

# A Qualitative Study of Application-level Caching

Jhonny Mertz and Ingrid Nunes

**Abstract**—Latency and cost of Internet-based services are encouraging the use of application-level caching to continue satisfying users' demands, and improve the scalability and availability of origin servers. Despite its popularity, this level of caching involves the manual implementation by developers and is typically addressed in an *ad-hoc* way, given that it depends on specific details of the application. As a result, application-level caching is a time-consuming and error-prone task, becoming a common source of bugs. Furthermore, it forces application developers to reason about a crosscutting concern, which is unrelated to the application business logic. In this paper, we present the results of a qualitative study of how developers handle caching logic in their web applications, which involved the investigation of ten software projects with different characteristics. The study we designed is based on comparative and interactive principles of grounded theory, and the analysis of our data allowed us to extract and understand how developers address cache-related concerns to improve performance and scalability of their web applications. Based on our analysis, we derived guidelines and patterns, which guide developers while designing, implementing and maintaining application-level caching, thus supporting developers in this challenging task that is crucial for enterprise web applications.

**Index Terms**—application-level caching, qualitative study, pattern, guideline.



## 1 INTRODUCTION

DYNAMICALLY generated web content represents a significant portion of web requests. This content changes according to a user request and due to this, the rate at which dynamic documents are delivered is often orders of magnitudes slower than static documents [1], [2]. Therefore, dynamic content generation and delivery place a significant burden on servers, often leading to performance bottlenecks [1]. To continue satisfying users' demands and also reducing the workload on content providers, over the past years many optimization techniques have been proposed [2]. The ubiquitous solution to this problem is some form of *caching*, which can reduce the user perceived latency and Internet traffic, and improve the scalability and availability of origin servers.

Recently, latency and cost of Internet-based services are driving the proliferation of *application-level caching*, which is placed on the server-side and typically uses a key-value in-memory store system to cache frequently accessed or expensive to compute data that remain not cached in other caching levels, lowering the load on database or services that are difficult to scale up [3], [4], [5]. Therefore, application-level caching has become a popular technique for enabling highly scalable web applications.

Despite its popularity, the implementation of application-level caching is not trivial and demands high effort, because it involves manual implementation by developers. Its design and maintenance involve four key challenging issues: determining *how* to cache the selected data, *what* data should be cached, *when* the selected data should be cached or evicted, and *where* the cached data

should be placed and maintained. Although application-level caching is commonly being adopted, it is typically implemented in an *ad hoc* way, and there are no available practical guidelines for developers to appropriately design, implement and maintain caching in their applications. To find caching best practices, developers can make use of conventional wisdom, consult development blogs, or simply search online for tips. Nevertheless, this empirical knowledge is unsupported by concrete data or theoretical foundations that demonstrate its effectiveness in practice. Despite there are existing tool-supported approaches that can help developers to implement caching with minimal impact on the application code [4], [5], [6], they do not consider all the aforementioned issues and do not take application specificities into account, letting most part of the reasoning, as well as the integration with the tool, to developers.

Given this context, we in this paper present a qualitative study performed to identify patterns and guidelines to support developers while designing, implementing and maintaining application-level caching in their applications. Previous qualitative studies have been conducted to gain a greater understanding of the behavior of developers during specific software development tasks [7], [8], [9], [10], [11]. Moreover, there is work that focused on proposing software development guidelines to help software developers while developing applications [12], [13], [14]. Our work is similar in nature, but focuses on the investigation of a development challenge that has not been previously addressed.

Besides helping developers while handling cache, the novel findings of this study provide insights to propose solutions to raise the abstraction level of caching or automate caching tasks, providing a better experience with caching for developers. The results of this study are thus informative and provide practical guidance for developers. It is important to note, however, that qualitative research is aimed at gaining a deep understanding of a particular target of study, rather than providing a general description

- J. Mertz is with the Departamento de Informática Aplicada, Instituto de Informática, Universidade Federal do Rio Grande do Sul, Brazil.  
E-mail: jnamertz@inf.ufrgs.br
- I. Nunes is with the Departamento de Informática Aplicada, Instituto de Informática, Universidade Federal do Rio Grande do Sul, Brazil and TU Dortmund, Germany.  
E-mail: ingridnunes@inf.ufrgs.br

Manuscript received April 19, 2005; revised August 26, 2015.

of a large sample of a population.

In short, this paper provides three main contributions: (i) the design and results of a qualitative study that provides an in-depth analysis of how developers employ application-level caching in web-based applications; (ii) useful guidelines for modeling an application-level caching component; and (iii) patterns for caching design, implementation and maintenance to be adopted under different circumstances.

The remainder of this paper is organized as follows. In Section 2 we give background and related work on application-level caching. We detail the performed study in Section 3, then present and discuss its results in Section 4. We describe the proposed guidelines and patterns derived from this study in Section 5. Finally, we discuss threats to the validity of our study in Section 6 and, in Section 7, we conclude.

## 2 BACKGROUND AND RELATED WORK

Focused on server-side, caching of dynamic web content can be implemented at several locations across a web-based system architecture [15]. Depending on where caching is implemented, different types of content can be stored, such as the final HTML page [16], intermediate HTML or XML fragments [17], [18], [19], database queries [20], [21], [22], [23], or database tables [24]. These alternatives are conceived out of application boundaries and can cache dynamic data automatically, which provide transparent caching components for developers. However, they do not take application specificities into account, providing good results in general but are less effective where complex logic and personalized web content are processed within the application [25]. Therefore, application-level caching is an appealing technique to improve performance, reduce workload and make the overall user experience more pleasurable by reducing communication delays of web-based systems.

Figure 1 illustrates an example where an application-level cache is used to lower the application workload. First, the application receives an HTTP request (Step 1) from the web infrastructure. This HTTP request is eventually treated by an application component called M1, which in turn depends on M2. However, M2 can be an expensive operation (i.e. request the database, call a service or process a large amount of data). Therefore, M1 implements a caching layer, which verifies whether the necessary processing is already in the cache before calling M2 (Step 2). Then, the cache component performs a look up for the requested data and returns either the cached result or a not found error (Step 3). If data is found, it means a cache *hit* and M1 can continue its execution without calling M2. If, however, a not found error is returned, it means a cache *miss* happened, then M2 needs to be invoked (Steps 4 and 5). The newly fetched result of M2 can be cached to serve future requests faster (Step 6) and eventually a response is sent to the user (Step 7). Furthermore, other caching layers can be deployed outside the application, at the web infrastructure.

The steps described in Figure 1 are those typically performed to implement cache at any level, and the key difference is the location where it is applied and its explicit manipulation by developers. Application-level cache allows

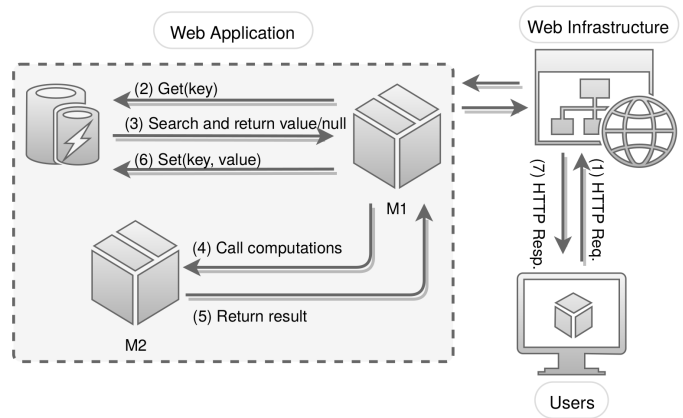


Fig. 1. Using Application-level Caching to Decrease the Application Workload.

caching at a granularity best suited to the application, providing a way to cache entire HTML pages, page fragments, database queries or even computed results; rather than laying in front of the application servers (e.g., a proxy-level cache) or between the application and the database (e.g., a query or database cache). For example, many websites have highly-personalized content, thus rendering whole-page web caches is mostly useless; application-level caches can be used to separate shared content from customized content, and then the shared content can be cached and reused among users [4]. Memcached [26] and Redis [27] are popular solutions and are a critical web infrastructure component for some big players such as Amazon, Facebook, LinkedIn, Twitter, Wikipedia, and YouTube [28].

As already discussed in the introduction, the development of application-level caching involves four key issues: determining *how*, *what*, *when*, and *where* to cache data. The main problem is that solutions for all of these issues usually depend on application-specific details, and are manually designed and implemented by developers, as shown in the example presented in Listing 1. In this example, we assume that an e-commerce application where a class named `ProductsRepository` is responsible for loading products from database. To reduce database workload, caching logic is inserted in this class within the methods `getProducts`, `updateProduct`, and `deleteProduct`, which retrieves, updates and deletes content from the database, respectively.

The first issue is related to the *cache-aside* implementation and the fact that the cache system and the underlying source of data are not aware of each other (as illustrated in Figure 1). The cache system itself is a passive component, and developers must implement, directly in the application code, ways to assign names to cached values, perform lookups, and keep the cache up to date. The extra logic also requires additional testing and debugging time, which can be expensive in some scenarios. The implementation and maintenance of application-level caching are a challenge because developers must refactor all the data access logic to encapsulate cache data into the proper application-level object, and its direct management leads to a concern spread all over the system—mixed with business code—which leads to increasing complexity and maintenance time. Thus, this on-demand and manual caching process can be very

Listing 1  
Example of Application-level Caching Implementation.

```

public class ProductsRepository {
    public List<Product> getProducts() {
        Cache cache = Cache.getInstance();
        List<Product> products =
            (List<Product>) cache.get("products");
        if (products == null) {
            products = getProductsFromDB();
            cache.put("products", products);
        }
        return products;
    }
    public void updateProduct(Product product) {
        Cache.getInstance().delete("products");
        updateProductIntoDB(product);
    }
    public void deleteProduct(Product product) {
        Cache.getInstance().delete("products");
        deleteProductFromDB(product);
    }
}

```

time-consuming, error-prone and, consequently, a common source of bugs [25]. Gupta et al. [6] and Ports et al. [4] both address these implementation issues by providing high-level caching abstractions, in which developers can simply designate application functions as cacheable, and the proposed system automatically caches their results and invalidates the cached data when the underlying source changes.

The second and third issues refer to deciding the right content to cache and the best moment of caching or clearing the cache in order to avoid cache thrashing and stale content. These decisions involve (i) choosing the granularity of cache objects, (ii) translating between raw data and cache objects, and (iii) maintaining cache consistency, which are all tasks to be accomplished by developers and might not be trivial in complex applications. Della Toffola et al. [5], Nguyen and Xu [29] and Infante [30] address part of these issues by identifying and suggesting caching opportunities. However, developers should still review the suggestions and refactor the code, integrating cache logic into the application.

The last issue is related to the maintenance of the cache system, which involves several aspects such as determining replacement policies and the size of the cache. According to Radhakrishnan [31], a traditional approach to maintain cache is to set up strategies based on initial assumptions or well-accepted characteristics of workload and access patterns. Thus, cache statistics such as hit ratio, the number of objects cached, and average object size can be collected by observing the application and cache at run-time. These data, compared in the context of desired values of these parameters, can be used to decide whether to change the first choices. Usually, this process is repeated until an acceptable trade-off between configuration and performance is reached, and then cache strategies are fixed for the lifetime of the product or at least for several release cycles. However, its performance may decay over time due to changes in the workload characteristics and access patterns. Therefore, achieving acceptable application performance requires constantly tuning cache decisions, which implies extra time

dedicated to maintenance [31].

To supply these limitations, recent studies have shown that caching approaches that deal with the changing dynamics of the web environment can perform better than a statically configured cache [31], [32]. These intelligent approaches try to optimize the cache behavior by dynamically configuring strategies like admission [22], [33], [34], [35], [36], consistency [21] and replacement [37], [38].

Despite many methods related to intelligent cache approaches have been published [15], adapting cache properties and decisions at application-level has only been tackled by few examples. They provide adaptive solutions for cache update algorithms [18], [39], [40], [41], or mapping of requests before repartitioning the cache [42]. Some approaches [31], [43], [44] focus on the infrastructure aspect of application-level caching by exploring the size and fault-tolerance adaptations in cache servers.

Even though these approaches focus on providing an adaptive behavior to application-level cache, none of them takes application specificities into account to autonomously manage their target. They attempt only to optimize cache decisions based on cache statistics like hit-ratio and access patterns at a higher level, thus, ignoring cache meta-data expressed by application-specific characteristics, which are closely related to application computations and could help to reach an optimal performance of the caching service on time. Addressing this issues, content-aware approaches have been proposed [45], [46] for admitting and replacing items in the cache by exploring descriptive caching hits in the application model or spatial locality (relationships among web objects). Though adaptive caching is in general new and innovative, it is far from being adopted as standard practice in software industry; many improvements must be achieved before this can happen, such as reducing the overhead in terms of resource consumption and processing time of the learning process, and providing easy ways to integrate application and caching method.

Thereby, substantial advances have been made in application-level cache. These design, implementation and maintenance shortcomings call for new methods and techniques, which can support developers to reach a good trade-off between cache development and application performance, leveraging application-specific details. An essential step towards such approaches, taken in this work, is to *understand, extract, structure* and *document* the application-level caching knowledge implicit and spread in existing web applications. Moreover, our results serve as guidance for the development of application-level caching, given that there is neither off-the-shelf solutions or systematic ways of designing, implementing, and maintaining it.

### 3 STUDY DESIGN

We performed a qualitative study to understand how application-level caching is implemented in existing software systems that rely on this kind of caching. We provide details of our study in next sections, starting by discussing the study approach in Section 3.1. We present our goal and research questions in Section 3.2, and then describe the study procedure in Section 3.3. We introduce the target systems of our study in Section 3.4, to later proceed to the analysis and interpretation of our obtained results.

TABLE 1  
Goal Definition (GQM Template).

Definition Element	Our Study Goal
Motivation	To identify patterns of application-level caching adopted by developers and understand what kinds of caching implementations and decisions can be automatically inferred,
Purpose	characterize and evaluate
Object	application-level caching-related design and implementation
Perspective	from a perspective of the researcher
Domain: web-based applications	as they are implemented in the source code and described in issues of web-based applications
Scope	in the context of 10 software projects, obtained from open-source repositories and software companies.

### 3.1 Study Approach

We chose a qualitative method rather than quantitative in order to take a holistic and comprehensive understanding of caching practices adopted by developers, exploring their inner experiences in application development.

Our study was designed based on comparative and interactive principles of *grounded theory* [47]. The purpose of grounded theory is to construct theory grounded in data by identifying general concepts, develop theoretical explanations that reach beyond the known, and offer new insights into the area of study. The systematic procedures of grounded theory enable qualitative researchers to generate ideas. In turn, these ideas can be later studied and verified through traditional quantitative forms of research.

### 3.2 Goal and Research Questions

As stated in the introduction, our primary objective while performing this study is to provide guidance to developers when adopting application-level caching in their applications. This study aims to provide such guidance using an *application-centric approach*, i.e. by identifying caching solutions that developers usually apply in their applications, taking into account application details. The study was guided by the framework proposed by Basili et al. [48], [49]. The paradigm includes the goal-question-metric (GQM) template, which was adopted to define the goal of the study, the research questions to be answered to achieve the goal and metrics for answering these questions.

Although the GQM approach has been extensively used as a way to structure and design empirical studies, it is focused primarily on *quantitative* studies. Given that our study is qualitative rather than quantitative, instead of metrics we specify an evaluation approach with sub-questions, which are used in a similar way as metrics, in the sense of being a way to guide the analysis and answer the mapped questions. This helped us to select the means of achieving our study goal. The description of the study, following the GQM template, is presented in Table 1.

To achieve our goal, we investigated different application-level caching concerns, which are associated with the three key research questions presented below.

**RQ1.** What and when is data cached at the application level?

**RQ2.** How is application-level caching implemented?

**RQ3.** Which design choices were made to maintain the application-level cache efficient?

Determining the cacheable content and the right moment of caching or clearing the cache content are a developer's responsibility and might not be trivial in complex applications, motivating RQ1. In this research question, we aim to identify *what* data is selected to be cached, and the criteria used to detect such cacheable data. Furthermore, this question also explores *when* data should be evicted, as well as constraints, consistency conditions, and the rationale for all these choices.

In addition to design issues, as shown in Figure 1, the cache system and the underlying source of data are not aware of each other, and the application must implement ways to interact with the cache. Therefore, our goal with RQ2 is to characterize patterns of *how* this implementation occurs in the application code; for example, ways to assigning names to cached values, performing lookups, and keeping the cache up to date.

Finally, determining maintenance strategies to manage efficiently *where* data is placed, such as replacement policies or size of the cache, requires additional knowledge and reasoning from application developers. Nevertheless, there are no foundations to make this choice adequately. Therefore, RQ3 complements the analysis above, questioning how application specificities are leveraged to provide the desired performance to the cache system.

By answering these three central questions, we can extract patterns and guidelines for caching design, implementation and maintenance, and also the context in which each pattern can be adopted.

### 3.3 Procedure

As previously described, this study is mainly based on development information of web applications. To investigate different caching constructs, we followed a set of steps to perform our qualitative study: (i) selection of a set of suitable web applications; (ii) specification of a set of questions to guide us in the data analysis and (iii) analysis of each web application using the specified questions. The collected data consists of six different sources of information, explained as follows.

**Information about the application.** Our goal is to identify caching patterns or decisions, which possibly depend on the application domain. Therefore, we collected general details of the applications to characterize them. The collected application data is (i) its description; (iii) programming languages and technologies involved; and (iii) size of the application in terms of the number of lines of code.

**Source code.** Application source code is our core source of information. Since we focus on application-level caching, our analysis is concentrated in the core of the application (i.e. the business logic), which is where the caching logic is typically implemented.

**Code comments.** Given that caching is an orthogonal concern in the application, unrelated to the business logic, but interleaved with its code, code comments are often

used to describe the reasons behind caching implementation and decisions.

**Issues.** An issue can represent a software bug, a project task, a help-desk ticket, a leave request form, or even user messages about the project in general. Usually, changes in the code are due to registered issues. Thus, implementation and design decisions are better explained by associated issues in issue platforms, such as GitHub Issue Tracker, JIRA, and Bugzilla.

**Developer documentation.** In general, developer documentation consists of useful resources, guides and reference material for developers to learn how the application was implemented. Thus, the available documentation about the project was also collected.

**Developers.** In addition to the manual inspection of the above data, we asked developers to which we had access about caching-related implementation, decisions, challenges, and problems.

To guide the extraction of the information needed from the above data for our qualitative analysis, we derived sub-questions for each research question, which convey characteristics that should be observed and evaluated while looking for answers to the main research questions, i.e. sub-questions were used to extract data. Results presented in Section 4 are derived from answers to those questions. Table 2 depicts the sub-questions derived, which serve as a checklist while analyzing our data.

We followed the analytical process of *coding* in our analysis [47], which makes it easier to search the data and identify patterns. This process combines the data for themes, ideas and categories, labeling similar passages of text with a code label, allowing them to be easily retrieved at a later stage for further comparison and analysis. We used what we learned from the analysis to adapt our evaluation approach and observation protocols. Insights we had while coding the data and clustering the codes were captured in memos. There are three coding phases in classical grounded theory: open coding, selective coding, and theoretical coding. Open coding generates concepts from the data that will become the building blocks for the theory [47]. The process of doing grounded theory is based on a concept indicator model of constant comparisons of incidents or indicators to incidents [50]. Indicators are actual data, such as behavioral actions and events observed or described in documents and interviews. In this case, an indicator may be an architectural style or design pattern adopted to implement the cache, a data structure, a class, a control flow logic, a comment, a discussion in the issues platform, a paragraph in the documentation, or any other evidence we can get from the data being analyzed.

By using grounded theory, we aimed to construct a well-integrated set of hypotheses that explain how the concepts operated. Thus, the selective coding involves identifying the core category that best explains how study data refers to a large portion of the variation in a pattern and is considered the primary concern or problem related to the study, integrating closely related concepts. Finally, theoretical codes conceptualize how the codes may relate to each other as hypotheses to be incorporated into the theory [47].

Figure 2 provides an example of such analytical process and illustrates how we collected the different evidences.

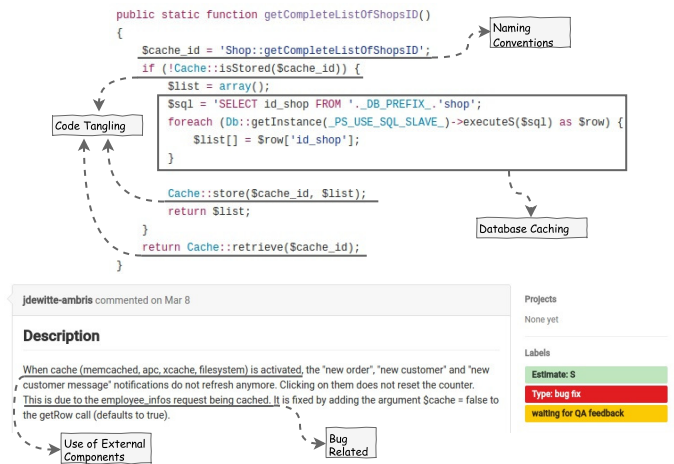


Fig. 2. Example of the Analytical Process of Coding.

First, we assigned concepts to pieces of extracted text (open coding), each representing application-level caching characteristics. Figure 2 exemplifies different codes identified during the analysis of an application. For instance, *Code Tangling* is created from the observation of cache-related implementation spread all over the application base code (underlined). Then, for each new concept, we verified whether they are connected somehow with existing ones, in order to generate categories (selective coding). Thus, the name assigned to a particular category aims at representing, at a higher abstraction level, all concepts related to it. Regarding the *Code Tangling* example, if *Code Scattered* were found afterwards, we could establish a relationship between the former and the latter to create a category, given that both are related to lack of separation of concerns.

Following the introduced phases of data analysis, we performed mainly a subjective analysis of the data, collecting: (i) typical caching design, implementation and maintenance strategies; (ii) motivations, challenges and problems behind caching, and (iii) characteristics of caching decisions. These collected data were evaluated to conceptualize how the open codes were related to each other as a set of hypotheses in accounting for resolving the primary concern. Furthermore, we also made a broad analysis of the target systems in order to investigate how application-level caching was conceived in them. All the research phases were performed manually, as the collected data (most of them expressed in natural language) analysis is associated with the interpretation of caching approaches.

### 3.4 Target Systems

In order to investigate different aspects of caching, it was important to select representative systems that make extensive use of application-level caching. To obtain applications that employ application-level caching, we searched through open-source repositories, from which information can be easily retrieved for our study. Based on text search mechanisms, we assessed how many occurrences of cache implementation and issues were present in applications to ensure they would contribute to the study. The more caching-related implementation and issues, the better, so

TABLE 2  
Evaluation Approach (Sub-questions).

#	Sub-question	RQ1. What and when is data cached at the application level?	RQ2. How is application-level caching implemented?	RQ3. Which design choices were made to maintain the application-level cache efficient?
1	What are the motivations to employ cache?	X		
2	What are the typical use scenarios?	X		
3	Where and when data is cached?	X		
4	What are the constraints related to data cached?	X		
5	What are the selection criteria adopted to detect cacheable content?	X		
6	Do developers adopt a pattern to decide which content should be cached?	X		
7	What kinds of bottlenecks are addressed by developers?	X		
8	What motivated the need for explicit caching manipulation?	X		
9	What is the granularity of the cached objects?	X		
10	What is the importance of the cached data to the application?	X		
11	Is there a relationship between the data cached and the application domain?	X		
12	How much memory does the cached data consume?	X		
13	What data is most frequently accessed?	X		
14	How often is the cached data going to be used and changed?	X		
15	What data is expensive?	X		
16	What data depends on user sessions?	X		
17	How up to date does the data need to be?	X	X	
18	How is consistency assurance implemented? Why was it chosen?	X	X	
19	Where and when is consistency assurance employed?	X	X	
20	Which kind of operation/behavior affects cache consistency?	X	X	
21	Do developers employ any technique to ease caching implementation, such as design patterns, third-party libraries or aspects?		X	
22	How is the caching logic mixed with the application code?	X	X	
23	Is this extra cache logic tested?		X	
24	What is the required format to cache?		X	X
25	How are objects translated to the cache?		X	X
26	How are names (keys) defined for cached objects?	X	X	
27	Do developers use another caching layer besides application-level?		X	
28	Is any transparent or automatic caching component being used?		X	X
29	Do developers rely on automatic caching components?		X	X
30	Is it necessary to explicitly manipulate a caching component that should be automatically managed?		X	X
31	Which application layers are more likely to have caching logic?	X	X	
32	Do developers perform analysis to measure cache efficiency?			X
33	What is the replacement policy adopted? Why was it chosen?			X
34	What is the size of the cache? Why was it defined?			X
35	What are the default parameter values?			X
36	Do developers use configurations different from to default?		X	X
37	Do developers take into account application-specific information when defining maintenance strategies?		X	X

that the rationale behind choices regarding application-level caching could be extracted. Moreover, we managed to obtain commercial applications with partner companies interested in the results of this study.

Therefore, we analyzed systems with a broad range of characteristics. Aiming at reducing the influence of particular software features on the results, we selected systems of different sizes (from 21 KLOC to 250 KLOC, without comments and blank lines), written in different programming languages, adopting different frameworks and architectural styles, and spanning different domains. We studied ten systems in total, of which nine are open-source software projects, and one is from the industry. Due to intellectual property constraints, we will refer to the commercial system as S1. Table 3 summarizes the general characteristics of each target system.

The open-source applications were selected from

GitHub, the widely known common code hosting service. We selected GitHub projects that match the following criteria: (i) projects with some popularity (at least 350 stars); (ii) projects containing application-level caching implementation and issues (at least 50 occurrences of cache-related aspects); (iii) projects written in different programming languages; and (iv) projects of different domains. The first criterion indicates that projects are interesting and were possibly evolved by people other than the initial developers. The second ensures that selected projects would present caching-related implementation and issues, which would contribute more to the study. The inclusion of applications was not restricted to any particular technology or programming language given that the study is focused on identifying language-independent caching patterns and guidelines expressed in the source code. Furthermore, we found applications that use cache in other architectural (e.g. database

TABLE 3  
Target Systems of our Study.

Project Name	Domain	Language	KLOC	Analysis Order
S1	Market trend analysis	Ruby	21	#8
Pencilblue	CMS and blogging platform	JavaScript	33	#2
Spree	e-Commerce	Ruby	50	#7
Shopizer	e-Commerce	Java	57	#3
Discourse	Platform for community discussion	Ruby	88	#5
OpenCart	e-Commerce	PHP	123	#4
OpenMRS API and web application	Patient-based medical record system	Java	127	#6
ownCloud core	File sharing solution for online collaboration and storage	PHP	193	#10
PrestaShop	e-Commerce	PHP	245	#1
Open edX	Online courses platform	Python	250	#9

systems) or infrastructure (e.g. proxy caching) levels. However, these caching approaches are not handled directly by the developers within the application; consequently, such applications do not fit the purposes of our study. The commercial system was selected by convenience. We selected the applications that better satisfied criteria (i) and (ii), without repeating programming languages and domains. As result, we achieved a set of *representative* systems, required for qualitative studies. In order to minimize bias, applications were analyzed according to a randomly generated order, which is presented in the last column of Table 3.

It is important to highlight that our goal is to capture the *state-of-practice* of application-level caching, by characterizing and deriving caching patterns. Consequently, we assume that the investigated applications—in which caching was introduced in previous releases and was already debugged—have adequate application-level implementations, i.e. we did not evaluate the quality of any caching implementation.

## 4 ANALYSIS AND RESULTS

This section details the results of our study and their analysis, according to the research questions we aim to answer. Our collected data consist mainly of source code and issues (expressed in natural language) and, as these are qualitative data, we have undertaken a subjective analysis of the application-level caching (hereafter referred to simply as “caching”) aspects represented in the target systems. Note that we labeled some findings with “*Evidence X*,” so that we can later refer to them to support the guidelines we derived from this analysis.

Before addressing each of our research questions, we show an objective analysis of the impact of caching in the investigated applications by identifying all code and issues related to them. This gives a broad sense of how caching is implemented in target systems.

### 4.1 Caching in Target Systems

To investigate how caching is present in target systems, we examined the number of files in which caching logic

TABLE 4  
Characteristics of Target Systems.

Project Name	#Cache Files	Cache LOC	#Cache Issues	Doc.	Dev.
S1	30 (4.72%)	446 (2.12%)	NA	No	Yes
Pencilblue	16 (7.04%)	526 (1.59%)	26 (4.99%)	No	No
Spree	27 (2.58%)	422 (0.84%)	205 (2.92%)	Yes	No
Shopizer	31 (4.20%)	1725 (3.02%)	0 (0%)	No	No
Discourse	53 (3.10%)	1190 (1.35%)	48 (1.34%)	Yes	No
OpenCart	44 (4.23%)	462 (0.37%)	77 (1.97%)	Yes	No
OpenMRS API and web application	22 (2.06%)	350 (0.27%)	19 (1.11%)	Yes	No
ownCloud core	240 (10.52%)	4344 (2.24%)	670 (2.99%)	Yes	No
PrestaShop	54 (4.78%)	2727 (1.11%)	95 (1.96%)	Yes	No
Open edX	198 (10.76%)	4693 (1.87%)	333 (2.95%)	No	No

is implemented, the number of LOC associated specifically with caching (without comments and blank lines), and the number of issues related to it. This analysis is shown in Table 4, in the columns *#Cache Files*, *Cache LOC* and *#Cache Issues*, respectively. This table also gives an overview of further information of each application to which we had access. They are some form of documentation and access to developers. It is important to note that available documentation about caching, in most cases, was limited to an abstract or general description of how caching was adopted. This documentation was available in some open source systems, and we only had access to developers of our system from the industry.

Based on data presented in Table 4, we can observe a significant amount of lines of code dedicated to implementing caching, ranging from 0.27% to 3.02%. It shows the importance of the caching in the project, considering that caching is a solution for a non-functional requirement (i.e. scalability and performance). Furthermore, caching logic is presented in a substantial amount of files, from 2.06% to 10.76%, which indicates the caching nature of being spread all over the application.

Moreover, we observed that all analyzed web applications did not adopt caching since their conception. As they increased in size, usage and performance analysis were performed and led to requests for improvements. Thus, developers had to refactor data access logic to encapsulate cache data into the proper application-level object, which is a task that can be very time-consuming, error-prone and, consequently, a common source of bugs. As result, we found a significant number of issues specifically related to caching, achieving the maximum of 4.99% of the *Pencilblue* issues, which can express in numbers the impact of caching in the entire project (*Evidence 15*).

We performed an analysis of the issues available in issue platforms to investigate the primary sources of cache-related problems, typically bugs, in the applications. Based on user messages, code reviews and commit messages that




Fig. 3. Classification of Caching Issues by Topic.

are described in the issues, we classified them into three different categories, which follow the main topic of our research questions: *Design* (e.g. changes in the selection of content to be put in the cache), *Implementation* (e.g. bugs in the implementation) and *Maintenance* (e.g. performance tests, adjustments in replacement algorithms or size limits of the cache). Results of the performed analysis are presented in Figure 3, in which applications are ordered ascending by the number of cache-related issues, which is shown next to each application name. Only open source projects with issues related to caching are detailed in this figure.

As can be seen in Figure 3, caching implementation and associated design decisions are much more discussed and revised by developers than maintenance decisions. Caching implementation, which is spread in the code and involves the choice for appropriate locations to add and remove elements from the cache, is error-prone and can compromise code legibility. Consequently, many issues are associated with bug fixes, technological details and code refactorings. Moreover, despite being less frequent, caching design is time-consuming and challenging, given that it requires understanding of the application behavior, as well as limitations, conditions and restrictions of content being cached. In applications analyzed, the mean (M) and standard deviation (SD) values of cache-related implementation issues are  $M = 47.45\%$  and  $SD = 5.76\%$ , while design issues achieve  $M = 37.92\%$  and  $SD = 7.11\%$ . Finally, because fine-grained configurations require empirical analysis such as cache profiling, and there is little evidence that this was performed in investigated applications, maintenance decisions often result in the choice for default settings. Consequently, a lower number of issues is associated with such decisions, specifically  $M = 14.61\%$  and  $SD = 3.44\%$ . This issue analysis allowed us to understand the aspects of caching that require more effort from developers.

After this broad analysis of our target systems, we further analyze them focusing on our research questions. For each research question, we detail coding results and explore the answers achieved while analyzing emerged codes.




Fig. 4. Amount of Emerged Categories based on the Analysis of Each System.

#### 4.2 RQ1: What and when is data cached at the application level?

In this question, we focus on going beyond the facts exposed by source code and analyze the reasoning behind caching decisions such as what and why data is cached or evicted, when and where caching is done, and what are the conditions and constraints involved with this. Therefore, issues, code comments, and documentation about the cache of applications were the primary source of information to find answers to this question, because they convey (in natural language) the rationale behind implementation decisions.

With the analysis of all provided information about application-level caching, we gathered and categorized it in an exhaustive way, observing fine-grained details. As said, we followed the principles of grounded theory, and its first phase consists of open coding. As a result, we identified 72 initial categories. Following the process, these categories were not predefined and were created and reviewed during the open coding process.

These initial categories that emerged during open coding were posteriorly analyzed and conceptualized, focusing on identifying relationships between them as hypotheses and grouping categories that may be theoretically coded as causal and associated with degrees. The open codes were directly used to derive more abstract categories during the theoretical coding. In the end, we identified 17 concepts, which better convey the implicit knowledge about application-level caching design, implementation, and maintenance acquired from the applications. After the analysis of the fourth application, codes stabilized, as shown in Figure 4.

The theoretical categories are described in Table 5, which presents their description, an original piece of content (example of evidence) and sources of data classified in each category. Moreover, categories are also shown in Figure 5a with the associated number of occurrences in the applications analyzed and classified according to each research question. Figure 5b shows the percentage of contribution of each application to the categories emerged from the study. These figures also present categories identified in RQ2 (described in Table 6) and RQ3 (described in Table 7), which are discussed in the following sections. Acronyms used in Figures 5a and 5b are introduced in Tables 5, 6, and 7.

We observed in eight from the ten analyzed applications that developers indicated that they had *uncertainty in cache design*, with respect to deciding what data should be cached and the right moment to do it, causing missed caching opportunities. There are comments in the source code and



TABLE 5  
Analysis of RQ1: Emerged Design Categories.

Category (Acronym)	Description	Example of Evidence	Evidence Sources					
			SC	COM	IS	DOC	DEV	
Uncertainty in Cache Design (UCD)	Indication, by developers, of uncertainty regarding cacheable content.	A code comment before an expensive operation: <i>TODO cache the global properties to speed this up??</i>		X	X			X
Explanations of Cache Design Choices (EDC)	Explanations provided by developers regarding cache design choices, such as specific criteria to determine cacheable content.	A code comment detailing the motivation for caching: <i>Fetch all dashboard data. This can be an expensive request when the cached data has expired, and the server must collect the data again.</i>	X	X	X	X		
Multiple Caching Solutions (MCS)	Occurrences of multiple caching solutions, which can involve different third-party components or different application layers. As a consequence, the same content can be cached at different places in varying forms.	A user announcement on issue platform: <i>Caching has now landed: Fragment caching for each product; Fragment caching for the lists of products in home/index and products/index; Caching in the ProductsController, using expires_in which caches for 15 minutes.</i>	X	X	X	X		
Caching in Application Boundaries (CAB)	Occurrences of caching content being added to the cache because they retrieve data from frameworks, libraries, or external applications.	A sentence in the documentation: <i>Most stores spend much time serving up the same pages over and over again. [...] In such cases, a caching solution may be appropriate and can improve server performance by bypassing time-consuming operations such as database access.</i>	X	X	X	X		X
Ensuring Consistency (ENC)	Design of some kind of consistency approach such as expiration policies and invalidation, preventing stale data.	A user message on issue platform detailing the invalidation approach adopted: <i>Ideally we cache until the topic changes (a new post is added, or a post changed) [...] less ideally we cache for N minutes</i>	X	X	X	X		X
Complex Cache Design (CCD)	Indication that a cache design choice is difficult to understand, confusing developers and requiring detailed comments.	A user message on issue platform: <i>i can see that it may speed up the query responses, but is the saved time substantial enough to be worth the effort?</i>			X			X
Choice for Simple Cache Design Solutions (CSD)	Indication that developers selected simple solutions for cache rather than complex ones, such as defining a time-to-live (TTL) instead of implementing manual invalidation, possibly in order to balance design effort and caching gains.	A code snippet defining a default setting: <pre>&lt;defaultCache maxElementsInMemory="1000" eternal="false" timeToIdleSeconds="60" timeToLiveSeconds="0" overflowToDisk="false" diskPersistent="false"/&gt;</pre>	X	X	X			X

Labels of Evidence Sources: **SC**-Source Code (without comments); **COM**-Code Comments; **IS**-Issues; **DOC**-Documentation; **DEV**-Developers.

issues about whether to cache some specific data or not, showing that the connection between observed bottlenecks in the application and opportunities for caching is not straightforward and requires a deeper analysis (Evidence 1).

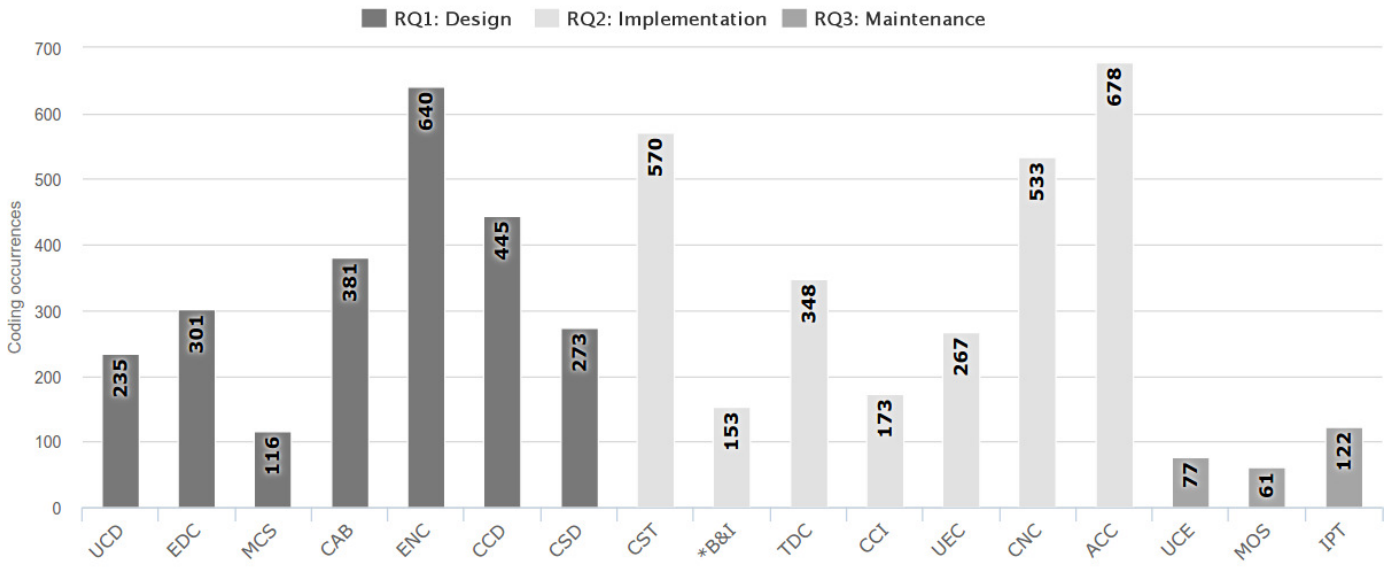
Therefore, selection criteria based on common sense or past experiences are initial assumptions to ease such decisions. Despite these criteria are usually unjustified, they guide developers while selecting content. We observed in 90% of the applications the definition of selection criteria to make the distinction of cacheable from uncacheable content easier. These criteria were observed in *explanations of cache design choices* in comments, issues, and documentation. We identified common criteria used to determine whether to cache or not a specific content, which are: (i) content change frequency (Evidence 2), (ii) content usage frequency (Evidence 3), (iii) content shareability (Evidence 4), (iv) content retrieval complexity (Evidence 5), (v) content size (Evidence 6), and (vi) size of the cache (Evidence 7).

Regarding design choices, we observed common practices. The first is associated with the lack of a specific approach to cache data. To process a client request, application components of distinct layers and other systems (databases, web services, and others) are invoked, and each interaction results in data transformation, which is likely cacheable. Due to this, nine analyzed applications present *multiple caching solutions* by not specifying cacheable layers, com-

ponents or data, and employing cache wherever can potentially provide performance and scalability benefits (Evidence 8), no matter which is the application layer, component or data (Evidence 9). As a consequence, the same content can be cached at different places, from the database to controllers, in varying forms such as query results, page fragments or lists of objects (Evidence 10).


The second design choice is associated with the concern of reducing the communication latency between the application and other systems, which increases the overall response time of a user request. Therefore, we noticed *caching in application boundaries* in nine of the analyzed applications, addressing remote server calls and requests to web services, database queries, and loads dependent on file systems, which are common bottlenecks (Evidence 11).

Despite choosing where and what to cache, cached values are valid only as long as the sources do not change, and when sources change, a consistency policy should be employed to ensure that the application is not serving stale data. Therefore, in nine analyzed applications, there are indications that developers demand an effort to design *consistency approaches*, reasoning about the lifetime of cached data, as well as eviction conditions and constraints. We identified common approaches to keep consistency, which are: (i) a less efficient and straightforward approach is to invalidate cached values based on mapping actions that



\*Only 9 applications had available sources to identify this category.

(a) Coding Occurrences.



(b) Coding Percentage by Application.

Fig. 5. Categories Emerged from the Study.

change its dependencies and, after a change is invoked, invalidate cached values and recompute them at the next request; and (ii) the use of a less intrusive alternative such as a time-to-live (TTL) or a replacement approach, which require a certain level of staleness (Evidence 12).

All these caching design options may become complex and difficult to understand. Indeed, identifying caching opportunities and ensuring consistency can add much code and may not be trivial to implement and understand. Due to the nature of application-level caching, such logic is spread all over the system. We noticed in 90% of the applications

the presence of pieces of evidence that *caching design achieves a high level of complexity*, requiring many comments explaining even less complex parts and issues related to the lack of understanding of caching decisions (Evidence 13).

Due to the increasing demand and complexity involved with caching, developers try to decrease cache-related effort by *adopting simple design solutions*. We observed in seven of the ten analyzed applications design choices such as managing consistency based on expiration time, keeping default parameters, selecting data without any criteria, or even adopting external solutions, which do not claim extra

reasoning, time, and modifications to the code (*Evidence 14*).

We discussed our findings regarding cache design decisions and we next discuss observations made associated with how to implement design decisions.

### 4.3 RQ2: How is application-level caching implemented?

To understand how developers integrate application-level caching logic in their applications, we analyzed the cache implementations from two perspectives. The first consists of examining the explicit caching logic present in the application code, focusing on analyzing where the caching logic is placed, for what this logic is responsible, and when it is executed. The second evaluates the integration and communication between the application and an external caching component, which is usually used to relieve the burden on developers, easing cache implementation, raising abstraction levels, or even taking full control of caching. For this question, the most valuable data comes from source code and comments, which express implementation details. The theoretical categories referring to RQ2 are described in Table 6 and shown in Figure 5.

We observed in eight from the ten analyzed applications that they present *code scattering and tangling*, on caching logic, causing low cohesion and high coupling in the code. Caching control code, responsible for caching data in particular places, was spread all over the base code, being invoked when application requests were processed. Consequently, there is a lack of separation of concerns, leading to increased maintenance effort (*Evidence 15*).

A possible cause for this is that caching was not part of all the applications since their conception. Applications were developed and, as they evolved, usage and performance problem reports led to requests for response time improvements. Thus, developers had to refactor existing business logic to include caching aspects. Interleaved caching logic can cause not only increased maintenance effort but also be a source of bugs—in eight (from nine, given that we had no access to an issue tracker of one of the applications) analyzed applications, there are issues associated with *bugs due to caching* (*Evidence 16*).

Although this problem is present in the code, there are indications that developers know about it and express they are willing to improve the provided solution. In order to reduce the impact of an infrastructure component to the system business logic, we identified cases where there are suggestions to design more extensible classes and modules, refactoring and reducing cache-related code, and reusing components (*Evidence 17*). This acknowledgment of *technical debt* was observed in 90% of the applications.

Regarding implementation choices, we observed common practices. The first is associated with how to name cached data. In order to use in-memory caching solutions, there is no prescribed way to organize data. Typically, unique names are assigned to each cached content, thus leading to a key-value model—and this was the case in all investigated applications. Given that cache stores lots of data, the set of possible names must be large; otherwise, two names (keys) can conflict with each other and, thus stale (or even entirely wrong) data can be retrieved from

the cache. Consequently, *complex naming conventions* were adopted (*Evidence 18*).

Therefore, in all applications, there are indications, in the form of source code, comments, and issues, that developers demand an effort to understand and improve cache keys, reasoning about alternative better ways to organize and identify content. We identified common content properties used to build cache keys, which are: (i) content *ids*, (ii) method signatures, and (iii) a tag-based identification, in which the type of content, the class, the module or hierarchy are used (*Evidence 19*).

The second implementation choice is associated with the concern of not increasing the throughput due to the introduction of communication between the cache and the application. Due to this, six of our applications make *use of non-trivial programming techniques* to ensure a high performance caching implementation, taking advantage of every improvement opportunity (*Evidence 20*). Solutions related to this include processing caching requests in the background, batching caching calls together, or even implementing a coarse-grained caching service that allows a single logical operation to be performed by using a single round trip. These prevent blocking client requests when a possibly large caching processing is being performed, e.g. a caching warm up or update (*Evidence 21*).

Given that implementing cache is challenging, we noticed that in six applications developers made *use of supporting libraries and frameworks*. This was done to prevent adding much cache-related code to the base code, because such components raise the abstraction level of caching, providing some ready-to-use features (*Evidence 22*). Examples of such external components are distributed cache systems, e.g. Redis [27] and Memcached [26], and libraries that can act locally, e.g. Spring Caching [51], EhCache [52], Infinispan [53], Rails low-level caching [54], Google Guava [55] and Caffeine [56].

Such supporting libraries and frameworks not only provide partial ready-to-use features but also reduce the amount of additional effort required to guarantee that the cache is working. We observed in all applications that the cache includes *code dedicated to test, debug and configure cache components*, which can be expensive in some scenarios (*Evidence 23*).

We next discuss observations made associated with how the cache is maintained and tuned.

### 4.4 RQ3: Which design choices were made to maintain the application-level cache efficient?

After designing and implementing the cache, maintenance issues are still open. Initial assumptions concerning caching design can become invalid as changes occur in the application domain, workload characteristics or access patterns, thus leading to a performance decay. Therefore, in order to keep the cache efficient, it is necessary to constantly tune cache settings, at least in between releases. Maintenance decisions involve: (i) detecting improvement opportunities in the cache design, (ii) dealing with specificities of the cache deployment scheme (e.g. local or remote, shared or dedicated, in memory or on disk storage), (iii) determining the appropriate size of the cache, (iv) defining how many

TABLE 6  
Analysis of RQ2: Emerged Implementation Categories.

Category (Acronym)	Description	Example of Evidence	Evidence Sources				
			SC	COM	IS	DOC	DEV
Code Scattering and Tangling (CST)	Presence of caching logic spread all over the application, indicating a lack of separation of concerns.	A developer quote: <i>At first glance, the cache code hinders the understanding of the business logic. Also, the cache logic itself it is not easy to get.</i>	X	X	X		X
Bugs and Issues (B&I)	Presence of bugs and problems due to caching reporting, for example, increased client response time, reduced throughput, increased server resource utilization.	A bug report on an issue platform: <i>When you import categories with a parent category which does not exist, it prevents from duplicate it because of the cache.</i>			X		
Technical Debt Concerns (TDC)	Indication of extra effort spent by developers to build an extensible and well designed caching component.	A user message on an issue platform: <i>my concern with this implementation is that it can lead to some confusing situations. [...] I think the caching needs to be more granular</i>	X	X	X	X	X
Complex Caching Implementation (CCI)	Presence of complex constructs such as batch processing and asynchronous communication, which require extra effort and reasoning from developers to be implemented.	A user message on an issue platform: <i>Can you give me some tips on how to safely implement an async step to save this? [...] I could have the end user polling and refreshing against that as needed.</i>	X	X	X	X	
Use of External Components (UEC)	Indication of use of third-party caching solutions to help the implementation, raising the level of abstraction of some caching aspects.	A sentence in documentation: <i>We use Redis [a third-party solution] as a cache and for transient data.</i>	X	X	X	X	X
Complex Naming Conventions (CNC)	Choice for complex keys of caching content, causing developers to spend time and effort to elaborate and understand such keys.	A code snippet: <code>cache_id = 'objectmodel_' . \$entity_defs['classname'] . '_' . (int)\$id . '_' . (int)\$id_shop . '_' . (int)\$id_lang;</code>	X	X	X		
Additional Caching Code (ACC)	Code implemented to support the caching logic, such as implemented caching tests, logic to monitor cache statistics and additional interfaces to support available caching providers.	A code snippet exposing a test of caching logic: <i>it "can set and get false values when return cache nil" do</i> <code>@store.set :test, false</code> <i>expect(@store.get(:test)).to be false</i> <code>end</code>	X	X	X		

Labels of Evidence Sources: **SC**-Source Code (without comments); **COM**-Code Comments; **IS**-Issues; **DOC**-Documentation; **DEV**-Developers.

objects are allowed to be cached at same time, and (v) what should be done in case of the cache is full. The theoretical categories identified while investigating how developers approach these issues and maintain cache efficient are described in Table 7 and shown in Figure 5.

As mentioned in Section 4.2, we observed that developers often rely on external components, which can raise the level of caching abstraction, providing some cache-related components ready-to-use. The most common tasks delegated to external components are associated with cache space allocation and management, since maintaining cache manually through lists or hashes inside the application involves issues such as dealing with concurrency, key management and size limits, which may not be trivial.

Furthermore, in cases where updates in the base code are not an option (due to time or technical restrictions), a transparent and automatic caching component can provide fast results. These solutions address layers before and after application boundaries, and require only a few adaptations to the application needs (Evidence 24).

However, we observed in three analyzed applications an *uncertainty regarding the quality, in terms of performance, of a transparent component* (Evidence 25). In fact, transparent caching solutions are out of the application control and are built and maintained by a third party community. Despite they provide ways to developers to manipulate and observe their behavior, it may become a problem when requests for improvements force developers to customize the behavior of the built-in cache component, thus creating a new and

particular behavior for the component, which is supposed to be generic and easy-to-use. All this manipulation should also be tested and evaluated, leading to the generation of test cases to assert the reliability of the external component being used.

The use of transparent and automatic caching solutions relieve developers also from two maintenance tasks. First, while it is true that the larger the cache size, the better the performance, very large caches are unrealistic. While cost is the primary reason for limiting the size of a cache, efficient invalidation approaches play a key role as well. We noticed in five of our analyzed applications indications that developers explicitly *deal with size limitations* (Evidence 26), by defining a specific size of the cache, replacement policies (Evidence 27) and the number of manageable objects allowed to be cached.

The second maintenance task refers to cache performance improvements. We observed in seven of the ten applications indications that developers *perform tests*, comparing the application behavior with and without a particular caching approach (Evidence 28). However, while focusing on tuning all the aspects of caching, developers perform trivial analyses, which are usually based on comparing logs or execution of several requests in a row and measurement of response times. These tests are inappropriately reported and documented, becoming irreproducible. As result, these tests can lead to false-positive improvements, based on poor and biased conclusions (Evidence 29).

Based on our analysis, we identified application-level

TABLE 7  
Analysis of RQ3: Emerged Maintenance Categories.

Category (Acronym)	Description	Example of Evidence	Evidence Sources				
			SC	COM	IS	DOC	DEV
Uncertainty in External Cache Component (UCE)	Indication of uncertainty in the behavior of an external cache component, when the applications rely on it instead of handling the cache manually.	A code comment: <i>Is hibernate taking care of caching and not hitting the db every time? (hopefully it is)</i>	X	X	X	X	
Maintenance of Cache Size (MOS)	Indication of choices made in order to control the cache size, such as definitions of eviction approaches to fit cache size limitations.	A user message on issue platform: <i>Be sure the local cache will not grow out of control (especially during big operations like product import)</i>	X	X	X		
Improvements based on Performance Tests (IPT)	Indication that improvements were made based on performance tests.	An user message on issue platform comparing two request log traces, with and without cache: <i>My results were something like: 1) Using rows: 30ms, 5mb RAM (peak usage); 2) Fetching all at once: 3ms, 5.75MB RAM</i>			X		X

Labels of Evidence Sources: **SC**-Source Code (without comments); **COM**-Code Comments; **IS**-Issues; **DOC**-Documentation; **DEV**-Developers.

caching decisions and behaviors adopted by developers. Our findings allowed us to propose a set of guidelines and patterns for the development of a caching approach, which are described in next section.

## 5 GUIDELINES AND PATTERNS

This study allowed us to understand how developers deal with application-level caching in their applications, by explaining design, implementation and maintenance choices. Our findings and observations were used as foundation to the provision of practical guidance for developers with respect to caching. In this section, we thus introduce guidelines, which are general rules to be followed by developers while developing application-level caching, and patterns, which are structured reusable solutions to recurring problems in cache design, implementation or maintenance. For each guideline, we indicate evidences that support it. Both guidelines and patterns derived from our study are classified into categories, explained below. Note that there are patterns associated with the proposed guidelines.

**Design.** Support to design decisions associated with application-level caching.

**Implementation.** Support to implementation issues of application-level caching, by providing solutions and guidance at the code level.

**Maintenance.** Support to performance analysis and improvement of application-level caching.

### 5.1 Design Guidelines

*Evaluate different abstraction levels to cache.* (Evidence 9) and (Evidence 11) It is important to cache data where it reduces the most processing power and round trips, choosing locations that support the lifetime needed for the cached items, despite where it is located in the application. Different levels of caching provide different behavior, and possibilities must be analyzed. For instance, caching in the model or database level offers higher hit ratios, while caching in presentation layer can reduce the application processing overhead significantly in the application in case of a hit. However, in the latter case, hit ratios are in general lower. It is possible to cache data at various layers of an application, according to the following layer-by-layer considerations.

**Controller layer.** Caching data at the controller layer should be considered when data needs to be frequently displayed to the user and is not cached on a per-user basis. At this level, controllers usually work by serving parameterized content, which can be used as an identifier in the cache. For example, if a list of states is presented to the user, the application can load these once from the database and then cache them, according to the parameters passed in the first request.

**Business or service layer.** Caching data at the business layer should be considered if an application needs to process requests from the presentation layer or when the data cannot be efficiently retrieved from the database or another service. It can be implemented by using hash tables, library or framework. However, at this level, a large amount of data tends to be manipulated and caching it can consume more memory and leads to memory bottlenecks.

**Model or database layer.** At the model or database layer, a large amount of data can be cached, for lengthy periods. It is useful to cache data from a database when it demands a long time to process queries, avoiding unnecessary round-trips.

*Stack caching layers.* (Evidence 9) and (Evidence 10) It is reasonable to say that the more data cached, the lower the chance of being hit without any content already loaded. Caching might be at the client, proxy server, inside the application in presentation, business, and model logics, or database. Despite the same data may be cached in multiple locations, when the cache expires in one of them, the application will not be hit with an entirely uncached content, avoiding processing and network round trips. However, it is important to keep in mind that caching layers imply a range of responsibilities, such as consistency conditions and constraints, and extra code and configuration. Due to this, it is important to consider many caching layers but, at the same time, achieve a good trade-off between caching benefits and implementation effort.

*Separate dynamic from static data.* (Evidence 3) and (Evidence 2) Content can be distinguished in static, dynamic, and user-specific. By partitioning the content, it is easier to select portions of the data to cache.

*Evaluate application boundaries.* (Evidence 10) and (Evidence 11) Communication between application and external

components is a common bottleneck and, consequently, an opportunity for caching. Consider caching for database queries, remote server calls and requests to web services, which are made across a network.

**Specify selection criteria.** (Evidence 1), (Evidence 2), (Evidence 3), (Evidence 5), (Evidence 6) and (Evidence 7) Selecting the right data to cache involves a great reasoning effort given that data manipulated by web applications range in dynamicity, from being completely static to changing constantly. To optimize this selection process, there are four primary selection criteria used by developers while detecting cacheable content, which should be used in decisions regarding whether to cache. These criteria are described below, ordered according to their importance; i.e. the higher the influence level, the earlier it is presented.

**Data change frequency.** Developers should seek for data that have some degree of stability, i.e. those that are more used than changed. Even if data are volatile and change in time intervals, caching still brings a benefit. This is the first factor to be considered since caching volatile data implies the implementation of consistency mechanisms, which is not trivial and requires an extra effort and reasoning from developers. In short, the cost of consistency approaches cannot be higher than the benefit of caching. Besides, when stale data is not a critical issue, an approach of weak consistency can be employed, such as time-to-live (TTL) eviction, where data is expired after a time in cache, regardless of possible changes.

**Data usage frequency.** Frequent requests, operations, queries and shared content (accessed by multiple users) must be identified, focusing on recomputation avoidance. Even if some processing can be fast enough at a glance, it can potentially become a bottleneck when being invoked many times. Despite being frequently used, user-specific data cannot be shared and may not bring the benefit of caching, being usually left out of the cache.

**Data retrieval complexity.** Data that is expensive to retrieve, compute, or render, regardless of its dynamicity, is always considered a good caching opportunity.

**Size of the data.** The size of the content being cached should be considered when using size-limited caches. In this case, an adequate trade-off between popularity (hits) and size of the items must be achieved. Keeping small popular items in the cache tends to optimize hit-ratio; however, a hit in a large item may be more beneficial for an application than many hits on small items. At the same time, filling the cache with few large items may turn the cache performance dependent on a good replacement policy.

**Evaluate staleness and lifetime of cached data.** (Evidence 2), (Evidence 3) and (Evidence 12) Every piece of cached data is already potentially stale, it is important to rethink the degree of integrity and potential staleness that the application can compromise for increased performance and scalability. Many cache implementations adopted an expiration policy to invalidate cached data based on a timeout since weak consistency is easier than defining a hard-to-maintain, but more robust, invalidation process. In short, developers must

ensure that the expiration policy matches the pattern of access to applications that use the data, which is based on determining how often the cached information is allowed to be outdated, and relaxing freshness when possible.

**Avoid caching per-user data.** (Evidence 4) and (Evidence 26) It is recommended to avoid caching per-user data unless the user base is small and the total size of the cached data does not require an excessive amount of memory; otherwise, it can cause a memory bottleneck. However, if users tend to be active for a while and then go away again, caching per-user data for short-time periods may be an appropriate approach. For instance, a search engine that caches query results by each user, so that it can page through results efficiently.

**Avoid caching volatile data.** (Evidence 2) and (Evidence 3) Data should be cached when it is frequently used and is not continually changing. Developers should remember that caching is most effective for relatively stable data, or data that is frequently read. Caching volatile data, which is required to be accurate or updated in real time, should be avoided.

**Do not discard small improvements.** (Evidence 3), (Evidence 8) and (Evidence 9) The user perceived latency is reduced by any caching solution employed. This means that even not obvious scenarios should be target of caching, i.e. it is not true that solely data that is frequently used and expensive to retrieve or create should be considered for caching. Furthermore, data that is expensive to retrieve and is modified on a periodic basis can still improve performance and scalability when properly managed. Caching data even for a few seconds can make a large difference in high volume sites. If the data is handled more often than it is updated, it is also a candidate for caching.

## 5.2 Implementation Guidelines

**Keep the cache API simple.** (Evidence 15), (Evidence 17), (Evidence 20) and (Evidence 22) Caching logic tends to be spread all over the application, and a good solution should be employed to avoid writing messy code at the cost of high maintenance efforts.

**Define naming conventions.** (Evidence 19) and (Evidence 18) To define appropriate names for cached data, it is important to assign a name that is related to its context, the data itself, and the caching location. It can provide two direct benefits: (a) prevention of key conflicts, and (b) guidance of cache actions such as updates and deletes of stale data in case of changes in the source of information.

**Perform cache actions asynchronously.** (Evidence 21) For large caches, it is adequate to load the cache asynchronously with a separate thread or batch process. Moreover, when an expired cache content is requested, it needs to be repopulated and doing so synchronously affects response time and blocks the request processing thread.

**Do not use cache as data storage.** (Evidence 27) An application can modify data held in a cache, but the cache should be considered as a transient data store that can disappear at any time. Therefore, developers should not save valuable data only in the cache, but keep the information where it should be as well, minimizing the chances of losing data if the cache unexpectedly becomes unavailable.

### 5.3 Maintenance Guidelines

**Perform measurements.** (Evidence 1), (Evidence 8), (Evidence 26), (Evidence 28) and (Evidence 29) Caching is an optimization technique and, as any optimization, it is important to perform measurements before making substantial changes, given that not all application performance and scalability problems can be solved with caching. Furthermore, if unnecessarily employed, caching can eventually decrease performance rather than improve it.

**Document and report measurements.** (Evidence 29) To compare and reproduce performance tests employed, it is important to document the setup used to perform the application analysis. It includes modules enabled, particular configurations and hardware settings.

**Consider using supporting libraries and frameworks.** (Evidence 22) and (Evidence 24) Supporting libraries and frameworks can raise the level of abstraction of cache implementation and provide useful features. In addition, they can scale up in a much easier and faster way than application-wide solutions.

**Tune default configurations.** (Evidence 22), (Evidence 25), (Evidence 26) and (Evidence 27) Default configurations provided by external components serve as a start point. However, because they are generic and may not fit the application specificities, other configurations must be evaluated and possibly adopted.

**Use of transparent caching components.** (Evidence 22) and (Evidence 25) The use of transparent caching solutions to address bottlenecks outside the application boundaries such as databases, final HTML pages or fragments, and static assets can provide fast results. These solutions do not explore application specificities, but can still provide performance benefits for typical usage scenarios.

All the proposed guidelines are summarized in Table 8 with the respective name, associated evidence and identification.

### 5.4 Patterns

Based on our study, we derived caching patterns, which can be used by developers to help them design, implement and manage cache. They can be used in combination with our guidelines. Moreover, we identified the components these patterns must have, which comprise a template for a caching pattern catalog. These components are (i) a *classification*, (ii) the pattern *intent*, (iii) the *problem* involved, (iv) the *solution* proposed, and (v) an *example*.

The proposed patterns are summarized in Table 9 along with the possible combination with guidelines. We, in this paper, limit ourselves to present in detail one of our patterns, namely the *Cacheability Pattern*, in Table 10, as an example. The complete description of the remaining patterns derived from our study is available elsewhere [57].

## 6 THREATS TO VALIDITY

We now analyze the possible threats to the validity of this study, and how we mitigated them. Researcher bias is a typical threat to qualitative research because the results are subject to the researcher interpretation of the data. In our study, prior knowledge might have influenced results. In

TABLE 8  
Caching Guidelines Derived from the Study.

Guideline	Associated Evidence	ID
Evaluate different abstraction levels to cache	(Evidence 9) and (Evidence 11)	DG-01
Stack caching layers	(Evidence 9) and (Evidence 10)	DG-02
Separate dynamic from static data	(Evidence 3) and (Evidence 2)	DG-03
Evaluate application boundaries	(Evidence 10) and (Evidence 11)	DG-04
Specify selection criteria	(Evidence 1), (Evidence 2), (Evidence 3), (Evidence 5), (Evidence 6) and (Evidence 7)	DG-05
Evaluate staleness and lifetime of cached data	(Evidence 2), (Evidence 3) and (Evidence 12)	DG-06
Avoid caching per-user data	(Evidence 4) and (Evidence 26)	DG-07
Avoid caching volatile data	(Evidence 2) and (Evidence 3)	DG-08
Do not discard small improvements	(Evidence 3), (Evidence 8) and (Evidence 9)	DG-09
Keep the cache API simple	(Evidence 15), (Evidence 17), (Evidence 20) and (Evidence 22)	IG-01
Define naming conventions	(Evidence 19) and (Evidence 18)	IG-02
Perform cache actions asynchronously	(Evidence 21)	IG-03
Do not use cache as data storage	(Evidence 27)	IG-04
Perform measurements	(Evidence 1), (Evidence 8), (Evidence 26), (Evidence 28) and (Evidence 29)	MG-01
Document and report measurements	(Evidence 29)	MG-02
Consider using supporting libraries and frameworks	(Evidence 22) and (Evidence 24)	MG-03
Tune default configurations	(Evidence 22), (Evidence 25), (Evidence 26) and (Evidence 27)	MG-04
Use of transparent caching components	(Evidence 22) and (Evidence 25)	MG-05

TABLE 9  
Caching Pattern Classification.

Pattern	Classification	Intent	Associated Guidelines
Asynchronous Loading	Implementation	Design a mediator to asynchronously deal with caching.	IG-03, MG-03
Cacheability	Design	Provide an intuitive process to decide whether to cache or not particular data.	DG-01, DG-03, DG-05, DG-07, DG-08, DG-09
Data Expiration	Design and Maintenance	Given cacheable content, provide an intuitive process to choose a consistency management approach based on data specificities.	DG-06, DG-08, MG-03, MG-04
Name Assignment	Implementation	Ensure a unique key and keep track of the content cached.	IG-02

TABLE 10  
Cacheability Pattern.

<b>Classification:</b> Design
<b>Intent:</b> provide an intuitive process to decide whether to cache or not particular data.
<b>Problem:</b> cache has limited size, so it is important to use the available space to cache data that maximizes the benefits provided to the application. Otherwise, it can end up reducing application performance instead of improving it, consuming more cache memory and at the same time suffering from cache misses, where the data is not getting served from cache but is fetched from the source.
<b>Solution:</b> even though there are many criteria that contribute for identifying the level of data cacheability, there is a subset that would confirm this decision regardless of the values of the other criteria. Changeability is the first criterion that should be analyzed while selecting cacheable data, then usage frequency, shareability, retrieval complexity, and cache properties should be considered.
Figure 6 expresses a flowchart of the reasoning process to decide whether to cache data, based on the observation of data and cache properties. All criteria are tightly related to the application specificities and should be specified by the developer.

```

graph TD
    Q1{Is the data completely static?} -- Yes --> C[Cacheable]
    Q1 -- No/Not Sure --> Q2{Does the data change more than it is used?}
    Q2 -- Yes --> UC[Uncacheable]
    Q2 -- No/Not Sure --> Q3{Is the data used by many requests?}
    Q3 -- Yes/Not Sure --> Q4{Is the data user-specific?}
    Q4 -- Yes --> UC
    Q4 -- No/Not Sure --> Q5{Is the data size large?}
    Q5 -- Yes/Not Sure --> Q6{Is the cache size large?}
    Q6 -- Yes --> C
    Q6 -- No/Not Sure --> Q7{Is the data expensive to compute?}
    Q7 -- Yes --> C
    Q7 -- No --> UC
    
```

Fig. 6. Cacheability flowchart: intuitive process to decide whether to cache or not particular data.

*Rules of thumb:*

- (a) Despite being frequently used, user-specific data are not shareable and may not bring the benefit of caching, being usually avoided by developers. In this case, a specific session component is used to keep and retrieve user sessions.
- (b) If the data changes frequently, it should not be immediately discarded from cache. An evaluation of the performance benefits of caching against the cost of building the cache should be done. Caching frequently changing data can provide benefits if slightly stale data is allowed.
- (c) Expensive spots (when much processing is required to retrieve or create data) are bottlenecks that directly affect application performance and should be cached, even though it can increase complexity and responsibilities to deal with. Methods with high latency or that consists of a large call stack are some examples of this situation and opportunities for caching.

In addition, we list content properties that should be avoided, which do not convey the influence factors in a good way and lead to problems such as cache trashing.

- (a) User-specific data. Avoid caching content that varies depending on the particularities of the request, unless weak consistency is acceptable. Otherwise, the cache can end up being fulfilled with small and less beneficial objects. As result, the caching component achieves its maximum capacity earlier and is flushed or replaced many times in a brief period, which is cache thrashing.
- (b) Highly time-sensitive data. Content that changes more than is used should not be cached given that it will not take advantage from caching. The cost of implementing and designing an efficient consistency policy may not be compensate.
- (c) Large-sized objects. Unless the size of the cache is large enough, do not cache large objects, it will probably result in a cache trashing problem, where the caching component is flushed or replaced many times in a short period.

**Example:** we list some typical scenarios where data should be cached and also give explanations based on the criteria presented.

- (a) Headlines. In most cases, headlines are shared by multiple users and updated infrequently.
- (b) Dashboards. Usually, much data need to be gathered across several application modules and manipulated to build a summarized information about the application.
- (c) Catalogs. Catalogs need to be updated at specific intervals, are shared across the application, and manipulated before sending the content to the client.
- (d) Metadata/configuration. Settings that do not frequently change, such as country/state lists, external resource addresses, logic/branching settings and tax definitions.
- (e) Historical datasets for reports. Costly to retrieve or create and does not need to change frequently.



order to avoid this, we followed a systematic analysis of our data, and conclusions are all founded on these data. Moreover, cross-checks were performed using our different sources of evidence.

Another source of bias in this study is the process adopted to select applications and the number of applications analyzed. We mitigated this problem by specifying criteria (described in Section 3.4) that allowed us to select representative systems with different characteristics to be part of this study. Moreover, we observed that our codes and patterns emerged after analyzing part of our systems and, in the remaining ones, we identified only recurrences of codes and patterns. This is an evidence that the investigation of a larger number of systems would not bring more information. Furthermore, given that the goal of qualitative research is to understand a particular event, rather than to provide a general description of a large sample of a population, selecting few representative applications is sufficient for our study.

The existence of such plateau reached after the analysis of the fourth application also gives evidence that the results may be widely applied to other web applications. However, we acknowledge that the selected applications may have influenced the results of this study as well as the order in which they were analyzed, thus being an external threat to the validity of this study. We mitigated this threat by selecting a sample of applications with different domains, programming languages, technologies, sizes and business models, and the use of a random order to analyze them.

Despite caching is associated with performance concerns and requirements, the focus of this study is on software engineering concerns of application-level caching, approaching design, implementation, and maintenance from a developer point of view. It is important to highlight that our goal is to capture the state-of-practice of application-level caching, by characterizing and deriving caching patterns. Consequently, we assume that the investigated applications have adequate caching implementations given that caching was introduced in previous releases and was already debugged and improved by developers, i.e. we assume that the design and implementation are all efficient; otherwise, there would be open issues in issue trackers to be resolved by developers. Therefore, another threat in our research is related to the possibility that we observed inadequately designed and implemented caching. Furthermore, the quality or performance improvement of design and implementation of application-level caching could be better explored in future work.

Finally, most of the web applications satisfy initial performance and scalability requirements using traditional forms of caching. The need for application-level caching usually appears after the application is released, with an increasing user demand. Due to this, all applications selected in our study did not adopt application-level caching since their conception—no application implemented with application-level caching in its first release was found. Consequently, it is not possible to assess if the overhead of adding cache to an application in later stages is higher than doing so during its initial development. For all analyzed systems, performed complexity or usage analysis led to requests for improvements, as systems became larger.

Consequently, developers had to refactor the application to insert caching logic. This is thus another external threat to validity, given that our results cannot be generalized to systems that use application-level caching since they are conceived.

## 7 DISCUSSION AND CONCLUSION

Application-level caching has been increasingly used in the development of web applications, in order to improve their response time given that they are becoming more complex and dealing with larger amounts of data over time. Caching has been used in different locations, such as proxy servers, often as seamless components. Application-level caching allows caching additional content taking into account application specificities not captured by off-the-shelf components. However, there is limited guidance to design, implement and manage application-level caching, which is often implemented in an *ad-hoc* way. In this paper, we presented a qualitative study performed to understand how developers approach application-level caching. The study consisted of the selection ten web applications and investigation of caching-related aspects, namely design, implementation and maintenance practices.

We observed a higher number of categories, which are classes of observations made based on the analysis of these applications, associated with caching design and implementation than those associated with maintenance. Furthermore, the number of occurrences of each category, design and implementation categories also have higher numbers. This phenomenon was expected since the most representative portion of our qualitative data consists of source code and issues. Moreover, simple maintenance tasks and configurations are already executed and provided by external components, being commonly adopted. However, the number of occurrences do not reflect the importance of a category.

The high number of occurrences related to *ensuring consistency* refers to the expiration process of cached content, which requires extra reasoning from developers, because they should track which changes cause data content to become outdated, and be aware for how long the cache can provide stale data, in case the data source has been updated. In fact, consistency approaches have been widely investigated [4], [6], and the typical way of dealing with it is to analyze data dependency, from which conditions and constraints for consistency can be derived.

Many emerged categories, such as *Bugs and Issues*, *Technical Debt Concerns* and *Complex Design and Implementation*, indicate the additional burden placed on developers while designing and implementing application-level caching, which can lead to a significant amount of time and effort added to software projects. Therefore, ideally, opportunities for caching should be identified as early as possible during the application design, avoiding the problems of employing cache only in late stages of the development cycle or even after it is in production, as an emergency solution. Furthermore, we observed that faster results can be achieved when a caching component is not implemented from scratch, but when libraries and frameworks are adopted to support cache implementation.

Although a well designed and implemented cache can achieve desired non-functional requirements, its performance may decay over time due to changes in the application domain, workload characteristics and access patterns. Therefore, achieving acceptable application performance requires constantly tuning cache decisions, which implies extra time dedicated to maintenance [31].

From all analyzed applications, we observed that none of them use a proactive approach to cache content. Content is always cached after it requested (i.e. reactive approach) and, as a consequence, the first request always results in a cache miss. Due to this, prefetching techniques can be used in order to populate the cache and prevent misses by predicting and caching data that will potentially be requested in the future. It can be based on heuristics, usage observations or even with the use of complex prediction algorithms [15], [32], [58], [59]. However, the design and implementation of a reactive cache component already requires significant effort and reasoning to be properly done, and a proactive approach increases the complexity of the caching solution even more.

Based on the results of our study, guidelines and patterns were derived to provide guidance to developers to deal with application-level caching. The former are high-level directions to develop caching solutions, and also serve as a checklist of points that must be analyzed while designing and implementing a caching component. The latter capture the reasoning used by developers to build a caching component and integrating it with a web application, providing a systematic way to develop a caching solution. Regardless of the use of our guidelines or patterns (or any other reusable component), it is important to determine exactly whether a particular caching approach is performing adequately. Therefore, it is essential to perform an application profiling, by measuring performance both with and without caching.

Our guidelines and patterns can be used as foundation to design, implement and manage a caching component for a particular application as well as to develop an application-independent caching component or framework. Future work involves the use of the guidance derived from our qualitative research to offer solutions to raise the abstraction level of application-level caching and, by distinguishing domain-neutral from domain-specific caching aspects, automate the reasoning captured in our patterns, providing a better overall experience with caching for developers.

## ACKNOWLEDGMENTS

Jhonny Mertz would like to thank CNPq grant 131550/2015-2. Ingrid Nunes thanks for research grants CNPq ref. 303232/2015-3, CAPES ref. 7619-15-4, and Alexander von Humboldt, ref. BRA 1184533 HFSTCAPES-P. We would also like to thank the editor of this journal and anonymous reviewers who have provided helpful comments and suggestions on the refinement of the paper.

## REFERENCES

- [1] S. Bouchenak, A. Cox, S. Dropsho, S. Mittal, and W. Zwaenepoel, "Caching Dynamic Web Content: Designing and Analysing an Aspect-Oriented Solution," *Proceedings of the ACM/IFIP/USENIX International Conference on Middleware*, vol. 4290, pp. 1–21, nov 2006. [Online]. Available: <http://dx.doi.org/10.1007/11925071>
- [2] J. Ravi, Z. Yu, and W. Shi, "A survey on dynamic Web content generation and delivery techniques," *Journal of Network and Computer Applications*, vol. 32, no. 5, pp. 943–960, sep 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.jnca.2009.03.005>
- [3] B. Atikoglu, Y. Xu, E. Frachtenberg, S. Jiang, and M. Paleczny, "Workload analysis of a large-scale key-value store," *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*, vol. 40, no. 1, p. 53, jun 2012. [Online]. Available: <http://dx.doi.org/10.1145/2318857.2254766>
- [4] D. R. K. Ports, A. T. Clements, I. Zhang, S. Madden, and B. Liskov, "Transactional Consistency and Automatic Management in an Application Data Cache," *Proceedings of the 9th USENIX Symposium on Operating Systems Design and Implementation*, pp. 279–292, oct 2010.
- [5] L. Della Toffola, M. Pradel, and T. R. Gross, "Performance problems you can fix: a dynamic analysis of memoization opportunities," in *Proceedings of the ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, vol. 50, no. 10. New York, New York, USA: ACM Press, oct 2015, pp. 607–622. [Online]. Available: <http://dx.doi.org/10.1145/2858965.2814290>
- [6] P. Gupta, N. Zeldovich, and S. Madden, "A trigger-based middleware cache for ORMs," in *Proceedings of the ACM/IFIP/USENIX International Middleware Conference*, vol. 7049 LNCS, 2011, pp. 329–349. [Online]. Available: [http://dx.doi.org/10.1007/978-3-642-25821-3\\_17](http://dx.doi.org/10.1007/978-3-642-25821-3_17)
- [7] M. P. Robillard, W. Coelho, and G. C. Murphy, "How effective developers investigate source code: An exploratory study," *IEEE Transactions on Software Engineering*, vol. 30, no. 12, pp. 889–903, 2004. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2016.2532873>
- [8] J. Sillito, G. C. Murphy, and K. De Volder, "Asking and answering questions during a programming change task," in *IEEE Transactions on Software Engineering*, vol. 34, no. 4. IEEE, jul 2008, pp. 434–451. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2008.26>
- [9] S. Nadi, T. Berger, C. Kästner, and K. Czarniecki, "Where do configuration constraints stem from? An extraction approach and an empirical study," *IEEE Transactions on Software Engineering*, vol. 41, no. 8, pp. 820–841, aug 2015. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2015.2415793>
- [10] J. Borstler and B. Paech, "The Role of Method Chains and Comments in Software Readability and Comprehension – An Experiment," *IEEE Transactions on Software Engineering*, pp. 1–1, 2016. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2016.2527791>
- [11] A. Namoun, A. Daskalopoulou, N. Mehandjiev, and Z. Xun, "Exploring Mobile End User Development: Existing Use and Design Factors," *IEEE Transactions on Software Engineering*, pp. 1–1, 2016. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2016.2532873>
- [12] M. Jorgensen, "Evidence-based guidelines for assessment of software development cost uncertainty," *IEEE Transactions on Software Engineering*, vol. 31, no. 11, pp. 942–954, nov 2005. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2005.128>
- [13] N. Juristo, A. M. Moreno, and M. I. Sanchez-Segura, "Guidelines for eliciting usability functionalities," *IEEE Transactions on Software Engineering*, vol. 33, no. 11, pp. 744–758, nov 2007. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2007.70741>
- [14] L. Carvajal, A. M. Moreno, M. I. Sánchez-Segura, and A. Seffah, "Usability through software design," *IEEE Transactions on Software Engineering*, vol. 39, no. 11, pp. 1582–1596, nov 2013. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2013.29>
- [15] S. Podlipnig and L. Böszörményi, "A Survey of Web Cache Replacement Strategies," *ACM Computing Surveys*, vol. 35, no. 4, pp. 374–398, dec 2003. [Online]. Available: <http://dx.doi.org/10.1145/954339.954341>
- [16] K. S. Candan, W.-S. Li, Q. Luo, W.-P. Hsiung, and D. Agrawal, "Enabling dynamic content caching for database-driven web sites," *Proceedings of the ACM SIGMOD International Conference on Management of Data*, vol. 30, no. 2, pp. 532–543, jun 2001. [Online]. Available: <http://dx.doi.org/10.1145/376284.375736>
- [17] L. Ramaswamy, L. Liu, A. Iyengar, and F. Douglass, "Automatic fragment detection in dynamic web pages and its impact on

- caching," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 6, pp. 859–874, jun 2005. [Online]. Available: <http://dx.doi.org/10.1109/TKDE.2005.89>
- [18] L. Li, C. Niu, H. Zheng, and J. Wei, "An adaptive caching mechanism for Web services," in *Proceedings of the International Conference on Quality Software*. IEEE, oct 2006, pp. 303–310. [Online]. Available: <http://dx.doi.org/10.1109/QSIC.2006.9>
- [19] C. Guerrero, C. Juiz, and R. Puigjaner, "Improving web cache performance via adaptive content fragmentation design," in *Proceedings of the IEEE International Symposium on Network Computing and Applications*. IEEE, aug 2011, pp. 310–313. [Online]. Available: <http://dx.doi.org/10.1109/NCA.2011.55>
- [20] G. Soundararajan and C. Amza, "Using semantic information to improve transparent query caching for dynamic content web sites," in *Proceedings of the International Workshop on Data Engineering Issues in E-Commerce*, vol. 2005. IEEE, 2005, pp. 132–138. [Online]. Available: <http://dx.doi.org/10.1109/DEEC.2005.25>
- [21] C. Amza, G. Soundararajan, and E. Cecchet, "Transparent caching with strong consistency in dynamic content web sites," in *Proceedings of the 19th annual international conference on Supercomputing*. New York, New York, USA: ACM Press, jun 2005, p. 264. [Online]. Available: <http://dx.doi.org/10.1145/1088149.1088185>
- [22] R. Baeza-Yate, F. Junqueira, V. Plachouras, and H. Witschel, "Admission Policies for Caches of Search Engine Results," in *Proceedings of the 14th International Conference on String Processing and Information Retrieval*, vol. 4726. Springer-Verlag, oct 2007, pp. 74–85. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-75530-2\\_7](http://dx.doi.org/10.1007/978-3-540-75530-2_7)
- [23] H. Ma, W. Liu, B. Wei, L. Shi, X. Bao, L. Wang, and B. Wang, "Paap," in *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*. New York, New York, USA: ACM Press, jul 2014, pp. 983–986. [Online]. Available: <http://dx.doi.org/10.1145/2600428.2609490>
- [24] P. Larson, J. Goldstein, and J. Zhou, "MTCache: Transparent mid-tier database caching in SQL server," *Proceedings of the International Conference on Data Engineering*, vol. 20, pp. 177–188, mar 2004. [Online]. Available: <http://dx.doi.org/10.1109/ICDE.2004.1319994>
- [25] W. Wang, Z. Liu, Y. Jiang, X. Yuan, and J. Wei, "EasyCache: a transparent in-memory data caching approach for internetware," in *Proceedings of the 6th Asia-Pacific Symposium on Internetware on Internetware*. New York, New York, USA: ACM Press, nov 2014, pp. 35–44. [Online]. Available: <http://dx.doi.org/10.1145/2677832.2677837>
- [26] "Memcached," 2016. [Online]. Available: <https://memcached.org/>
- [27] "Redis," 2016. [Online]. Available: <http://redis.io/>
- [28] D. Waddington, J. Colmenares, J. Kuang, and F. Song, "KV-Cache: A scalable high-performance web-object cache for manycore," in *Proceedings of the IEEE/ACM 6th International Conference on Utility and Cloud Computing*. IEEE, dec 2013, pp. 123–130. [Online]. Available: <http://dx.doi.org/10.1109/UCC.2013.34>
- [29] K. Nguyen and G. Xu, "Cachetor: detecting cacheable data to remove bloat," in *Proceedings of the 9th Joint Meeting on Foundations of Software Engineering*. New York, New York, USA: ACM Press, aug 2013, p. 268. [Online]. Available: <http://dx.doi.org/10.1145/2491411.2491416>
- [30] A. Infante, "Identifying caching opportunities, effortlessly," in *Companion Proceedings of the 36th International Conference on Software Engineering*. New York, New York, USA: ACM Press, may 2014, pp. 730–732. [Online]. Available: <http://dx.doi.org/10.1145/2591062.2591198>
- [31] G. Radhakrishnan, "Adaptive application caching," *Bell Labs Technical Journal*, vol. 9, no. 1, pp. 165–175, may 2004. [Online]. Available: <http://dx.doi.org/10.1002/bltj.20011>
- [32] W. Ali, S. M. Shamsuddin, and A. S. Ismail, "A survey of Web caching and prefetching," *International Journal of Advances in Soft Computing and Its Applications*, vol. 3, no. 1, pp. 18–44, mar 2011. [Online]. Available: [http://eprints.utm.my/36316/2/IJASCA15\\_A-Survey-of-Web-Caching.pdf](http://eprints.utm.my/36316/2/IJASCA15_A-Survey-of-Web-Caching.pdf)
- [33] S. Surapaneni, V. K. S. Nerella, S. K. Madria, and T. Weigert, "Exploring caching for efficient collection operations," in *Proceedings of the 26th IEEE/ACM International Conference on Automated Software Engineering*. IEEE, 2011, pp. 468–471. [Online]. Available: <http://dx.doi.org/10.1109/ASE.2011.6100101>
- [34] M. Fiore, F. Mininni, C. Casetti, and C. F. Chiasserini, "To cache or not to cache?" in *Proceedings of the IEEE INFOCOM*. New York, New York, USA: ACM Press, may 2009, pp. 235–243. [Online]. Available: <http://dx.doi.org/10.1109/INFCOM.2009.5061926>
- [35] C. Guerrero, I. Lera, and C. Juiz, "Performance improvement of web caching in Web 2.0 via knowledge discovery," *Journal of Systems and Software*, vol. 86, no. 12, pp. 2970–2980, dec 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.jss.2013.04.060>
- [36] G. Einziger and R. Friedman, "TinyLFU : A Highly Efficient Cache Admission Policy," in *Proceedings of the 22nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*. IEEE, feb 2014, pp. 146–153. [Online]. Available: <http://dx.doi.org/10.1109/PDP.2014.34>
- [37] Nimrod Megiddo and Dharmendra S. Modha, "ARC: A Self-Tuning, Low Overhead Replacement Cache," *Proceedings of the 2nd USENIX conference on File and Storage Technologies*, vol. 3, pp. 115–130, mar 2003.
- [38] Q. Yang and H. H. Zhang, "Web-log mining for predictive web caching," *IEEE Transactions on Knowledge and Data Engineering*, vol. 15, no. 4, pp. 1050–1053, jul 2003. [Online]. Available: <http://dx.doi.org/10.1109/TKDE.2003.1209022>
- [39] R. Subramanian, Y. Smaragdakis, and G. H. Loh, "Adaptive caches: Effective shaping of cache behavior to workloads," in *Proceedings of the Annual International Symposium on Microarchitecture*. IEEE, dec 2006, pp. 385–396. [Online]. Available: <http://dx.doi.org/10.1109/MICRO.2006.7>
- [40] S. Ghandeharizadeh, S. Irani, J. Lam, and J. Yap, "Camp," in *Proceedings of the 15th International Conference on Middleware*. New York, New York, USA: ACM Press, dec 2014, pp. 289–300. [Online]. Available: <http://dx.doi.org/10.1145/2663165.2663317>
- [41] C. Li and A. L. Cox, "GD-Wheel," in *Proceedings of the 10th European Conference on Computer Systems*. New York, New York, USA: ACM Press, apr 2015, pp. 1–15. [Online]. Available: <http://dx.doi.org/10.1145/2741948.2741956>
- [42] X. Hu, X. Wang, Y. Li, L. Zhou, Y. Luo, C. Ding, S. Jiang, C. Ding, and Z. Wang, "LAMA : Optimized Locality-aware Memory Allocation for Key-value Cache," in *Proceedings of the USENIX Annual Technical Conference*, 2015, pp. 57–69.
- [43] X. Qin, W. Wang, W. Zhang, J. Wei, X. Zhao, H. Zhong, and T. Huang, "PRESC 2: Efficient self-reconfiguration of cache strategies for elastic caching platforms," *Computing*, vol. 96, no. 5, pp. 415–451, nov 2014. [Online]. Available: <http://dx.doi.org/10.1007/s00607-013-0365-6>
- [44] J. H. Wood and Timothy, "Adaptive Performance-Aware Distributed Memory Caching," in *Proceedings of the 10th International Conference on Autonomic Computing*, 2013, pp. 33–43.
- [45] C. Pohl, "Adaptive caching of distributed components," Ph.D. dissertation, Dresden University of Technology, jan 2005.
- [46] A. P. Negrão, C. Roque, P. Ferreira, and L. Veiga, "An adaptive semantics-aware replacement algorithm for web caching," *Journal of Internet Services and Applications*, vol. 6, no. 1, p. 4, feb 2015. [Online]. Available: <http://dx.doi.org/10.1186/s13174-015-0018-4>
- [47] B. Glaser, *Basics of grounded theory analysis.*, 1992.
- [48] V. R. Basili, R. W. Selby, and D. H. Hutchens, "Experimentation in software engineering," *IEEE Transactions on Software Engineering*, vol. SE-12, no. 7, pp. 733–743, jul 1986. [Online]. Available: <http://dx.doi.org/10.1109/TSE.1986.6312975>
- [49] V. R. Basili, G. Caldiera, and H. D. Rombach, "The goal question metric approach," *Encyclopedia of Software Engineering*, vol. 2, pp. 528–532, 1994. [Online]. Available: <http://dx.doi.org/10.1.1.104.8626>
- [50] B. G. Glaser and A. L. Strauss, "The discovery of grounded theory," *International Journal of Qualitative Methods*, vol. 5, pp. 1–10, 1967. [Online]. Available: <http://dx.doi.org/10.2307/588533>
- [51] "Spring Caching," 2016. [Online]. Available: <https://spring.io/guides/gs/caching/>
- [52] "EhCache," 2016. [Online]. Available: <http://www.ehcache.org/>
- [53] "Infinispan," 2016. [Online]. Available: <http://infinispan.org/>
- [54] "Rails low-level caching," 2016. [Online]. Available: [http://guides.rubyonrails.org/caching\\_with\\_rails.html](http://guides.rubyonrails.org/caching_with_rails.html)
- [55] "Google Guava," 2016. [Online]. Available: <https://github.com/google/guava>
- [56] "Caffeine," 2016. [Online]. Available: <https://github.com/ben-manes/caffeine>
- [57] J. Mertz and I. Nunes, "A Qualitative Study of Application-level Caching," 2016. [Online]. Available: <http://www.inf.ufrgs.br/~jmamertz/caching-patterns/>

- [58] J. Domènech, J. A. Gil, J. Sahuquillo, and A. Pont, "Web prefetching performance metrics: A survey," *Performance Evaluation*, vol. 63, no. 9-10, pp. 988–1004, oct 2006. [Online]. Available: <http://dx.doi.org/10.1016/j.peva.2005.11.001>
- [59] M. B. Pal and D. C. Jain, "Web service enhancement using web pre-fetching by applying markov model," in *Proceedings of the 4th International Conference on Communication Systems and Network Technologies*. IEEE, apr 2014, pp. 393–397. [Online]. Available: <http://dx.doi.org/10.1109/CSNT.2014.84>



**Jhonny Mertz** is an MSc student in computer science at the Federal University of Rio Grande do Sul (UFRGS), in Brazil. He obtained his BSc from the Universidade Estadual do Oeste do Paraná (UNIOESTE), Brazil, in 2012. His research interests include software maintenance and evolution, application-level caching and autonomic computing. See <http://www.inf.ufrgs.br/~jmamertz> for more information.



**Ingrid Nunes** is an Associate Professor at the Institute of Informatics, Universidade Federal do Rio Grande do Sul (UFRGS), Brazil, currently in a sabbatical year at TU Dortmund in Germany. She obtained her PhD in Informatics at the Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Brazil. Her PhD was in cooperation with King's College London (UK) and University of Waterloo (Canada). She is the head of the Prosoft research group, and her main research areas are software maintenance and evolution and agent-oriented software engineering.