mathcal{T}$ . This set includes 156,777 questions and 249,662 answers, where 88,764 (36%) of these answers are accepted answers. To reduce noise, following previous work [7, 11, 17], we add questions and their accepted answers from this set to our set of concurrency post  $\mathcal{C}$  and discard unaccepted answers. The set of concurrency posts  $\mathcal{C}$  includes 156,777 questions and 88,764 accepted answers. In total,  $\mathcal{C}$  includes 245,541 questions and answers.

Using the tag set  $\mathcal{T}$  to extract concurrency questions does not mean that concurrency questions cannot have other tags in addition to the tags in  $\mathcal{T}$ . In fact, our concurrency questions have 11,703 extra tags, such as *asynchronous*, *blocking*, and *atomic*.

In addition, our concurrency questions are related to a broad spectrum of programming languages ranging from C languages, such as C, C++, C#, and Objective C, to Java and JavaScript and from Visual Basic to Go and Haskell. Figure 2 shows the number of our concurrency questions for different programming languages. To extract these languages, we manually inspect top 200 most used tags in  $\mathcal{C}$  and select tags that identify a programming language.

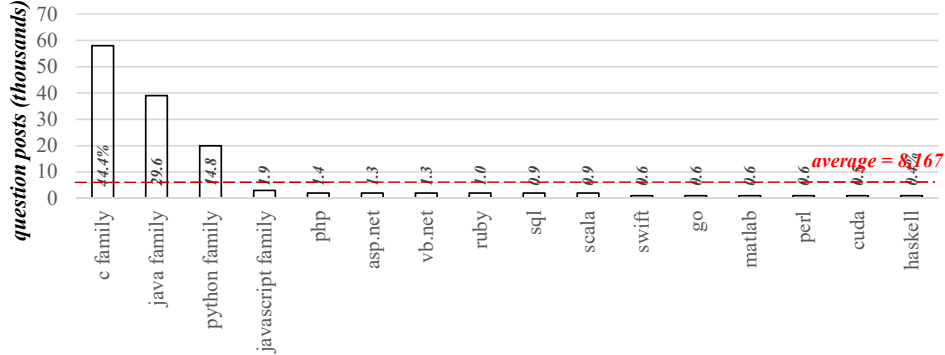


Figure 2: Programming languages of our concurrency questions and their numbers.

Table 1: Concurrency tags for select relevance and significance threshold values. Our concurrency tag set  $\mathcal{T}$  is in gray.

$(\mu, \nu)$	Tag set $\mathcal{T}$	No.
(0.3, 0.015)	concurrency multithreading pthreads thread-safety threadpool	5
(0.3, 0.01)	concurrency multithreading mutex pthreads python-multithreading thread-safety threadpool	7
(0.3, 0.005)	backgroundworker concurrency executorservice multithreading mutex pthreads python-multithreading runnable semaphore synchronized thread-safety threadpool	12
(0.2, 0.015)	concurrency locking multithreading pthreads synchronization thread-safety threadpool	7
(0.2, 0.01)	concurrency locking multithreading mutex pthreads python-multithreading synchronization task-parallel-library thread-safety threadpool	10
(0.2, 0.005)	atomic backgroundworker concurrency deadlock executorservice grand-central-dispatch locking multithreading mutex openmp pthreads python-multithreading runnable semaphore synchronization synchronized task-parallel-library thread-safety threadpool wait	20
(0.1, 0.015)	concurrency locking multithreading parallel-processing pthreads synchronization thread-safety threadpool	8
(0.1, 0.01)	concurrency locking multiprocessing multithreading mutex parallel-processing pthreads python-multithreading synchronization task-parallel-library thread-safety threadpool	12
(0.1, 0.005)	atomic backgroundworker concurrency deadlock executorservice grand-central-dispatch locking multiprocessing multithreading mutex openmp parallel-processing pthreads python-multithreading queue runnable semaphore synchronization synchronized task task-parallel-library thread-safety threadpool wait	24
(0.05, 0.015)	asynchronous c++11 concurrency locking multithreading parallel-processing pthreads sockets synchronization thread-safety threadpool	11
(0.05, 0.01)	android-asynctask asynchronous boost c++11 concurrency locking multiprocessing multithreading mutex parallel-processing pthreads python-multithreading sockets synchronization task-parallel-library thread-safety threadpool timer	18
(0.05, 0.005)	android-asynctask async-await asynchronous atomic backgroundworker boost c++11 concurrency deadlock executorservice grand-central-dispatch locking multiprocessing multithreading mutex openmp parallel-processing process pthreads python-multithreading queue runnable semaphore sockets synchronization synchronized task task-parallel-library thread-safety threadpool timer wait	32

**Step 4: Develop race tool keyword set** In this step, we develop a set of keywords  $\mathcal{R}$  to identify and extract concurrency questions that are about tools and techniques that are used to find and fix data race bugs. We develop a



keyword set instead of a tag set because of the following reasons. First, there are no Stack Overflow tags that are specifically designed to represent data race tools and techniques. Second, most Stack Overflow questions with the generic data race tags [32], such as *data-race* and *race-condition*, are not about the usage of data race tools and techniques.

To construct the keyword set  $\mathcal{R}$ , we include the names of both industrial as well as academic data race tools and techniques that are discussed in the previous work. For industrial tools, we include Google’s *ThreadSanitizer* (*TSan*) and *ThreadSafety* (*annotalysis*) [33], Intel’s *Inspector* [34], and Valgrind’s *Helgrind* and *DRD* [35, 36]. For academic tools and techniques, we add 29 tools and techniques from a recent and comprehensive survey on data race tools and techniques [37]. We exclude tools and techniques with no specific names or overly generic names such as *HAVE*.

$$\mathcal{R} = \left\{ \begin{array}{l} \textit{ThreadSanitizer} \textit{TSan} \textit{ThreadSafety} \textit{annotalysis} \textit{Helgrind} \textit{DRD} \textit{'Intel} \\ \textit{Inspector'} \textit{Acculock} \textit{Eraser} \textit{FastTrack} \textit{LiteRace} \textit{Racer} \textit{Race-} \\ \textit{Track} \textit{RACEZ} \textit{SOS} \textit{TRaDe} \textit{Atom-Aid} \textit{AtomRace} \textit{AVIO} \textit{CTrigger} \\ \textit{DefUse} \textit{FALCON} \textit{McPatom} \textit{PENELOPE} \textit{MultiRace} \textit{MUVI-Eraser} \\ \textit{AtomTracker} \textit{ColorSafe} \textit{COPPER} \textit{Marathon} \textit{MUVI-AVIO} \textit{Veldrome} \\ \textit{Chord} \textit{RacerX} \textit{RccJava} \textit{RELAY} \end{array} \right\}$$

The keyword set  $\mathcal{R}$  with 36 keywords covers a broad set of static and dynamic tools and techniques for finding and fixing data race bugs. To illustrate, *Chord* and *RacerX* are static whereas *ThreadSanitizer* and *ThreadSafety* are dynamic. A static race detection technique analyzes a program without running it, whereas a dynamic technique analyzes the runtime information of the program.

**Step 5: Extract race tool posts** In this step, we first extract concurrency questions from  $\mathcal{C}$  that their title or body include a keyword in  $\mathcal{R}$ . This set includes 446 questions. Second, we manually study these questions to exclude the questions that are not about data race tools. This is because a keyword in  $\mathcal{R}$  may refer to different tools and techniques used for different purposes. To illustrate, *FALCON* is not only a data race tool but also a Python web development framework. Similarly, *Helgrind* can not only be used for race detection but also deadlock. Our final set  $\mathcal{U}$  of race tool posts includes 82 questions. Interestingly, the questions in  $\mathcal{U}$  are related to the following 6 keywords in  $\mathcal{R}$  and there are no questions in  $\mathcal{U}$  for the rest of keywords.

$$\{ \textit{ThreadSanitizer} \textit{TSan} \textit{Helgrind} \textit{DRD} \textit{'Intel} \textit{Inspector'} \textit{CTrigger} \textit{RacerX} \}$$

**Step 6: Preprocess concurrency posts** In this step, we preprocess the set of our concurrency posts  $\mathcal{C}$  to reduce the noise for sentiment analysis and topic modeling in the next steps. For sentiment analysis, we preprocess  $\mathcal{C}$  by removing code snippets, enclosed in  $\langle \textit{code} \rangle \backslash \langle \textit{code} \rangle$  tags, HTML tags, such as  $\langle p \rangle \langle /p \rangle$ , numbers, punctuation marks, non-alphabetical characters, and URLs. For topic modeling, we further preprocess  $\mathcal{C}$  by removing stop words, such as “a”, “the”, and “is”, and reduce words to their base representations. For example,

“reading”, “read” and “reads” all reduce to their base “read”. Reducing words to their bases allows grouping of words with similar meanings together. For stop words and word reduction, we use MALLET’s popular list of stop words [38] and Porter stemming algorithm [38], respectively. We do not reduce words for sentiment analysis because an unreduced word may convey important sentiment information [39, 4, 27].

### 2.1.1. Data collection for concurrency questions and answers

Our data collection steps are in line with the best practices of previous work. Previous work uses tags, keywords, and significance and relevance heuristics often to identify security [12, 40], chatbot [41], mobile [11], big data [17], and deep learning [42, 43] questions and answers from Stack Overflow.

## 2.2. Data Analysis

**Step 7: Model and label concurrency topics** In this step, we use MALLET [44] with latent Dirichlet allocation (LDA) topic modeling [23] to group our concurrency posts  $\mathcal{C}$  into topics. In our topic model, a concurrency post is a probabilistic distribution of several topics with different proportions. A post has a dominant topic that covers the biggest proportion of the text of the post. A topic is a set of frequently co-occurring words that approximates a real-world concept. To produce the model, we treat each individual question and accepted answer as an individual document. MALLET processes 245,541 questions and answers documents.

MALLET groups posts into  $K$  topics after  $I$  grouping iterations and returns a set of topics and their proportions for each post along with a set of words for each topic. Our experiments with a broad range of values show that  $K = 30$  and  $I = 1,000$  allow for sufficiently granular topics. These values are in line with values used by previous work [8, 11, 12, 17]. MALLET uses hyperparameters  $\alpha$  and  $\beta$  to control the distribution of words in topics and distributions of topics in posts. Following previous work [8, 11, 12, 17], we use standard values  $50/K$  and  $0.01$  for these hyperparameters. We use this standard values because as Bigger et al. [45] show “ $\alpha$  and  $\beta$  have little influence on the accuracy of the LDA [topic modeling]”.

MALLET groups concurrency posts  $\mathcal{C}$  into topics  $t$  that are a set of words. However, MALLET cannot decide about the meaning of the topics and label them with names useful for human beings. To illustrate, the word set  $w = \{task, execute, async, complete, run, cancel, wait, asynchronous, parallel, schedule\}$  represents a topic in MALLET. Following previous work [11, 12, 7, 8, 25, 17], we use an open card sort technique [24] to label the topics. In an open card sort, there are no predefined topics and participants develop their own topics during the sorting and labeling process. To label a topic, we find it sufficient to manually inspect the top 20 words in the set of words for the topic and read through 15 random posts for the topic to decide for a label that best explains the words and posts of the topic. To illustrate, we label the word set  $w$  with *task parallelism*. The numbers of topic words and random posts used to label a

topic are in line with previous work [17, 8, 7, 14, 18, 11, 12]. Table 2 shows our 27 concurrency topics, their names, and top 10 reduced words.

During the labeling process, we merge topics 1 and 15 into *basic concepts* topic due to close similarities between their topic words and questions. Similarly, we merge topics 5 and 28 into *object-oriented concurrency*. We remove topic 12 that is about synchronization between local and remote repositories in version control systems such as Git and is not about concurrency.

During the card sorting process for concurrency topics, the first and second authors individually assign labels to topics and reiterate and refine topics until they agree on topic labels. The first author is a Programming Languages and Software Engineering professor with extensive expertise in concurrent and event-based systems [16, 46, 47, 17, 30, 48, 49, 50] and the second author is a graduate student with coursework in concurrent and distributed systems. The first author also has several years of industrial experience as a Software Engineer.

**Step 8: Construct concurrency topic hierarchy** In this step, we use a card sorting process, similar to the step 7 and with the same authors, to construct the topic hierarchy by repeated grouping of similar topics into categories, and lower-level categories into higher-level categories. There are no predefined categories and participants develop their own categories during the sorting and labeling process. Table 2 and Figure 4 show the textual and pictorial representations of the topic hierarchy. To illustrate, *thread life cycle management* and *thread scheduling* topics are grouped into the lower-level category *multithreading*, where *multithreading* itself is grouped into the higher-level category *concurrency models*. The category *concurrency models* includes other lower-level categories such as *multiprocessing* and *parallel computing*.

**Step 9: Determine concurrency topic popularity** In this step, we use 3 well-known metrics that are used often by the previous work to measure the popularity of a concurrency topic. The first metric, that is used by the previous work [8, 11, 12, 25, 17, 41], is the average number of views for questions with the topic as their dominant topic. This metric includes views by both registered users and visitors of Stack Overflow. The inclusion of views by visitors is important because in Stack Overflow there are many more visitors than there are registered users [51]. The second metric, that is used by the previous work [12, 18, 25, 8, 17, 41] is the average number of questions of the topic marked as favorite by users. The third metric, that is used by the previous work [12, 18, 25, 8, 17, 41], is the average score of questions of the topic. Intuitively, a topic with a higher number of views and favorites and a higher score is more popular. Table 3 shows the popularity measurements of concurrency topics.

**Step 10: Determine concurrency topic difficulty** We use 2 well-known metrics that are used often by previous work to measure the difficulty of a concurrency topic. The first metric, that is used by the previous work [12, 11, 9, 17, 41], is the percentage of questions of the topic that have no accepted answers. And the second metric, that is used by the previous work [12, 11, 17, 41], is the average median time needed for a question to receive an accepted answer. Intuitively, a topic with higher percentage of questions with no accepted answers and taking longer to receive accepted answers is more

difficult. Table 4 shows the difficulty measurements of our concurrency topics.

**Step 11: Determine concurrency topic sentiment** We use Senti4SD [52] to determine the sentiment of our concurrency topics. Senti4SD uses a combination of lexical and semantic features to determine the sentiment polarity of a concurrency post and classify it as positive, negative, or neutral. For example, Senti4SD classifies the sentence “I am so used to thinking about solutions in a linear/serial/OOP/functional way that I am struggling to cast any of my domain problems in a way that merits using concurrency” from a Stack Overflow post as negative whereas it classifies the sentence “what challenges promote the use of parallel/concurrent architectures?” as neutral. Senti4SD is specifically trained to analyze communications of software developers. Previous work [53] shows that Senti4SD provides the best performance, in comparison to other sentiment analysis tools such SentiStrength [54], SentiCR [55], and SentiStrengthSE [56], when analyzing Stack Overflow questions and answers. The sentiment of a concurrency topic is the percentages of the positive, negative, and neutral polarities of its questions and answers. Table 5 shows the sentiment measurements of concurrency topics.

**Step 12: Determine popularity, difficulty, and sentiment correlations** In this step, we use Kendall correlation to identify if there are any correlations between popularity, difficulty, and sentiment of our concurrency topics. We use Kendall because it is less sensitive to outliers and is relatively more stable. In total, we investigate 21 correlations between 3 popularity, 2 difficulty, and 3 sentiment metrics of our concurrency topics.

**Step 13: Label race tool topics** In this step, we use the card sorting technique, similar to steps 7 and 8, to group race tool questions in  $\mathcal{U}$  into race tool topics. During the card sorting process, the first and fourth authors individually assign labels to topics and reiterate and refine the topics until they agree on topic labels. The fourth author is a Software Engineering professor with extensive expertise in parallel systems [46, 2]. The fourth author also has several years of industrial experience as a Software Engineer.

**Step 14: Investigate the relation of our findings with the findings of previous work** Throughout steps 7–13 of our analysis, we investigate the relation of our findings in each of these steps with the findings of previous work.

### 2.2.1. Data analysis of concurrency questions and answers

Our data analysis steps are in line with the best practices of previous work. Previous work uses card sorting and manual analysis often to group security [12, 40], chatbot [41], mobile [11], big data [17], and deep learning [42, 43] questions and answers from Stack Overflow. In addition, previous work uses the average number of views, favorites, and scores as popularity metrics and the percentage of questions with no answers and median time to receive accepted answers as difficulty metrics to measure the popularity and difficulty of security [12, 40], chatbot [41], mobile [11], and big data [17] topics.

### 3. Results

In this section, we present and discuss the results of our study for research questions **RQ1-RQ7**, investigate the relation of our findings with the findings of the previous work, and present a set of example questions that developers ask for each of our concurrency and race tool topics and categories.

#### 3.1. RQ1: Concurrency Topics

Table 2 shows the concurrency topics that developers ask questions about on Stack Overflow. These topics are determined in step 7 of our analysis. According to Table 2, developers ask questions about a broad spectrum of concurrency topics. These questions can be categorized into 27 concurrency topics ranging from *thread pool* to *parallel computing*, *mobile concurrency* to *web concurrency*, and *memory consistency* to *runtime speedup*.

The meaning of a concurrency topic is best understood by looking at the questions that belong to the topic. To illustrate, the following is a question in the *thread pool* topic. In this question, the developer is asking how to implement a thread pool where the size of the pool can change based on the number of its jobs. The Stack Overflow identifier for this question is *11249342* and it can be accessed at <https://stackoverflow.com/questions/11249342>. This question has been viewed 13,448 times, more than 6 times the average number of views for a Stack Overflow question. The average number of views for a Stack Overflow question is 2,154 [17]. A thread is an abstraction for the execution of a computation that allows to control its concurrency by running, suspending, and resuming the execution of the computation. A thread pool is an abstraction for a set of threads that allows for their creation, allocation, and deallocation.

*Q.11249342* **Creating a dynamic (growing/shrinking) thread pool** I need to implement a thread pool in Java (`java.util.concurrent`) whose number of threads is at some minimum value when idle, grows up to an upper bound (but never further) when jobs are submitted into it faster than they finish executing, and shrinks back to the lower bound when all jobs are done ... How would you implement something like that? ...

Similarly, the following is a question in the *web concurrency* topic. In this question, the developer is asking how to send emails using background threads in their web application where the background thread prevents blocking of the main thread and therefore does not force the user to wait until the email is sent. Classic ASP (Active Server Pages) is a scripting language to write server-side web applications. Ajax is a client-side web development technique for asynchronous communication between a web client and server. In an asynchronous request, the client does not block to wait for the response from the server and can immediately continue its computation after the request is made.

*Q.17052243* **How to perform multithreading/background process in classic ASP** I need to send emails via a background job on a classic-ASP app so the user doesn't have to wait for a slow webserver to complete sending the email. I know I can use Ajax to generate two separate requests, but I'd

Table 2: Names, categories and top 10 reduced words of our 27 concurrency topics.

<i>No.</i>	<i>Topic name</i>	<i>Category</i>	<i>Topic words</i>
1	basic concepts	basic concepts	code question work answer understand read find edit issu make
2	thread life cycle management	multithreading\ concurrency models	thread main creat run start execut back-ground separ join finish
3	concurrent collections	correctness	list arrai map element collect iter number item kei add
4	entity management	persistence	concurr session spring entiti transact actor collect updat model version
5	object-oriented concurrency	programming paradigms	object class method instanc creat static access refer variabl synchron
6	task parallelism	concurrency models	task execut async complet run cancel wait asynchron parallel schedul
7	thread sharing	multithreading\ concurrency models	function variabl pass call pointer pthread argument type return global
8	process life cycle management	multiprocessing\ concurrency models	process child parent termin exit fork creat kill share start
9	thread scheduling	multithreading\ concurrency models	loop time wait stop run start sleep set check finish
10	thread safety	correctness	thread java safe multipl multi multithread concurr singl implement applic
11	runtime speedup	performance	time core cpu run perform memori number system process machin
12	REMOVED		context user set sync local chang synchron save work updat
13	web concurrency	programming paradigms	request servic web server applic user app net respons http
14	database management systems	persistence	databas tabl queri updat row record lock insert sql transact
15	locking	correctness	lock mutex wait condit releas semaphor acquir deadlock synchron resourc
16	basic concepts	basic concepts	implement make gener librari approach oper case requir solut good
17	thread pool	multithreading\ concurrency models	worker pool job number work process task creat limit threadpool
18	data scraping	performance	data time problem load download work page structur solut url
19	unexpected output	debugging	code program work run output problem print line result function
20	irreproducible behavior	debugging	error test code run problem applic issu window compil crash
21	event-based concurrency	programming paradigms	event signal handler timer callback call handl fire receiv slot
22	memory consistency	correctness	read memori write variabl oper atom cach share synchron access
23	producer-consumer concurrency	concurrency models	queue messag consum produc process item buffer block wait empti
24	GUI	GUI	updat form gui window applic button control user progress click
25	parallel computing	concurrency models	parallel node loop comput calcul openmp result algorithm mpi function
26	python multiprocessing	multiprocessing\ concurrency models	python script run multiprocessing process function modul command parallel php
27	file management	persistence	file read write log line open stream directori folder text
28	mobile concurrency	programming paradigms	android imag app activ game view updat frame async task devic
29	object-oriented concurrency	programming paradigms	call method block return code execut synchron function make time
30	client-server concurrency	programming paradigms	server client connect send socket messag receiv data port read

rather not require Javascript. Plus, I suspect there's a better way to pull this off. Ideas?

Finally, the following is a question in the *basic concepts* topic. In this question, the developer asks about the basic motivations behind the need for concurrency and concurrency architectures.

Q.541344 **What challenges promote the use of parallel/concurrent architectures?** .. I am so used to thinking about solutions in a linear/serial/OOP [Object-oriented Programming]/functional way that I am struggling to cast any of my domain problems in a way that merits using concurrency.

**Finding 1:** Questions that concurrency developers ask can be grouped into 27 concurrency topics ranging from *thread pool* to *parallel computing*, *mobile concurrency* to *web concurrency*, and *memory consistency* to *runtime speedup*.




Figure 3: Concurrency topics and numbers of their questions.

Figure 3 shows the number of questions that developers ask about our concurrency topics. According to Figure 3, the numbers of questions that developers ask for concurrency topics are not uniform. Developers ask the most number of questions (8%) about *basic concepts* and the least number of questions (1%) about *event-based concurrency*.

Pinto et al. [18] study the 250 most popular concurrency questions on Stack Overflow and group them into several themes. Our observation that the concurrency topic *basic concepts* has the most number of questions is similar to Pinto et al.’s observation that “most of them [concurrency questions] are related to basic concepts”. Our observation is also in line with the general understanding that the concurrency and even its basics remains challenging for developers. Bagherzadeh and Khatchadourian [17] study and categorize Stack Overflow

questions and answers related to big data software development into several big data topics and categories including the *basic concepts* topic. Our *basic concepts* topic overlaps with Bagherzadeh and Khatchadourian’s *basic concepts*. Interestingly, both concurrency and big data developers ask more than average number of questions about *basic concepts*.

**Finding 2:** Developers ask the most questions about *basic concepts* inline with the general understanding that the concurrency and even its basics remains challenging for developers. Developers ask the least about *event-based concurrency*.

### 3.2. RQ2: Concurrency Topics Hierarchy

Figure 4 shows the hierarchy of our concurrency topics. This hierarchy is constructed in step 8 of our study. In this figure, concurrency topics are in gray and their categories are in white with the percentages of their questions. The hierarchy expands outwards from higher-level to lower-level categories to topics.

According to Figure 4, questions that concurrency developers ask can be grouped into a hierarchy with 8 high-level categories: *concurrency models*, *programming paradigms*, *correctness*, *debugging*, *basic concepts*, *persistence*, *performance*, and *graphical user interface (GUI)*. In addition, the number of questions that developers ask in each category is not uniform. Developers ask the most (28%) about the *concurrency models* and the least (5%) about *GUI*. Interestingly, developers ask more (12%) about *correctness* of their concurrent programs than *performance* (7%). This is in line with the general tradeoff between the performance advantages of concurrency and its correctness issues [48, 50, 49, 57].

**Finding 3:** Questions that developers ask can be grouped into a hierarchy with 8 high-level categories: *concurrency models*, *programming paradigms*, *correctness debugging*, *basic concepts*, *persistence*, *performance*, and *GUI*.

**Finding 4:** Developers ask the most questions about *concurrency models* category and the least about *GUI*.

**Finding 5:** Developers ask more about *correctness* than *performance*, in line with the correctness and performance tradeoffs of concurrent software.

In the following, we discuss each higher-level concurrency category and its lower-level categories and topics. In addition, we present a set of example questions for each concurrency topic and category.



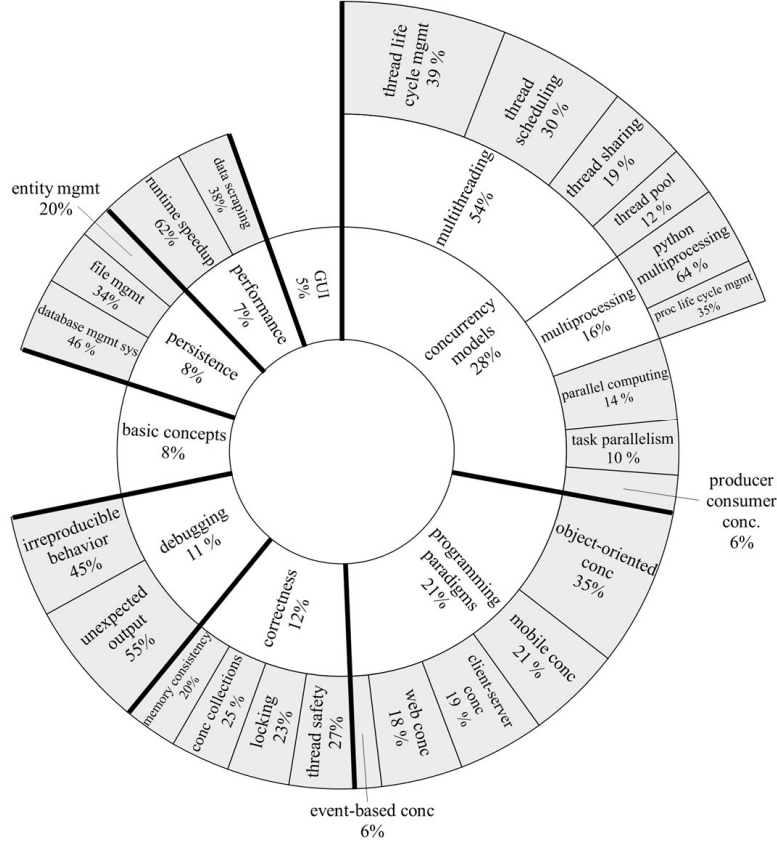


Figure 4: Hierarchy of concurrency topics, in gray, their categories, in white, and percentages of their questions.

### 3.2.1. Concurrency Models

Concurrency models are concerned about concurrency abstractions (e.g. threads and processes) and execution models (e.g. multicore and single core executions). The *concurrency models* category includes five lower-level categories among which *multithreading* alone contains more than half of the questions that developers ask in this category. This is in line with the general understanding that multithreading is the defacto dominant concurrency model. Developers ask the most (54%) about *multithreading* and the least (6%) about *producer-consumer concurrency*, when asking questions about concurrency models.

In *multithreading*, developers ask questions with titles like (thread life cycle management): “How to safely destruct Posix thread pool in C++?” and (thread pool): “Is there a way to create a pool of pools using the Python work-erpool module?”. The topic of a question is included inside parentheses. In *producer-consumer concurrency*, developers ask questions with titles like “pro-ducer Consumer with BlockingQueues in Java EE as background task”.

Our *multithreading* category and its *thread life cycle management* topic overlap with Pinto et al.’s [18] *threading* and *thread life cycle* themes. Rosen and Shihab [11] study and categorize Stack Overflow questions related to mobile development into several mobile topics including *threading*. Our *multithreading* category overlaps with their mobile *threading* topic.

**Finding 6:** Developers ask the most about *multithreading*, when asking about *concurrency models*, in line with the general understanding that multithreading is the defacto dominant concurrency model. Developers ask the least about *producer-consumer concurrency*.

### 3.2.2. Programming Paradigms

Programming paradigms are concerned about programming abstractions (e.g. objects and events), platforms (e.g. web and mobile), and patterns (e.g. producer and consumer). The *programming paradigms* category includes five lower-level categories among which *object-oriented concurrency* alone contains more than a third of questions that developers ask in this category. This is in line with the general understanding that object-orientation is a dominant programming paradigm. Developers ask the most (35%) about *object-oriented concurrency* and the least (6%) about *event-based concurrency*, when asking about programming paradigms.

In the *object-oriented concurrency* and *event-based concurrency* topics, developers ask questions with titles like (object-oriented concurrency): “*Is it better to synchronize object from inside of the class that encapsulates access it or from outside?*” and (event-based concurrency): “*What type of timer event should I use for a background process when my timer fires very quickly?*”.

Barua et al. [7] study and categorize all Stack Overflow questions and answers into several general topics including web and mobile development. Our concurrency topics *mobile concurrency* and *web concurrency* overlap with their general web development and mobile development topics. Our *mobile concurrency* topic overlaps with Pinto et al.’s [18] observation that “concurrent programming has reached mobile developers”. In Pinto et al.’s [18] study, 9% of Stack Overflow questions are related to mobile development. In contrast, according to Figure 3, only 4.6% of our concurrency questions are in *mobile concurrency* topic.

**Finding 7:** Developers ask the most about *object-oriented concurrency*, when asking about *programming paradigms*, in line with the general understanding that object-orientation is a dominant concurrent programming paradigm. Developers ask the least about *event-based concurrency*.

### 3.2.3. Correctness

Correctness is concerned with the prevention of data corruption for concurrently accessed (e.g. read and write) data using mechanisms like locking,

consistent memory models, and thread-safe data structures and programming patterns. *Correctness* questions are divided almost evenly among its *thread safety*, *locking*, *concurrent collections*, and *memory consistency* topics.

In *correctness* category, developers ask questions like (thread safety): “How to make factory [pattern] thread safe?”, (locking): “Is there a way to lock 2 or more locks or monitors atomically?”, (concurrent collections): “Threadsafe dictionary that does lookups with minimal locking”, and (memory consistency): “Atomic read-modify-write in C#”.

Lu et al. [1] categorize concurrency bug patterns and their fixes. Our *memory consistency* topic overlaps with their observations that most concurrency bugs are atomicity violation bugs where the “desired serializability among multiple memory accesses is violated” and order violation bugs where the “desired order between two (groups of) memory accesses is flipped.” Our *locking* topic overlaps with their designation that locking is one of the main fixes for concurrency bugs to ensure correctness. In addition, our *correctness* category and its *locking* topic overlap with Pinto et al.’s [18] *correctness* and *locking* themes. Similarly, our *concurrent collections* topic overlaps with Pinto et al.’s *concurrent libraries* theme.

**Finding 8:** *Correctness* questions are divided almost evenly between *thread safety*, *locking*, *concurrent collections*, and *memory consistency* topics.

#### 3.2.4. Basic Concepts

*Basic concepts* are about both theoretical and practical aspect of concurrency and includes questions with titles like “How many threads are involved in deadlock?”, “What is a race condition?”, “Lock, mutex, semaphore... what’s the difference?”, and “Java: notify() vs. notifyAll() all over again”.

Our *basic concepts* category overlaps with Pinto et al.’s [18] *theoretical concepts* and *practical concepts* themes.

#### 3.2.5. Debugging

Debugging is concerned about finding and fixing concurrency bugs which manifest either in the behavior or output of programs. Debugging questions are almost evenly divided among its topics *irreproducible behavior* and *unexpected output*, which includes question with titles like “Trace non-reproducible bug in C++” and “Synchronized codes with unexpected outputs”.

Our *debugging* category overlaps with Bagherzadeh and Khatchadourian’s [17] big data *debugging* topic. Our *irreproducible behavior* topic in *debugging* category overlaps with Lu et al.’s [1] observation that some “concurrency bugs are very difficult to repeat”.

**Finding 9:** *Debugging* questions are divided almost evenly between *irreproducible behavior* and *unexpected output* topics.

### 3.2.6. Persistence

Persistence is about storing and retrieving of data using persistence management systems (e.g. databases management systems, entity/object persistence<sup>1</sup> or file systems). *Persistence* includes three topics among which *database management systems* includes near half of persistence questions. The *persistence* category includes question with titles like (database management system): “How do I lock read/write to MySQL tables so that I can select and then insert without other programs reading/writing to the database?”, (file management): “Is this is correct use of mutex to avoid concurrent modification to file?”, and (entity management): “Save entity using threads with JPA [Java Persistence API] (synchronized)”.

Our *database management systems* topic overlaps with Barua et al.’s [7] general *MySQL* topic.

**Finding 10:** Nearly half of *persistence* questions are about *database management systems*.

### 3.2.7. Performance

Performance is about speeding up execution of programs (e.g. data scraping programs<sup>2</sup>). *Performance* includes two topics. These topics includes questions with titles like (runtime speedup): “Poor multithreading performance in .Net” and (data scraping): “Using multithreading to speed up web crawler written by beautifulsoup4 and python”.

Pinto et al. [18] observe that they “did not find questions that ask for advices on how to use concurrent programming constructs to improve application performance, which is surprising, since performance is one of the most important motivations for the use of concurrency and parallelism”. In contrast, according to Figure ??, our *performance* topic includes more than 7% of all of our concurrency questions. In addition, our *performance* topic overlaps with Bagherzadeh and Khatchadourian’s [17] big data *performance* topic.

### 3.2.8. GUI

Graphical user interface (GUI) allows for the graphical interaction between a software and its users. *GUI* is the smallest category with regard to the number of questions and includes question with titles like “Force GUI update from UI Thread” and “Object synchronization with GUI Controls”.

Our concurrency topic *GUI* overlaps with Rosen and Shihab’s [11] *UI* topic for mobile programming.

---

<sup>1</sup>An entity management system automates the serialization of objects (entities) for storage in a database.

<sup>2</sup>A data scraping program downloads data from remote web URLs and stores it locally.

### 3.3. RQ3: Popularity of Concurrency Topics

Table 3 shows the popularities of our concurrency topics. These popularities are measured using 3 metrics average number of views, favorites and score, in step 9 of our analysis. Table 3 is sorted by the average number of views for topics and the highest and the lowest values of popularity metrics are in gray. A topic with higher number of views, favorites, and score is more popular.

According to Table 3, the *thread safety* topic has the highest views, third highest favorites, and second highest score whereas *client-server concurrency* has the third lowest views and lowest favorites and score.

Table 3: Popularity of concurrency topics. The highest and lowest values of popularity metrics are in gray.

<i>Topic</i>	<i>Avg. views</i>	<i>Avg. favorites</i>	<i>Avg. score</i>
thread safety	2848	1.5	4.6
basic concepts	2222	1.6	4.3
task parallelism	2216	1.3	4.0
locking	2152	1.3	3.5
thread life cycle management	2130	0.7	2.4
thread scheduling	2032	0.7	2.2
process life cycle management	2004	1.0	2.6
thread pool	1841	0.9	2.7
object-oriented concurrency	1773	0.8	2.6
database management systems	1727	0.6	1.8
thread sharing	1671	0.6	2.0
GUI	1664	0.5	1.6
irreproducible behavior	1647	0.6	2.3
event-based concurrency	1636	0.7	2.3
python multiprocessing	1587	0.9	2.5
entity management	1583	0.8	2.3
memory consistency	1531	1.7	4.8
file management	1458	0.6	1.9
producer-consumer concurrency	1311	0.8	2.2
unexpected output	1304	0.5	1.6
mobile concurrency	1292	0.5	1.3
runtime speedup	1276	0.9	2.7
web concurrency	1252	0.8	1.9
concurrent collections	1155	0.5	2.0
client-server concurrency	1083	0.4	1.1
data scraping	1003	0.6	1.4
parallel computing	899	0.6	1.9
<b>Average</b>	<b>1641</b>	<b>0.8</b>	<b>2.5</b>

**Finding 11:** *Thread safety* topic, in *correctness* category, and *client-server concurrency*, in *programming paradigms*, are among the most and least popular concurrency topics, respectively.

### 3.4. RQ4: Difficulty of Concurrency Topics

Table 4 shows the difficulty of concurrency topics. These difficulties are measured using 2 metrics percentage of questions with no accepted answers and average median time to get an accepted answer, in step 10 of our analysis. Table 4 is sorted by the percentage of questions with no accepted answers and the highest and the lowest values of difficulty metrics are in gray. A topic with a higher percentage of questions with no accepted answers and median time to receive accepted answers is more difficult.

According to Table 4, the *irreproducible behavior* topic has the second highest percentage of questions with no accepted answers and highest time to accepted answers whereas *memory consistency* has the lowest percentage of questions with no accepted answers and the second lowest time to accepted answers.

Table 4: Difficulty of concurrency topics. The highest and lowest values of difficulty metrics are in gray.

<i>Topic</i>	<i>% w/o acc. answer</i>	<i>Hrs to acc. answer</i>
database management systems	51.2	1.0
irreproducible behavior	51.1	2.1
web concurrency	50.7	0.9
mobile concurrency	50.4	0.8
client-server concurrency	50.4	0.9
python multiprocessing	50.3	0.9
parallel computing	50.1	2.1
data scraping	48.9	1.0
file management	48.8	0.6
entity management	48.0	1.8
runtime speedup	48.0	0.7
thread pool	47.0	0.7
process life cycle management	44.9	0.6
producer-consumer concurrency	43.3	0.7
unexpected output	41.8	0.7
GUI	41.1	0.4
thread scheduling	40.8	0.4
thread life cycle management	40.4	0.3
locking	40.1	0.3
event-based concurrency	39.7	0.6
thread safety	39.6	0.3
concurrent collections	38.6	0.4
basic concepts	37.0	0.7
thread sharing	35.6	0.3
task parallelism	35.3	0.4
object-oriented concurrency	35.2	0.3
memory consistency	33.2	0.4
<i>Average</i>	<i>43.8</i>	<i>0.7</i>

**Finding 12:** *Irreproducible behavior* topic, in *debugging* category, and *memory consistency*, in *correctness*, are among the most and least difficult concurrency topics, respectively.

### 3.5. RQ5: Sentiment of Concurrency Topics

Table 5 shows the sentiment of concurrency topics. These sentiments are measured using percentages of their positive, neutral, and negative polarities, in step 11 of our study. Table 5 is sorted by the positive polarity and the highest and lowest values of polarities are in gray. A topic is neutral if its neutral polarity is higher than its positive and negative polarities. The positive and negative topics are defined similarly.

According to Table 5, the neutrality of all concurrency topics is higher than their positivity and negativity. That is, all concurrency topics are neutral in sentiment. *Process life cycle management* and *irreproducible behavior* are the most and least neutral topics. *Data scraping* and *process life cycle management* topics are the most and least positive whereas *irreproducible behavior* and *process life cycle management* are the most and least negative.

Interestingly, in sentiments, the most positive is not the same as the least negative. For example, *data scraping* is the most positive topic but not the least negative. The same is true for the least positive and most negative sentiments.

Guzman et al. [26] analyze the emotions of commit logs in 29 general GitHub projects. Our observation that all concurrency topics are neutral is similar to Guzman et al.’s observation that “the average emotion score of the commit comments for each of the projects tended to neutrality”. Similarly, Sinha et al. [27] analyze the sentiment of 2,251,585 commit logs in 28,466 general GitHub projects. Our observation is also similar to Sinha et al.’s observation that “a majority [74.74%] of the sentiment in GitHub projects are categorized as neutral”. Tourani et al. [28] analyze emotions of 595,673 developer and user emails from Apache’s Ant and Tomcat mailing lists. Finally, our observation is similar to Tourani et al.’s observation that “19.77% of the sampled emails were positive, 11.27% negative. [And, 68.95% were neutral]”.

**Finding 13:** All concurrency topics are neutral in sentiment.

**Finding 14:** *Data scraping* and *process life cycle management* topics are the most and least positive whereas *irreproducible behavior* and *process life cycle management* are the most and least negative.

### 3.6. RQ6: Correlations of Popularity, Difficulty, and Sentiment

According to Tables 3-5, a topic like *client-server concurrency* has lower popularity, higher difficulty, and a more negative sentiment. Such anecdotal evidence could suggest an intuition that there may be correlations between the

Table 5: Sentiment of concurrency topics. The highest and lowest sentiment polarities are in gray.

<i>Topic</i>	<i>Positive%</i>	<i>Neutral%</i>	<i>Negative%</i>
data scraping	26.7	45.5	27.7
database management systems	21.7	56.8	21.4
entity management	21.5	58.2	20.1
thread safety	20.4	60.1	19.3
concurrent collections	20.3	62.8	16.8
mobile concurrency	19.4	49.4	31.1
basic concepts	19.3	59.7	20.9
parallel computing	18.8	60.3	20.8
unexpected output	18.3	44.7	36.9
web concurrency	18.2	59.1	22.6
python multiprocessing	17.3	55.6	26.9
client-server concurrency	16.9	52.1	30.8
GUI	16.7	52.4	30.7
runtime speedup	16.6	62.4	20.9
object-oriented concurrency	16.4	63.4	20.1
irreproducible behavior	16.3	40.6	43.0
memory consistency	15.0	68.1	16.8
producer-consumer concurrency	14.4	62.4	23.1
thread scheduling	14.3	57.1	28.4
task parallelism	13.5	66.3	20.0
file management	12.1	69.3	18.5
event-based concurrency	11.7	67.9	20.2
thread pool	11.6	71.7	16.6
thread life cycle management	10.7	66.5	22.6
thread sharing	10.5	64.9	24.5
locking	9.6	71.6	18.6
process life cycle management	7.6	78.2	14.0
<i>Average</i>	<i>16.4</i>	<i>59.8</i>	<i>24.0</i>

popularity, difficulty, and sentiment of concurrency topics where a less popular topic is more difficult and more negative. To investigate, we calculate 21 correlations between 3 popularity, 2 difficulty, and 3 sentiment metrics of our concurrency topics. These correlations are calculated in step 12 of our analysis.

Table 6 shows the p-values and directions (positive or negative) of these correlations. According to Table 6, there is a statically significant negative correlation with 95% confidence between the popularity and difficulty metrics of concurrency topics, except for the average number of favorites popularity metric and the hours to accepted answer difficulty metric. That is, more popular concurrency topics are usually less difficult to find answers to their questions. Similarly, there is a statically significant negative correlation between the popularity metrics and negative sentiment of concurrency topics, except for the average number of views popularity metric and the negative sentiment. That is, more popular topics are usually less negative. Note that these correlation do not imply causality.



Pinto et al. [14] study 300 questions and 550 answers related to software energy consumption on Stack Overflow and group them into several themes. Our observation that the more popular concurrency questions are less difficult to find answers for is in contrast to Pinto et al.’s observation that their *Measurement* theme is not only their most difficult theme but also the most popular. Our observation is also in contrast to Rosen and Shihab’s [11] observation that finds Barua et al.’s [7] popular *mobile development* topic to be difficult. Wang et al. [29] study the relation between time to answer a question with 46 factors of questions, answers, askers, and answerers in 55853, 70336, 7134, and 10776 questions on 4 question and answer websites Stack Overflow, Mathematics, Super User, and Ask Ubuntu. Our observation is also similar to Wang et al.’s observation that “slow-answered [more difficult] questions are usually associated with rarer [less popular] tags than fast-answered questions across the four studied websites”.

Table 6: Correlations of popularity, difficulty, and sentiment of concurrency topics. The statically significant correlations are in gray.

<i>direction/p-value</i>		<i>Popularity</i>		
		<i>Avg. views</i>	<i>Avg. favorites</i>	<i>Avg. score</i>
<i>Difficulty</i>	<i>% w/o acc. answer</i>	- / 0.0044	- / 0.01166	- / 0.001073
	<i>Hrs to acc. answer</i>	- / 0.01404	- / 0.09196	- / 0.02726

<i>direction/p-value</i>		<i>Sentiment</i>		
		<i>Positive</i>	<i>Neutral</i>	<i>Negative</i>
<i>Difficulty</i>	<i>% w/o acc. answer</i>	+ / 0.21851	- / 0.15604	+ / 0.57319
	<i>Hrs to acc. answer</i>	+ / 0.01700	- / 0.05727	+ / 0.67260

<i>direction/p-value</i>		<i>Popularity</i>		
		<i>Avg. views</i>	<i>Avg. favorites</i>	<i>Avg. score</i>
<i>Sentiment</i>	<i>Positive</i>	- / 0.02860	- / 0.28198	- / 0.19498
	<i>Neutral</i>	+ / 0.24294	+ / 0.00469	+ / 0.00657
	<i>Negative</i>	- / 0.78629	- / 0.01915	- / 0.04043

**Finding 15:** More popular concurrency topics are usually less difficult and less negative.

### 3.7. RQ7: Race Tool Topics

Figure 5 shows the race tool topics that concurrency developers ask questions about on Stack Overflow. These topics are determined in step 13 of our analysis. According to Figure 5, the questions that developers ask about race tools can be grouped into 4 topics: *root cause identification*, *general*, *configuration & execution* and *alternative use*. In addition, the number of questions that developers ask in each topic is not uniform. Developers ask the most questions (68%) about *root cause identification* and the least (9%) about *alternative use* of race tools.

**Finding 16:** Questions that concurrency developers ask about race tools can be grouped into 4 topics *root cause identification*, *general*, *configuration & execution* and *alternative use*.

**Finding 17:** Developers ask the most number of race tool questions about *root cause identification* and the least about *alternative use*.




Figure 5: Race tool topics and numbers of their questions.

In the following, we discuss race tool topics and present a set of example questions for each topic.

### 3.7.1. Root Cause Identification

*Root cause identification* is about using the output of race detection tools to identify the source of the reported races in the code. *Root cause identification* could become challenging because of the high false positive rate, buggy behavior of the tools, and the difficulty in understanding their outputs. More than two thirds of race tool questions are about *root cause identification*. This is in line with the general understanding that root cause identification is a major step when debugging software systems. *Root cause identification* include questions like “ThreadSanitizer says my Atomic Inc/Dec has data races, false positive?”, “Why does valgrind drd think pthread\_barrier\_wait is buggy?”, “intepretation of Valgrind output to figure out the location of data race”, and “Can’t figure out where race condition is occuring”.

Sadowski and Yi [33] study the use of widely used ThreadSafety and Thread Sanitizer (TSan) concurrency tools at Google and identify tool usage themes by interviewing 7 developers. Our observation that developers are concerned about false positives in *root cause identification* is similar to Sadowski and Yi’s observation that “another theme that repeatedly emerged was the importance of a low false positive rate”.

**Finding 18:** More than two thirds of race tool questions are about *root cause identification*, in line with the general understanding that root cause identification is a major debugging step.

### 3.7.2. General

*General* topic is about fundamental guarantees, features, and best practices of race tools and includes question with titles like “Does one run of Helgrind suffice to guarantee that the given multithreaded implementation is data-race free and deadlock-free?”, “What features does threadsanitizer lack that are supported by helgrind, and vice-versa?”, and “What threading analysis tools do you recommend?”. Interestingly, about 1 in 7 (13.5%) of questions that developers ask about race tool are *general* questions. This means that usage of race tools and even their basics remains challenging for developers.

**Finding 19:** Near 1 in 7 questions that developers ask about race tools are *general* questions which means that the usage of race tools and even their basics remains challenging for concurrency developers.

### 3.7.3. Configuration and execution

*Configuration and execution* is about the proper configuration and execution of race detection tools for different execution environments (e.g. operating systems and frameworks). *Configuration and execution* includes question like “Is there a way to run helgrind/drd in android?”, “valgrind/helgrind gets killed on stress test”, and “Valgrind hanging to profile a multi threaded program”. About 1 in 10 race tool questions are about *configuration & execution*.

Our *configuration and execution* topic overlaps with Bagherzadeh and Khatchadourian’s [17] big data *configuration* category.

### 3.7.4. Alternative Use

*Alternative use* is about using race tools not only to detect races but also to monitor, trace, and profile concurrent software for its behavior and performance. *Alternative use* includes questions like “How to monitor each thread behavior of a multithread (pthread) C++ program on Linux?”, “Trace non-reproducible bug in c++”, “How to measure lock contention?”, and “How can I debug what portion of a multi-threaded C++ program is taking excessive time?”. About 1 in 10 race tool questions are about *alternative use* of these tools.

**Finding 20:** Developers ask questions about *alternative use* of race tools for purposes such as profiling and performance monitoring of their concurrent software.

#### 4. Implications

The results of our study can not only help concurrency developers and their practice but also the research and education of concurrent software development to better decide where to focus their efforts. Members of these communities could use tradeoffs between our concurrency and tool usage topics as one of several factors they use in their decision making process. One way to tradeoff a topic against another is using their popularities, difficulties, sentiments, or the number of their questions. Obviously, there are many factors that go into tradeoffs that practitioners, researchers, and educators make to decide where to focus their efforts. However, we believe our findings can contribute to inform and improve these decisions.

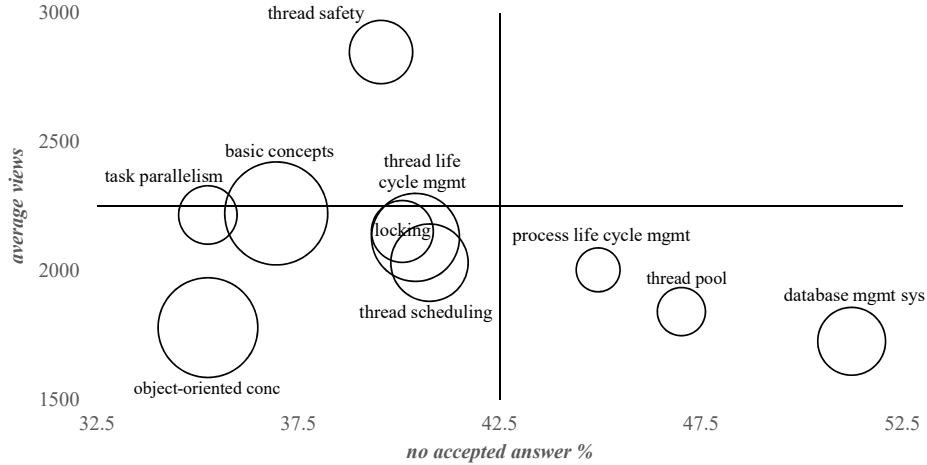


Figure 6: Trading off concurrency topics based on their popularity and difficulty.

To illustrate, Figure 6 shows the difficulty of our top 10 most popular concurrency topics. For simplicity, we use the average number of views as the popularity metric and the percentage of questions without accepted answers as the difficulty metrics of these topics. The popularity of topics on the vertical axis increases from bottom to top. And the difficulty of topics on the horizontal axis increases from left to right. The quadrants halve the ranges for the values of the popularity and difficulty metrics. In Figure 6, a topic is represented by a circle where the size of the circle is proportional to the number of questions in the topic, according to Figure 3. Similarly, Figure 7 shows the sentiment of our top 10 most popular concurrency topics. Again, for simplicity, we use the percentage of positive sentiment as the metric for the sentiment of topics. The sentiment of topics on the horizontal axis increases from left to right.

**Practice** According to Figures 6 and 7, a developer with interest in concurrency may make a tradeoff between the popularities, difficulties, and sentiments of concurrency topics to decide to start their learning from the more popular,

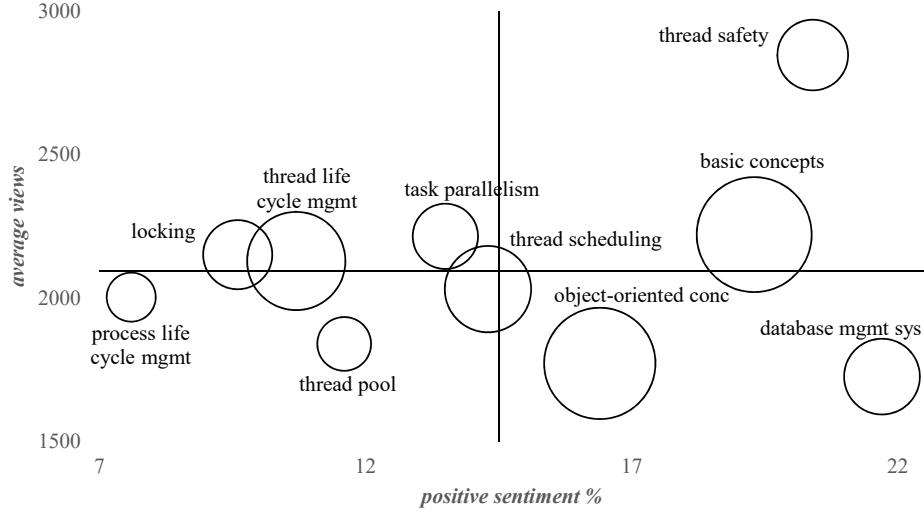


Figure 7: Trading off concurrency topics based on their popularity and sentiment.

less difficult, and more positive *task parallelism* topic before the less popular, more difficult, and less positive *thread pool* topic. Similarly, a manager of a concurrency development team may decide to assign a task related to the less difficult and more positive *object-oriented concurrency* topic to a less experienced developer and a task related to the more difficult and less positive *thread pool* topic to a more experienced developer [12].

**Research** According to Figures 6 and 7, an experienced concurrency researcher may make a tradeoffs between difficulties, popularities, and sentiments of concurrency topics to decide to focus their research on the more difficult, less popular, and less positive *process life cycle management* rather than the less difficult, more popular, and more positive *task parallelism* in the hope of making more contributions in a less crowded area. Conversely, a researcher looking to transition anew to the concurrency research may decide to start their research from the more popular, less difficult, and more positive *thread safety* topic.

**Education** According to Figures 6 and 7, a concurrency educator may decide to teach the material related to the less difficult, more popular, and more positive *thread safety* topic before the more difficult, less popular, and less positive *thread pool*, accounting for the dependencies between these topics. The educator may also decide to prepare more material and spend more time in both the class and lab on the more difficult and less positive *thread scheduling* topic than the less difficult and more positive *thread safety*. This decision is further supported by the observation that there are more questions about *thread scheduling* than *thread safety*. In addition, according to Figure 5, the educator teaching about data race tools may decide to spend more time on *root cause identification* topic with more number of questions than *configuration and execution* topic for these data race tools with less questions.

**Finding 21:** Concurrency practitioners, researchers, and educators can use our findings to make decisions about where to focus their efforts.

## 5. Threats to Validity

Our use of concurrency tags to identify concurrency questions and answers in Stack Overflow is a threat to the validity of our study. This is because concurrency tags may not be able to identify the complete set of questions and answers related to concurrency. To minimize this threat, we use well-known techniques used by previous work [11, 12] in developing our concurrency tags and perform solid experiments with a broad range of tag relevance and significance thresholds for  $\mu$  and  $\nu$ . Similarly, our use of race tool keywords to identify race tool questions is a threat. To minimize this threat, we include industrial race tools that are in use by companies like Google [33] and academic race tools and techniques surveyed by previous work [37].

Our use of Stack Overflow as the sole dataset is another threat. This is because Stack Overflow questions and answers may not be representative of interests, difficulties, sentiments, and race tool usages of concurrency developers. However, the large number of participant developers and questions and answers in Stack Overflow along with its wide-spread popularity among developers could mitigate this risk. Also, unlike some previous work that only use the title of Stack Overflow questions in their study [11], we use both the title and the body of the questions as well as their accepted answers to mitigate this risk.

Our metrics used to measure the popularity, difficulty, and sentiment of topics is another threat. This is because other factors, such as the technology, time and day of a question, length of the title, body, and code snippets [29], of the question can affect its popularity [9], difficulty [29], and sentiment [27, 26]. In fact, Wang et al. [29] study 46 factors that can affect the time to answer for a Stack Overflow question. To minimize this threat, we use well-known metrics and tools used by previous work to measure popularity [11, 12, 8], difficulty [11, 12, 9], and sentiment [52, 53, 58, 59].

Our manual labeling of topics is another threat. To minimize this threat, we use a well-known approach used by previous work [8] to label topics using their top 10 words and 15 random questions. Determining an optimal value for  $K$  when modeling topics is another threat. To minimize this threat, we use a well-known approach used by previous work [7, 8, 11] to find a reasonable value for  $K$  using experiments with a broad range of values for  $K$ . It is well-known that determining an optimal value for  $K$  is difficult [7].

Parsing Stack Overflow dataset, labeling topics from textual contents of questions and answers, and reduction of words to their bases is another threat. To minimize this threat, we use well-known tools used by previous work. We parse Stack Overflow posts using Python elementTree XML API [11], model topics using MALLET [7, 8, 11], and reduce words using the Porter stemming algorithm [8].

## 6. Related Work

Previous works that are closer to our work study software knowledge repositories, such as Stack Overflow, to understand interest [60, 61, 62, 9, 10, 7, 11], difficulty [9, 11, 8, 12], and sentiment [5, 63, 26, 64, 28, 3, 27, 4] of developers.

### 6.1. Concurrency

Closest to our work is the work of Pinto et al. [18] that uses the 250 most popular concurrency questions on Stack Overflow to study difficulties that developers face when writing concurrent programs. They categorize these difficulties into a set of themes including theoretical and practical concepts, threading, and first steps themes.

In other previous work, Pinto et al. [65] analyze the code for 2,227 projects to understand the usage of Java’s concurrent programming constructs and libraries and the evolution of their usage. Lin and Dig [66] study a corpus of 611 widely used Android apps to understand how developers use Android constructs for asynchronous concurrency. Blom et al. [67] study the usage of *java.util.concurrent* library in Qualitas corpus. Godefroid and Nagappan [68] survey 684 developers to study the spread and popularity of concurrency platforms and models at Microsoft.

In contrast, in this work, we develop a set of concurrency tags to extract concurrency questions and answers from Stack Overflow; group these posts into concurrency topics, categories, and a topic hierarchy; and analyze popularity, difficulty, and sentiment of these concurrency topics and their correlations. We also develop a set of race tool keywords to extract concurrency questions about data race tools and techniques; and group these questions into race tool topics. Finally, we discuss the implications of our findings for the practice, research, and education of concurrent software development, investigate the relation of our findings with the findings of previous work and present a sample set of questions for each of our concurrency topics and categories.

### 6.2. Non-concurrency

Rosen and Shihab [11] use LDA to model mobile development topics on Stack Overflow. They study popularity and difficulty of their mobile topics and categorize developers’ questions based on platforms for mobile development and type of questions that developers ask (why, what, and how). Yang et al. [12] use LDA tuned with a genetic algorithm to model security topics on Stack Overflow and manually organize their topics into 5 categories. They study popularity and difficulty of their security topics. Bajaj et al. [8] use LDA to model client-side web development topics using Stack Overflow and study interest of developers in these topics and challenges they face when working with these topics. Barua et al. [7] use LDA to model general topics on Stack Overflow. They study relations of questions and answers of these topics and evolution of developers’ interest in these topics both in general and for specific technologies.

Gyöngyi et al. [60] and Adamic et al. [61] study Yahoo!Answers posts to determine developers’ interests in a set of predefined categories. Hindle et

al. [62] use LDA to model topics related to development tasks from commit messages of a standalone software project and study evolution of developers' interest in these topics. Treude et al. [9] and Allamanis and Sutton [10] study Stack Overflow questions and answers to infer types of questions that developers ask and determine their difficulties with these question types. Bajracharya and Lopes [69] study logs of Koders, a code search engine, to learn about general topic of interest to developers in code search.

Guzman et al. [26] study the sentiment of commit logs and its relation with programming languages, team distribution, and approvals in 90 GitHub projects. They find that commit logs for Java projects tend to be more negative, projects with more distributed teams tend to have more positive commit logs, and Monday commit logs are more negative. Similarly, Sinha et al. [27] study the sentiment of commit logs in 28K GitHub projects and find that a majority of commit logs are neutral, Tuesday commit logs are more negative, and a strong correlation exists between sentiment of commit logs and the average number of files changed in the commits.

Ortu et al. [3] study the sentiment, emotion, and politeness of issue comments and their relations with the time to fix the issues in 506K issue comments of 14 projects in the Apache issue tracking system (JIRA). They find that issues with positive sentiment tend to take less time to be fixed and there is a weak relation between sentiment, emotion, and politeness. Souza and Silva [4] study the sentiment of commit logs in 1,262 GitHub projects and its relation with builds performed by Travis CI continuous integration server. They find that commit logs with negative sentiment are more likely to result in broken builds. Garcia et al. [5] study the relation between emotion and activity of contributors in the open source project GENTOO using its bug tracking issues and mail archives. They find that contributors with strong positive or negative emotions are more likely to leave the project. Bazelli et al. [63] study the relation between the sentiment of developers' answers on Stack Overflow and their reputation and find that top reputed authors express less negative emotions in their answers. Muriga et al. [64] study whether software artifacts such as issue reports carry emotional information and find that issue reports include emotional information about design choices, maintenance activity, and colleagues.

In contrast, in this work, we develop a set of concurrency tags to extract concurrency questions and answers from Stack Overflow; group these posts into concurrency topics, categories, and a topic hierarchy; and analyze popularity, difficulty, and sentiment of these concurrency topics and their correlations. We also develop a set of race tool keywords to extract concurrency questions about data race tools and techniques; and group these questions into race tool topics. Finally, we discuss the implications of our findings for the practice, research, and education of concurrent software development, investigate the relation of our findings with the findings of previous work and present a sample set of questions for each of our concurrency topics and categories.



## 7. Conclusions

In this work, we conduct a large-scale study on the entirety of Stack Overflow to understand interests, difficulties, sentiments, and tool usages of concurrency developers. To conduct this study, we develop a set of concurrency tags to extract concurrency questions and answers from Stack Overflow; group these questions and answers into concurrency topics, categories, and a topic hierarchy; analyze popularities, difficulties, and sentiments of these concurrency topics and their correlations; develop a set of tool keywords to extract concurrency questions about data race tools and group these questions into race tool topics; and finally discuss the implications of our findings for the practice, research, and education of concurrent software development, investigate the relation of our findings with the findings of the previous work, and present a set of examples questions that developers ask for each of our concurrency and tool topics.

## Acknowledgements

Bagherzadeh, Ahmed, and Sripathi were supported in part by the Department of Computer Science and Engineering and the Research Office at Oakland University. Khatchadourian was supported by PSC-CUNY Awards #617930049 and #638010051, jointly funded by The Professional Staff Congress and The City University of New York. We would also like to thank Nicholas Fireman for his help.

## References

- [1] S. Lu, S. Park, E. Seo, Y. Zhou, Learning from mistakes: A comprehensive study on real world concurrency bug characteristics, in: Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS XIII, ACM, New York, NY, USA, 2008, pp. 329–339. doi:10.1145/1346281.1346323.  
URL <http://doi.acm.org/10.1145/1346281.1346323>
- [2] R. Khatchadourian, Y. Tang, M. Bagherzadeh, S. Ahmed, Safe automated refactoring for intelligent parallelization of Java 8 streams, in: International Conference on Software Engineering, ICSE '19, ACM/IEEE, IEEE Press, Piscataway, NJ, USA, 2019, pp. 619–630. doi:10.1109/ICSE.2019.00072.  
URL [http://academicworks.cuny.edu/hc\\_pubs/489](http://academicworks.cuny.edu/hc_pubs/489)
- [3] M. Ortu, B. Adams, G. Destefanis, P. Tourani, M. Marchesi, R. Tonelli, Are bullies more productive? empirical study of affectiveness vs. issue fixing time, in: 2015 IEEE/ACM 12th Working Conference on Mining Software Repositories, 2015, pp. 303–313. doi:10.1109/MSR.2015.35.
- [4] R. Souza, B. Silva, Sentiment analysis of Travis CI builds, in: Proceedings of the 14th International Conference on Mining Software Repositories, MSR '17, IEEE Press, Piscataway, NJ, USA, 2017, pp. 459–462. doi:10.1109/MSR.2017.27.  
URL <https://doi.org/10.1109/MSR.2017.27>
- [5] D. Garcia, M. S. Zanetti, F. Schweitzer, The role of emotions in contributors activity: A case study on the GENTOO community, in: Proceedings of the 2013 International Conference on Cloud and Green Computing, CGC '13, IEEE Computer Society, Washington, DC, USA, 2013, pp. 410–417. doi:10.1109/CGC.2013.71.  
URL <https://doi.org/10.1109/CGC.2013.71>
- [6] M. De Choudhury, S. Counts, Understanding affect in the workplace via social media, in: Proceedings of the 2013 Conference on Computer Supported Cooperative Work, CSCW '13, ACM, New York, NY, USA, 2013, pp. 303–316. doi:10.1145/2441776.2441812.  
URL <http://doi.acm.org/10.1145/2441776.2441812>
- [7] A. Barua, S. W. Thomas, A. E. Hassan, What are developers talking about? an analysis of topics and trends in Stack Overflow, Empirical Softw. Engg. 19 (3) (2014) 619–654. doi:10.1007/s10664-012-9231-y.  
URL <http://dx.doi.org/10.1007/s10664-012-9231-y>
- [8] K. Bajaj, K. Pattabiraman, A. Mesbah, Mining questions asked by web developers, in: Proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014, ACM, New York, NY, USA, 2014, pp. 112–121. doi:10.1145/2597073.2597083.  
URL <http://doi.acm.org/10.1145/2597073.2597083>

- [9] C. Treude, O. Barzilay, M.-A. Storey, How do programmers ask and answer questions on the web? (NIER track), in: Proceedings of the 33rd International Conference on Software Engineering, ICSE '11, ACM, New York, NY, USA, 2011, pp. 804–807. doi:10.1145/1985793.1985907. URL <http://doi.acm.org/10.1145/1985793.1985907>
- [10] M. Allamanis, C. Sutton, Why, when, and what: Analyzing Stack Overflow questions by topic, type, and code, in: Proceedings of the 10th Working Conference on Mining Software Repositories, MSR '13, IEEE Press, Piscataway, NJ, USA, 2013, pp. 53–56. URL <http://dl.acm.org/citation.cfm?id=2487085.2487098>
- [11] C. Rosen, E. Shihab, What are mobile developers asking about? a large scale study using Stack Overflow, Empirical Softw. Engg. 21 (3) (2016) 1192–1223. doi:10.1007/s10664-015-9379-3. URL <http://dx.doi.org/10.1007/s10664-015-9379-3>
- [12] X.-L. Yang, D. Lo, X. Xia, Z.-Y. Wan, J.-L. Sun, What security questions do developers ask? a large-scale study of Stack Overflow posts, Journal of Computer Science and Technology 31 (5) (2016) 910–924. doi:10.1007/s11390-016-1672-0. URL <https://doi.org/10.1007/s11390-016-1672-0>
- [13] S. Neuhaus, T. Zimmermann, Security trend analysis with CVE topic models, in: 2010 IEEE 21st International Symposium on Software Reliability Engineering, 2010, pp. 111–120. doi:10.1109/ISSRE.2010.53.
- [14] G. Pinto, F. Castor, Y. D. Liu, Mining questions about software energy consumption, in: Proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014, ACM, New York, NY, USA, 2014, pp. 22–31. doi:10.1145/2597073.2597110. URL <http://doi.acm.org/10.1145/2597073.2597110>
- [15] Stack Overflow, <http://www.stackoverflow.com/> (June 2017).
- [16] F. N. S. A. Bagherzadeh, Mehdi, R. Khatchadourian, Actor concurrency bugs: A comprehensive study on symptoms, root causes, api usages, and differences, in: Proceedings of 2020 Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2020, 2020.
- [17] M. Bagherzadeh, R. Khatchadourian, Going big: A large-scale study on what big data developers ask, in: Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2019, ACM, New York, NY, USA, 2019, pp. 432–442. doi:10.1145/3338906.3338939. URL <http://doi.acm.org/10.1145/3338906.3338939>
- [18] G. Pinto, W. Torres, F. Castor, A study on the most popular questions about concurrent programming, in: Proceedings of the 6th Workshop on Evaluation and Usability of Programming Languages and Tools,

- PLATEAU 2015, ACM, New York, NY, USA, 2015, pp. 39–46. doi: 10.1145/2846680.2846687.  
URL <http://doi.acm.org/10.1145/2846680.2846687>
- [19] S. A. Asadollah, D. Sundmark, S. Eldh, H. Hansson, Concurrency bugs in open source software: a case study, *Journal of Internet Services and Applications* 8 (1) (2017) 1–15.
  - [20] N. G. Leveson, C. S. Turner, An investigation of the Therac-25 accidents, *Computer* 26 (7) (1993) 18–41. doi:10.1109/MC.1993.274940.  
URL <https://doi.org/10.1109/MC.1993.274940>
  - [21] U.-C. P. S. O. T. Force, Final report on the August 14th blackout in the United States and Canada, Tech. rep., U.S. Department of Energy (2004).
  - [22] E. Pozniansky, A. Schuster, Multirace: efficient on-the-fly data race detection in multithreaded C++ programs, *Concurrency and Computation: Practice and Experience* 19 (3) (2007) 327–340. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.1064>, doi:10.1002/cpe.1064.  
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.1064>
  - [23] D. M. Blei, A. Y. Ng, M. I. Jordan, Latent Dirichlet allocation, *J. Mach. Learn. Res.* 3 (2003) 993–1022.  
URL <http://dl.acm.org/citation.cfm?id=944919.944937>
  - [24] S. Fincher, J. Tenenbergs, Making sense of card sorting data, *Expert Systems* 22 (3) (2005) 89–93. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0394.2005.00299.x>, doi:10.1111/j.1468-0394.2005.00299.x.  
URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0394.2005.00299.x>
  - [25] S. Nadi, S. Krüger, M. Mezini, E. Bodden, "Jumping through hoops": Why do Java developers struggle with cryptography APIs?, in: 2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE), 2016, pp. 935–946. doi:10.1145/2884781.2884790.
  - [26] E. Guzman, D. Azócar, Y. Li, Sentiment analysis of commit comments in GitHub: An empirical study, in: Proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014, ACM, New York, NY, USA, 2014, pp. 352–355. doi:10.1145/2597073.2597118.  
URL <http://doi.acm.org/10.1145/2597073.2597118>
  - [27] V. Sinha, A. Lazar, B. Sharif, Analyzing developer sentiment in commit logs, in: Proceedings of the 13th International Conference on Mining Software Repositories, MSR '16, ACM, New York, NY, USA, 2016, pp. 520–523. doi:10.1145/2901739.2903501.  
URL <http://doi.acm.org/10.1145/2901739.2903501>

- [28] P. Tourani, Y. Jiang, B. Adams, Monitoring sentiment in open source mailing lists: Exploratory study on the Apache ecosystem, in: Proceedings of 24th Annual International Conference on Computer Science and Software Engineering, CASCON '14, IBM Corp., Riverton, NJ, USA, 2014, pp. 34–44.  
URL <http://dl.acm.org/citation.cfm?id=2735522.2735528>
- [29] S. Wang, T.-H. Chen, A. E. Hassan, Understanding the factors for fast answers in technical Q&A websites: An empirical study of four Stack Exchange websites, in: Proceedings of the 40th International Conference on Software Engineering, ICSE '18, ACM, New York, NY, USA, 2018, pp. 884–884. doi:10.1145/3180155.3182521.  
URL <http://doi.acm.org/10.1145/3180155.3182521>
- [30] S. Ahmed, M. Bagherzadeh, What do concurrency developers ask about?: A large-scale study using Stack Overflow, in: Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM '18, ACM, New York, NY, USA, 2018, pp. 30:1–30:10. doi:10.1145/3239235.3239524.  
URL <http://doi.acm.org/10.1145/3239235.3239524>
- [31] Stack Exchange Dump, <https://archive.org/details/stackexchange> (June 2017).
- [32] Stack Overflow questions with race-condition and data-race tags, <https://data.stackexchange.com/stackoverflow/query/edit/1089810>.
- [33] C. Sadowski, J. Yi, How developers use data race detection tools, in: Proceedings of the 5th Workshop on Evaluation and Usability of Programming Languages and Tools, PLATEAU '14, ACM, New York, NY, USA, 2014, pp. 43–51. doi:10.1145/2688204.2688205.  
URL <http://doi.acm.org/10.1145/2688204.2688205>
- [34] Intel Corporation, <https://software.intel.com/en-us/inspector>.
- [35] Valgrind Developers, <http://valgrind.org/docs/manual/hg-manual.html>.
- [36] Valgrind Developers, <http://valgrind.org/docs/manual/drd-manual.html>.
- [37] S. Hong, M. Kim, A survey of race bug detection techniques for multi-threaded programmes, *Softw. Test. Verif. Reliab.* 25 (3) (2015) 191–217. doi:10.1002/stvr.1564.  
URL <http://dx.doi.org/10.1002/stvr.1564>
- [38] M. F. Porter, Readings in information retrieval, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1997, Ch. An Algorithm for Suffix Stripping, pp. 313–316.  
URL <http://dl.acm.org/citation.cfm?id=275537.275705>

- [39] M. V. Mäntylä, N. Novielli, F. Lanubile, M. Claes, M. Kuuttila, Bootstrapping a lexicon for emotional arousal in software engineering, in: Proceedings of the 14th International Conference on Mining Software Repositories, MSR '17, IEEE Press, Piscataway, NJ, USA, 2017, pp. 198–202. doi:10.1109/MSR.2017.47. URL <https://doi.org/10.1109/MSR.2017.47>
- [40] N. Meng, S. Nagy, D. D. Yao, W. Zhuang, G. A. Argoty, Secure coding practices in Java: Challenges and vulnerabilities, in: Proceedings of the 40th International Conference on Software Engineering, ICSE '18, Association for Computing Machinery, New York, NY, USA, 2018, pp. 372–383. doi:10.1145/3180155.3180201. URL <https://doi.org/10.1145/3180155.3180201>
- [41] A. Abdellatif, D. Costa, K. Badran, R. Abdalkareem, E. Shihab, Challenges in chatbot development: A study of stack overflow posts, in: Proceedings of the 17th International Conference on Mining Software Repositories, MSR '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 174–185. doi:10.1145/3379597.3387472. URL <https://doi.org/10.1145/3379597.3387472>
- [42] M. J. Islam, G. Nguyen, R. Pan, H. Rajan, A comprehensive study on deep learning bug characteristics, in: Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2019, ACM, New York, NY, USA, 2019, pp. 510–520. doi:10.1145/3338906.3338955. URL <http://doi.acm.org/10.1145/3338906.3338955>
- [43] M. J. Islam, R. Pan, G. Nguyen, H. Rajan, Repairing deep neural networks fix patterns and challenges, in: Proceedings of the 42nd International Conference on Software Engineering - Volume 2, ICSE '20, IEEE Press, 2020.
- [44] A. K. McCallum, MALLET: A machine learning for language toolkit, <http://mallet.cs.umass.edu> (2002).
- [45] L. R. Biggers, C. Bocovich, R. Capshaw, B. P. Eddy, L. H. Etzkorn, N. A. Kraft, Configuring latent Dirichlet allocation based feature location, Empirical Softw. Engg. 19 (3) (2014) 465–500. doi:10.1007/s10664-012-9224-x. URL <http://dx.doi.org/10.1007/s10664-012-9224-x>
- [46] R. Khatchadourian, Y. Tang, M. Bagherzadeh, B. Ray, An empirical study on the use and misuse of java 8 streams, in: H. Wehrheim, J. Cabot (Eds.), Fundamental Approaches to Software Engineering, Springer International Publishing, Cham, 2020, pp. 97–118.
- [47] R. Khatchadourian, Y. Tang, M. Bagherzadeh, S. Ahmed, A tool for optimizing Java 8 stream software via automated refactoring, in: International Working Conference on Source Code Analysis and Manipulation, SCAM

- '18, IEEE, IEEE Press, 2018, pp. 34–39, engineering Track. (9/17; 53% acceptance rate). **Distinguished Paper Award**. doi:10.1109/SCAM.2018.00011.  
URL [http://academicworks.cuny.edu/hc\\_pubs/429](http://academicworks.cuny.edu/hc_pubs/429)
- [48] M. Bagherzadeh, H. Rajan, Order types: Static reasoning about message races in asynchronous message passing concurrency, in: Proceedings of the 7th ACM SIGPLAN International Workshop on Programming Based on Actors, Agents, and Decentralized Control, AGERE 2017, ACM, New York, NY, USA, 2017, pp. 21–30. doi:10.1145/3141834.3141837.  
URL <http://doi.acm.org/10.1145/3141834.3141837>
  - [49] Y. Long, M. Bagherzadeh, E. Lin, G. Upadhyaya, H. Rajan, On ordering problems in message passing software, in: Proceedings of the 15th International Conference on Modularity, MODULARITY 2016, ACM, New York, NY, USA, 2016, pp. 54–65. doi:10.1145/2889443.2889444.  
URL <http://doi.acm.org/10.1145/2889443.2889444>
  - [50] M. Bagherzadeh, H. Rajan, Panini: A concurrent programming model for solving pervasive and oblivious interference, in: Proceedings of the 14th International Conference on Modularity, MODULARITY 2015, ACM, New York, NY, USA, 2015, pp. 93–108. doi:10.1145/2724525.2724568.  
URL <http://doi.acm.org/10.1145/2724525.2724568>
  - [51] L. Mamykina, B. Manoim, M. Mittal, G. Hripcsak, B. Hartmann, Design lessons from the fastest Q&A site in the west, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, ACM, New York, NY, USA, 2011, pp. 2857–2866. doi:10.1145/1978942.1979366.  
URL <http://doi.acm.org/10.1145/1978942.1979366>
  - [52] F. Calefato, F. Lanubile, F. Maiorano, N. Novielli, Sentiment polarity detection for software development, Empirical Softw. Engg. 23 (3) (2018) 1352–1382. doi:10.1007/s10664-017-9546-9.  
URL <https://doi.org/10.1007/s10664-017-9546-9>
  - [53] N. Novielli, D. Girardi, F. Lanubile, A benchmark study on sentiment analysis for software engineering research, in: Proceedings of the 15th International Conference on Mining Software Repositories, MSR '18, ACM, New York, NY, USA, 2018, pp. 364–375. doi:10.1145/3196398.3196403.  
URL <http://doi.acm.org/10.1145/3196398.3196403>
  - [54] M. Thelwall, K. Buckley, G. Paltoglou, Sentiment strength detection for the social web, J. Am. Soc. Inf. Sci. Technol. 63 (1) (2012) 163–173. doi:10.1002/asi.21662.  
URL <http://dx.doi.org/10.1002/asi.21662>
  - [55] T. Ahmed, A. Bosu, A. Iqbal, S. Rahimi, SentiCR: A customized sentiment analysis tool for code review interactions, in: Proceedings of the 32Nd

IEEE/ACM International Conference on Automated Software Engineering, ASE 2017, IEEE Press, Piscataway, NJ, USA, 2017, pp. 106–111.  
URL <http://dl.acm.org/citation.cfm?id=3155562.3155579>

- [56] M. R. Islam, M. F. Zibran, Leveraging automated sentiment analysis in software engineering, in: Proceedings of the 14th International Conference on Mining Software Repositories, MSR '17, IEEE Press, Piscataway, NJ, USA, 2017, pp. 203–214. doi:10.1109/MSR.2017.9.  
URL <https://doi.org/10.1109/MSR.2017.9>
- [57] E. A. Lee, The problem with threads, Computer 39 (5) (2006) 33–42. doi:10.1109/MC.2006.180.  
URL <https://doi.org/10.1109/MC.2006.180>
- [58] N. Novielli, F. Calefato, F. Lanubile, A gold standard for emotion annotation in Stack Overflow, in: Proceedings of the 15th International Conference on Mining Software Repositories, MSR '18, ACM, New York, NY, USA, 2018, pp. 14–17. doi:10.1145/3196398.3196453.  
URL <http://doi.acm.org/10.1145/3196398.3196453>
- [59] R. Jongeling, P. Sarkar, S. Datta, A. Serebrenik, On negative results when using sentiment analysis tools for software engineering research, Empirical Softw. Engg. 22 (5) (2017) 2543–2584. doi:10.1007/s10664-016-9493-x.  
URL <https://doi.org/10.1007/s10664-016-9493-x>
- [60] Z. Gyöngyi, G. Koutrika, J. Pedersen, H. Garcia-Molina, Questioning Yahoo! Answers, in: workshop on question answering on the Web, 2008.
- [61] L. A. Adamic, J. Zhang, E. Bakshy, M. S. Ackerman, Knowledge sharing and Yahoo answers: Everyone knows something, in: Proceedings of the 17th International Conference on World Wide Web, WWW '08, ACM, New York, NY, USA, 2008, pp. 665–674. doi:10.1145/1367497.1367587.  
URL <http://doi.acm.org/10.1145/1367497.1367587>
- [62] A. Hindle, M. W. Godfrey, R. C. Holt, What's hot and what's not: Windowed developer topic analysis, in: 2009 IEEE International Conference on Software Maintenance, 2009, pp. 339–348. doi:10.1109/ICSM.2009.5306310.
- [63] B. Bazelli, A. Hindle, E. Stroulia, On the personality traits of StackOverflow users, in: Proceedings of the 2013 IEEE International Conference on Software Maintenance, ICSM '13, IEEE Computer Society, Washington, DC, USA, 2013, pp. 460–463. doi:10.1109/ICSM.2013.72.  
URL <http://dx.doi.org/10.1109/ICSM.2013.72>
- [64] A. Murgia, P. Tourani, B. Adams, M. Ortu, Do developers feel emotions? an exploratory analysis of emotions in software artifacts, in: Proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014, ACM, New York, NY, USA, 2014, pp. 262–271. doi:10.1145/2597073.



2597086.

URL <http://doi.acm.org/10.1145/2597073.2597086>

- [65] G. Pinto, W. Torres, B. Fernandes, F. Castor, R. S. Barros, A large-scale study on the usage of Java's concurrent programming constructs, *J. Syst. Softw.* 106 (C) (2015) 59–81. doi:10.1016/j.jss.2015.04.064.  
URL <http://dx.doi.org/10.1016/j.jss.2015.04.064>
- [66] Y. Lin, S. Okur, D. Dig, Study and refactoring of Android asynchronous programming, in: *Proceedings of the 30th IEEE/ACM International Conference on Automated Software Engineering, ASE '15*, IEEE Press, Piscataway, NJ, USA, 2015, pp. 224–235. doi:10.1109/ASE.2015.50.  
URL <https://doi.org/10.1109/ASE.2015.50>
- [67] S. Blom, J. Kiniry, M. Huisman, How do developers use APIs? a case study in concurrency, in: *Proceedings of the 2013 18th International Conference on Engineering of Complex Computer Systems, ICECCS '13*, IEEE Computer Society, Washington, DC, USA, 2013, pp. 212–221. doi:10.1109/ICECCS.2013.39.  
URL <http://dx.doi.org/10.1109/ICECCS.2013.39>
- [68] P. Godefroid, N. Nagappan, Concurrency at Microsoft: An exploratory survey, in: *CAV Workshop on Exploiting Concurrency Efficiently and Correctly*, 2008.
- [69] S. K. Bajracharya, C. V. Lopes, Analyzing and mining a code search engine usage log, *Empirical Softw. Engg.* 17 (4-5) (2012) 424–466. doi:10.1007/s10664-010-9144-6.  
URL <http://dx.doi.org/10.1007/s10664-010-9144-6>