

Performance Analysis for Object-Oriented Software: A Systematic Mapping

David Maplesden, Ewan Tempero, *Member, IEEE*, John Hosking, *Member, IEEE*,
and John C. Grundy, *Member, IEEE*

Abstract—Performance is a crucial attribute for most software, making performance analysis an important software engineering task. The difficulty is that modern applications are challenging to analyse for performance. Many profiling techniques used in real-world software development struggle to provide useful results when applied to large-scale object-oriented applications. There is a substantial body of research into software performance generally but currently there exists no survey of this research that would help identify approaches useful for object-oriented software. To provide such a review we performed a systematic mapping study of empirical performance analysis approaches that are applicable to object-oriented software. Using keyword searches against leading software engineering research databases and manual searches of relevant venues we identified over 5000 related articles published since January 2000. From these we systematically selected 253 applicable articles and categorised them according to ten facets that capture the intent, implementation and evaluation of the approaches. Our mapping study results allow us to highlight the main contributions of the existing literature and identify areas where there are interesting opportunities. We also find that, despite the research including approaches specifically aimed at object-oriented software, there are significant challenges in providing actionable feedback on the performance of large-scale object-oriented applications.

Index Terms—Systematic review, survey, performance, object-oriented

1 INTRODUCTION

SOFTWARE performance has been of interest to researchers and practitioners since the earliest days of computing but it is still highly relevant today. The performance of object-oriented software is of particular interest because object-oriented languages and techniques are ubiquitous in industry, underlined by the popularity of object-oriented languages such as Java, C++ and C#. There is a significant collection of published literature relating to software performance, a large subset of which is applicable to object-oriented software. However, to the best of our knowledge, there exists no previous attempt to survey the literature in this field. Such a survey would be a valuable contribution for researchers and practitioners looking to understand the existing research in the field either for the purposes of leveraging or contributing to that research. To that end we have undertaken an in-depth systematic mapping study, a form of systematic literature review [1], of the field.

Software efficiency is still important for today's applications despite the prodigious improvement in hardware

performance over the last three decades. There are several trends in application development that contribute to a on-going need for efficient software. The size and complexity of software has increased at a similar or even greater rate than hardware has advanced [2]. There are increasing numbers of mobile applications being created which must run on devices with limited resources. Also many applications are delivered using cloud deployment models where running costs are directly impacted by software efficiency. The challenge is that the growing scale and complexity of the software under development means that analysing and improving the performance of these systems has become increasingly difficult.

There are specific challenges for object-oriented applications because they have characteristics that make performance analysis difficult. Following object-oriented principles tends to lead to applications with inter-procedural rather than intra-procedural control flow and a great number of methods. Additionally many object-oriented methodologies focus on developer productivity, producing maintainable and flexible software, and promoting componentisation and reuse. As a result most applications are built from reusable generalised frameworks and leverage established design patterns, making them very layered and complex. This approach means that the handling of even the simplest request in these framework-based applications goes through many layers and will require hundreds, maybe thousands, of method calls to complete [3]. This excessive activity to achieve seemingly simple results is a problem that has become

- David Maplesden and Ewan Tempero are with the Department of Computer Science, The University of Auckland, Private Bag 92019, Auckland 1142, New Zealand. E-mail: dmap001@aucklanduni.ac.nz, e.tempero@cs.auckland.ac.nz.
- John Hosking is with the Faculty of Science, The University of Auckland, Private Bag 92019, Auckland 1142, New Zealand. E-mail: j.hosking@auckland.ac.nz.
- John C. Grundy is with the Faculty of Information and Communication Technologies, Swinburne University of Technology, PO Box 218, Hawthorn, Victoria 3122, Australia. E-mail: jgrundy@swin.edu.au.

Manuscript received May 23, 2014, Revised Nov 2014, Accepted Jan 2015.

known as *runtime bloat* [4], and it has led many large scale object-oriented applications to suffer from chronic performance problems [5].

The scope of this survey then is performance analysis approaches that are *applicable* to object-oriented software. We have not limited the scope to just approaches specific to object-oriented software but have also included approaches that are applicable to most applications and are therefore likely to be useful when analysing a typical object-oriented application. The aim is provide a survey of techniques that are relevant both for practitioners who develop object-oriented applications as well as researchers interested in large-scale object-oriented applications and runtime bloat.

The survey covers empirical performance analysis approaches described in the literature since January 2000. We have chosen to cover empirical approaches only, and exclude model-based approaches (discussed in Section 2.1), to keep the survey to a manageable size and because empirical approaches are prevalent in industry [6]. We chose January 2000 as the start date as this represented a practical starting point for the review, it was far enough in the past to include the majority of relevant approaches but recent enough that all relevant published literature was indexed and available within the major electronic databases.

The remainder of this article is organised as follows: in the next section we describe in more detail the domain of our survey and summarise the most closely related work that we did *not* include in our formal survey. Section 3 describes our methodology for the review, including our specific research questions and the details of the categorisation we used for the data extraction. Section 4 contains the results of our data extraction and conclusions for our research questions. Section 5 is a discussion of the results. Section 6 evaluates the threats to validity of our survey and we conclude this article in section 7.

2 BACKGROUND AND MOTIVATION

2.1 Software Performance Engineering

Many different approaches are used when seeking to improve system performance. These include research into microprocessor or computer architecture design (i.e. hardware performance), static optimisation techniques (usually compiler based), dynamic binary translation or dynamic optimisation systems (which look to automatically tune the executing code at runtime) and software engineering approaches (which target improvements in the software architecture or source code implementation). Generally speaking each of these approaches is complementary to the others. They tend to target different types of inefficiency, and each can produce useful performance improvements independently of the others.

The realm of software engineering approaches to improving performance is known as *software performance*

engineering. Woodside, Franks & Petriu [6] give the following definition:

Software Performance Engineering (SPE) represents the entire collection of software engineering activities and related analyses used throughout the software development cycle, which are directed to meeting performance requirements.

Woodside et al also distinguish between two distinct approaches to SPE: model-based predictive performance engineering and empirical performance analysis. Often the term software performance engineering is used to refer only to model-based approaches rather than the broader definition given by Woodside et al.

Model-based approaches, made popular by Smith [7], consist of constructing and solving theoretical models of a system to predict its performance characteristics. There are a variety of different mathematical models used to support model-based performance prediction including queueing networks, petri nets and stochastic process algebras as surveyed by Balsamo et al [8].

By contrast empirical performance analysis approaches are based on the analysis of concrete measurements taken from running software and include the traditional application profiling and tuning approaches.

Model-based approaches and empirical analysis are generally accepted as complementary as they tend to target different categories of performance problem i.e. structural or architectural versus implementation respectively.

2.2 Challenges of Empirical Performance Analysis

Empirical performance analysis approaches record detailed data concerning the dynamic behaviour of a running system. With the increasing scale of dynamic behaviour in modern systems the amount of data being produced has become overwhelming. D'Elia, Demetrescu & Finocchi [9] report results for short runs of a variety of off-the-shelf applications in a typical Linux distribution, some of which yield statistics for millions of different active code paths. This amount of data is very difficult to interpret manually therefore we are interested in the form of feedback provided by empirical performance analysis approaches and the techniques they use to aid with interpreting performance data. In particular we are looking for approaches that aim to provide *actionable* feedback, information that provides specific advice on what to change, and even how to make the change to improve the performance of an application.

It is generally held that performance analysis is challenging and consequently a significant number of projects suffer from performance problems [5], [6]. Therefore we are interested in how effective a performance analysis approach is at helping to improve performance. To that end we are interested in how the approaches described in the literature are evaluated.

It is also relevant how practical it is to apply an approach, as this can impact the utility of the approach. In order to obtain useful results it is important that the dynamic behaviour being measured is representative of typical behaviour. Much of today's software runs in *live production environments*. These are systems (such as web applications) which are required to be always online and service users' requests in real-time. Approaches that can be directly used in these live production systems have an advantage because they can measure the target behaviour directly. Approaches that cannot be used in production environments, because they rely on specialised customisations or impose too much overhead, must instead use a representative test system to gather their measurements. This can be a challenge because for many systems replicating realistic environment, data and load conditions is very difficult.

2.3 Related Surveys

There are three existing survey papers published since the year 2000 that discuss literature in some area of software performance [4], [8], [10]. None of the three survey papers we found were systematic literature reviews. Two were traditional expert reviews of a particular specialised field and one was a position paper that included a survey of current research into runtime bloat.

Koziolok [10] presented a survey on the performance evaluation of component-based software systems. The survey was limited to discussing approaches for component-based software, that is software based on frameworks such as Java EJB, Microsoft COM or CORBA CCM, but did cover both empirical measurement based and predictive model-based approaches. The survey contains a detailed breakdown and analysis of the literature covered with useful information and conclusions for both practitioners and researchers, however with its focus on component-based software and inclusion of model-based approaches it only covers a small subset of the research we are interested in.

Xu et al [4] summarised the state of research into runtime bloat. The focus of this position paper was to describe software bloat and to argue why it is primarily a software engineering problem. The paper was motivated by many of the same concerns that have motivated our research. It surveys existing research into runtime bloat and outlines possible future research directions. However currently there is only a small body of work that identifies itself as investigating software bloat and therefore the paper only covers a small amount of the literature we are interested in that is applicable to object-oriented software performance more generally.

Balsamo et al [8] surveyed model-based performance prediction in software development. The survey completes a detailed description, classification and analysis of 16 different integrated methods for model-based software performance prediction. The goal for the survey is to assess the maturity of the research in the field and

point out promising research directions. The survey has a different focus from ours and they do not consider empirical performance analysis approaches.

2.4 Related Literature

Besides the related surveys described above there is a great deal of literature related to software performance we did not include in our review. Here we provide a description of the major research areas that we excluded from our survey.

As we noted in section 2.1 our focus on empirical performance engineering has meant we have excluded literature related specifically to model-based SPE approaches. These are performance prediction approaches based on constructing and solving mathematical models of the running system. There is a large body of literature describing how to create or improve useful mathematical models, discussing new more efficient algorithms for solving the created models or describing tools and methodologies for applying the approaches in practice.

Because of our focus on object-oriented software there is literature that we excluded because it had a specific focus on some other specialised field. This included performance analysis approaches for:

- Real-time systems
- Embedded systems
- Distributed systems
- High Performance Computing (HPC) and massively parallel systems

Performance analysis approaches focussed in these areas are concerned with aspects of performance that are not applicable to object-oriented software generally. For example embedded systems tend to have severely constrained hardware requirements and specialised runtime environments that make performance monitoring difficult, and performance analysis for distributed systems usually focuses on remote communication patterns. We also excluded a substantial amount of research into performance for HPC and massively parallel systems that focuses on the challenges unique to supercomputing and cluster computing scenarios. This includes research into understanding and improving parallelisation and inter-node communication, computer architecture related inefficiencies such as non-uniform memory access, and approaches for processing, analysing and visualising the enormous quantities of performance data that are generated by high-end systems.

Finally we excluded a range of research that leveraged dynamic analysis of software but had goals other than performance understanding. These goals included:

- General purpose program comprehension or reverse engineering
- Automated verification or validation, defect detection
- Workload characterisation, normally for automated test workload generation
- Monitoring for intrusion detection

- Monitoring for quality of service violations or capacity planning purposes

All of these use runtime monitoring or profiling systems that are similar in nature to the dynamic data capture systems that are used for performance analysis.

3 METHODOLOGY

3.1 Overview

A systematic literature review (SLR) is described by Kitchenham & Charters [1] as *"a methodologically rigorous review of research results"*. A systematic mapping study is a form of SLR that aims to give a broader overview of a particular field. It does not evaluate the articles in as much depth as an SLR, with the advantage that a broader range of primary studies is covered [11].

We have undertaken a systematic mapping study of empirical performance analysis approaches applicable to object-oriented software. In general we followed the SLR guidelines produced by Kitchenham & Charters [1], [12] and also incorporated some of the recommendations given by Petersen et al [11].

The high level steps in our review process were:

- 1) Define research questions
- 2) Perform manual search to pilot inclusion/exclusion criteria and generate a reference set of articles
- 3) Develop automated search strategy
- 4) Formalise review protocol
- 5) Conduct search for relevant studies
- 6) Screen and select studies for inclusion
- 7) Data extraction
- 8) Analysis

The manual search at step 2 was a necessary addition for us to the typical SLR process. As described in section 3.3 it assisted us with the development of the inclusion criteria and the automated search phrases that were required to create the formal review protocol.

The other departure from recommended SLR procedure was that we did not perform any quality assessment on the primary studies. This is customary for mapping studies aiming to structure the literature in an entire field in order to include as much relevant literature as possible [11].

3.2 Research Questions

The focus of our systematic review is empirical software engineering approaches to performance analysis for object-oriented software. Our research questions are motivated by this focus and by the challenges in empirical performance analysis discussed in section 2.2:

- 1) What approaches to empirical performance analysis have been proposed that are applicable to object-oriented software?
- 2) How can these approaches be characterised?
- 3) What form of feedback does the approach provide?
- 4) How practical to apply are the approaches?
 - a) What is required to apply the approach?

- b) Can the approaches be applied to live production environments?

- 5) How are approaches evaluated and validated?

3.3 Initial Manual Search

In order to develop our review protocol we needed to develop both the automated search phrase and the inclusion criteria used to select the primary studies. To do this we undertook a manual search phase against selected venues where we could not only pilot and test our inclusion criteria but also generate a reference set of results with which to test and refine our automated searches. Ideally such a reference set would come from an independent source, such as a previously published literature review, but we had not been able to identify a suitable existing source.

We conducted our initial manual search against:

- **ICPE** - ACM/SPEC International Conference on Performance Engineering - 2012, 2011, 2010
- **WOSP** - International Workshop on Software and Performance - 2008, 2007
- **Performance Evaluation** - Journal - 2012, 2011, 2010, 2009

We then expanded our reference set by screening the papers referenced from the selected papers (reference snowballing) using our inclusion/exclusion criteria until we had an initial reference set of 46 papers.

3.4 Search Phrase Development

The search phrase was developed based on search terms extracted from the abstracts, titles and keywords from the reference set of articles and expanded with synonyms.

The development of the search phrase was challenging due to the generic nature of the obvious individual search terms: 'software', 'performance' - using these terms resulted in searches that returned many thousands of results i.e. low precision. There is also a plethora of synonyms for software performance, meaning that each individual search term has only a very low recall.

After extensive experimentation we found an effective approach was to use a search phrase made up of two parts. The first part aimed at returning approaches to finding or detecting performance problems:

```
( "detecting" OR
  "finding" OR
  "profiling" )
AND
( "performance problems" OR
  "performance issues" OR
  "performance bugs" OR
  "performance bottlenecks" OR
  "performance antipatterns" OR
  "performance anti-patterns" OR
  "object churn" OR
  "memory bloat" OR
```

```
"runtime bloat" )
```

The second part of the search criteria is aimed finding approaches to profiling code, because profiling based approaches are such a ubiquitous method of understanding program performance:

```
( "profiling" OR "profiler" )
AND
( "calling context" OR
  "path profiling" OR
  "call-path" OR
  "call-graph" OR
  "call-tree" OR
  "calltree" OR
  "instrumentation sampling" OR
  "sampling based" OR
  "instrumentation based" OR
  "performance analysis" )
```

We developed the search phrase to be executed against the title, abstract and keywords of the articles in each electronic database. The development of our search phrase was completed using the Scopus database. Scopus has comprehensive searching functionality and indexes all of the venues from the initial manual search, making it convenient to evaluate the effectiveness of our searches by comparison against our initial reference set.

This search returned 1086 results and contained 38/46 = 82% of our reference set. Ideally we would have preferred to achieve a higher recall than 82%. We were only able to improve the recall by either including much more generic search terms, which vastly expanded the number of returned results, or by using artificially specific search terms that were derived from the remaining papers and did not return any other relevant results.

3.5 Venues

To select the electronic venues in which to conduct our search we performed a survey of the search engines used in previous systematic literature reviews and screened them for suitability for our review. Zhang *et al* [13] reported the results of a similar survey. From these surveys and experiences reported by Dybå *et al* [14] and Petersen [15] the final list of electronic databases we chose to search was:

- ACM digital library
- IEEE Xplore Digital Library
- ScienceDirect
- SpringerLink
- SCOPUS
- ISI Web of Science

The other notable electronic databases that were potential candidates for inclusion that we did not include were:

- Wiley InterScience (now Wiley Online)
- Kluwer Online
- EI Compendex
- Inspec

These were excluded because of their high degree of overlap with one or more of the included search venues.

- Kluwer Online has been merged with Springer and is now indexed through the SpringerLink database.
- EI Compendex and Inspec are indexing services that have a high degree of overlap with the SCOPUS and ISI Web of Science indexing services. For example statistics from the JISC Academic Database Assessment Tool (<http://adat.crl.edu/>) indicate that 3194 of the 3610 titles covered by Compendex and 3142 of the 4748 titles covered by Inspec are also covered by SCOPUS.

3.6 Article Screening

To build our list of included articles we screened the automated search results from each of the venues searched and applied the inclusion and exclusion criteria listed in Table 1.

Articles were screened by following an adaptive reading depth approach [11], that is the article title and abstract were read first and if possible the article was included or excluded based on the information therein. Otherwise the full text of the article was retrieved and further sections of the article were read as necessary in order to make the inclusion determination.

All articles were screened by the lead author. To help insure the consistency of the screening process two of the other authors independently tested the inclusion and exclusion criteria on an initial set of 100 papers randomly selected from the automated search results. This initial testing showed a generally high level of consistency on the screening results, with the two other authors agreeing with the lead author on 81 and 86 of the 100 papers respectively. The discrepancies were discussed and used to refine the inclusion criteria used for the full screening.

Duplicate papers returned by the searches were excluded. Many papers are indexed in multiple electronic databases and some appear in multiple venues. We detected duplicate papers by manually comparing any papers where the title and lead author matched. There were some examples of similar, but not identical, papers from the same authors appearing in different venues. In this situation we included both papers. Often systematic reviews avoid this by reporting on the underlying primary studies rather than the papers that have been published. However in a mapping study of this size it was not practical to identify the underlying primary study for each paper, so we report simply on the number of papers we found.

3.6.1 Inclusion and Exclusion Criteria

Our inclusion criteria reflect the emphasis of our research questions from section 3.2, in particular:

- We are interested in empirical rather than model-based performance engineering

TABLE 1
Inclusion and Exclusion Criteria

Inclusion Criteria	Exclusion Criteria
<ul style="list-style-type: none"> • Scientific literature written in English and published in peer-reviewed venues in the year 2000 or later • Literature that describes empirical approaches to analysing and understanding performance applicable to object-oriented software • Papers about finding performance problems, bugs or anti-patterns applicable to object-oriented software • Papers about profiling for <i>performance understanding</i> applicable to object-oriented software 	<ul style="list-style-type: none"> • Non-empirical (i.e. model-based) performance engineering techniques • Performance optimisations applied by compilers, virtual machines, middleware or hardware • Profiling for purposes other than performance understanding e.g. <ul style="list-style-type: none"> - reverse engineering software architecture models - program comprehension - workload characterisation and generation - bug detection e.g. deadlocks, memory leaks - intrusion detection • Performance monitoring approaches not focussed on improvement e.g. <ul style="list-style-type: none"> - detecting quality of service violations - performance prediction - service provisioning - capacity planning • Approaches specific to domains other than object-oriented software e.g. <ul style="list-style-type: none"> - embedded systems - real-time systems - massively parallel HPC systems - big data systems such as MapReduce applications

- We are not interested in automated performance optimisations applied by compilers, virtual machines, middleware or hardware
- We are focussed on approaches applicable to object-oriented software in general and so excluded approaches specific to specialised domains such as embedded, real-time or massively parallel systems

We regard each of these excluded areas of research as being distinct from and complementary to the research we are focussed on. We discussed these areas of research and how they relate to our area of interest in section 2.4.

Note that this exclusion of approaches *specific* to other domains does not exclude approaches from those domains that are more general in nature and therefore still applicable to object-oriented software.

3.7 Completing article selection

To complete article selection, after the screening of the automated searches, we undertook a final manual search phase. This consisted of two parts, manual searches of selected venues and *reference snowballing* — manually checking paper references for other relevant literature. We felt this final manual search phase was necessary to mitigate the risk of missing relevant literature due to the only moderate recall of the automated searches. We discuss this further in section 6.

Before the manual search phase we had just over 200 included articles from 100 different venues, including 66 venues with only one article. We chose to perform manual searches against the top 5 venues which each had contributed more than 10 articles. No other venue had contributed more than 6. These venues were:

- **OOPSLA** - ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications
- **CGO** - International Symposium on Code Generation and Optimization
- **PLDI** - ACM SIGPLAN Conference on Programming Language Design and Implementation
- **ECOOP** - European Conference on Object-Oriented Programming
- **Euro-Par** - International Euro-Par Conference on Parallel Processing

For each of these venues we completed a manual search on all published proceedings from 2000 to 2013. The CGO conference only began in 2003 but otherwise each of these venues had one published proceedings each year. In total our manual venue search covered 66 published proceedings.

3.8 Search Results

The results of our systematic search process are summarised in Table 2. We ran the same logical query against each electronic database but the syntax used is different for each venue due to their different search engine implementations. We report the exact queries we used to allow them to be easily reproduced. For each search we completed the table lists:

- The exact search phrase that was used
- The date the final search was executed
- The total number of results returned by the search
- The number of results that were selected for inclusion after screening

- The number of included results that were unique to that search

Note that because of the overlap in the results returned by the searches (many papers were returned by more than one search) the total number of included results is not equal to the sum of the included results from the individual searches.

The initial automated searches were completed from May to August 2013 but we refreshed all our automated search results in December 2013 when we completed the final manual phases of our search.

There were some particular challenges that we faced with the SpringerLink search. The SpringerLink database search engine did not support searching against the title, abstract and keywords only. It supported either searching against title only (which returned very few results) or against the full text of the articles, including references. This meant that the SpringerLink search returned 26,677 results. In order to practically screen such a large number of results we reverted to screening them in relevance order as ranked by the SpringerLink search engine and stopping once the rate at which we were finding new articles for inclusion had reduced to the point where we felt finding any further articles was unlikely. In practice we screened the first 2000 results from the search and then stopped because there had been no new included articles in the last 500 results.

Excluding SpringerLink, we had 2779 automated search results. Including the 2000 results we screened from the SpringerLink search we screened 4779 automated search results and from these included 202 articles. The manual searches covered 66 individual conference proceedings and resulted in 32 additional articles. The reference snowballing involved screening many hundreds of references and resulted in 19 additional articles. The final count was 253 included papers, 202 from the screened automated searches, 32 from the manual search and 19 from reference snowballing.

3.9 Search Results Discussion

We will leave the majority of our analysis of the literature until we discuss the results of the data extraction and the threats to validity for our study but there are some interesting conclusions that can be drawn from the search results we have reported in Table 2.

The first is that each phase of the search (the automated searches, the manual search and the reference snowballing) was a valuable part of the search process. Each led us to discover and include literature that would not otherwise have been covered. That this was true for our mapping study does not necessarily mean it would be true for all systematic reviews. Reviews with a narrower focus may find that either the automated search phase or the manual search phase is superfluous, but that was certainly not our experience.

What was interesting was the very small number of unique search results returned by some of the auto-

mated searches. Indeed the ACM, IEEE and ScienceDirect searches together only returned 6 unique results not returned by either the SCOPUS or Web of Science index searches. The implication of this is that the manual search and reference snowballing phases were more productive and more critical to the coverage of our literature review than the automated searches against those databases.

3.10 Data extraction

Data extraction was completed by reviewing each included article and categorising it according to a classification scheme that we developed to allow us to structure and map the literature in our review. We used an adaptive reading depth approach, similar to our approach when screening articles for inclusion, beginning with reading the title, abstract and introduction and then reading further sections as required to complete our classification.

3.11 Classification Scheme

Our classification scheme allows us to structure the literature in our study to map the literature in general and answer our research questions in particular. The classification scheme we have developed is a novel contribution in its own right, providing a framework for categorising and describing performance engineering approaches applicable to object-oriented software. Our scheme consists of 10 facets, shown in Figure 1.

Our facets were selected to capture both attributes that are useful to help us answer our research questions (such as Feedback and Probe Implementation) and attributes that represent commonly recognised aspects of performance analysis approaches (e.g. Analysis Type and Probe Type). We did not predetermine the categories for each facet ahead of time, instead using a *keywording* approach similar to that described by Petersen et al [11]. We classified each paper's approach using key words or phrases for each facet that describe the characteristics of the approach. We then consolidated the list of key phrases that had been used for each facet into a coherent set of categories, generalising sets of phrases into a single category where necessary.

3.11.1 Contribution

The contribution of a paper is the type of novel contribution it is providing to the field of research. Capturing the contribution allows us to compare papers that broadly had the same objectives. Our contribution facet contained 8 categories in total that were grouped into 3 broader categories:

- Profiling infrastructure papers describe a general purpose approaches relevant to the collection of performance data e.g. binary instrumentation techniques

TABLE 2
Systematic Search Results

Venue	Search Details	Date	Results	Included	Unique
ACM digital library - performance search	((Title:("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks") OR Abstract:("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks") OR Keywords:("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks")) AND (Title:("detecting" OR "finding" OR "profiling" OR "performance understanding") OR Abstract:("detecting" OR "finding" OR "profiling" OR "performance understanding") OR Keywords:("detecting" OR "finding" OR "profiling" OR "performance understanding"))))	18 Dec 2013	106	24	0
ACM digital library - profiling search	((Title:("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis") OR Abstract:("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis") OR Keywords:("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis")) AND (Title:(profiling OR profiler) OR Abstract:(profiling OR profiler) OR Keywords:(profiling OR profiler)))	18 Dec 2013	150	46	0
IEEE Xplore - performance search	(("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks") AND ("detecting" OR "finding" OR "profiling" OR "performance understanding"))	18 Dec 2013	105	12	1
IEEE Xplore - profiling search	((("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis") AND ("profiling" OR "profiler")))	18 Dec 2013	297	38	4
ScienceDirect	TITLE-ABSTR-KEY((("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks") AND ("detecting" OR "finding" OR "profiling" OR "performance understanding")) OR (("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis") AND (profiling OR profiler)))	18 Dec 2013	80	9	1
SpringerLink	(("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks") AND ("detecting" OR "finding" OR "profiling" OR "performance understanding")) OR (("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis") AND (profiling OR profiler))	18 Dec 2013	26677 ^a	61	40
Scopus	TITLE-ABS-KEY((("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks") AND ("detecting" OR "finding" OR "profiling" OR "performance understanding")) OR (("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis") AND (profiling OR profiler)))	18 Dec 2013	1086	137	16
Web of Science	TS=((("object churn" OR "memory bloat" OR "runtime bloat" OR "performance bugs" OR "performance antipatterns" OR "performance anti-patterns" OR "performance problems" OR "performance issues" OR "performance bottlenecks") AND ("detecting" OR "finding" OR "profiling" OR "performance understanding")) OR (("calling context" OR "path profiling" OR "call-path" OR "call-graph" OR "call-tree" OR calltree OR "instrumentation sampling" OR "sampling based" OR "instrumentation based" OR "performance analysis") AND (profiling OR profiler)))	19 Dec 2013	955	102	7
Manual Search	Euro-Par 2000-2013, OOPSLA 2000-2013, PLDI 2000-2013, CGO 2003-2013, ECOOP 2000-2013	Dec 2013		98	32
Reference Snowballing		Dec 2013			19
Total (duplicates removed)			2779^b	253	

a. See discussion on SpringerLink in section 3.8

b. Excluding SpringerLink and Manual Searches

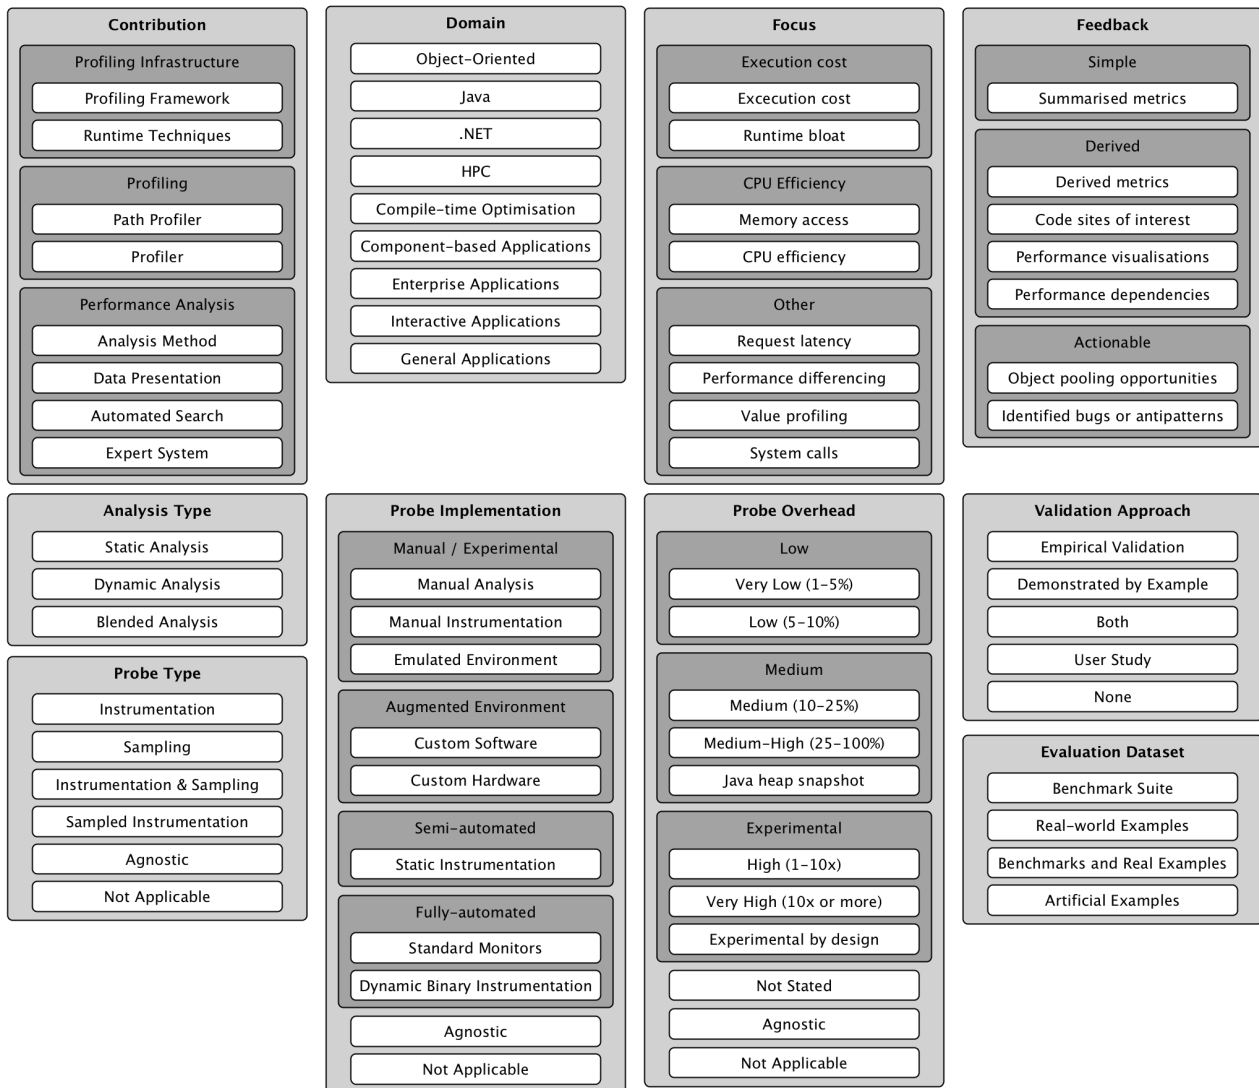


Fig. 1. Facets and Categories

- Profiling papers are focussed on a technique for collecting a specific type of performance data e.g. path, edge or calling context profilers
- The analysis approaches are focussed on how performance data is processed or analysed. This includes techniques for graphing, navigating and searching performance data as well automated search and pattern matching approaches that seek to highlight interesting information in the performance data.

3.11.2 Domain

The domain is the particular research domain that motivated the approach. Mostly this is the type of software that the approach is applied to e.g. object-oriented software or massively parallel HPC applications, but there are also approaches specific to a particular technology

framework such as Java or .NET (that are not generalisable to all object-oriented software). Having this facet allows us to contrast and compare results from different research domains or communities.

3.11.3 Focus

The focus of an approach is the particular type of performance information that the approach investigates. Performance of software is a wide-ranging concern and this facet allows us to classify more specifically the types of performance that the literature was interested in. Approaches that are focussed on the use of runtime resources (using measures such as CPU time or invocation counts) we classify as execution cost approaches. The other major group of approaches are those that focus on understanding and improving the efficiency of CPU execution, either by specifically tracking the instructions

per cycle rate for an execution or by monitoring the rate of events that reduce efficiency, such as memory accesses, memory cache misses or branch mispredictions. Additionally there are a number of other more specialised focii such as request latency profiling or profiling the use and cost of system calls. Request latency profiling is investigating the time taken for an application to respond to an incoming request (such as an HTTP request or a UI trigger event) and is important to improve the responsiveness of interactive applications.

3.11.4 Feedback

The feedback provided by an approach is the style of information or data generated for the user. This is one of our primary motivations for our study, to investigate the style and richness of the feedback that is generated for the user about the performance of software. The simplest feedback returns straight-forward raw or summarised metrics, essentially a long list of measurements often collated by instruction, method or calling context (the hierarchy of active methods calls leading to the current call). The derived feedback approaches perform some level of calculation or interpretation of results. This includes those that generate an advanced visualisation of the performance data (over and above a simple graphing of the summarised metrics), those that calculate new derived metrics from the raw measurements and those that are able to highlight specific code sites of interest. The approaches we regard as providing actionable feedback are those that are able to identify concrete opportunities where a known solution can be applied to improve performance. This includes the identification of object pooling opportunities and finding known coding bugs or anti-patterns.

3.11.5 Analysis Type

Analysis type is the style of analysis used by the approach i.e. whether the approach used static analysis, dynamic analysis or blended (both static and dynamic) analysis. Static analysis approaches are those that rely purely on an analysis of the source code or statically generated artifacts and do not observe the software during its execution. Dynamic analysis approaches are those that rely purely on measurements and analysis of data gathered from software that is executing. Blended analysis, described as a paradigm by Dufour, Ryder and Sevitsky in 2007 [16], uses both static and dynamic analysis information during their application. There are approaches that use basic static information (for example basic block or method information) to insert instrumentation to collect dynamic measurements. Given the simple nature of this static information we still regard these as pure dynamic analysis approaches.

3.11.6 Probe Type

This facet and the two following all capture some aspect of the mechanism used to record dynamic measurements. As such they are not applicable to pure static

analysis approaches (and hence each have a not applicable category). Additionally they each have an agnostic category to categorise performance analysis approaches that are applied to data after it is captured and are independent of how the data was captured.

Probe type describes the type of mechanism used to capture dynamic measurements. Broadly speaking these are either sampling based or instrumentation based. Sampling based approaches observe the state of the system at sampled points in time and build a statistical profile of the dynamic behaviour based upon these samples e.g. java stack sampling. Instrumentation based approaches use probes inserted at specific points in the system to record exact measurements every time the probe is activated e.g. recording every method invocation. Some approaches use sampled instrumentation, this is when instrumented probes are used to record measurements but only during certain sampling periods i.e. they are not permanently active.

3.11.7 Probe Implementation

The probe implementation is the actual mechanism used to capture dynamic measurements, of which there are a wide variety. For example the application source code may be manually modified to include code for logging specific measurements (manual instrumentation), modifications might be made to a kernel library to record how often it is invoked (custom software), or there may be standard APIs that can be used to retrieve standard measures (standard monitors).

Capturing the probe implementation allows us to investigate which mechanisms are popular and gives some insight into how challenging to apply the approach is in practice, as some mechanisms are more accessible than others. We can group the different categories we have for probe implementation into:

- those that can only be applied manually or in an experimental fashion
- those that require an augmented environment with either customised hardware or patched system software, standard libraries or middleware
- static instrumentation approaches that require an extra build-time step but are otherwise automated
- fully automated approaches based on standard monitors or dynamic binary instrumentation that can be activated dynamically with already running applications

3.11.8 Probe Overhead

The probe overhead is the amount of execution cost overhead imposed by the dynamic measurement capture. Capturing overhead allows us to classify the impact of the approaches to judge how practical they are to apply.

3.11.9 Validation Approach

This facet and the next are concerned with how the approaches proposed in the literature are evaluated.

The validation approach is the type of evidence that is put forward to establish the effectiveness of a technique. This can include empirical evaluation, where some measurable aspect of the technique (such as its overhead or accuracy) is analysed to prove its effectiveness, demonstration by example, where a case study gives a detailed report on using the technique to achieve a useful outcome, or a user study, where human subjects use the technique and report on its utility.

3.11.10 Evaluation Dataset

The evaluation dataset for an approach is the type of input used for the described evaluation. Many evaluations are completed using a standard suite of benchmarks (often specific to the particular domain of the research). Others use real-world applications as input, either large open-source frameworks or real-world industrial applications with realistic attributes. Some approaches are evaluated over artificial examples that have been constructed specifically for the purpose of demonstrating or evaluating the approach.

4 RESULTS

4.1 Overview

In total there are 253 papers we included in our completed review. With ten different facets and a number of other pieces of data (such as the publication year and publication venue) defined for each paper there is a large amount of data available for analysis. The complete list of included papers is given in Appendix A and the complete categorisation for all 253 included papers is given in Appendix B. Here we focus on presenting the results that address our specific research questions.

We begin by giving an overview of the included literature in addressing research question 1. The categorisation results for the facets shown in Figure 3 and Figure 8 are discussed in answering the other research questions.

4.2 Research Question 1

What approaches to empirical performance analysis have been proposed that are applicable to object-oriented software?

This narrative summary and Figure 2 describes the literature using selected facets in a way that aims to provide a useful overview of the literature. For this purpose we have primarily used the results from the Contribution and Focus facets as these best capture the particular problem or area the paper was addressing. Note that in order to improve readability Figure 2 is not drawn precisely to scale, but the relative sizes of the labelled areas are approximately in proportion to the number of papers they represent.

We have first divided the papers into 3 groups representing the type of novel contribution they provide:

- 138 Profiling papers
- 77 Performance Analysis papers

- 38 Profiling Infrastructure papers

The profiling papers are focussed on a technique for collecting a specific type of performance data. The performance analysis papers describe an approach to analysing or interpreting the performance data. In most cases the distinction is clear, indeed many profilers essentially do no data analysis and many analysis approaches are completely agnostic of how the data is collected, but in some cases where the approach includes both the collection and analysis of data the distinction can be a fine one.

The profiling infrastructure papers are those that describe tools and runtime techniques that are useful when creating profilers or other performance understanding tools but do not in-of-themselves describe a profiling or performance understanding approach. These papers are still applicable to object-oriented software performance, and therefore included in our review, because they directly support creating new or improving existing performance understanding approaches. There are 18 runtime technique papers. These include the probabilistic [s40]¹, inferred [s161] and precise [s206], [s207] calling context encoding approaches, which describe low overhead approaches to establishing the current full calling context, and papers on instrumentation approaches [s35], [s72], [s160] or aspect oriented programming approaches [s8], [s32], [s222] aimed at supporting dynamic analysis in general and profiling in particular. 20 papers describe tools and frameworks designed to support dynamic analysis in general and profiling in particular. Some of these are instrumentation tools such as Pin [s133] or Valgrind [s165] and others are tools which facilitate the parallelisation of dynamic analysis [s92], [s158], [s224], [s247], [s251] to reduce its cost.

There are 138 papers that describe a form of profiler. Of these there are a small number that have a highly specialised focus, such as the 3 papers concerned with measuring the amount of time spent in kernel code or making operating system calls [s34], [s145], [s182] and the 3 papers concerned with value profiling – tracking the most commonly returned values from individual functions or statements [s41], [s167], [s226]. Ten papers are focussed on understanding request latency rather than overall application performance or throughput. These are either specifically looking at unresponsive UIs in interactive desktop applications [s71], [s109], [s110], [s111] or investigating the reasons for slow responses from multi-tier web-based enterprise applications [s5], [s11], [s46], [s143], [s144], [s243]. The 37 CPU efficiency papers are focussed on monitoring and understanding the code that cause a CPU to run at less than maximum performance. All of these leveraged hardware performance monitors to track low level events that cause CPU pipeline stalls and a drop in overall computation speed. The majority of these (24 papers e.g. [s129], [s171],

¹ Paper references beginning with s refer to the List of Included Papers in Appendix A

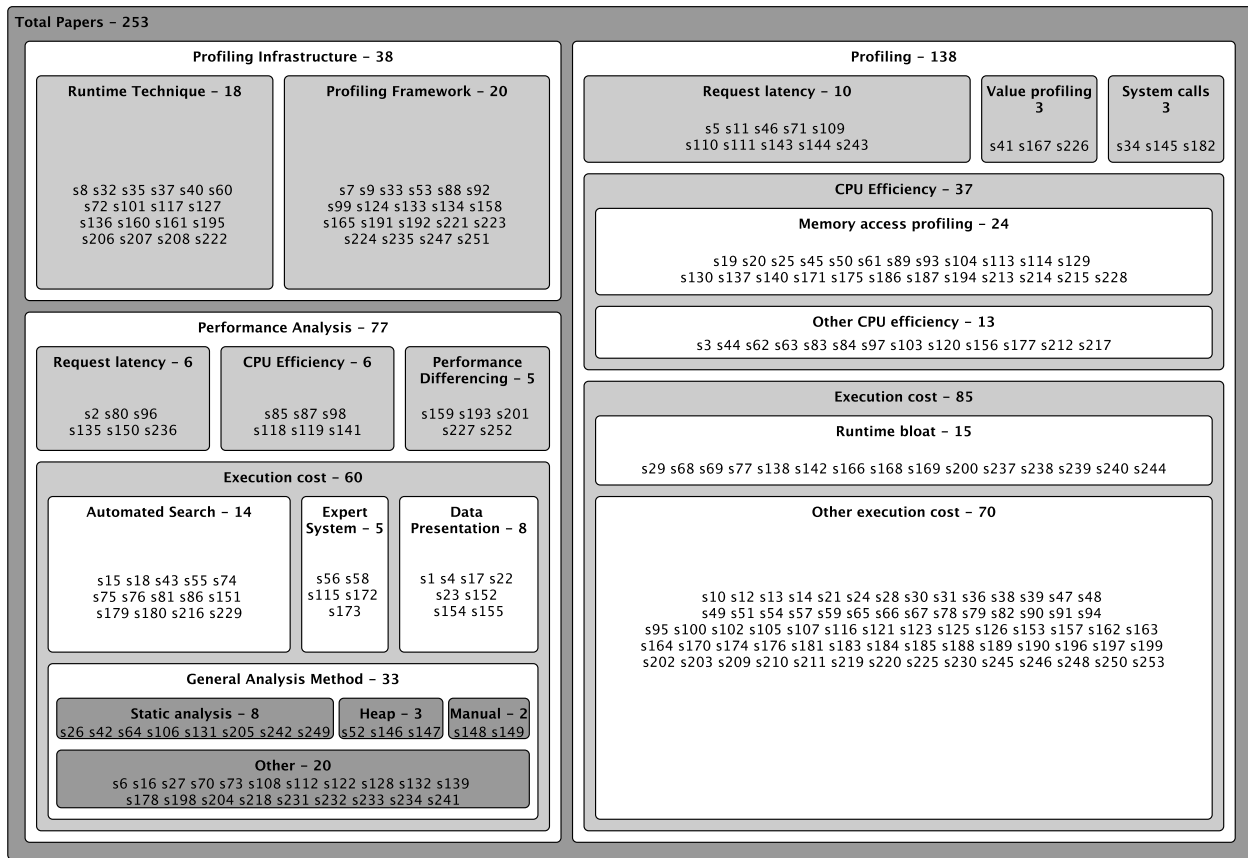


Fig. 2. Literature Overview by Contribution and Focus

[s175], [s213], [s214]) are specifically interested in profiling memory accesses to understand where and why memory loads, stores and cache misses are occurring as these are often the main culprits in stalling the CPU. The remaining 85 profiling papers focus on execution cost generally, that is understanding which code regions are consuming the most computation resources. A small subset of these (15 papers e.g. [s68], [s69], [s166], [s168], [s169], [s237], [s238], [s239], [s240], [s244]) are interested in runtime bloat in object-oriented software.

There are 77 papers concerned with analysis of performance data. As with the profiling approaches we can classify these by focus and again we find that there are a small number of papers with a specialised focus. 5 papers look at performance differencing approaches [s159], [s193], [s201], [s227], [s252] that aim to highlight changes in performance between two captured profiles. There are 6 request latency papers [s159], [s193], [s201], [s227], [s252] with the same goal as the profiling approaches that focussed on request latency - to understand the reasons for unresponsive UIs or slow responses from web-applications. There are only 6 analysis papers that look to understand CPU efficiency [s85], [s87], [s98], [s118], [s119], [s141], a much lower ratio than for the

profiling papers, and the remaining 60 papers focus on execution cost generally. The 14 automated search papers describe approaches where the performance analysis system automatically searches for code regions that exhibit certain pre-defined performance characteristics. Most are systems similar in approach to Periscope [s18], [s81], [s85], [s86] and Paradyn [s43], [s55], [s179], [s180] which run a series of systematic experiments, monitoring progressively finer grained regions of code and measuring and evaluating the performance metrics of interest to narrow down the regions of interest. 5 papers [s56], [s58], [s115], [s172], [s173] are approaches that employ an expert system, rules based, approach to analysing a given set of performance measurements and identifying known patterns of undesirable behaviour. There are 8 papers that propose new approaches to presenting performance data in order to aid the interpretation or comprehension of that data. Most of these relate to visualising or presenting the hierarchical data contained in large calling context trees, for example using calling context ring charts [s152], [s154], [s155]. Execution profiling blueprints [s22], [s23] is an approach to displaying multiple performance metrics or characteristics in the one visualisation to aid in drawing performance infer-

ences.

The remaining 33 papers are general analysis papers. 8 of these are static analysis approaches to understanding performance. Several of these static analysis papers are a type of feedback generated from static optimisers [s64], [s205] (normally but not necessarily associated with compilers) that indicate to the user where code changes could be made to enable further optimisations. Others statically estimate loop, path or method execution frequency [s42], [s131], [s249]. One approach looks for code patterns that equate to known performance bugs [s106]. The patterns have been derived from an in-depth study of performance bugs in open source bug databases. 3 of the other general analysis approaches are based on an analysis of Java heaps to understand memory performance [s52], [s146], [s147], generally looking for patterns that indicate wasteful or inefficient data structures. There are two general analysis papers that rely on customised manual experiments to investigate the way complex framework-based applications transform data [s148], [s149], to understand the reasons for inefficient computation. The remaining 20 papers are a collection of different approaches that focus on the analysis of dynamic performance data to understand the execution cost of applications. These include techniques for aggregating or summarising performance across user-defined code regions [s128], [s204], [s231], [s232] or using user-defined metrics [s73], [s108], [s198].

4.3 Research Question 2

How can these approaches be characterised?

We have already given an overall description of the literature in our review in the previous section. Here we describe some selected results and trends that are apparent in our classification that are not covered in more detail elsewhere.

4.3.1 Venue

One of the traditional results to report in a mapping study is a breakdown of the literature by the publication venue (conference or journal) as this provides an indication for researchers as to which venues are likely to have the most interest in similar research. Our breakdown, shown in Table 3, highlights a number of interesting things:

- Literature related to performance has been published in a very broad range of venues, our survey includes literature from 100 different publication venues, including 66 venues that only provided a single paper.
- Although the performance focussed venues such as the International Conference on Performance Engineering (ICPE) and International Symposium on Performance Analysis of Systems and Software (ISPASS) did contribute a number of articles (5 from each venue) the most important venues for the empirical approaches to performance analysis that

TABLE 3
Categorisation by Venue

Venue	Description	Count
Euro-Par	International Euro-Par Conference on Parallel Processing	26
OOPSLA	Conference on Object-Oriented Programming, Systems, Languages, and Applications	22
CGO	International Symposium on Code Generation and Optimization	21
PLDI	Conference on Programming Language Design and Implementation	18
ECOOP	European Conference on Object-Oriented Programming	11
PPPJ	International Symposium on Principles and Practice of Programming in Java	6
ICSE	International Conference on Software Engineering	6
Tools for HPC	International Workshop on Parallel Tools for High Performance Computing	5
SPE	Software: Practice and Experience	5
ISPASS	International Symposium on Performance Analysis of Systems and Software	5
ICPE/WOSP	International Conference on Performance Engineering	5
SIGMETRICS	International Conference on Measurement and Modeling of Computer Systems	4
SC	Supercomputing Conference	4
CCPE	Concurrency and Computation: Practice and Experience	4
5 more venues with 3 papers		
15 more venues with 2 papers		
66 more venues with 1 paper		

we were interested in were the programming and languages focussed conferences such as OOPSLA, PLDI and ECOOP (which contributed over 50 articles between them).²

4.3.2 Domain

The categorisation of the articles by research domain (shown in Figure 3) again gives some insight into the breadth of the research that has been conducted into software performance. Despite our inclusion criteria restricting us only to articles applicable to object-oriented software and excluding approaches specific to other domains we still had a large contribution from the HPC and compile-time optimisation domains. We made a distinction between articles that were focussed on object-oriented applications in general and those that focussed on applications built with the Java or .NET frameworks. Of course Java and .NET are object-oriented language frameworks, but some of the approaches were focussed not on the object-oriented nature of these frameworks but something specific to the framework design, such as

² Excluded ICPE papers focussed on areas, such as model-based performance prediction, which are related to but not covered by our survey. We discussed these areas in section 2.4.

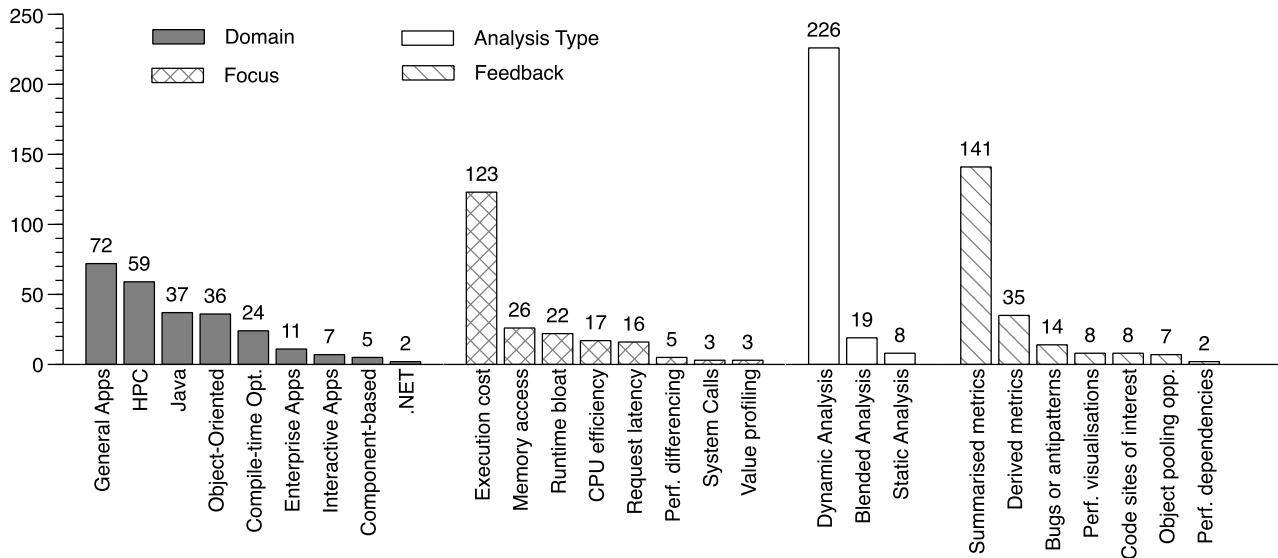


Fig. 3. Categorisation Results - Domain, Focus, Analysis Type and Feedback

techniques for efficient profiling in the JVM. One of the interesting things we noticed here was the great contrast between the number of papers relating specifically to Java (37) compared to .NET (2). Given the importance of .NET as a commercial development platform it would seem to be severely underrepresented in the literature.

4.3.3 Focus

The Focus categorisation results shown in Figure 3 include only the 215 profiling and performance analysis papers. We excluded the 38 profiling infrastructure papers as they were all general purpose approaches applicable when investigating a number of different performance concerns i.e. they had no specific focus.

The most obvious trend here is that by far the majority of the literature (66% - 142 out of 215 once the runtime bloat literature is included) is focussed on analysing the

execution cost of applications. Perhaps the surprising thing though is not that this number is so high but that it is so low, that there are fully a third of the papers that investigate other aspects of performance besides the raw time and resources they consume. In particular 20% (43 out of 215) of the literature focuses on CPU efficiency and memory access profiling, marking it as a significant form of performance analysis that attracts little attention in some domains. A bubble plot of Focus vs Domain (see Figure 5) shows the relative lack of interest in CPU efficiency from the Object-Oriented and Java domains.

There are some other interesting points illustrated in the bubble plot. Unsurprisingly, given its nature, we can see that the concept of runtime bloat is one that mostly exists in the Object-Oriented domain. Finally we can see the particular focus on request latency profiling for Enterprise and Interactive applications.

There are also some noticeable trends over time for the Focus facet, as shown in Figure 4. Whilst the amount of activity in some focus categories remains relatively steady over time, such as execution cost, memory access profiling and CPU efficiency, there is a clear upward trend for both request latency profiling and runtime bloat.

4.3.4 Analysis Type

The clear trend here (Figure 3) is the dominance of pure dynamic analysis approaches. The lack of pure static analysis approaches is not surprising, given that performance is a runtime property. However the small number of blended analysis approaches indicates there are likely opportunities to leverage more static analysis in the performance analysis process.

When we look at the trends over time for analysis type (Figure 6) we can see that both blended and static analysis have become more common in recent years.

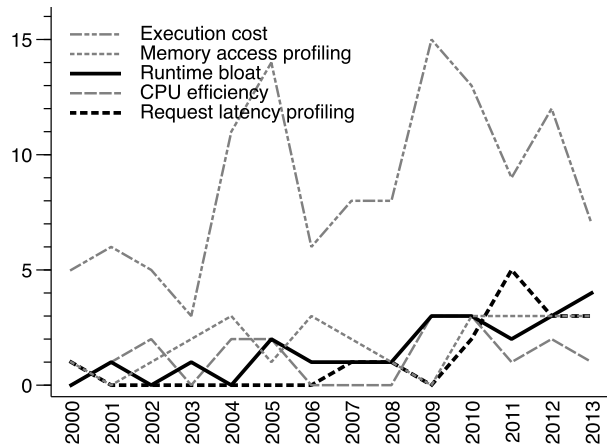


Fig. 4. Focus over Time




Fig. 5. Focus vs Domain




Fig. 6. Analysis Type over Time

4.4 Research Question 3

What form of feedback does the approach provide?

The focus of this research question is to try to understand how helpful the approaches are to the end goal of improving software performance. We captured the type of feedback provided by each approach in the Feedback Type facet (Figure 3). As with the earlier Focus facet we excluded the 38 profiling infrastructure papers from this categorisation as they did not describe approaches that generated performance understanding feedback.

Of the remaining 215 papers, 141 of them provide only raw or summarised metrics. These are metrics that may be helpful in understanding simple properties of an application, such as the absolute cost of a particular method or execution path, but generally require a significant amount of interpretation to find useful performance improvements. In particular they tend to only indicate where an application is consuming a large number of resources but not whether this consumption is wasteful

or useful or how to reduce the consumption.

There are 53 approaches that generate some form of derived feedback by performing additional analysis or transformation on the raw performance data. These include:

- 35 derived metrics
- 8 approaches to advanced visualisation or presentation of performance information expressly designed to aid its interpretation
- 8 code sites of interest
- 2 performance dependencies

In general the derived feedback approaches aim to make the interpretation of the performance data easier. They generally create or define a higher-level abstraction for which they are able to provide a metric that not only helps indicate where an application is consuming resources but why this may be occurring or helps imply a potential solution. Examples include approaches for finding inefficiently used data structures, automated search approaches that search for execution paths that exhibit particular undesirable behaviours and various other novel performance metrics.

Finally there are 21 approaches that generate concrete actionable feedback. 7 of these approaches identify concrete locations where object pooling was both possible and desirable. These use a variety of object lifetime and object usage profiling techniques to establish which allocation sites create objects that are suitable for being pooled and reused. The other 14 actionable feedback approaches identify known performance bugs or antipatterns in concrete source code locations. These cover a wide range of different types of performance problems such as inefficient memory access patterns, inappropriate data structures and never-used or rarely-used object allocations. Several of the antipattern identification approaches make use of configurable detection rules allowing expert users to define new antipatterns.

Overall the majority of approaches (141/215) still only




Fig. 7. Feedback over Time

provide simple feedback.

There are also some interesting trends over time (Figure 7). The obvious trend is the increase in approaches looking at more sophisticated feedback in recent years, object pooling opportunities and identified bugs or antipatterns in particular.

4.5 Research Question 4

How practical to apply are the approaches?

- 1) *What is required to apply the approach?*
- 2) *Can the approaches be applied to live production environments?*

This research question was motivated by the interest practitioners would have in the practicality of applying any performance monitoring or analysis approach. Approaches that are complex to apply can still be useful; there is opportunity within the development life-cycle of many projects to use time-consuming approaches if they provide valuable information. However it is advantageous if approaches can be simply and rapidly applied because it means they can be used more often and in more situations.

We had two facets that were primarily aimed at helping us answer this question, Probe Implementation and Probe Overhead. The categorisation results for these facets are shown in Figure 8. They exclude the 8 pure static analysis approaches, which do not use dynamic probes and therefore do not have an implementation or impose any overhead. It is worth noting that in most situations the static analysis approaches are the most practical to apply because they do not require a running environment. This is one of their great advantages and means that they can be fully integrated into development environments and processes with relative ease. However there are only a few such approaches in the literature we reviewed.

The results of the Probe Implementation categorisation indicate that there are a significant number of

approaches (39% - 96 out of 245) that rely on either dynamic binary instrumentation or standard monitors. These are approaches that we regard as being relatively straight-forward to apply as they do not require a customised environment or build-time instrumentation, the approach can be used in any standard environment with pre-existing binaries. There are a further 34 approaches (~14%) that are analysis approaches that can work with data obtained from any suitable performance monitor and therefore should also be practical to use. The other approaches mostly rely either on custom hardware or software, or require static or manual instrumentation of the application before they can be used. This doesn't preclude them being used in real-world environments but does require extra effort to facilitate their use.

The Probe Overhead categorisation is more difficult to interpret for several reasons. The high number of papers (58 out of 245) that don't address the overhead of their dynamic analysis approach introduces significant uncertainty about any conclusions that are drawn. It also is not clear precisely what the cut-off point is for 'acceptable' levels of overhead, generally this will depend on the requirements of the system being analysed. Another complication is that interpreting some of the reported overheads for the different approaches is difficult due to the variety of approaches taken to measuring and reporting overhead. Whilst many papers take a rigorous approach involving assessing overhead against a variety of standard benchmark and realistic real-world applications others simply used a single case study or report only comparative overhead results relative to other profiling techniques, rather than a comparison with the unprofiled application. Also many of the approaches (especially those using execution sampling methods) could be tuned to reduce overhead at the cost of accuracy. In these cases the literature often suggested a particular overhead/accuracy point as the best-case scenario and those are the figures we used for our categorisation. We have categorised 60 papers where overhead is not applicable. These are either profiling infrastructure papers where it is impossible to assess overhead without knowing the actual analysis implemented using the infrastructure, or analysis approaches that focus on the post-processing of previously collected data and so the approach itself is independent of the style of data collection.

What is clear from the results is that we do have a significant number of approaches, about 25%, that have an emphasis on minimising overhead, those with Very Low or Low reported overhead. The majority of the remaining approaches either have a high overhead or do not report on overhead at all, there are relatively few approaches that fall into the medium or medium-high overhead areas.

When we look at the categorisation of approaches by both implementation approach and overhead we see that a significant proportion of the low overhead approaches are achieved by using custom hardware or software




Fig. 8. Categorisation Results - Implementation and Evaluation

implementations. Overall only around 10% (25 out of 245) of the approaches are both low overhead and fully automated.

The other trend that is apparent in our data is the type of feedback generated for low overhead approaches is dominated by simple summarised metrics. Nearly 90% of low overhead approaches return only summarised metrics as opposed to around 65% of all profiling and analysis method approaches.

4.6 Research Question 5

How are approaches evaluated and validated?

We were interested in how approaches to improving performance engineering were evaluated and validated. This followed from our interest in the utility of performance analysis approaches, to understand how effective reported performance analysis approaches were we wanted to understand how they had been evaluated.

We used two facets to classify how approaches were evaluated, the Evaluation Dataset and the Validation Approach. The categorisation for these facets is shown in Figure 8. There are two notable findings from the Validation Approach classification:

- 1) A significant proportion of the approaches (79/253 ~ 31%) did no validation of their approach.
- 2) Only a tiny minority (4/253 ~ 1.6%) validated their approach with real users in a user study.

The majority of the approaches that were validated were done so via either empirical validation or a case study that demonstrated using the approach to achieve a desirable outcome. The empirical validations were generally quantitative comparisons of accuracy or overhead compared to some established baseline. In some cases, where it was possible to have an oracle that could identify true positives, the empirical validation

included measuring false positive and false negative rates generated by the tool. Most of the 79 papers that we have classified as doing no validation did demonstrate applying their approach to real-world or benchmark applications but did not validate the output in any way to demonstrate the accuracy or utility of the output. They relied solely on the description of their approach to establish that it was useful. This is in contrast to the 51 papers that demonstrated using their approach with a case study that highlighted how using the approach resulted in information that led to a measurable performance improvement.

The Evaluation Dataset facet classifies the approaches by the type of input they were applied to. Here we see an fairly equal split between using real-world examples (usually large open source or industrial applications) or a standard benchmark suite. There were some approaches that were evaluated over both standard benchmarks and real-world examples. The approaches that did not use either real-world examples or benchmarks used artificial examples. Note that this classification includes all 253 papers, including the 79 that we classified as having not completed any validation. This is because all 253 papers used illustrative examples, even if they were only artificial examples, either when describing or evaluating their approach.

5 DISCUSSION

The intent of our survey was to map and analyse the existing literature pertaining to empirical performance analysis approaches for object oriented software, with a particular emphasis on reviewing the type of performance feedback provided by these approaches. In the previous sections we have described the results of our categorisation of the relevant literature, here we

highlight some of those results and discuss some key conclusions we have drawn from our analysis. In particular we highlight where we think there are opportunities for future research.

5.1 Very little static or blended analysis, virtually all dynamic

The analysis type facet gave us the most skewed distribution of results in our entire categorisation. It is clear that not only do pure dynamic analysis approaches dominate in our study but that much of the work done using static or blended analysis (20 out of 27 papers) has been published since 2009. This trend in recent times towards using more static analysis, in conjunction with dynamic analysis or alone, indicates to us that this is an emerging area of work where there are very likely still substantial opportunities. Our definition of blended analysis for performance analysis is not novel, it is specifically described by Dufour, Ryder & Sevitsky [16], but the results of our survey have allowed us to quantify just how infrequently it has been leveraged.

5.2 Cross-Domain approaches

We found our study covered literature from a very broad range of venues and domains, as we quantified with the categorisation results we covered in section 4.3. We found this interesting because the focus of our study was literature applicable to object-oriented software and yet we found literature from a wide range of areas that was applicable to object-oriented software. The implication of this is that researchers or practitioners working with object-oriented software need to be aware that there is potentially related work being done within these other domains, such as HPC or compile-time optimisation. We feel the reverse is likely to be true as well. For example object-oriented software frameworks and concepts are being adopted in some HPC applications and this makes it likely that some of the performance characteristics of object-oriented software will be inherited as well, therefore making the research into object-oriented application performance relevant to the HPC community.

We also found that there were a number of techniques or approaches that had been investigated by researchers from different domains that were similar in nature and therefore likely had some relevance to each other. For example there were expert system approaches to performance classification described by researchers in both the HPC and EJB application domains. Request latency profiling was of interest to researchers from both the enterprise web application and interactive desktop application domains. There were many performance data visualisation or aggregation techniques from different domains that had fundamentally similar goals. This was notable because we found examples where the research in one domain did not reference or compare their approach to work with similar aspects that was from another domain e.g. the expert system approaches [s56], [s58], [s115], [s172], [s173].

5.3 Evaluating the utility of performance analysis approaches

Based upon our analysis in section 4.6 of the evaluation approaches used in the literature we reviewed we have concluded that there has been remarkably little work done to prove the utility of empirical performance analysis approaches, particularly for large scale object-oriented software. By this we mean very few papers addressed the question of whether the performance feedback being generated for the user was actually useful in helping to improve software performance. In all the papers we included in our review, only 4 included an evaluation involving human subjects to assess whether the approach was useful for ultimately improving performance. A further 70 gave concrete examples of how using a particular approach led to real performance improvements but few of these examples could convincingly be extrapolated to provide evidence that the approach is useful for object-oriented software in general. Only 42 used real-world examples and only 7 of these were from the object-oriented domain (as classified by our Domain facet).

Note that this is not a commentary on the quality of the evaluations done in the individual research papers in our survey. Many papers were presenting a new approach to obtaining a standard set of data with better accuracy or lower overhead, a situation where the utility of the data is assumed to already be proven, and therefore the focus in the evaluation for those papers was in validating the improved accuracy or reduced overheads. Nevertheless there were approaches describing novel types of performance feedback that relied solely on the description of the approach to establish that it was useful. We feel there is a clear gap in the research when it comes to evaluating the utility of performance analysis approaches, particularly for large scale object-oriented software.

5.4 Improving feedback on the performance of large-scale object-oriented software

The overriding conclusion we have drawn from our study is that there are still opportunities to improve feedback on the performance of large-scale object-oriented software. Whilst there has been research done in this area, which we describe below, there are still possible extensions and innovations to contribute.

We had 253 papers in our study. This included:

- 38 profiling infrastructure papers
- 74 papers focussed on some characteristic of performance other than execution cost e.g. memory accesses, request latency or value profiling
- 82 that returned only simple feedback i.e. summarised metrics

Therefore only 59 papers describe profiling or analysis approaches investigating execution cost that provide a more advanced (derived or actionable) form of feedback. We describe these papers in more detail below

to highlight where we think there are further research opportunities.

30 of these papers, whilst applicable to object-oriented software, do not specifically address the challenges provided by large-scale object-oriented software. This means they are likely to struggle to produce useful feedback when used with large-scale object-oriented software. These included:

- 5 expert system approaches. These use a rules-based encoding of expert knowledge to classify whether observed behaviour matches known, pre-defined, undesirable patterns. The major challenge with these approaches is creating suitable expert-system rules. The examples we have seen have been focussed either on HPC applications, with rules mostly focussed around parallel scalability concerns, or component-based (essentially EJB) applications, with rules mostly focussed around poor distributed computing patterns. We did not find any expert system approaches focussed on analysing large-scale object-oriented applications.

The other approaches provide feedback at the method level or basic block level of granularity and so can fail to provide useful feedback when faced with the overwhelming number of fine-grained methods that exist in large-scale object-oriented applications. These consisted of:

- 14 papers that perform a type of automated search for code regions that exhibit certain pre-defined performance characteristics. These approaches help to rapidly identify specific problem areas, often with a relatively low overhead because of the iterative and targeted search process. Only pre-defined characteristics, such as a poor rate of memory cache hits, can be identified.
- 8 papers covering a range of different approaches to deriving new metrics from the raw measurements that are captured during profiling. Some of these new metrics are built-in advanced metrics that expose a particular characteristic of performance and some allow the user to define their own metrics using calculations combining existing measurements.
- 3 static analysis approaches. Two of these are a type of feedback generated from static optimisers (normally but not necessarily associated with compilers) that indicate where code changes could be made that would enable further static optimisations. The other approach looks for code patterns that equate to known performance bugs. All of these approaches have the typical benefits and weaknesses of static approaches. They can be applied at build time without needing a running environment but are unable to distinguish which code is performance critical.

The remaining 29 papers are those that most directly addressed performance concerns specific to large-scale object-oriented software, and runtime bloat in particular:

- 4 papers analyse memory efficiency in Java pro-

grams by analysing the allocated objects and structures on the heap. These help give insight into memory bloat but do not directly address issues of execution time performance.

- 7 papers discuss performance data presentation approaches. These techniques assist with interpreting performance data but only help visualise or present the existing data. New insights into the performance behaviour and how the performance could be improved still need to be drawn by the user.
- 1 static analysis approach, which combines concern input with program analysis for bloat detection. This helps analyse when optional concerns (features) have an impact on the performance of mandatory features. It requires an external source of concern input and does not help identify when the implementation of a mandatory feature is inefficient.
- 1 paper presents algorithmic profiling, an approach that automatically partitions a program into multiple algorithms and determines a cost function for each of these algorithms in terms of the input size for the algorithm. The focus is on understanding how these algorithms will behave with increasing input sizes. The current implementation is only able to work on algorithms that operate on data structures and has a very high runtime overhead (several orders of magnitude).

The last 16 papers are those that most directly tackle runtime execution bloat:

- 7 approaches focus on identifying object pooling opportunities, these identify object allocation sites that produce many objects that are either very short-lived or very similar or both and as such create repeated and potentially unnecessary work.
- 5 papers that focus on tracking possible symptoms of runtime bloat in a data-flow centric manner. They do this by monitoring data copies, reference propagations or the usage of temporary objects and relating them to the code locations that induce them.
- 4 approaches focus on the way data structures are created and used, either identifying situations where an implementation is inappropriate for the way it is used or where a data-structure has a low value (is used very little) relative to the cost of its creation.

All of these last 16 papers are focussed on understanding bloat in terms of the inefficient use of data-structures, inefficient data-flows or inefficient creation of objects. In other words they all took a data-flow centric view of the problem, rather than a control-flow centric one. Whilst the majority of the more traditional profilers did take a control-flow centric view of performance, they provided only simple method or calling context summarised metrics. Consequently there is opportunity for providing more advanced control-flow centric feedback on performance of object-oriented software. For example, feedback approaches based on aggregating performance information over the repeated calling patterns that are

induced by the design of the software. Additionally the object pooling and data structure based approaches address specific identified performance problems, i.e. object churn and inappropriate data structure respectively, rather than providing general performance feedback. This means that whilst these approaches are useful they are not comprehensive and there are further opportunities for complementary approaches.

6 THREATS TO VALIDITY

The primary threats to the validity of our results and conclusions are the two classic problems faced by many literature reviews, namely the possibility of researcher bias and the possibility of missing relevant articles. The primary defense against these threats has been to follow a systematic review methodology as we have described in section 3. This process is designed to be rigorous and objective and relies primarily upon the discipline of defining the protocol to be followed throughout the review before undertaking the review. However it does still rely on subjective interpretation and in our case the risk of researcher bias is increased by the fact that the majority of the work was undertaken by a single researcher, the primary author of this paper. Therefore we took additional steps, which we describe below, at several stages to mitigate this risk.

One of our biggest challenges in conducting this systematic review was the search phase. We experienced a number of challenges similar to those discussed by Brereton et al [17], in particular the differences in functionality and sophistication between the different mainstream software engineering search engines. As documented in our search results in section 3.8 each search engine required a different search expression syntax. They also each had different approaches to both stemming (the process of identifying the root of a search term to find relevant partial matches e.g. plurals) and identification of key words for articles, meaning that it was possible to get different results from different search engines even when the same paper was being indexed. However these challenges were minor compared to our biggest challenge, which was generating suitable key phrases for our initial searches. As we discussed when we described our search phrase development in section 3.4 the key terms for our domain suffered from having both low precision (terms like ‘software’ and ‘performance’ return tens of thousands of results) and low recall (there are many synonyms for terms like ‘performance analysis’). We lacked that one key unique domain specific term that would have helped both the precision and recall of our searches.

To mitigate the challenges of the search phase we resorted to being as thorough as possible. We included multiple electronic database searches, manual searches targeted to the venues highlighted as key venues by the results of our electronic searches, and finally referencing checking of included papers. In total we performed

automated searches against 6 major electronic databases, manual searches across 66 conference proceedings, and reference checks against many hundreds of papers. In total we manually screened many thousands of papers for inclusion in our survey. Despite this thoroughness we feel it is likely that some relevant articles have been missed. This was highlighted by our manual searches, where a detailed search of 5 key venues that were already contributing over 60 articles found by our electronic searches, found another 32 articles suitable for inclusion. This worryingly represents an approximately 50% miss rate from the electronic searches. More reassuring were the results of the reference snowballing, which resulted in only an additional 19 papers being included, meaning that comparatively few papers of referenceable quality were missed by the earlier searches. Due to the completeness of our searches we believe the number of possibly missed articles is small enough to not substantially affect our conclusions.

Our decision to limit our mapping study only to literature published since January 2000 does mean that we have potentially excluded relevant literature from before this time. However given our particular focus on object-oriented applications, which were still relatively new in 2000, the rapid pace of change in computing since then, and the large numbers of papers we have included in our study, we feel it is unlikely that the excluded papers from before 2000 would significantly impact our conclusions.

To reduce the risk of researcher bias during article selection and data extraction we undertook a process where a random selection of 100 papers returned from the automatic searches was screened for inclusion and then categorised according to the developed review protocol by each of the four authors of this paper. Once each author had completed this process independently the results were compared and any discrepancies were discussed and resolved. The intent of this process was to clarify any ambiguous aspects of the inclusion/exclusion criteria and categorisation process to make them less subjective and easier to apply more consistently. This then allowed us to fine-tune our review protocol before it was finalised and applied to the full review.

The results interpretation process is by its nature a subjective process and is vulnerable to bias or limitations at the earlier stages of the process as missing papers or ambiguous classifications during the data extraction will potentially change the results that are being discussed. One of our primary defenses against these threats was to limit ourselves to analysing and drawing conclusions on clear trends in the data, trends that are not invalidated by small numbers of papers being missed or classified inappropriately. Our results interpretation was also an area where the conclusions were reviewed and discussed amongst the authors of the paper to mitigate the risk of bias from a single researcher.

7 CONCLUSION

We have presented the details of a systematic mapping study on empirical performance analysis approaches applicable to object-oriented software. We performed a wide-ranging systematic literature search including automated searches of 6 major software engineering electronic databases and detailed manual searches of 5 key venues since 2000, screening many thousands of papers, to eventually select 253 articles related to empirical approaches for performance analysis applicable to object-oriented software. We developed a classification scheme (see Figure 1) to structure and describe literature in the field of empirical software performance analysis and subsequently applied it to our 253 selected articles. This allowed us to give an overview of the current research in the field (see Figure 2) and examine the trends and gaps in that research. The resulting mapping is useful both for researchers looking to identify new opportunities and related work in this field, and practitioners from industry interested in discovering state-of-art approaches to performance analysis they might wish to apply.

After analysis of our categorisation we came to 4 key conclusions. Firstly there are relatively few approaches that incorporate the use of any in-depth static information or analysis when analysing performance. There seems to be great scope for leveraging more static analysis to augment the traditional dynamic analysis based approaches. Secondly we found that there are performance analysis approaches from a wide range of different domains that are applicable to object-oriented software. The reverse is likely to be true as well, with object-oriented approaches being increasingly adopted in diverse areas such as HPC and mobile computing the performance analysis approaches developed for object-oriented software are likely to be applicable in those areas in future. Thirdly, there has been little work done evaluating the utility of performance analysis approaches, particularly for large-scale object-oriented software.

Finally we found there are still opportunities to improve on the feedback generated on the performance of large-scale object-oriented software. The majority of existing approaches provide only simple metrics, which are challenging to interpret in large-scale object-oriented software. There is some more recent research, particularly that into runtime bloat, that addresses this problem but to date it has been limited to data-flow, rather than control-flow, based approaches. Despite the importance and prevalence of object-oriented software there are significant challenges in understanding and improving its performance.

REFERENCES

- [1] B. Kitchenham and S. Charters, "Guidelines for performing systematic literature reviews in software engineering," *EBSE Technical Report, EBSE-2007-01*, 2007.

- [2] J. Larus, "Spending Moore's Dividend," *Communications of the ACM*, vol. 52, no. 5, pp. 62–69, 2009.
- [3] N. Mitchell, G. Sevitsky, and H. Srinivasan, "Modeling Runtime Behavior in Framework-Based Applications," *Lecture Notes in Computer Science*, vol. 4067, pp. 429–451, 2006.
- [4] G. Xu, N. Mitchell, M. Arnold, A. Rountev, and G. Sevitsky, "Software Bloat Analysis: Finding, Removing, and Preventing Performance Problems in Modern Large-Scale Object-Oriented Applications," *Proceedings of the FSE/SDP Workshop on the Future of Software Engineering Research - FoSER 2010*, pp. 421–425, 2010.
- [5] N. Mitchell, E. Schonberg, and G. Sevitsky, "Four Trends Leading to Java Runtime Bloat," *IEEE Software*, vol. 27, no. 1, pp. 56–63, 2010.
- [6] M. Woodside, G. Franks, and D. Petriu, "The future of software performance engineering," *Future of Software Engineering*, pp. 171–187, 2007.
- [7] C. U. Smith, *Performance Engineering of Software Systems*. Addison Wesley, 1990.
- [8] S. Balsamo, A. Di Marco, P. Inverardi, and M. Simeoni, "Model-based performance prediction in software development: a survey," *IEEE Transactions on Software Engineering*, vol. 30, no. 5, pp. 295–310, May 2004.
- [9] D. C. D'Elia, C. Demetrescu, and I. Finocchi, "Mining hot calling contexts in small space," *Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation - PLDI '11*, pp. 516–527, 2011.
- [10] H. Koziolok, "Performance evaluation of component-based software systems: A survey," *Performance Evaluation*, vol. 67, no. 8, pp. 634–658, Aug. 2010.
- [11] K. Petersen, R. Feldt, S. Mujtaba, and M. Mattsson, "Systematic mapping studies in software engineering," in *12th International Conference on Evaluation and Assessment in Software Engineering*, 2008, pp. 68–77.
- [12] B. Kitchenham, R. Pretorius, D. Budgen, P. Brereton, M. Turner, M. Niazi, and S. Linkman, "Systematic literature reviews in software engineering – A tertiary study," *Information and Software Technology*, vol. 52, no. 8, pp. 792–805, Aug. 2010.
- [13] H. Zhang, M. A. Babar, and P. Tell, "Identifying relevant studies in software engineering," *Information and Software Technology*, vol. 53, no. 6, pp. 625–637, Jun. 2011.
- [14] T. Dyba, T. Dingsoyr, and G. K. Hanssen, "Applying Systematic Reviews to Diverse Study Types: An Experience Report," *First International Symposium on Empirical Software Engineering and Measurement - ESEM 2007*, no. 7465, pp. 225–234, Sep. 2007.
- [15] K. Petersen, "Measuring and predicting software productivity: A systematic map and review," *Information and Software Technology*, vol. 53, no. 4, pp. 317–343, Apr. 2011.
- [16] B. Dufour, B. G. Ryder, and G. Sevitsky, "Blended analysis for performance understanding of framework-based applications," *Proceedings of the 2007 international symposium on Software testing and analysis - ISSTA '07*, pp. 118–128, 2007.
- [17] P. Brereton, B. a. Kitchenham, D. Budgen, M. Turner, and M. Khalil, "Lessons from applying the systematic literature review process within the software engineering domain," *Journal of Systems and Software*, vol. 80, no. 4, pp. 571–583, Apr. 2007.



David Maplesden David Maplesden completed a MSc (Hons) in Computer Science at The University of Auckland, New Zealand, in 2000. He spent 12 years as a professional software developer both in NZ and overseas mostly building enterprise Java web applications. He is currently working towards a PhD in the Department of Computer Science at The University of Auckland. His current research interests are the performance of large scale object-oriented software and tool support for software engineering.



Ewan Tempero Ewan Tempero is an Associate Professor in the Department of Computer Science at The University of Auckland, New Zealand. He graduated from the University of Otago, New Zealand, with a B.Sc., (Honours) in Mathematics in 1983 and received his Ph.D. in Computer Science from the University of Washington, USA, in 1990. He has published over 140 papers in journals and internationally-refereed conferences, mainly in the areas of software reuse, software tools, and software metrics. His

current research is developing metrics for measuring the quality of software designs. He is the developer and maintainer of the Qualitas Corpus.



John C. Grundy John C. Grundy received the BSc (Hons), MSc, and PhD degrees in computer science from the University of Auckland, New Zealand. Previously, he was a lecturer and senior lecturer at the University of Waikato, New Zealand, and a professor of software engineering and the head of electrical and computer engineering at the University of Auckland, New Zealand. He is currently a professor of software engineering and the head of computer science and software engineering at Swinburne University of Technology, Melbourne, Australia. He is an associate editor of the

IEEE Transactions on Software Engineering, the Automated Software Engineering Journal, and IEEE Software. His current interests include domain-specific visual languages, model-driven engineering, large-scale systems engineering, and software engineering education. He is a member of the IEEE and the IEEE Computer Society. More details about his research can be found at <http://www.ict.swin.edu.au/ictstaff/jgrundy>.



John Hosking John Hosking is Dean of Science at the University of Auckland. John obtained his BSc and PhD in Physics from the University of Auckland then joined the Department of Computer Science at the University of Auckland working his way through to full Professor before taking up the position of Dean of the College of Engineering and Computer Science at the Australian National University. He has recently returned to his new role in Auckland. John's research interests are in software engineering,

visual languages, and knowledge visualisation. He is a member of the IEEE and a Fellow of the Royal Society of New Zealand.

APPENDIX A

LIST OF INCLUDED PAPERS

- [s1] A. Adamoli and M. Hauswirth, "Trevis: A Context Tree Visualization & Analysis Framework and its Use for Classifying Performance Failure Reports," *Proceedings of the 5th international symposium on Software visualization - SOFTVIS '10*, pp. 73–82, 2010.
- [s2] A. Adamoli, M. Jovic, and M. Hauswirth, "LagAlyzer: A latency profile analysis and visualization tool," *2010 IEEE International Symposium on Performance Analysis of Systems & Software - ISPASS '10*, pp. 13–22, Mar. 2010.
- [s3] L. Adhianto, S. Banerjee, M. Fagan, M. Krentel, G. Marin, J. Mellor-Crummey, and N. R. Tallent, "HPCTOOLKIT: tools for performance analysis of optimized parallel programs," *Concurrency and Computation: Practice and Experience*, vol. 22, no. 6, pp. 685–701, 2010.
- [s4] L. Adhianto, J. Mellor-Crummey, and N. R. Tallent, "Effectively Presenting Call Path Profiles of Application Performance," *2010 39th International Conference on Parallel Processing Workshops*, pp. 179–188, Sep. 2010.
- [s5] E. Altman, M. Arnold, S. Fink, and N. Mitchell, "Performance analysis of idle programs," *Proceedings of the ACM international conference on Object oriented programming systems languages and applications - OOPSLA '10*, vol. 45, no. 10, pp. 739–753, Oct. 2010.
- [s6] G. Ammons, J.-d. Choi, M. Gupta, and N. Swamy, "Finding and Removing Performance Bottlenecks in Large Systems," *Proceedings of the 18th European Conference on Object-Oriented Programming - ECOOP '04*, pp. 172–196, 2004.
- [s7] D. Ansaloni, W. Binder, A. Heydarnoori, and L. Y. Chen, "Deferred Methods: Accelerating Dynamic Program Analysis on Multicores," *Proceedings of the Tenth International Symposium on Code Generation and Optimization - CGO '12*, pp. 242–251, 2012.
- [s8] D. Ansaloni, W. Binder, A. Villazón, and P. Moret, "Rapid development of extensible profilers for the Java virtual machine with aspect-oriented programming," *Proceedings of the first joint WOSP/SIPEW international conference on Performance engineering - WOSP/SIPEW '10*, pp. 57–62, 2010.
- [s9] D. Ansaloni, S. Kell, Y. Zheng, L. Bulej, W. Binder, and P. Tüma, "Enabling Modularity and Re-use in Dynamic Program Analysis Tools for the Java Virtual Machine," *Lecture Notes in Computer Science*, vol. 7920, pp. 352–377, 2013.
- [s10] T. Apiwattanapong and M. J. Harrold, "Selective path profiling," *Proceedings of the 2002 ACM SIGPLAN-SIGSOFT workshop on Program analysis for software tools and engineering - PASTE '02*, pp. 35–42, 2002.
- [s11] A. Arkles and D. Makaroff, "MT-WAVE: Profiling Multi-Tier Web Applications," *Proceedings of the second joint WOSP/SIPEW international conference on Performance engineering - ICPE '11*, pp. 247–258, 2011.
- [s12] M. Arnold and D. Grove, "Collecting and Exploiting High-Accuracy Call Graph Profiles in Virtual Machines," *International Symposium on Code Generation and Optimization - CGO '05*, pp. 51–62, 2005.
- [s13] M. Arnold and B. G. Ryder, "A framework for reducing the cost of instrumented code," *Proceedings of the ACM SIGPLAN 2001 conference on Programming language design and implementation - PLDI '01*, pp. 168–179, 2001.
- [s14] M. Arnold and P. Sweeney, "Approximating the calling context tree via sampling," *IBM TJ Watson Research Center*, 2000.
- [s15] M. Attariyan, M. Chow, and J. Flinn, "X-ray: Automatic Root-cause Diagnosis of Performance Anomalies in Production Software," *Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation - OSDI'12*, pp. 307–320, 2012.
- [s16] G. Ausiello, C. Demetrescu, I. Finocchi, and D. Firmani, "k-Calling context profiling," *Proceedings of the ACM international conference on Object oriented programming systems languages and applications - OOPSLA '12*, pp. 867–878, Nov. 2012.
- [s17] F. Beck, O. Moseler, S. Diehl, and G. D. Rey, "In situ understanding of performance bottlenecks through visually augmented code," *2013 21st International Conference on Program Comprehension (ICPC)*, pp. 63–72, May 2013.
- [s18] S. Benedict, M. Brehm, M. Gerndt, C. Guillen, W. Hesse, and V. Petkov, "Automatic Performance Analysis of Large Scale Simulations," *Lecture Notes in Computer Science*, vol. 6043, pp. 199–207, 2010.
- [s19] E. Berg and E. Hagersten, "SIP: Performance Tuning through Source Code Interdependence," *Lecture Notes in Computer Science*, vol. 2400, pp. 177–186, 2002.
- [s20] —, "Fast Data-Locality Profiling of Native Execution," *Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems - SIGMETRICS '05*, pp. 169–180, 2005.
- [s21] A. Bergel, "Counting Messages as a Proxy for Average Execution Time in Pharo," *Lecture Notes in Computer Science*, vol. 6813, pp. 533–557, 2011.
- [s22] A. Bergel, F. Bañados, R. Robbes, and W. Binder, "Execution profiling blueprints," *Software: Practice and Experience*, vol. 42, no. 9, pp. 1165–1192, 2012.
- [s23] A. Bergel, R. Robbes, and W. Binder, "Visualizing Dynamic Metrics with Profiling Blueprints," *Lecture Notes in Computer Science*, vol. 6141, pp. 291–309, 2010.
- [s24] A. R. Bernat and B. P. Miller, "Incremental call-path profiling," *Concurrency and Computation: Practice and Experience*, vol. 19, no. 11, pp. 1533–1547, 2007.
- [s25] K. Beyls and E. H. D'Hollander, "Discovery of Locality-Improving Refactorings by Reuse Path Analysis," *Lecture Notes in Computer Science*, vol. 4208, pp. 220–229, 2006.
- [s26] S. Bhattacharya, K. Gopinath, and M. G. Nanda, "Combining concern input with program analysis for bloat detection," *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications - OOPSLA '13*, pp. 745–764, 2013.
- [s27] S. Bhattacharya, M. G. Nanda, K. Gopinath, and M. Gupta, "Reuse, Recycle to De-bloat Software," *Lecture Notes in Computer Science*, vol. 6813, pp. 408–432, 2011.
- [s28] W. Binder, "A Portable and Customizable Profiling Framework for Java Based on Bytecode Instruction Counting," *Lecture Notes in Computer Science*, vol. 3780, pp. 178–194, 2005.
- [s29] —, "Portable Profiling of Memory Allocation in Java," *Proceeding of the 6th International Conference on Net.ObjectDays 2005 - NODE '05*, pp. 110–128, 2005.
- [s30] —, "Portable and accurate sampling profiling for Java," *Software: Practice and Experience*, vol. 36, no. 6, pp. 615–650, May 2006.
- [s31] W. Binder, D. Ansaloni, A. Villazón, and P. Moret, "Parallelizing calling context profiling in virtual machines on multicores," *Proceedings of the 7th International Conference on Principles and Practice of Programming in Java - PPPJ '09*, pp. 111–120, 2009.
- [s32] —, "Flexible and efficient profiling with aspect-oriented programming," *Concurrency Computation: Practice and Experience*, vol. 23, pp. 1749–1773, 2011.
- [s33] W. Binder and J. Hulaas, "Flexible and Efficient Measurement of Dynamic Bytecode Metrics," *Proceedings of the 5th international conference on Generative programming and component engineering - GPCE '06*, pp. 171–180, 2006.
- [s34] W. Binder, J. Hulaas, and P. Moret, "A Quantitative Evaluation of the Contribution of Native Code to Java Workloads," *2006 IEEE International Symposium on Workload Characterization - IISWC '06*, pp. 201–209, Oct. 2006.
- [s35] —, "Advanced Java bytecode instrumentation," *Proceedings of the 5th international symposium on Principles and practice of programming in Java - PPPJ '07*, p. 135, 2007.
- [s36] W. Binder, J. Hulaas, P. Moret, and A. Villazón, "Platform-independent profiling in a virtual execution environment," *Software - Practice and Experience*, vol. 39, no. 1, pp. 47–79, 2009.
- [s37] M. D. Bond, G. Z. Baker, and S. Z. Guyer, "Breadcrumbs: efficient context sensitivity for dynamic bug detection analyses," *Proceedings of the 2010 ACM SIGPLAN conference on Programming language design and implementation - PLDI '10*, pp. 13–24, 2010.
- [s38] M. D. Bond and K. S. McKinley, "Continuous Path and Edge Profiling," *Proceedings of the 38th Annual IEEE/ACM International Symposium on Microarchitecture, 2005 - MICRO-38*, pp. 130–140, 2005.
- [s39] —, "Practical Path Profiling for Dynamic Optimizers," *Proceedings of the 2005 International Symposium on Code Generation and Optimization - CGO '05*, pp. 205–216, 2005.
- [s40] M. D. Bond and K. S. McKinley, "Probabilistic calling context," *Proceedings of the 22nd annual ACM SIGPLAN conference on Object-oriented programming systems and applications - OOPSLA '07*, pp. 97–112, Oct. 2007.
- [s41] M. Burrows, U. Erlingsson, S.-T. a. Leung, M. T. Vandevoorde, C. a. Waldspurger, K. Walker, and W. E. Wehl, "Efficient and flexible value sampling," *Proceedings of the ninth international*

- conference on Architectural support for programming languages and operating systems - ASPLOS IX, pp. 160–167, Nov. 2000.
- [s42] R. P. L. Buse and W. Weimer, “The road not taken: Estimating path execution frequency statically,” in *Proceedings - International Conference on Software Engineering*, 2009, pp. 144–154.
- [s43] H. W. Cain, B. P. Miller, and B. J. N. Wylie, “A Callgraph-Based Search Strategy for Automated Performance Diagnosis,” *Lecture Notes in Computer Science*, vol. 1900, pp. 108–122, 2000.
- [s44] P. Calafiura, S. Eranian, D. Levinthal, S. Kama, and R. a. Vitillo, “GOoDA: The Generic Optimization Data Analyzer,” *Journal of Physics: Conference Series*, vol. 396, no. 5, p. 052072, Dec. 2012.
- [s45] M. Chabbi and J. Mellor-Crummey, “DeadSpy: A Tool to Pinpoint Program Inefficiencies,” *Proceedings of the Tenth International Symposium on Code Generation and Optimization - CGO '12*, pp. 124–134, 2012.
- [s46] A. Chanda, A. L. Cox, and W. Zwaenepoel, “Whodunit: Transactional Profiling for Multi-Tier Applications,” *Proceedings of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems 2007 - EuroSys '07*, pp. 17–30, 2007.
- [s47] S. Chaudhuri, V. Narasayya, and M. Syamala, “Bridging the application and DBMS divide using static analysis and dynamic profiling,” *Proceedings of the 35th SIGMOD international conference on Management of data - SIGMOD '09*, p. 1039, 2009.
- [s48] D. Chen, N. Vachharajani, R. Hundt, X. Li, S. Eranian, W. Chen, and W. Zheng, “Taming Hardware Event Samples for Precise and Versatile Feedback Directed Optimizations,” *Computers, IEEE Transactions on*, vol. 62, no. 2, pp. 376–389, 2013.
- [s49] D. Chen, N. Vachharajani, R. Hundt, S.-w. Liao, V. Ramasamy, P. Yuan, W. Chen, and W. Zheng, “Taming hardware event samples for FDO compilation,” *Proceedings of the 8th annual IEEE/ACM international symposium on Code generation and optimization - CGO '10*, p. 42, 2010.
- [s50] L. Chen, Z. Cui, Y. Bao, M. Chen, Y. Huang, and G. Tan, “A lightweight hybrid hardware/software approach for object-relative memory profiling,” *2012 IEEE International Symposium on Performance Analysis of Systems & Software*, pp. 46–57, Apr. 2012.
- [s51] S. Chilingarova and V. O. Safonov, “Sampling profiler for Rotor as part of optimizing compilation system,” *Proceedings - 4th International Conference. NET Technologies 2006*, pp. 43–50, 2006.
- [s52] A. E. Chis, N. Mitchell, E. Schonberg, G. Sevitsky, P. O. Sullivan, T. Parsons, and J. Murphy, “Patterns of Memory Inefficiency,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6813, pp. 383–407, 2011.
- [s53] H. K. Cho, T. Moseley, R. Hank, D. Bruening, and S. Mahlke, “Instant Profiling : Instrumentation Sampling for Profiling Data-center Applications,” *International Symposium on Code Generation and Optimization - CGO '13*, p. Article in print, 2013.
- [s54] R. Chouhan, S. Roy, and S. Baswana, “Pertinent Path Profiling : Tracking Interactions among Relevant Statements,” *International Symposium on Code Generation and Optimization (CGO 2013)*, p. Article in print, 2013.
- [s55] E. D. Collins and B. P. Miller, “A Loop-Aware Search Strategy for Automated Performance Analysis,” *Lecture Notes in Computer Science*, vol. 3726, pp. 573–584, 2005.
- [s56] G. Cong, I.-h. Chung, H. Wen, D. Klepacki, H. Murata, Y. Negishi, and T. Moriyama, “A Holistic Approach towards Automated Performance Analysis and Tuning,” *Lecture Notes in Computer Science*, vol. 5704, pp. 33–44, 2009.
- [s57] E. Coppa, C. Demetrescu, and I. Finocchi, “Input-sensitive profiling,” *Proceedings of the 33rd ACM SIGPLAN conference on Programming Language Design and Implementation - PLDI '12*, pp. 89–98, 2012.
- [s58] M. Crasso, A. Zunino, L. Moreno, and M. Campo, “JEETuning-Expert: A software assistant for improving Java Enterprise Edition application performance,” *Expert Systems with Applications*, vol. 36, no. 9, pp. 11 718–11 729, Nov. 2009.
- [s59] D. C. D’Elia and C. Demetrescu, “Ball-Larus path profiling across multiple loop iterations,” *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications - OOPSLA '13*, pp. 373–390, 2013.
- [s60] D. C. D’Elia, C. Demetrescu, and I. Finocchi, “Mining hot calling contexts in small space,” *Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation - PLDI '11*, pp. 516–527, 2011.
- [s61] L. DeRose, K. Ekanadham, and S. Sbaraglia, “An Approach for Symbolic Mapping of Memory References,” *Lecture Notes in Computer Science*, vol. 3149, pp. 141–148, 2004.
- [s62] L. Derose and F. Wolf, “CATCH - A Call-Graph Based Automatic Tool for Capture of Hardware Performance Metrics for MPI and OpenMP Applications,” *Lecture Notes in Computer Science*, vol. 2400, pp. 167–176, 2002.
- [s63] L. A. DeRose, “The Hardware Performance Monitor Toolkit,” *Lecture Notes in Computer Science*, vol. 2150, pp. 122–132, 2001.
- [s64] D. V. Dincklage and A. Diwan, “Optimizing Programs with Intended Semantics,” *Proceedings of the 24th ACM SIGPLAN conference on Object oriented programming systems languages and applications - OOPSLA '09*, pp. 409–424, 2009.
- [s65] M. Dmitriev, “Profiling Java Applications Using Code Hotswapping and Dynamic Call Graph Revelation,” *Proceedings of the 4th international workshop on Software and performance - WOSP '04*, pp. 139–150, 2004.
- [s66] —, “Selective Profiling of Java Applications Using Dynamic Bytecode Instrumentation,” *2004 IEEE International Symposium on Performance Analysis of Systems and Software - ISPASS '04*, pp. 141–150, 2004.
- [s67] E. Duesterwald and M. Street, “Software Profiling for Hot Path Prediction: Less is More,” *Proceedings of the ninth international conference on Architectural support for programming languages and operating systems - ASPLOS '00*, pp. 202–211, 2000.
- [s68] B. Dufour, B. G. Ryder, and G. Sevitsky, “Blended analysis for performance understanding of framework-based applications,” *Proceedings of the 2007 international symposium on Software testing and analysis - ISSTA '07*, pp. 118–128, 2007.
- [s69] —, “A scalable technique for characterizing the usage of temporaries in framework-intensive Java applications,” *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of software engineering - SIGSOFT '08/FSE-16*, pp. 59–70, 2008.
- [s70] M. Dunlavey, “Performance Tuning with Instruction-Level Cost Derived from Call-Stack Sampling,” *ACM SIGPLAN Notices*, vol. 42, no. 8, pp. 4–8, 2007.
- [s71] Y. Endo and M. Seltzer, “Improving interactive performance using TIPME,” *Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems - SIGMETRICS '00*, pp. 240–251, Jun. 2000.
- [s72] M. Factor, A. Schuster, and K. Shagin, “Instrumentation of standard libraries in object-oriented languages: the twin class hierarchy approach,” *Proceedings of the 19th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications - OOPSLA '04*, pp. 288–300, 2004.
- [s73] T. Fahringer, M. Gerndt, G. Riley, and J. L. Träff, “On Performance Modeling for HPF Applications with ASL,” *Lecture Notes in Computer Science*, vol. 1940, pp. 191–204, 2000.
- [s74] T. Fahringer and C. Seragiotta, “Modeling and Detecting Performance Problems for Distributed and Parallel Programs with JavaPSL,” *Supercomputing Conference, SC2001*, no. November, p. 38, 2001.
- [s75] —, “Automatic Search for Performance Problems in Parallel and Distributed Programs by Using Multi-experiment Analysis,” *Lecture Notes in Computer Science*, vol. 2552, pp. 151–162, 2002.
- [s76] —, “Aksum: A Performance Analysis Tool for Parallel and Distributed Applications,” *Performance Analysis and Grid Computing*, pp. 189–208, 2004.
- [s77] U. Finkler, “An analytic framework for detailed resource profiling in large and parallel programs and its application for memory use,” *IEEE Transactions on Computers*, vol. 59, no. 3, pp. 358–370, 2010.
- [s78] N. Froyd, J. Mellor-Crummey, and R. Fowler, “Low-overhead call path profiling of unmodified, optimized code,” *Proceedings of the 19th annual international conference on Supercomputing - ICS '05*, p. 81, 2005.
- [s79] N. Froyd, N. Tallent, J. Mellor-Crummey, and R. Fowler, “Call path profiling for unmodified, optimized binaries,” *Proceedings of the GCC Developers’ Summit 2006*, pp. 21–35, 2006.
- [s80] Q. Fu, J.-G. Lou, Q.-W. Lin, R. Ding, D. Zhang, Z. Ye, and T. Xie, “Performance Issue Diagnosis for Online Service Systems,” *2012 IEEE 31st Symposium on Reliable Distributed Systems*, pp. 273–278, Oct. 2012.
- [s81] K. Fuerlinger and M. Gerndt, “Automated Performance Analysis Using ASL Performance Properties,” *Lecture Notes in Computer Science*, vol. 4699, pp. 390–397, 2007.

- [s82] K. Fuerlinger, M. Gerndt, and J. Dongarra, "On Using Incremental Profiling for the Performance Analysis of Shared Memory Parallel Applications," *Lecture Notes in Computer Science*, vol. 4641, pp. 62–71, 2007.
- [s83] K. Fuerlinger, D. Terpstra, H. You, P. Mucci, and S. Moore, "Enabling Data Structure Oriented Performance Analysis with Hardware Performance Counter Support," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 5415, pp. 263–272, 2009.
- [s84] A. Georges, D. Buytaert, L. Eeckhout, and K. De Bosschere, "Method-Level Phase Behavior in Java Workloads," *Proceedings of the 19th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications - OOPSLA '04*, pp. 270–286, 2004.
- [s85] M. Gerndt and E. Kereku, "Automatic Memory Access Analysis with Periscope," *Lecture Notes in Computer Science*, vol. 4488, pp. 847–854, 2007.
- [s86] —, "Search Strategies for Automatic Performance Analysis Tools," *Lecture Notes in Computer Science*, vol. 4641, pp. 129–138, 2007.
- [s87] J. Giménez, J. Labarta, F. X. Pegenau, H.-F. Wen, D. Klepacki, I.-H. Chung, G. Cong, F. Voigtländer, and B. Mohr, "Guided Performance Analysis Combining Profile and Trace Tools," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 6586, pp. 513–521, 2011.
- [s88] S. Goldsmith, R. O'Callahan, and A. Aiken, "Relational queries over program traces," *Proceedings of the 20th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications - OOPSLA '05*, vol. 40, no. 10, pp. 385–402, Oct. 2005.
- [s89] C. Gottbrath, "Cache-Related Performance Analysis Using Rogue Wave Software's ThreadSpotter," *Tools for High Performance Computing 2012*, pp. 75–93, 2013.
- [s90] G. Gousios and D. Spinellis, "Java Performance Evaluation Using External Instrumentation," *2008 Panhellenic Conference on Informatics*, pp. 173–177, Aug. 2008.
- [s91] S. L. Graham, P. B. Kessler, and M. K. Mckusick, "gprof: a Call Graph Execution Profiler," *ACM SIGPLAN Notices - Best of PLDI 1979-1999*, vol. 39, no. 4, pp. 49–57, 2004.
- [s92] J. Ha, M. Arnold, S. M. Blackburn, and K. S. McKinley, "A concurrent dynamic analysis framework for multicore hardware," *Proceedings of the 24th ACM SIGPLAN conference on Object oriented programming systems languages and applications - OOPSLA '09*, pp. 155–174, Oct. 2009.
- [s93] E. Hagersten, M. Nilsson, and M. Vesterlund, "Improving Cache Utilization Using Accumem VPE," *Tools for High Performance Computing*, pp. 115–135, 2008.
- [s94] R. J. Hall, "CPPROF: Aspect-capable Call Path Profiling of Multi-threaded Java Applications," *Proceedings of the 17th IEEE international conference on Automated software engineering - ASE '02*, pp. 107–116, 2002.
- [s95] —, "Aspect-Capable Call Path Profiling of Multi-Threaded Java Applications," *Automated Software Engineering*, vol. 12, no. 4, pp. 415–446, Oct. 2005.
- [s96] S. Han, Y. Dang, S. Ge, D. Zhang, and T. Xie, "Performance debugging in the large via mining millions of stack traces," *2012 34th International Conference on Software Engineering (ICSE)*, pp. 145–155, Jun. 2012.
- [s97] M. Hauswirth, P. F. Sweeney, A. Diwan, and M. Hind, "Vertical Profiling: Understanding the Behavior of Object-Oriented Applications," *Proceedings of the 19th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications - OOPSLA '04*, pp. 251–269, 2004.
- [s98] M. Hauswirth, P. F. Sweeney, A. Diwan, and M. C. Mozer, "Automating Vertical Profiling," *Proceedings of the 20th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications - OOPSLA '05*, pp. 281–296, 2005.
- [s99] O. Hernandez, H. Jin, and B. Chapman, "Compiler Support for Efficient Instrumentation," *Advances in Parallel Computing*, vol. 15, pp. 661–668, 2008.
- [s100] M. Hirzel and T. Chilimbi, "Bursty tracing: A framework for low-overhead temporal profiling," *4th ACM Workshop on Feedback-Directed and Dynamic Optimization - FDDO-4*, pp. 117–126, 2001.
- [s101] J. Huang and M. D. Bond, "Efficient context sensitivity for dynamic analyses via calling context uptrees and customized memory management," *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications - OOPSLA '13*, pp. 53–72, Nov. 2013.
- [s102] S.-H. Hung, S.-J. Huang, and C.-H. Tu, "New Tracing and Performance Analysis Techniques for Embedded Applications," *2008 14th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, pp. 143–152, Aug. 2008.
- [s103] H. Inoue and T. Nakatani, "How a Java VM can get more from a hardware performance monitor," *ACM SIGPLAN Notices*, vol. 44, no. 10, p. 137, Oct. 2009.
- [s104] M. Itzkowitz, B. J. N. Wylie, C. Aoki, and N. Kosche, "Memory Profiling using Hardware Counters," *2003 ACM/IEEE Supercomputing Conference - SC '03*, p. 17, 2003.
- [s105] R. Jalan and A. Kejariwal, "Trin-Trin: Who's Calling? A Pin-Based Dynamic Call Graph Extraction Framework," *International Journal of Parallel Programming*, vol. 40, no. 4, pp. 410–442, May 2012.
- [s106] G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu, "Understanding and Detecting Real-World Performance Bugs," in *Proceedings of the 33rd ACM SIGPLAN conference on Programming Language Design and Implementation - PLDI '12*, 2012, pp. 77–87.
- [s107] R. Joshi, M. D. Bond, and C. Zilles, "Targeted Path Profiling: Lower Overhead Path Profiling for Staged Dynamic Optimization Systems," *International Symposium on Code Generation and Optimization*, 2004. CGO 2004., pp. 239–250, 2004.
- [s108] G. Jost, O. Mazurov, and D. Mey, "Adding New Dimensions to Performance Analysis Through User-Defined Objects," *Lecture Notes in Computer Science*, vol. 4315, pp. 255–266, 2008.
- [s109] M. Jovic, A. Adamoli, and M. Hauswirth, "Catch Me If You Can: Performance Bug Detection in the Wild," in *Proceedings of the 2011 ACM international conference on Object oriented programming, systems, languages and applications - OOPSLA '11*, 2011, pp. 155–170.
- [s110] M. Jovic and M. Hauswirth, "Measuring the performance of interactive applications with listener latency profiling," *Proceedings of the 6th international symposium on Principles and practice of programming in Java - PPPJ '08*, pp. 137–146, 2008.
- [s111] —, "Listener latency profiling: Measuring the perceptible performance of interactive Java applications," *Science of Computer Programming*, vol. 76, no. 11, pp. 1054–1072, Nov. 2011.
- [s112] C. Jung, S. Rus, B. P. Railing, N. Clark, and S. Pande, "Brainy: Effective Selection of Data Structures," *Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation - PLDI '11*, pp. 86–97, 2011.
- [s113] E. Kereku, T. Li, M. Gerndt, and J. Weidendorfer, "A Data Structure Oriented Monitoring Environment for Fortran OpenMP Programs," *Lecture Notes in Computer Science*, vol. 3149, pp. 133–140, 2004.
- [s114] A. Ketterlin and P. Clauss, "Efficient memory tracing by program skeletonization," (*Lee Ispass*) *IEEE International Symposium on Performance Analysis of Systems and Software*, pp. 97–106, Apr. 2011.
- [s115] S. W. Kim, I. Park, and R. Eigenmann, "A Performance Advisor Tool for Shared-Memory Parallel Programming," *Lecture Notes in Computer Science*, vol. 2017, pp. 274–288, 2001.
- [s116] P. Kłaczewski and J. Wytrebowicz, "J2eeprof - a tool for testing multithreaded applications," *IFIP International Federation for Information Processing*, vol. 227, pp. 199–210, 2007.
- [s117] A. Knüpfer and W. E. Nagel, "Construction and Compression of Complete Call Graphs for Post-Mortem Program Trace Analysis," *International Conference on Parallel Processing*, 2005 - ICPP 2005., pp. 165–172, 2005.
- [s118] S. Koliai, Z. Bendifallah, M. Tribalat, C. Valensi, J.-T. Acquaviva, and W. Jalby, "Quantifying Performance Bottleneck Cost Through Differential Analysis," *Proceedings of the 27th International conference on supercomputing - ICS '13*, pp. 263–272, 2013.
- [s119] S. Koliai, S. Zuckerman, E. Oseret, M. Ivascot, T. Moseley, D. Quang, and W. Jalby, "A Balanced Approach to Application Performance Tuning," *Lecture Notes in Computer Science*, vol. 5898, pp. 111–125, 2010.
- [s120] R. Kufirin, "Perfsuite: An accessible, open source performance analysis environment for linux," *The 6th International Conference on Linux Clusters: The HPC Revolution 2005 (LCI-05)*, vol. 2005, no. April, 2005.
- [s121] T. Küstner, J. Weidendorfer, and T. Weinzierl, "Argument Controlled Profiling," *Lecture Notes in Computer Science*, vol. 6043, pp. 177–184, 2010.
- [s122] J. Labarta, "New Analysis Techniques in the CEPBA-Tools Environment," *Tools for High Performance Computing 2009*, pp. 125–143, 2010.

- [s123] J. Lambert and A. Podgurski, "xdProf: A tool for the capture and analysis of stack traces in a distributed Java system," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 4521, pp. 96–105, 2001.
- [s124] C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis & transformation," *Proceedings of the international symposium on Code generation and optimization: feedback-directed and runtime optimization - CGO '04*, pp. 75–86, 2004.
- [s125] B. Lee, K. Resnick, M. D. Bond, and K. S. McKinley, "Correcting the Dynamic Call Graph Using Control-Flow Constraints," *Lecture Notes in Computer Science*, vol. 4420, pp. 80–95, 2007.
- [s126] B. Li, L. Wang, H. Leung, and F. Liu, "Profiling all paths: A new profiling technique for both cyclic and acyclic paths," *Journal of Systems and Software*, vol. 85, no. 7, pp. 1558–1576, Jul. 2012.
- [s127] Z. Li and J. Tian, "A New Approach of Capturing System Call Context for Software Behavior Automaton Model," *Communications in Computer and Information Science*, vol. 159, pp. 148–153, 2011.
- [s128] S. Lin, F. Taiani, T. C. Ormerod, and L. J. Ball, "Towards Anomaly Comprehension: Using Structural Compression to Navigate Profiling Call-Trees," *Proceedings of the 5th international symposium on Software visualization - SOFTVIS '10*, pp. 103–112, 2010.
- [s129] X. Liu and J. Mellor-Crummey, "Pinpointing data locality problems using data-centric analysis," *International Symposium on Code Generation and Optimization (CGO 2011)*, pp. 171–180, Apr. 2011.
- [s130] —, "A data-centric profiler for parallel programs," *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis on - SC '13*, p. Article No. 28, 2013.
- [s131] P. Lokuciejewski, D. Cordes, H. Falk, and P. Marwedel, "A Fast and Precise Static Loop Analysis Based on Abstract Interpretation, Program Slicing and Polytope Models," *Proceedings of the 7th annual IEEE/ACM International Symposium on Code Generation and Optimization - CGO '09*, pp. 136–146, Mar. 2009.
- [s132] D. Lorenz, D. Böhme, B. Mohr, A. Strube, and Z. Szebenyi, "Extending Scalasca's Analysis Features," *Tools for High Performance Computing 2012*, pp. 115–126, 2013.
- [s133] C.-k. Luk, R. Cohn, R. Muth, H. Patil, A. Klauser, S. Wallace, V. J. Reddi, G. Lowney, and K. Hazelwood, "Pin: Building Customized Program Analysis Tools with Dynamic Instrumentation," *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation - PLDI '05*, pp. 190–200, 2005.
- [s134] J. Maebe, D. Buytaert, L. Eeckhout, and K. De Bosschere, "Javana: A System for Building Customized Java Program Analysis Tools," *Proceedings of the 21st annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications - OOPSLA '06*, pp. 153–168, 2006.
- [s135] J. a. P. Magalhães and L. M. Silva, "Root-cause analysis of performance anomalies in web-based applications," *Proceedings of the 2011 ACM Symposium on Applied Computing - SAC '11*, p. 209, 2011.
- [s136] A. D. Malony and S. S. Shende, "Overhead Compensation in Performance Profiling," *Lecture Notes in Computer Science*, vol. 3149, pp. 119–132, 2004.
- [s137] J. Marathe, F. Mueller, T. Mohan, B. R. de Supinski, S. A. McKee, and A. Yoo, "METRIC: Tracking Down Inefficiencies in the Memory Hierarchy via Binary Rewriting * C9," *Proceedings of the international symposium on Code generation and optimization: feedback-directed and runtime optimization - CGO '03*, pp. 289–300, 2003.
- [s138] D. Marinov and R. O'Callahan, "Object equality profiling," *Proceedings of the 18th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications - OOPSLA '03*, pp. 313–325, Nov. 2003.
- [s139] M. Marron, C. Sanchez, Z. Su, and M. Fahndrich, "Abstracting Runtime Heaps for Program Understanding," *IEEE Transactions on Software Engineering*, vol. 39, no. 6, pp. 774–786, Jun. 2013.
- [s140] C. McCurdy and J. Vetter, "Memphis: Finding and fixing NUMA-related performance problems on multi-core platforms," *2010 IEEE International Symposium on Performance Analysis of Systems & Software (ISPASS)*, pp. 87–96, Mar. 2010.
- [s141] J. Mellor-Crummey, R. J. Fowler, G. Marin, and N. Tallent, "HPCView : A Tool for Top-down Analysis of Node Performance," *The Journal of Supercomputing*, vol. 23, no. 1, pp. 81–104, 2002.
- [s142] M. Meyerhöfer and B. Volz, "EJBMProf - A Memory Profiling Framework for Enterprise JavaBeans," *Lecture Notes in Computer Science*, vol. 3489, pp. 17–32, 2005.
- [s143] H. Mi, H. Wang, H. Cai, Y. Zhou, M. R. Lyu, and Z. Chen, "P-Tracer: Path-Based Performance Profiling in Cloud Computing Systems," *2012 IEEE 36th Annual Computer Software and Applications Conference*, pp. 509–514, Jul. 2012.
- [s144] H. Mi, H. Wang, Y. Zhou, M. R.-T. Lyu, H. Cai, and G. Yin, "An online service-oriented performance profiling tool for cloud computing systems," *Frontiers of Computer Science*, vol. 7, no. 3, pp. 431–445, Feb. 2013.
- [s145] A. V. Mirgorodskiy and B. P. Miller, "CrossWalk: A Tool for Performance Profiling Across the User-Kernel Boundary," *Advances in Parallel Computing*, vol. 13, pp. 745–752, 2004.
- [s146] N. Mitchell, "The Runtime Structure of Object Ownership," *Lecture Notes in Computer Science*, vol. 4067, pp. 74–98, 2006.
- [s147] N. Mitchell, E. Schonberg, and G. Sevitsky, "Making Sense of Large Heaps," *Lecture Notes in Computer Science*, vol. 5653, pp. 77–97, 2009.
- [s148] N. Mitchell, G. Sevitsky, and H. Srinivasan, "The diary of a datum: an approach to modeling runtime complexity in framework-based applications," *Library-Centric Software Design - LCS'D'05*, p. 85, 2005.
- [s149] —, "Modeling Runtime Behavior in Framework-Based Applications," *Lecture Notes in Computer Science*, vol. 4067, pp. 429–451, 2006.
- [s150] N. Mitchell and P. F. Sweeney, "On-the-fly capacity planning," *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications - OOPSLA '13*, pp. 849–866, Nov. 2013.
- [s151] B. Mohr and F. Wolf, "KOJAK - A Tool Set for Automatic Performance Analysis of Parallel Programs," *Lecture Notes in Computer Science*, vol. 2790, pp. 1301–1304, 2003.
- [s152] P. Moret, W. Binder, D. Ansaloni, and A. Villazón, "Visualizing Calling Context Profiles with Ring Charts," *5th IEEE International Workshop on Visualizing Software for Understanding and Analysis VISSOFT '09*, pp. 33–36, 2009.
- [s153] P. Moret, W. Binder, and A. Villazón, "CCCP - Complete Calling Context Profiling in Virtual Execution Environments," *Proceedings of the 2009 ACM SIGPLAN workshop on Partial evaluation and program manipulation - PEPM '09*, pp. 151–160, 2009.
- [s154] P. Moret, W. Binder, A. Villazón, and D. Ansaloni, "Exploring large profiles with calling context ring charts," *Proceedings of the first joint WOSP/SIPEW international conference on Performance engineering - WOSP/SIPEW '10*, pp. 63–68, 2010.
- [s155] P. Moret, W. Binder, A. Villazón, D. Ansaloni, and A. Heydamoori, "Visualizing and exploring profiles with calling context ring charts," *Software: Practice and Experience*, vol. 40, no. 9, pp. 825–847, 2010.
- [s156] A. Morris, A. D. Malony, S. Shende, and K. Huck, "Design and Implementation of a Hybrid Parallel Performance Measurement System," *2010 39th International Conference on Parallel Processing*, pp. 492–501, Sep. 2010.
- [s157] A. Morris, W. Spear, A. D. Malony, and S. Shende, "Observing performance dynamics using parallel profile snapshots," *Lecture Notes in Computer Science*, vol. 5168, pp. 162–171, 2008.
- [s158] T. Moseley, A. Shye, V. J. Reddi, D. Grunwald, and R. Peri, "Shadow Profiling: Hiding Instrumentation Costs with Parallelism," *International Symposium on Code Generation and Optimization (CGO'07)*, pp. 198–208, Mar. 2007.
- [s159] N. Mostafa and C. Krintz, "Tracking performance across software revisions," *Proceedings of the 7th International Conference on Principles and Practice of Programming in Java - PPPJ '09*, pp. 162–171, 2009.
- [s160] J. Mußler, D. Lorenz, and F. Wolf, "Reducing the Overhead of Direct Application Instrumentation Using Prior Static Analysis," *Lecture Notes in Computer Science*, vol. 6852, pp. 65–76, 2011.
- [s161] T. Mytkowicz, D. Coughlin, and A. Diwan, "Inferred call path profiling," *ACM SIGPLAN Notices*, vol. 44, no. 10, p. 175, Oct. 2009.
- [s162] T. Mytkowicz, A. Diwan, M. Hauswirth, and P. F. Sweeney, "Evaluating the accuracy of Java profilers," *Proceedings of the 2010 ACM SIGPLAN conference on Programming language design and implementation - PLDI '10*, pp. 187–197, 2010.
- [s163] S. Nandy, X. Gao, and J. Ferrante, "TFP: Time-Sensitive, Flow-Specific Profiling at Runtime," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 2958, pp. 32–47, 2004.

- [s164] H. Narayanappa, M. S. Bansal, and H. Rajan, "Property-aware program sampling," *Proceedings of the 9th ACM SIGPLAN-SIGSOFT workshop on Program analysis for software tools and engineering - PASTE '10*, pp. 45–52, 2010.
- [s165] N. Nethercote and J. Seward, "Valgrind: A Framework for Heavyweight Dynamic Binary Instrumentation," *Proceedings of the 2007 ACM SIGPLAN conference on Programming language design and implementation - PLDI '07*, pp. 89–100, 2007.
- [s166] K. Nguyen and G. Xu, "Cachetor: detecting cacheable data to remove bloat," *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering - ESEC/FSE 2013*, pp. 268–278, 2013.
- [s167] A. Nistor, L. Song, D. Marinov, and S. Lu, "Toddler: Detecting performance problems via similar memory-access patterns," *2013 35th International Conference on Software Engineering (ICSE)*, pp. 562–571, May 2013.
- [s168] G. Novark, E. D. Berger, and B. G. Zorn, "Efficiently and precisely locating memory leaks and bloat," *Proceedings of the 2009 ACM SIGPLAN conference on Programming language design and implementation - PLDI '09*, pp. 397–407, 2009.
- [s169] R. Odaira and T. Nakatani, "Continuous object access profiling and optimizations to overcome the memory wall and bloat," *Proceedings of the seventeenth international conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS '12*, p. 147, 2012.
- [s170] D. Oren, Y. Matias, and M. Sagiv, "Online Subpath Profiling," *Lecture Notes in Computer Science*, vol. 2304, pp. 78–94, 2002.
- [s171] S. A. Ostadzadeh, M. Corina, C. Galuzzi, and K. Bertels, "Runtime extraction of memory access information from the application source code," *2011 International Conference on High Performance Computing & Simulation*, pp. 647–655, Jul. 2011.
- [s172] T. Parsons, "A Framework for Detecting Performance Design and Deployment Antipatterns in Component Based Enterprise Systems," *Proceedings of the 2nd international doctoral symposium on Middleware - DSM '05*, pp. 1–5, 2005.
- [s173] T. Parsons and J. Murphy, "Detecting Performance Antipatterns in Component Based Enterprise Systems," *Journal of Object Technology*, vol. 7, no. 3, pp. 55–90, 2008.
- [s174] E. Perelman, T. Chilimbi, and B. Calder, "Variational path profiling," *14th International Conference on Parallel Architectures and Compilation Techniques (PACT'05)*, pp. 7–16, 2005.
- [s175] A. Pesterev, N. Zeldovich, and R. T. Morris, "Locating cache performance bottlenecks using data profiling," *Proceedings of the 5th European conference on Computer systems - EuroSys '10*, p. 335, 2010.
- [s176] S. P. Reiss, "Controlled dynamic performance analysis," *Proceedings of the 7th international workshop on Software and performance - WOSP '08*, pp. 43–54, 2008.
- [s177] G. Ren, E. Tune, T. Moseley, Y. Shi, S. Rus, and R. Hundt, "Google-Wide Profiling: A Continuous Profiling Infrastructure for Data Centers," *IEEE Micro*, vol. 30, no. 4, pp. 65–79, 2010.
- [s178] M. Rohr, A. Van Hoorn, S. Giesecke, J. Matevska, W. Haselbring, and S. Alekseev, "Trace-context sensitive performance profiling for enterprise software applications," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 5119 LNCS, pp. 283–302, 2008.
- [s179] P. C. Roth and B. P. Miller, "Deep Start: A Hybrid Strategy for Automated Performance Problem Searches," *Lecture Notes in Computer Science*, vol. 2400, pp. 86–96, 2002.
- [s180] —, "Deep Start: a hybrid strategy for automated performance problem searches," *Concurrency and Computation: Practice and Experience*, vol. 15, no. 11-12, pp. 1027–1046, Sep. 2003.
- [s181] S. Roy and Y. Srikant, "Profiling k-Iteration Paths: A Generalization of the Ball-Larus Profiling Algorithm," *2009 International Symposium on Code Generation and Optimization*, pp. 70–80, Mar. 2009.
- [s182] Y. Ruan and V. Pai, "Making the 'Box' Transparent: System Call Performance as a First-class Result," *Proceedings of the annual conference on USENIX Annual Technical Conference - ATEC '04*, pp. 1–1, 2004.
- [s183] N. Rutar and J. K. Hollingsworth, "Assigning Blame : Mapping Performance to High Level Parallel Programming Abstractions," *Lecture Notes in Computer Science*, vol. 5704, pp. 21–32, 2009.
- [s184] —, "Data Centric Techniques for Mapping Performance Measurements," *2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*, pp. 1274–1281, May 2011.
- [s185] —, "Data centric techniques for mapping performance data to program variables," *Parallel Computing*, vol. 38, no. 1-2, pp. 2–14, Jan. 2012.
- [s186] —, "Software techniques for negating skid and approximating cache miss measurements," *Parallel Computing*, vol. 39, no. 3, pp. 120–131, Mar. 2013.
- [s187] J. Sánchez and A. González, "Analyzing Data Locality in Numeric Applications," *IEEE Micro*, vol. 20, no. 4, pp. 58–66, 2000.
- [s188] A. Sarimbekov, P. Moret, W. Binder, A. Sewe, and M. Mezini, "Complete and Platform-Independent Calling Context Profiling for the Java Virtual Machine," *Electronic Notes in Theoretical Computer Science*, vol. 279, no. 1, pp. 61–74, Dec. 2011.
- [s189] A. Sarimbekov, A. Sewe, W. Binder, P. Moret, and M. Mezini, "JP2: Call-site aware calling context profiling for the Java Virtual Machine," *Science of Computer Programming*, vol. 79, pp. 146–157, Jan. 2014.
- [s190] A. Sarimbekov, A. Sewe, W. Binder, P. Moret, M. Schoeberl, and M. Mezini, "Portable and accurate collection of calling-context-sensitive bytecode metrics for the Java virtual machine," *Proceedings of the 9th International Conference on Principles and Practice of Programming in Java - PPPJ '11*, p. 11, 2011.
- [s191] A. Sarimbekov, Y. Zheng, D. Ansaloni, L. Bulej, L. Marek, W. Binder, P. Tuma, and Z. Qi, "Productive Development of Dynamic Program Analysis Tools with DiSL," *2013 22nd Australian Software Engineering Conference*, pp. 11–19, Jun. 2013.
- [s192] S. S. Sastry, R. Bodik, and J. E. Smith, "Rapid Profiling via Stratified Sampling," *Proceedings - 28th Annual International Symposium on Computer Architecture - ISCA '01*, pp. 278–289, 2001.
- [s193] M. Schulz and B. R. de Supinski, "Practical Differential Profiling," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 4641, pp. 97–106, 2007.
- [s194] A. Sembrant, D. Black-Schaffer, and E. Hagersten, "Phase guided profiling for fast cache modeling," *Proceedings of the Tenth International Symposium on Code Generation and Optimization - CGO '12*, pp. 175–185, 2012.
- [s195] M. Serrano and X. Zhuang, "Building Approximate Calling Context from Partial Call Traces," *2009 International Symposium on Code Generation and Optimization*, pp. 221–230, Mar. 2009.
- [s196] H. Servat, G. Llort, J. Giménez, K. Huck, and J. Labarta, "Folding: Detailed Analysis with Coarse Sampling," *Tools for High Performance Computing 2011*, pp. 105–118, 2012.
- [s197] H. Servat, G. Llort, J. Giménez, and J. Labarta, "Detailed Performance Analysis Using Coarse Grain Sampling," *Lecture Notes in Computer Science*, vol. 6043, pp. 185–198, 2010.
- [s198] G. Sevitsky, W. de Pauw, and R. Konuru, "An information exploration tool for performance analysis of Java programs," *Proceedings Technology of Object-Oriented Languages and Systems. TOOLS 38*, pp. 85–101, 2001.
- [s199] O. Shacham, M. Vechev, and E. Yahav, "Chameleon: Adaptive Selection of Collections," *Proceedings of the 2009 ACM SIGPLAN conference on Programming language design and implementation - PLDI '09*, pp. 408–418, 2009.
- [s200] R. Shaham, E. K. Kolodner, and M. Sagiv, "Heap profiling for space-efficient Java," *Proceedings of the ACM SIGPLAN 2001 conference on Programming language design and implementation - PLDI '01*, pp. 104–113, 2001.
- [s201] K. Shen, C. Stewart, C. Li, and X. Li, "Reference-driven performance anomaly identification," *Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems - SIGMETRICS '09*, pp. 85–96, 2009.
- [s202] A. Shye, M. Iyer, T. Moseley, D. Hodgdon, D. Fay, V. J. Reddi, and D. A. Connors, "Analysis of Path Profiling Information Generated with Performance Monitoring Hardware," *Proceedings - Annual Workshop on Interaction between Compilers and Computer Architectures - INTERACT 2005*, pp. 34–45, 2005.
- [s203] J. M. Spivey, "Fast, accurate call graph profiling," *Software: Practice and Experience*, vol. 34, no. 3, pp. 249–264, Mar. 2004.
- [s204] K. Srinivas and H. Srinivasan, "Summarizing application performance from a components perspective," *Proceedings of the 10th European software engineering conference held jointly with 13th ACM SIGSOFT international symposium on Foundations of software engineering - ESEC/FSE-13*, pp. 136–145, Sep. 2005.

- [s205] V. St-Amour, S. Tobin-Hochstadt, and M. Felleisen, "Optimization Coaching: Optimizers Learn to Communicate with Programmers," *Proceedings of the ACM international conference on Object oriented programming systems languages and applications - OOPSLA '12*, pp. 163–178, 2012.
- [s206] W. N. Sumner, Y. Zheng, D. Weeratunge, and X. Zhang, "Precise calling context encoding," *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - ICSE '10*, vol. 1, pp. 525–534, 2010.
- [s207] —, "Precise Calling Context Encoding," *IEEE Transactions on Software Engineering*, vol. 38, no. 5, pp. 1160–1177, Sep. 2012.
- [s208] Z. Szebenyi, F. Wolf, and B. J. N. Wylie, "Space-efficient time-series call-path profiling of parallel applications," *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis - SC '09*, p. Article no. 37, 2009.
- [s209] Z. Szebenyi, T. Gamblin, M. Schulz, B. R. D. Supinski, F. Wolf, and B. J. Wylie, "Reconciling Sampling and Direct Instrumentation for Unintrusive Call-Path Profiling of MPI Programs," *2011 IEEE International Parallel & Distributed Processing Symposium*, pp. 640–651, May 2011.
- [s210] Z. Szebenyi, F. Wolf, and B. J. Wylie, "Performance Analysis of Long-Running Applications," *2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*, pp. 2105–2108, May 2011.
- [s211] S. Tallam, X. Zhang, and R. Gupta, "Extending path profiling across loop backedges and procedure boundaries," *International Symposium on Code Generation and Optimization, 2004. CGO 2004.*, pp. 251–262, 2004.
- [s212] N. R. Tallent, J. M. Mellor-Crummey, and M. W. Fagan, "Binary analysis for measurement and attribution of program performance," *Proceedings of the 2009 ACM SIGPLAN conference on Programming language design and implementation - PLDI '09*, pp. 441–452, May 2009.
- [s213] J. Tao, T. Gauger, and W. Karl, "A Profiling Tool for Detecting Cache-Critical Data Structures," *Lecture Notes in Computer Science*, vol. 4641, pp. 52–61, 2007.
- [s214] J. Tao and W. Karl, "Detailed cache simulation for detecting bottleneck, miss reason and optimization potentialities," *Proceedings of the 1st international conference on Performance evaluation methodologies and tools - ValueTools '06*, p. Article No. 62, 2006.
- [s215] —, "Supporting Cache Locality Optimization with a Toolset," *Lecture Notes in Computer Science*, vol. 4128, pp. 25–34, 2006.
- [s216] A. Traeger, I. Deras, and E. Zadok, "DARC: Dynamic Analysis of Root Causes of Latency Distributions," *Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems - SIGMETRICS '08*, pp. 277–288, 2008.
- [s217] J. Treibig, G. Hager, and G. Wellein, "LIKWID: Lightweight Performance Tools," *Competence in High Performance Computing 2010*, pp. 165–175, 2012.
- [s218] H.-L. Truong and T. Fahringer, "Soft computing approach to performance analysis of parallel and distributed programs," in *Lecture Notes in Computer Science*, vol. 3648, 2005, pp. 50–60.
- [s219] K. Vaswani, A. V. Nori, and T. M. Chilimbi, "Preferential Path Profiling: Compactly Numbering Interesting Paths," *Proceedings of the 34th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages - POPL '07*, pp. 351–362, 2007.
- [s220] K. Vaswani, M. J. Thazhuthaveetil, and Y. N. Srikant, "A Programmable Hardware Path Profiler," *Proceedings of the 2005 International Symposium on Code Generation and Optimization - CGO '05*, pp. 217–228, 2005.
- [s221] A. Villazón, W. Binder, and P. Moret, "Aspect Weaving in Standard Java Class Libraries," *Proceedings of the 6th international symposium on Principles and practice of programming in Java - PPPJ '08*, pp. 159–167, 2008.
- [s222] —, "Flexible calling context reification for aspect-oriented programming," *Proceedings of the 8th ACM international conference on Aspect-oriented software development - AOSD '09*, pp. 63–74, 2009.
- [s223] A. Villazón, W. Binder, P. Moret, and D. Ansaloni, "Comprehensive aspect weaving for Java," *Science of Computer Programming*, vol. 76, no. 11, pp. 1015–1036, 2011.
- [s224] S. Wallace and K. Hazelwood, "SuperPin: Parallelizing Dynamic Instrumentation for Real-Time Performance," *International Symposium on Code Generation and Optimization 2007 - CGO '07*, pp. 209–220, Mar. 2007.
- [s225] L. Wang and B. Li, "A Technology of Profiling Inter-procedural Paths," *Proceedings of the 23rd International Conference on Software Engineering and Knowledge Engineering - SEKE 2011*, pp. 31–36, 2011.
- [s226] S. Watterson and S. Debray, "Goal-Directed Value Profiling," *Lecture Notes in Computer Science*, vol. 2027, pp. 319–333, 2001.
- [s227] M. Weber, K. Mohror, M. Schulz, B. R. de Supinski, H. Brunst, and W. E. Nagel, "Alignment-Based metrics for trace comparison," *Lecture Notes in Computer Science*, vol. 8097, pp. 29–40, 2013.
- [s228] J. Weidendorfer, M. Kowarschik, and C. Trinitis, "A Tool Suite for Simulation Based Analysis of Memory Access Behavior," *Lecture Notes in Computer Science*, vol. 3038, pp. 440–447, 2004.
- [s229] A. Wert, J. Happe, and L. Happe, "Supporting swift reaction: Automatically uncovering performance problems by systematic experiments," *2013 35th International Conference on Software Engineering (ICSE)*, pp. 552–561, May 2013.
- [s230] J. Whaley, "A portable sampling-based profiler for Java virtual machines," *Proceedings of the ACM 2000 conference on Java Grande - JAVA '00*, pp. 78–87, 2000.
- [s231] D. Wiese, "Utilizing Data Marts for Performance Analysis," *18th GI-Workshop on the Foundations of Databases*, pp. 150–154, 2006.
- [s232] —, "Multidimensional Performance Management," *World Academy of Science, Engineering and Technology*, vol. 58, pp. 710–719, 2009.
- [s233] F. Wolf, B. Mohr, J. Dongarra, and S. Moore, "Efficient Pattern Search in Large Traces Through Successive Refinement," *Lecture Notes in Computer Science*, vol. 3149, pp. 47–54, 2004.
- [s234] B. Wu, M. Zhou, X. Shen, Y. Gao, R. Silvera, and G. Yiu, "Simple Profile Rectifications Go a Long Way," *Lecture Notes in Computer Science*, vol. 7920, pp. 654–678, 2013.
- [s235] Y. Wu and Y.-f. Lee, "Exploiting Free Execution Slots on EPIC Processors for Efficient and Accurate Runtime Profiling," *Lecture Notes in Computer Science*, vol. 3189, pp. 226–240, 2004.
- [s236] X. Xiao, S. Han, D. Zhang, and T. Xie, "Context-sensitive delta inference for identifying workload-dependent performance bottlenecks," *Proceedings of the 2013 International Symposium on Software Testing and Analysis - ISSTA 2013*, pp. 90–100, 2013.
- [s237] G. Xu, "Finding reusable data structures," *Proceedings of the ACM international conference on Object oriented programming systems languages and applications - OOPSLA '12*, p. 1017, 2012.
- [s238] —, "Resurrector: A Tunable Object Lifetime Profiling Technique for Optimizing Real-World Programs," *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications - OOPSLA '13*, pp. 111–130, 2013.
- [s239] G. Xu, M. Arnold, N. Mitchell, A. Rountev, and G. Sevitsky, "Go with the Flow: Profiling Copies To Find Runtime Bloat," *Proceedings of the 2009 ACM SIGPLAN conference on Programming language design and implementation - PLDI '09*, pp. 419–430, 2009.
- [s240] G. Xu, N. Mitchell, M. Arnold, A. Rountev, E. Schonberg, and G. Sevitsky, "Finding low-utility data structures," *Proceedings of the 2010 ACM SIGPLAN conference on Programming language design and implementation - PLDI '10*, pp. 174–186, 2010.
- [s241] G. Xu and A. Rountev, "Detecting inefficiently-used containers to avoid bloat," *Proceedings of the 2010 ACM SIGPLAN conference on Programming language design and implementation - PLDI '10*, pp. 160–173, May 2010.
- [s242] G. Xu, D. Yan, and A. Rountev, "Static Detection of Loop-Invariant Data Structures," *Proceedings of the 26th European conference on Object-Oriented Programming - ECOOP '12*, pp. 738–763, 2012.
- [s243] J. Xu, H. Zhang, and Q. Li, "An Execution Tracing Tool for Multi-tier Web," *Communications in Computer and Information Science*, vol. 153, pp. 244–250, 2011.
- [s244] D. Yan, G. Xu, and A. Rountev, "Uncovering performance problems in Java applications with reference propagation profiling," *2012 34th International Conference on Software Engineering (ICSE)*, pp. 134–144, Jun. 2012.
- [s245] T. Yasue, T. Suganuma, H. Komatsu, and T. Nakatani, "An efficient online path profiling framework for Java just-in-time compilers," *Proceedings of the 12th International Conference on Parallel Architectures and Compilation Techniques, 2003 - PACT '03.*, pp. 148–158, 2003.
- [s246] —, "Structural path profiling: An efficient online path profiling framework for just-in-time compilers," *Journal of Instruction-Level Parallelism*, vol. 6, no. April 2004, pp. 1–28, 2004.
- [s247] Z. Yu, W. Zhang, and X. Tu, "MT-Profiler: A Parallel Dynamic Analysis Framework Based on Two-Stage Sampling," *Lecture Notes in Computer Science*, vol. 6965, pp. 172–185, 2011.

- [s248] D. Zapanu and M. Hauswirth, "Algorithmic profiling," *Proceedings of the 33rd ACM SIGPLAN conference on Programming Language Design and Implementation - PLDI '12*, pp. 67–76, 2012.
- [s249] C. Zhang, H. Xu, S. Zhang, J. Zhao, and Y. Chen, "Frequency Estimation of Virtual Call Targets for Object-Oriented Programs," *Lecture Notes in Computer Science*, vol. 6813, pp. 510–532, 2011.
- [s250] H. Zhang, J. Ji, X. Zhou, H. Ma, and C. Wang, "Design and implementation of a configurable hardware profiler supporting path profiling and sampling," *2009 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery*, pp. 325–330, Oct. 2009.
- [s251] Q. Zhao, I. Cutcutache, and W.-f. Wong, "PiPA: Pipelined Profiling and Analysis on Multi-core Systems," *Proceedings of the 6th annual IEEE/ACM international symposium on Code generation and optimization - CGO '08*, pp. 185–194, 2008.
- [s252] X. Zhuang, S. Kim, M. Serrano, and J.-D. Choi, "PerfDiff: A Framework For Performance Difference Analysis In A Virtual Machine Environment," *Proceedings of the 6th annual IEEE/ACM international symposium on Code generation and optimization - CGO '08*, pp. 4–13, 2008.
- [s253] X. Zhuang, M. J. Serrano, H. W. Cain, and J.-D. Choi, "Accurate, efficient, and adaptive calling context profiling," *Proceedings of the 2006 ACM SIGPLAN conference on Programming language design and implementation - PLDI '06*, pp. 263–271, Jun. 2006.

APPENDIX B SYSTEMATIC MAPPING CATEGORISATION

TABLE 1
Key to Results Abbreviations

Contribution (Con)		Domain (Dm)	
Profiling Framework	PF	Object-Oriented	OO
Runtime Technique	RT	Java	J
Path Profiler	PP	.NET	N
Profiler	PR	HPC	H
Analysis Method	AM	Compile-time Optimisation	CO
Data Presentation	DP	Component-based Applications	CB
Automated search	AS	Enterprise Applications	E
Expert System	ES	Interactive Applications	I
		General Applications	G
Focus (Fs)		Feedback (FT)	
Execution cost	E	Summarised metrics	S
Runtime bloat	R	Derived metrics	M
Memory access	M	Code sites of interest	C
CPU efficiency	C	Performance visualisations	V
Request latency	L	Performance dependencies	D
Performance differencing	D	Object pooling opportunities	O
Value profiling	V	Identified bugs or antipatterns	B
System calls	S		
Not Applicable	-		
Analysis Type (AT)		Probe Implementation (PI)	
Static	S	Manual Analysis	MA
Dynamic	D	Manual Instrumentation	MI
Blended	B	Emulated Environment	EE
		Custom Software	CS
		Custom Hardware	CH
Probe Type (PT)			
Instrumentation	I	Static Instrumentation	SI
Sampling	S	Standard Monitors	MO
Instrumentation & Sampling	B	Dynamic Binary Instrumentation	DI
Sampled Instrumentation	SI	Agnostic	A
Agnostic	A	Not Applicable	-
Not Applicable	-		
Probe Overhead (PO)		Validation Approach (VA)	
Very Low (1-5%)	VL	Empirical Validation	E
Low (5-10%)	L	Demonstrated by Example	D
Medium (10-25%)	M	Both	B
Medium-High (25-100%)	MH	User Study	U
Java heap snapshot	JH	None	N
High (1-10x)	H		
Very High (10x or more)	VH		
Experimental by design	E	Evaluation Dataset (ED)	
Not stated	NS	Benchmark Suite	S
Agnostic	A	Real-world Examples	R
Not Applicable	-	Benchmarks & Real Examples	B
		Artificial Examples	A

TABLE 2. Systematic Mapping Results

Key	Year	Con	Dm	Fs	FT	AT	PT	PI	PO	VA	ED
[s1]	2010	DP	G	E	V	D	A	A	A	N	R
[s2]	2010	AM	I	L	M	D	B	DI	A	E	R
[s3]	2010	PR	H	C	S	D	S	DI	NS	N	A
[s4]	2010	DP	H	E	S	D	A	A	A	D	R
[s5]	2010	PR	OO	L	M	D	S	MO	VL	N	R
[s6]	2004	AM	OO	E	S	D	A	A	A	D	R
[s7]	2012	PF	J	-	-	D	I	DI	A	N	S
[s8]	2010	RT	J	-	-	D	I	DI	A	N	A
[s9]	2013	PF	J	-	-	D	I	DI	A	E	B
[s10]	2002	PP	G	E	S	D	I	SI	NS	E	R
[s11]	2011	PR	E	L	S	D	I	CS	MH	D	R
[s12]	2005	PR	OO	E	S	D	S	CS	VL	E	S
[s13]	2001	PR	G	E	S	D	SI	CS	L	E	S
[s14]	2000	PR	G	E	S	D	S	CS	NS	E	R
[s15]	2012	AS	G	E	M	D	I	CS	L	E	R
[s16]	2012	AM	G	E	M	D	A	A	A	E	R
[s17]	2013	DP	G	E	V	D	A	A	A	U	R
[s18]	2010	AS	H	E	M	D	A	MO	NS	N	R
[s19]	2002	PR	H	M	S	D	I	EE	VH	D	S
[s20]	2005	PR	G	M	S	D	S	MO	MH	E	S
[s21]	2011	PR	OO	E	S	D	I	CS	MH	E	R
[s22]	2012	DP	OO	E	V	D	A	A	A	D	R
[s23]	2010	DP	OO	E	V	D	A	A	A	D	R
[s24]	2007	PR	G	E	S	D	I	DI	L	D	R
[s25]	2006	PR	H	M	M	D	SI	SI	H	D	S
[s26]	2013	AM	OO	R	M	S	-	-	-	B	R
[s27]	2011	AM	OO	R	O	B	I	MO	NS	E	S
[s28]	2005	PR	J	E	S	D	B	DI	MH	E	S
[s29]	2005	PR	J	R	S	D	I	DI	H	N	S
[s30]	2006	PR	J	E	S	D	B	DI	MH	E	S
[s31]	2009	PR	J	E	S	D	I	DI	H	E	S
[s32]	2011	RT	J	-	-	D	I	DI	A	N	S
[s33]	2006	PF	J	-	-	D	I	SI	A	N	A
[s34]	2006	PR	J	S	S	D	I	MO	M	N	S
[s35]	2007	RT	J	-	-	D	I	DI	A	N	S
[s36]	2009	PR	J	E	S	D	I	DI	H	E	S
[s37]	2010	RT	G	-	-	D	I	CS	M	E	S
[s38]	2005	PP	CO	E	S	D	SI	CS	VL	E	S
[s39]	2005	PP	CO	E	S	D	I	SI	L	E	S
[s40]	2007	RT	OO	-	-	D	I	CS	VL	E	S
[s41]	2000	PR	CO	V	S	D	S	MO	L	D	B
[s42]	2009	AM	G	E	S	S	-	-	-	E	S
[s43]	2000	AS	H	E	M	D	I	DI	NS	E	R
[s44]	2012	PR	H	C	S	D	S	MO	NS	N	R
[s45]	2012	PR	G	M	S	D	I	DI	VH	D	S
[s46]	2007	PR	E	L	S	D	B	CS	L	D	R
[s47]	2009	PR	G	E	S	D	I	SI	NS	N	A
[s48]	2013	PR	CO	E	S	D	S	MO	VL	B	S
[s49]	2010	PR	CO	E	S	D	S	MO	VL	B	S
[s50]	2012	PR	G	M	S	D	I	CH	VL	D	S
[s51]	2006	PR	N	E	S	D	S	CS	VL	E	R
[s52]	2011	AM	OO	R	B	D	S	MO	JH	E	R
[s53]	2013	PF	CO	-	-	D	SI	DI	L	E	B
[s54]	2013	PP	G	E	S	D	I	SI	NS	E	S
[s55]	2005	AS	H	E	M	D	I	DI	NS	E	R
[s56]	2009	ES	H	E	B	B	A	A	A	D	R
[s57]	2012	PR	G	E	M	D	I	DI	VH	B	B
[s58]	2009	ES	CB	E	B	D	I	CS	MH	B	R
[s59]	2013	PP	G	E	S	D	I	CS	H	E	S
[s60]	2011	RT	G	-	-	D	A	A	A	E	R
[s61]	2004	PR	H	M	S	D	I	SI	NS	N	A
[s62]	2002	PR	H	C	S	D	I	DI	NS	N	A
[s63]	2001	PR	H	C	S	D	I	MI	NS	N	A
[s64]	2009	AM	OO	E	C	S	-	-	-	B	R
[s65]	2004	PR	J	E	S	D	I	CS	VH	E	B
[s66]	2004	PR	J	E	S	D	I	CS	VH	E	B
[s67]	2000	PP	CO	E	S	D	I	DI	NS	E	S
[s68]	2007	PR	OO	R	M	B	I	MO	NS	N	R
[s69]	2008	PR	OO	R	M	B	I	MO	NS	N	R
[s70]	2007	AM	G	E	S	D	S	MO	NS	D	A
[s71]	2000	PR	I	L	S	D	I	CS	VL	D	R
[s72]	2004	RT	J	-	-	D	I	DI	A	N	S
[s73]	2000	AM	H	E	M	D	A	A	A	N	A
[s74]	2001	AS	H	E	M	D	I	SI	E	N	R
[s75]	2002	AS	H	E	M	D	I	SI	E	N	R
[s76]	2004	AS	H	E	M	D	I	SI	E	N	R
[s77]	2010	PR	G	R	S	D	SI	CS	L	N	R
[s78]	2005	PR	G	E	S	D	S	MO	L	E	S
[s79]	2006	PR	G	E	S	D	S	MO	L	N	S
[s80]	2012	AM	E	L	D	D	S	A	A	E	R
[s81]	2007	AS	H	E	M	D	A	SI	NS	N	S
[s82]	2007	PR	H	E	S	D	I	SI	NS	N	S
[s83]	2009	PR	H	C	S	D	I	MI	VL	D	A

Key	Year	Con	Dm	Fs	FT	AT	PT	PI	PO	VA	ED
[s84]	2004	PR	OO	C	S	D	I	CS	L	N	S
[s85]	2007	AS	H	M	M	D	A	SI	NS	N	R
[s86]	2007	AS	H	E	M	D	I	SI	NS	N	S
[s87]	2011	AM	H	C	S	D	B	A	A	N	R
[s88]	2005	PF	OO	-	-	D	I	SI	VH	N	B
[s89]	2013	PR	H	M	B	D	S	MO	NS	N	A
[s90]	2008	PR	J	E	S	D	I	DI	NS	N	A
[s91]	2004	PR	G	E	S	D	B	SI	MH	N	A
[s92]	2009	PF	G	-	-	D	B	CS	A	E	S
[s93]	2008	PR	H	M	B	D	S	MO	NS	D	S
[s94]	2002	PR	J	E	S	D	S	MO	L	D	R
[s95]	2005	PR	J	E	S	D	S	MO	L	D	R
[s96]	2012	AM	I	L	S	D	S	A	A	U	R
[s97]	2004	PR	OO	C	S	D	I	CS	M	D	S
[s98]	2005	AM	OO	C	S	D	I	MO	A	E	S
[s99]	2008	PF	H	-	-	B	I	SI	A	N	S
[s100]	2001	PR	CO	E	S	D	SI	SI	M	E	S
[s101]	2013	RT	G	-	-	D	I	CS	A	E	S
[s102]	2008	PR	G	E	S	D	I	SI	NS	N	A
[s103]	2009	PR	J	C	S	D	S	CS	VL	E	B
[s104]	2003	PR	H	M	S	D	I	MO	NS	D	R
[s105]	2012	PR	G	E	S	D	I	DI	H	B	B
[s106]	2012	AM	G	E	C	S	-	-	-	E	R
[s107]	2004	PP	CO	E	S	D	I	SI	M	E	S
[s108]	2008	AM	H	E	M	D	A	A	NS	D	R
[s109]	2011	PR	I	L	S	D	B	MO	VL	D	R
[s110]	2008	PR	I	L	S	D	I	SI	VL	E	R
[s111]	2011	PR	I	L	S	D	I	SI	VL	D	R
[s112]	2011	AM	G	E	B	D	I	SI	NS	B	R
[s113]	2004	PR	H	M	S	D	I	CH	NS	D	A
[s114]	2011	PR	G	M	S	B	I	DI	H	E	S
[s115]	2001	ES	H	E	B	B	A	A	A	D	R
[s116]	2007	PR	CB	E	S	D	I	DI	VH	N	R
[s117]	2005	RT	H	-	-	D	I	A	A	E	R
[s118]	2013	AM	H	C	S	B	I	MO	NS	N	B
[s119]	2010	AM	H	M	S	B	I	MO	NS	D	R
[s120]	2005	PR	H	C	S	D	B	MO	NS	N	A
[s121]	2010	PR	G	E	S	D	I	DI	NS	D	R
[s122]	2010	AM	H	E	V	D	B	SI	A	N	R
[s123]	2001	PR	J	E	S	D	S	MO	L	N	S
[s124]	2004	PF	G	-	-	B	I	SI	A	N	S
[s125]	2007	PR	OO	E	S	B	S	CS	NS	E	S
[s126]	2012	PP	G	E	S	D	I	SI	H	E	B
[s127]	2011	RT	G	-	-	D	A	A	A	E	R
[s128]	2010	AM	G	E	S	D	A	A	A	U	R
[s129]	2011	PR	G	M	S	D	B	DI	M	D	R
[s130]	2013	PR	H	M	S	D	B	DI	L	D	S
[s131]	2009	AM	G	E	S	S	-	-	-	E	S
[s132]	2013	AM	H	E	M	D	I	SI	A	N	S
[s133]	2005	PF	G	-	-	D	I	DI	A	E	S
[s134]	2006	PF	J	-	-	D	I	CS	A	N	S
[s135]	2011	AM	E	L	D	D	SI	DI	VL	E	A
[s136]	2004	RT	H	-	-	D	I	SI	A	E	S
[s137]	2003	PR	G	M	S	D	SI	DI	NS	D	A
[s138]	2003	PR	OO	R	O	D	I	SI	VH	B	S
[s139]	2013	AM	N	R	M	D	SI	SI	VH	D	S
[s140]	2010	PR	H	M	S	D	S	MI	NS	D	B
[s141]	2002	AM	H	C	S	D	B	MO	A	D	R
[s142]	2005	PR	CB	R	S	D	I	MO	VH	N	R
[s143]	2012	PR	E	L	S	D	SI	MI	VL	D	R
[s144]	2013	PR	E	L	S	D	SI	MI	VL	D	R
[s145]	2004	PR	G	S	S	D	I	DI	M	D	R
[s146]	2006	AM	OO	R	M	D	S	MO	JH	E	B
[s147]	2009	AM	OO	R	M	D	S	MO	JH	N	R
[s148]	2005	AM	OO	E	M	D	B	MA	E	N	R
[s149]	2006	AM	OO	E	S	D	B	MA	E	N	S
[s150]	2013	AM	OO	L	M	D	S	MO	VL	E	R
[s151]	2003	AS	H	E	B	D	I	SI	A	N	A
[s152]	2009	DP	OO	E	V	D	A	A	A	N	A
[s153]	2009	PR	J	E	S	D	I	DI	H	E	B
[s154]	2010	DP	G	E	V	D	A	A	A	D	S
[s155]	2010	DP	J	E	V	D	A	A	A	D	B
[s156]	2010	PR	H	C	S	D	B	SI	L	N	R
[s157]	2008	PR	H	E	S	D	I	SI	NS	N	R
[s158]	2007	PF	CO	-	-	D	SI	DI	VL	E	S
[s159]	2009	AM	G	D	C	D	I	CS	E	E	B
[s160]	2011	RT	H	-	-	B	I	SI	A	E	S
[s161]	2009	RT	G	-	-	D	A	A	VL	E	S
[s162]	2010	PR	J	E	S	D	S	MO	NS	B	S
[s163]	2004	PP	CO	E	S	D	I	SI	L	D	B
[s164]	2010	PR	G	E	S	B	SI	SI	NS	E	B
[s165]	2007	PF	G	-	-	D	I	DI	A	N	B
[s166]	2013	PR	J	R	O	D	I	CS	VH	B	S
[s167]	2013	PR	G	V	M	D	I	SI	VH	E	R
[s168]	2009	PR	OO	R	S	D	S	CS	VL	E	B

Key	Year	Con	Dm	Fs	FT	AT	PT	PI	PO	VA	ED
[s169]	2012	PR	OO	R	O	D	SI	CS	VL	E	S
[s170]	2002	PP	CO	E	S	D	SI	SI	MH	E	B
[s171]	2011	PR	H	M	S	D	I	DI	NS	N	R
[s172]	2005	ES	CB	E	B	D	I	DI	NS	N	A
[s173]	2008	ES	CB	E	B	D	I	DI	NS	E	R
[s174]	2005	PP	G	E	M	D	SI	SI	VL	E	B
[s175]	2010	PR	G	M	S	D	S	CS	VL	D	R
[s176]	2008	PR	G	E	S	D	B	DI	VL	N	R
[s177]	2010	PR	E	C	S	D	S	MO	VL	N	A
[s178]	2008	AM	E	E	S	D	SI	SI	NS	E	R
[s179]	2003	AS	H	E	C	D	B	DI	NS	E	R
[s180]	2003	AS	H	E	C	D	B	DI	NS	E	R
[s181]	2009	PP	CO	E	S	D	I	SI	NS	N	S
[s182]	2004	PR	G	S	S	D	I	CS	VL	D	R
[s183]	2009	PR	H	E	S	B	S	SI	NS	N	R
[s184]	2011	PR	H	E	S	B	S	SI	NS	N	R
[s185]	2012	PR	H	E	S	B	S	SI	NS	N	R
[s186]	2013	PR	G	M	S	B	S	SI	NS	E	S
[s187]	2000	PR	G	M	S	B	I	SI	L	E	R
[s188]	2011	PR	J	E	S	D	I	DI	VH	N	S
[s189]	2012	PR	J	E	S	D	I	DI	H	N	S
[s190]	2011	PR	J	E	S	D	I	DI	VH	N	S
[s191]	2013	PF	J	-	-	D	I	DI	A	U	A
[s192]	2001	PF	CO	-	-	D	S	CH	VL	E	S
[s193]	2007	AM	G	D	S	D	A	A	A	N	R
[s194]	2012	PR	G	M	S	D	S	MO	MH	E	S
[s195]	2009	RT	G	-	-	D	A	A	A	E	S
[s196]	2012	PR	H	E	S	D	B	MO	VL	E	B
[s197]	2010	PR	H	E	S	D	B	MO	VL	N	R
[s198]	2001	AM	J	E	S	D	A	A	A	N	R
[s199]	2009	PR	OO	E	B	D	B	CS	H	B	B
[s200]	2001	PR	J	R	S	D	I	CS	NS	D	B
[s201]	2009	AM	G	D	M	D	A	A	A	D	R
[s202]	2005	PP	CO	E	S	D	S	MO	VL	E	S
[s203]	2004	PR	G	E	S	D	I	SI	MH	N	R
[s204]	2005	AM	G	E	S	D	A	A	A	E	R
[s205]	2012	AM	G	E	C	S	-	-	-	D	S
[s206]	2012	RT	G	-	-	D	I	SI	VL	E	B
[s207]	2012	RT	G	-	-	D	I	SI	VL	E	B
[s208]	2009	RT	H	-	-	D	A	A	A	E	S
[s209]	2011	PR	H	E	S	D	B	SI	VL	B	S
[s210]	2011	PR	H	E	S	D	B	SI	NS	N	S
[s211]	2004	PP	CO	E	S	D	I	SI	MH	E	S
[s212]	2009	PR	G	C	S	D	B	DI	L	E	B
[s213]	2007	PR	H	M	S	D	S	MO	L	N	S
[s214]	2006	PR	G	M	S	D	I	SI	NS	N	R
[s215]	2006	PR	G	M	M	D	I	SI	NS	D	A
[s216]	2008	AS	G	E	C	D	I	DI	VL	N	R
[s217]	2012	PR	H	C	S	D	I	MO	VL	D	R
[s218]	2005	AM	H	E	M	D	A	A	A	N	R
[s219]	2007	PP	CO	E	S	D	I	SI	M	E	S
[s220]	2005	PP	CO	E	S	D	I	CH	VL	E	S
[s221]	2008	PF	J	-	-	D	I	DI	A	E	S
[s222]	2009	RT	J	-	-	D	I	DI	A	E	B
[s223]	2011	PF	J	-	-	D	I	DI	A	E	A
[s224]	2007	PF	G	-	-	D	I	DI	A	E	S
[s225]	2011	PP	G	E	S	D	I	SI	NS	E	A
[s226]	2001	PR	CO	V	S	D	I	SI	H	B	S
[s227]	2013	AM	H	D	S	D	I	A	A	N	R
[s228]	2004	PR	G	M	S	D	I	DI	NS	D	R
[s229]	2013	AS	J	E	B	D	I	DI	E	B	R
[s230]	2000	PR	J	E	S	D	S	CS	VL	E	S
[s231]	2006	AM	E	E	S	D	A	A	A	N	A
[s232]	2009	AM	E	E	S	D	A	A	A	N	R
[s233]	2004	AM	H	E	B	D	I	SI	A	E	R
[s234]	2013	AM	CO	E	S	D	S	A	A	E	S
[s235]	2004	PF	CO	-	-	D	I	CH	VL	N	S
[s236]	2013	AM	I	L	M	D	I	SI	E	E	R
[s237]	2012	PR	OO	R	O	D	I	CS	MH	B	S
[s238]	2013	PR	OO	R	O	D	I	CS	VH	B	S
[s239]	2009	PR	OO	R	M	D	I	CS	VH	D	S
[s240]	2010	PR	OO	R	M	D	I	CS	VH	D	S
[s241]	2010	AM	OO	R	C	B	I	SI	NS	B	S
[s242]	2012	AM	OO	E	O	S	-	-	-	D	S
[s243]	2011	PR	E	L	S	D	I	CS	L	N	R