# <sup>1</sup> Access Permission-based Program Verification: A survey

Ayesha Sadiq<sup>\*</sup>, Yuan-Fang Li<sup>\*</sup>, Sea Ling<sup>\*</sup>

Faculty of Information Technology, Monash University, Clayton, Australia

# 4 Abstract

2

3

Verifying the correctness and reliability of imperative and object-oriented 5 programs is one of the grand challenges in computer science. In imperative programming models, programmers introduce concurrency manually by us-7 ing explicit concurrency constructs such as multi-threading. Multi-threaded 8 programs are prone to synchronization problems such as data races and dead-9 locks, and verifying API protocols in object-oriented programs is a non-trivial 10 task due to the improper and unexpected state transition at run-time. This 11 is in part due to the unwanted sharing of program states in such programs. 12 With these considerations in mind, access permissions have been investigated 13 as a means to reasoning about the correctness of such programs. Access per-14 missions are abstract capabilities that characterize the way a shared resource 15 can be accessed by multiple references. 16

This paper provides a comprehensive survey of existing access permission-17 based verification approaches. We describe different categories of permissions 18 and permission-based contracts. We elaborate how permission-based speci-19 fications have been used to ensure compliance of API protocols and to avoid 20 synchronization problems in concurrent programs. We compare existing ap-21 proaches based on permission usage, analysis performed, language and/or 22 tool supported, and properties being verified. Finally, we provide insight 23 into the research challenges posed by existing approaches and suggest future 24 directions. 25

Keywords: Access permissions, program verification, concurrency, protocol
 verification, permission inference, survey

yuanfang.li@monash.edu (Yuan-Fang Li), chris.ling@monash.edu (Sea Ling)

Preprint submitted to The Journal of Systems and Software

<sup>\*</sup>Corresponding author

Email addresses: ayesha.sadiq@monash.edu (Ayesha Sadiq),

### 28 1. Introduction

Correctness and reliability of software programs written in imperative and 29 object-oriented languages such as Java and C++ have always been a major 30 challenge for the IT industry. This is because of the implicit dependencies 31 that exist between the code and the program states. In imperative programs 32 different program parts may access the same mutable state, without expos-33 ing this information to each other and, consequently, may cause unwanted 34 interference or inconstant states. Preventing such errors is important to 35 ensure the compliance of API (Application Programming Interface) proto-36 cols in object-oriented programs and to verify the correctness of increasingly 37 ubiquitous multi-threaded applications. 38

Modern object-oriented programs are highly reliant on reusable APIs that 39 often define usage protocols i.e., the desired sequence of method calls that 40 API clients must follow for underlying objects to work properly. Typestates 41 (Strom and Yemini, 1986) have been designed to specify usage protocols and 42 verify their behavior. A typestate abstractly defines an object's state at ex-43 ecution time. However, statically tracking object state is a non-trivial task 44 because of unexpected transitions between states during program execution. 45 In multi-threaded programs, managing synchronization between threads is 46 a complicated and challenging task for programmers due to thread inter-47 leaving and heap interference, which can lead to concurrency problems such 48 as deadlocks, data races. The situation becomes worse in the presence of 49 unrestricted aliasing, the hallmark feature of imperative and object-oriented 50 languages (Bierhoff et al., 2009b). 51

Earlier work on verifying correctness of sequential programs dates back 52 to Hoare's logic (Hoare, 1969) that defines a set of axioms (calculus) and 53 formal inference rules to specify and verify desired properties and static anal-54 ysis techniques such as (Floyd, 1967). Handling thread non-interference for 55 Java-like concurrent programs dates back to Owicki-Gries' axiomatic method, 56 (Owicki and Gries, 1976), and Jones' rely-guarantee principle (Jones, 1983). 57 These approaches are considered the traditional and general ways of perform-58 ing shared-memory program verification. Since the seminal work of Hoare 59 (Hoare, 1969) and Floyd (Floyd, 1967), many logic-based verification ap-60 proaches and type-effect systems have been developed to avoid problems such 61 as data races, deadlocks, atomicity violations in shared-memory programs. 62 The commonly used verification approaches either perform deductive ver-63

<sup>64</sup> ification or employ theorem proving techniques where formal correctness

proofs are used to verify program properties based on the input specifications 65 (Huisman, 2001; Flanagan et al., 2003; Abadi et al., 2006; OHearn, 2007; Vil-66 lard et al., 2010; Caires and Seco, 2013). Others conduct dynamic analysis of 67 the input program through model checking (Visser et al., 2003; Chaki et al., 68 2004; Chaki and Gurfinkel, 2018) or perform static analysis (Boyapati et al., 69 2002; Engler and Ashcraft, 2003; Voung et al., 2007; Naik et al., 2009; Dias 70 et al., 2013) to approximate runtime behavior of the program to verify its 71 correctness. 72

However, in the last decades, static contract checking based on Hoare 73 logic and the development of advanced, simplified, automated program veri-74 fiers (Fähndrich and Logozzo, 2011; Pradel et al., 2012; Fillitre and Paskevich, 75 2013; Carr et al., 2017) had been an active areas of research for the verifica-76 tion of program behavior. Although these approaches are usable and quite 77 promising, the support for concurrency and aliasing is limited. As most of the 78 real-world applications are inherently multi-threaded, the next step was to 79 develop tools and techniques that can reason about shared-memory programs 80 and control aliasing in a sound and efficient way. With these considerations in 81 mind, permission-based program logics and tools became influential because 82 of their practicality, expressiveness and strong reasoning power to handle 83 both aliasing and concurrency. 84

Access permission, formally called Boyland's fractional permission (Boy-85 land, 2003), is a formalism inspired by Linear Logic (Girard, 1987) and Sep-86 aration Logic (O'Hearn et al., 2001; Reynolds, 2002). The former treats per-87 missions as linear resources and the latter simplifies the specifications and 88 verification of shared-memory programs in an efficient way. Fractional per-89 missions were originally proposed to verify non-interference of the program 90 states in parallel programs, using either read or write accesses on the refer-91 enced objects. The formalism was later extended by Bornat et al. (2005) to 92 allow read sharing of the shared program states. Bierhoff and Aldrich (2007) 93 extended fractional permissions as symbolic permissions to model both the 94 read/write operations and aliasing information of a program state at one 95 place. 96

A study of the literature shows that access permissions have been investigated to address different concerns related to security, concurrency, and protocol verification. Notable threads of research include Plural (Bierhoff and Aldrich, 2007, 2008; Beckman, 2009), Chalice (Leino et al., 2009; Leino and Müller, 2009), VeriFast (Jacobs et al., 2011), VPerm (Le et al., 2012), Pulse (Siminiceanu et al., 2012), Sample (Ferrara and Müller, 2012), Plaid (Aldrich et al., 2012), Æminium (Stork et al., 2014), VerCors (Amighi et al., 2012,
2014), and Viper (Müller et al., 2017), to name a few.

This survey summarizes the pragmatics of different access permission sharing and accounting models, discussed in detail in Section 3 and 5.1, that have been used in the literature to verify the correctness and behavior of sequential and concurrent programs, based on access permissions. To the best of our knowledge, this survey is the first attempt to provide readers a comprehensive overview of existing access permission-based verification approaches and tools since their introduction.

We categorize the research work covered in this survey along the following three dimensions:

- Verification of API protocols in typestate-based sequential and concurrent programs.
- 2. Verification of common concurrency problems such as race conditions
   and deadlock etc., in concurrent (multi-threaded) programs.
- Automatic inference of access permissions for sequential and concurrent
   programs.

Within each of the above categories, the existing approaches are compared and contrasted based on the following criteria:

- The type of underlying program (Prog) such as sequential or concurrent program.
- The programming language (Language) used or developed as a specification language.
- The realization of the proposed technique in the form of a tool (Tool ).
- The type of analysis (Analysis) i.e. static or dynamic performed.
- The kind of permission abstraction (Perm-Kind) such as fractional, 130 counting or symbolic permissions, supported as a part of specifications.
- The access notations or contracts (Perm-Specs) specified as program annotations.
- The permissions or access notations (Perm-Infer) inferred.

- The annotation overhead (Anno) (if any) posed by the approach.
- 135 136

• The functional or behavioral properties (**Properties**) verified, based on the permission-based specifications.

Section 2 briefly discusses the related formal theories and type systems in 137 the literature as seminal and background work to program verification. Sec-138 tion 3 provides an overview of access permissions. Sections 4, 5, 6 covers the 139 three categories of the state of access permission-based program verification. 140 mentioned previously. We elaborate some approaches with sample code snip-141 pets to show how different types of access permissions are manually added 142 in the program to specify design intents and verify program behavior gainst 143 specifications. The objective is to analyze the annotation overhead posed by 144 these approaches. Table ?? provides the more detailed statistics, in terms 145 of annotation overhead, of the existing approaches from the surveyed papers 146 where possible. Further, in every section, a brief summary of the studies work 147 in chronological order (where possible), following the above mentioned crite-148 ria, is given in tabular form. Finally, Section 7 provides an insight into the 149 use of access permission-based specifications, the research challenges posed 150 by the existing approaches and suggest future research directions. 151

# <sup>152</sup> 2. Related Formalisms to Program Verification and Parallelization

This section briefly presents other formal type theories for program verification and parallelization. A type system can verify the desired interaction between system components as types can classify program entities and the permissible results of the computations. The beauty of type systems is the assurance that if a program is type-checked then it is guaranteed to be free from errors.

Earlier work on type systems mainly focused on the results of a computation in a program in terms of its correctness. Then the study of type discipline and concurrency theory inspired the development of formal type systems that can statically formulate and verify the intended properties of a program behavior, along with the permissible results of the computations.

# 164 Behavioral Type

Behavioral type theory is one such formalism that was originally proposed to verify concurrent programs based on process algebra (Nielson and Nielson,

1993, 1996). A behavioural type system uses behavioral types, a type-based 167 abstraction, to formally describe the software entities such as communication 168 protocols, interfaces, web services and contracts, as a sequence of operations 169 in a concurrent and distributed environment. Formally speaking, a behav-170 ioral type system is a compositional type system that can directly model 171 the interaction between system components as a notion of choice, causality 172 and resource usage. Session types and behavioural contracts are two notions 173 related to behavioural types. 174

Since the introduction of behavioral types, many type-based effect and 175 proof systems, using the concepts of behavioural types, session types, spatial 176 logic and processes as types, have been developed to study various correctness 177 and behavioral properties such as unique receptiveness, race freedom, dead-178 locks, livelocks in large-scale concurrent and distributed systems (Sangiorgi, 179 1999; Chaki et al., 2002; Igarashi and Kobayashi, 2005; Dezani-Ciancaglini 180 et al., 2005; Kobayashi and Sangiorgi, 2010). However, in these approaches, 181 type-based specifications are explicitly added to the model and verify the 182 usage patterns of resources and communication objects. The behavioral 183 type theory then subsequently integrated with session types (Igarashi and 184 Kobayashi, 2001; Chaki et al., 2002; Igarashi and Kobayashi, 2005; Dezani-185 Ciancaglini et al., 2005) where type-based specifications are explicitly added 186 to model and verify the usage patterns of communication objects in concur-187 rent and distributed environment. 188

Recent work on behavioral separation in a distributed environment can be 189 attributed to Caires (2008) who developed a spatial-behavioural typing sys-190 tem, to model the resource independence and synchronization in a distributed 191 (concurrent) environment, based on the Spatial logic (Caires and Cardelli, 192 2003, 2002). The type system using parallel and sequential composition op-193 erators and resource ownership are handled using the type modality. Later, 194 Caires and Seco (2013) developed a behavioral separation programming lan-195 guage based on  $\lambda$ -calculus to ensure the disciplined interference between re-196 sources in higher-order concurrent programs having fork/join parallelism. 197 The language is based on the behavioral type systems (Honda et al., 1998; 198 Chaki et al., 2002) that incorporates a behavioral view of the program prop-190 erties and employs Concurrent Separation Logic (Reynolds, 2002; OHearn, 200 2007), to separate the dynamic behavior of run-time values rather than sepa-201 rating program states itself. However, program effects are computed based on 202 the explicitly specified assertions to ensure the safety of concurrent programs. 203 A detailed study of behavioral type systems and related methodologies can 204

#### $_{205}$ be found in Ancona et al. (2016).

#### 206 Session Type

Session type, a notion of behavioral types, was originally introduced by Honda et al. (1998) to ensure the disciplined interaction between two partners in a distributed environment and later extended in the work of Honda (Honda et al., 2008) to incorporate the arbitrary number of participants in the same environment.

Recently, session types were extended with formal type system to control 212 aliasing and to enforce usage protocols in a concurrent environment for Java 213 and Java-like languages, such as Session J (Hu et al., 2008, 2010), Yak (Milito 214 and Caires, 2009) and Mungo (Gay et al., 2015a), to name a few. Further, 215 the linearity of resources is an important and recurring theme in concurrency 216 that studies behavioral type systems, for process calculi and as mentioned 217 by Honda (1993), Linear Logic can be considered as a source of inspiration 218 for some aspects of session types. 219

Recent work on Linear Logic-based session types includes the work of 220 Caires and Pfenning (2010), who proposed a Curry-Howard style interpreta-221 tion of binary session types in intuitionistic Linear Logic, to expose a deep 222 relationship between both concepts. Recently, with this line of work, Gay 223 and Vasconcelos (2010) and Wadler (2014) developed a new calculus CP and 224 a linear functional language GV, to establish a connection between session 225 types and the Linear Logic, that can yield a process calculus, free from data 226 races and deadlocks. A comprehensive study of program verification ap-227 proaches based on the session and behavioral types can be found in (Hüttel 228 et al., 2016). 229

#### 230 Typestate

Yet another important formalism, a notion of behavioural types, is type-231 state. Typestate was first defined by Strom and Yemini (1986) as a new 232 programming language concept that determines the operations permitted on 233 objects in a given context. According to Garcia et al. (2014) "typestate re-234 flects how legal operations on objects can change at run-time as their internal 235 state changes". Typestate associates state information with a variable of a 236 given type, which is then subsequently used to decide the valid operations to 237 be called on an instance of that type. Typestate is suitable to represent re-238 sources that follow state transition systems that follow the 'open then close' 239

semantics. For example, a database connection can only execute a database
command if it is in the open state.

Typestates were originally developed for imperative languages without 242 the notion of objects, but later extended with the behavioural-type discipline, 243 to support verification of object-oriented languages such as Vault (DeLine 244 and Fähndrich, 2002), Fugue (DeLine and Fähndrich, 2004). Furthermore, 245 typestates were integrated as a first-class language construct in typestate-246 oriented language (Aldrich et al., 2009; Sunshine et al., 2011) to verify the 247 correctness of the usage protocols in state-based sequential programs. In the 248 new language, objects are being modeled not only as classes, but also the 249 changing abstract states, where the correctness of the program is determined 250 by tracking state transitions between different objects at execution time, 251 thereby ensuring the correct usage of protocols. Later, the typestate-oriented 252 programming led the development of a gradual typestate system (Garcia 253 et al., 2014) that integrates access permissions with gradual types (Siek and 254 Taha, 2007) to control aliasing in a more robust way. 255

Recent developments in state-based protocol checking and verification in-256 cludes approaches (Caires and Seco, 2013; Garcia et al., 2014; Militão et al., 257 2014b; Gay et al., 2015b), that handle aliasing in a more robust way and ver-258 ify communication protocols in distributed and concurrent object-oriented 259 languages. These approaches ensure the basic memory safety conservatively, 260 by associating typestate invariants with each referenced location and by en-261 suring that the type invariants hold for every store of this location. However, 262 all the approaches discussed above focus on identifying violations of protocols 263 but none of them check the higher-level (behavioral) characteristics of the 264 protocols themselves, for example, their usage in practice and the complex-265 ity associated with their definition, that are vital in verifying correctness of 266 many program properties (Beckman et al., 2011). 267

# 268 Ownership Type

Another stream of type-based systems is ownership type, a mechanism to express the sharing of program references (Noble et al., 1998; Clarke et al., 271 2013), in the way that allows controlled aliasing between objects by mitigat-272 ing the undesirable effects to other objects.

Since its introduction, ownership types have been used in many formal approaches to provide safe aliasing control mechanisms (Boyapati and Rinard, 275 2001; Clarke and Wrigstad, 2003; Müller and Rudich, 2007; Cohen et al., 276 2009), and to control object deallocation explicitly (Matsakis et al., 2014). However, in ownership-based verification approaches, programmers explicitly define ownership invariants as locking or access information to specify
dependencies at the object level and use this information to avoid data races.
A complete stream of research on ownership types can be found in the work
of Clarke et al. (2013).

#### 282 Atomic Set

Atomic sets have been used to detect atomicity violation and to avoid data races in concurrent programs. An atomic set defines a set of memory locations that share some consistency property and needs to be updated atomically. The atomic set can be viewed as a generalization of Hoare's monitors (Hoare, 1974) to multiple objects. They can be better integrated into the Java language.

Earlier work using atomic sets dates back to Vaziri et al. (2006)'s data cen-280 tric programming model that defines atomic set serializability. Atomic 290 set serializability is a disciplined interference criterion to avoid the problem-291 atic interleaving scenarios in the shared-memory programs based on atomic 292 Vaziri extended his previous work to support multiple object intersets. 293 ference (Vaziri et al., 2010). Subsequent work used atomic sets to detect 294 atomicity violation statically (Kidd et al., 2011) and dynamically (Xu et al., 295 2005; Hammer et al., 2008; Lai et al., 2010). The trend is followed by many 296 local data-centric concurrency control mechanism and type systems (Dolby 297 et al., 2012; Marino et al., 2013), to verify program behavior based on atomic 298 sets. 299

Data-centric concurrency control is one alternative to the explicit locking 300 mechanism. In contrast to the control-centric synchronization approaches, 301 (Artho et al., 2003; Lu et al., 2008), where each program instruction is pro-302 tected by synchronization constraints and then changes to the program states 303 are tracked for every execution path of the program flow. The local data-304 centric approaches combine all fields of an object that require consistency, 305 for all the control flow paths of program execution, into an atomic set and 306 updates them atomically to avoid data races. 307

Contrary to data-centric concurrency control and the use of atomic sets, the development of type systems (Flanagan et al., 2003; Abadi et al., 2006; Flanagan et al., 2008) has been influential to ensure the atomicity and datarace freedom in concurrent programs, and to reduce the annotation overhead associated with manually adding synchronization primitives at code level. However, unlike atomic sets, in type systems, programmers provide explicit synchronization primitives, as locking and guarded specifications at field orclass level, required by the code.

Recently, the notion of atomic sets was replaced by atomic variables by 316 Paulino et al. (2016). The objective was to handle the complexity asso-317 ciated with the use of memory structures in atomic sets. The proposed 318 approach applies a resource-centered view of the data-centric concurrency 319 control. However, in the proposed type system, programmers explicitly de-320 fine the synchronization primitives on the individual data items that require 321 atomic updates, to guarantee the progress of synchronization for all program 322 execution scenarios. 323

#### 324 Uniqueness and Immutability

Another area of interest has been the controlled sharing and interference of object references in imperative object-oriented programs. Sharing is a situation when a piece of memory is accessed by more than one reference, say x and y, so that a change to x affects y as well. Therefore, changes to one object may leave other objects in an inconsistent state, causing unwanted interference and subsequently, data races.

Work has been done to restrict the usage of references notably using 331 access-based type annotations such as uniqueness, immutability and read-only 332 (Clarke and Wrigstad, 2003; Boyland, 2006, 2010; Gordon et al., 2012; Cleb-333 sch et al., 2015). The objective was to identify isolated states that can be 334 safely handled by one or more threads, thereby avoiding the unwanted inter-335 ference and alternatively data races. However, the type system infers sharing 336 effects such as uniqueness and immutability for the object references by 337 computing an equivalence relationship for a set of free variables by evaluating 338 the input expression. The inferred effects are then used to determine which 339 part of the code can be safely shared between multiple threads to maintain 340 the integrity of the data. 341

Recent development in this area is the Giannini's type and effect system (Giannini et al., 2018a,b) that expresses sharing in imperative programs based on the **pure** calculus (Capriccioli et al., 2016), where memory stores are modeled by rewriting the source code terms rather than by modifying the auxiliary storage.

### 347 Rely-guarantee Protocols

The logic-based program verification that employs rely-guarantee reasoning has been another active area of research that verifies the correctness of usage protocols and avoids inter-thread interference in concurrent programs
(Parkinson and Bierman, 2005; Vafeiadis and Parkinson, 2007; DinsdaleYoung et al., 2010). However, in these approaches, rely-guarantee specifications are explicitly added at the state or thread level to model and control
the concurrent interactions safely.

The most recent is a sub-structural type system (Militão et al., 2014a) 355 that uses **rely-gurantee** protocol abstraction to model the interfering in-356 teraction of aliases to the shared states. The type system then ensures that 357 the aliases always get a determined value regardless of the potential changes 358 made by the program context during interleaving. However, the system ex-359 plicitly assigns a separate role to each alias with a "rely  $\Rightarrow$  guarantee" re-360 lation between aliases. Later, Militão et al. (2016) extended the approach 361 and developed a composition procedure based on linear capabilities (Mor-362 risett et al., 2005), to address the decidability of protocol composition and 363 its integration with the protocol abstraction. 364

#### 365 Separation Logic

Among other logic-based approaches for data race freedom, Separation Logic that is based on Hoare logic, attained much attention for controlling aliasing and verifying program behavior.

Hoare's logic for conditional critical regions (Hoare, 1972) and monitors 369 (Hoare, 1974) was widely adopted because of the simplicity and practicality 370 of their use in Java-like programs. Hoare's logic limits thread interference to 371 a few synchronization points. However, it cannot syntactically enforce a safe 372 monitor synchronization. This is because of the potential risk of aliasing in a 373 Java program where multiple threads can manipulate shared-memory data in 374 an unsafe manner. O'Hearn (O'Hearn et al., 2001) and Reynolds (Reynolds, 375 2002) extended the Hoare's logic to Classic Separation Logic, a new program 376 logic with new connectives and separation conjunction (\*), to reason about 377 the sequential programs that manipulate pointer data structures. 378

Hoare logic uses triples  $\{P\}S\{Q\}$  where P and Q are predicates over 379 program states that define the required and ensured properties of an ex-380 pression statement S. However, in Separation Logic, the idea is to explicitly 381 divide each program state, related to a current method call, into a heap and 382 a store part, to allow explicit local reasoning about the heap memory. In this 383 approach, the heap h is divide into two disjoint parts say  $h_1$  and  $h_2$  using 384 separation formula of the form  $\phi_1 * \phi_2$  where  $\phi_1$  is a pointer valid for the 385 part  $h_1$  and  $\phi_2$  is a pointer valid for the part  $h_2$ . The conjunction \* operator 386

combines two disjoint parts of the same heap. The idea of using separation formula to verify heap structure, dates back to the seminal work of Reynolds (Reynolds, 1978) who proposed a syntactic interference control mechanism to constrain the effects of interference in Algol-like languages using the Separation Logic. The separation formula ensures that two threads accessing the same location do not interfere to verify program behavior

Eventually, Separation Logic was realized as a new program logic called 393 Concurrent Separation Logic (CSL) (Brookes, 2004; OHearn, 2007) to rea-394 son about multi-threaded programs, with an assumption that if two threads 395 can operate on disjoint parts of the same heap location without interfering 396 with each other, they can be verified in a safe and isolated way. CSL enforces 397 correct synchronization of the shared-memory data logically, rather than syn-398 tactically. The idea of CSL was then extended in several sub-structural type 399 systems and concurrent approaches (Gotsman et al., 2007; Appel and Blazy, 400 2007: Hobor et al., 2008). to guarantee data race freedom in the shared-401 memory concurrent programs, and have recently been applied in high-order 402 imperative concurrent languages and type systems (Schwinghammer et al., 403 2011; Jensen and Birkedal, 2012). However, in the separation logics-based 404 approaches, the separation predicates are explicitly specified in the program 405 to define the access rights on the memory locations. 406

# 407 A Move to Permission-based Specifications

Among all formal approaches to the verification of shared-memory programs such as atomic set, behavioral type, session-type, relay-guarantee reasoning and Separation Logic, is access permission. Access permission is an abstract capability that combines type (effect) systems and, provides more advanced support for reasoning about the heap resources (Boyland, 2003).

The notion of access permissions is built on Linear Logic (Girard, 1987). 413 that treats permissions as linear resources, and the Classic Separation Logic 414 (O'Hearn et al., 2001) that performs local reasoning of program behavior 415 against specifications. However, Classic Separation Logic does not support 416 the concurrent read access of a memory location by multiple references or 417 threads. Therefore, Boyland (2003) and Bornat et al. (2005) combined Sep-418 aration Logic with abstract capabilities, called access permissions, to allow 419 concurrent reading of a program state. 420

421 Compared to classic verification methods such as Owicki-Gries (Owicki 422 and Gries, 1976) for concurrent programs, the permission-based Separation 423 Logic ensures that: a) only one reference (thread) can write on a particular

location at any given time, thereby mutating data in a safe way; b) if a 424 location is read by a thread, all other threads can only have read permission 425 for that location, thereby implying data race freedom without the need to 426 explicitly check for the interference between threads in concurrent programs. 427 Plural (Bierhoff and Aldrich, 2007; Beckman et al., 2008), a permission-428 based program verifier, was the next development phase where access per-429 missions were combined with typestate abstractions to statically ensure the 430 mutability of object's states, following the Separation Logic, and to verify 431 protocol compliance in Java-like sequential and concurrent programs. 432

Access permission was then subsequently used in many formal approaches (Huisman and Mostowski, 2015; ?) to identify and ensure the mutability of object's states and to verify program behavior in shared-memory programs.

# 436 3. Access Permissions: An Overview

Access permissions are abstract capabilities that can encode effects (read
and write) and aliasing information of a referenced object at one place. There
are three main categories of access permissions:

**Fractional Permissions** (Boyland, 2003). A fractional (share) permission 440 say s is a concrete mathematical value that defines a shared ownership 441 (access) of referenced objects  $\circ$  in a concurrent setting where s is a 442 fraction between 0 and 1 inclusive. A value 0 represents absence of 443 permission, while a value 1 represents full permission and any value 444 greater than zero represents a shared read-only access on the referenced 445 object. Fractional permissions can be used to split a full permission 446 (with value 1) into number of fractions and then to distribute these 447 fractions among multiple references, and so on. The splitting function 448 for a full permission, say s, when divided between two references, say 449 s1 and s2, can be written as  $s_1 + s_2 = s$  with each reference having a 450 fraction of  $\mathbf{s}$  in the range (0, 1). 451

**Counting Permissions** (Bornat et al., 2005). A counting permission is a special fractional permission where **s** is an integer value between 0 and a maximum constant value, where zero represents absence of permission and the maximum value represents **full** permission on the referenced object **o**. The read-only access on the referenced object **o** is represented by a non-zero integer value such that  $0 < s \leq max$ .

Symbolic permissions (Bierhoff and Aldrich, 2007). Symbolic permis-458 sions, simple called access permissions, are extension of Boyland's frac-459 tional permission sharing model but, instead of using concrete frac-460 tional value to represent and split permissions among multiple refer-461 ences, symbolic permissions represent and track permission flow through 462 the system using permission types such as unique or immutable etc. 463 Access permissions splitting and joining rules for symbolic permissions 464 are given in Section 3.2. 465

There are five types of symbolic permissions that can be assigned to a 466 reference x, for a referenced object o, in the presence of its alias y. 467

unique(x): This permission provides reference **x** an exclusive read and modify access on the referenced object o at any given time.

468 No other reference (e.g. y) to the same object can co-exist while x has unique permission on an object. 469





full(x): This permission grants reference x with read and write access to the referenced object o, and at the same time o may also be

471 read, but not written, by other references such as y. 472

473

share(x): This permission is the same as full(x), except that now other references 474 such as y can also write on the referenced object o. 475

476

pure(x): This permission gives reference x a read but not a write access on a referenced object o. Moreover, other references such as 477 y may have both read and write access on the same object. 478

479

480

immutable(x): This permission grants a nonmodifying access on the referenced object o

to both the current reference  $\mathbf{x}$  and any other reference such as y.









Table 1 summarises how access permissions can co-exist on a referenced object o by the current reference (*This Reference* x), and by other reference (*Other Reference* y).

This Reference <b>x</b>	Access Rights	Other Reference y
unique	read/write	none
full	read/write	pure
share	read/write	share, pure
pure	read	full, pure, immutable
immutable	read	immutable, pure

Table 1: Co-existing access permissions (Bierhoff and Aldrich, 2007)

# 484 3.1. Access Permission Contracts in the Spirit of Linear Logic and the Design 485 by Contract Principle

Linear logic (Girard, 1987) traditionally treats access permissions as re-486 sources that cannot be duplicated (discarded). Access permission contracts 487 in Linear Logic are specified using the Linear Logic implication connective 488  $(-\infty)$ . The connective  $(-\infty)$  operator is used to specify a method's pre- and 489 post-conditions in Linear Logic. As indicated by  $P \rightarrow Q$ , permissions in the 490 pre-conditions P are consumed before a method runs, and it produces Q as 491 post-conditions when the method completes its execution. Once a method 492 consumes its permissions they are no longer available to other methods until 493 the method returns the same permissions again. 494

In the Design by Contract Principle (Meyer, 1988), contracts are obligations and rights of the client and the implementing class itself. Contracts are specified using the **requires** and the **ensures** clauses that represent a method's pre and post-conditions respectively (Meyer, 1992; Leavens et al., 2006).

In the spirit of the Design by Contract Principle, permission-based specifications at method level represent contracts where permission-based obligations are defined as pre-conditions P that client of a class must guarantee before calling methods of the class, and permission-based rights represent post-conditions Q that must hold for both the client and the implementing class after executing the specified method. The idea of specifying pre- and post-conditions as contracts dates back to Hoare's work (Hoare, 1969) on formal verification of software applications and has recently been applied to
permission-based verification and parallelization approaches (Cataño et al.,
2014; Militao et al., 2010; Huisman and Mostowski, 2015; ?).

# 510 3.2. Access Permission Splitting and Joining Rules

Access permission can be split into one or more relaxed permissions (frac-511 tions of original permission, using fractional values in the range (0, 1) and 512 then merged back into more restrictive or original permission. This phe-513 nomenon is known as fractional permission analysis where fractions keep 514 tracks of the way the permissions were split and joined back. This informa-515 tion can be used to verify system properties based on certain specific criteria 516 and to parallelise execution of a program by tracking permission-based de-517 pendencies in the program. Table 2 shows access permissions splitting and 518 joining rules.

Splitting and joining Rules	Rule #
unique(x;o;k) $\Leftrightarrow$ full( $x_1$ ;o; $k_1$ ) $\bigotimes$ pure( $x_2$ ;o; $k_2$ )	Rule I
unique(x;o;k) $\Leftrightarrow$ immutable( $x_1$ ;o; $k_1$ ) $\bigotimes$ immutable( $x_2$ ;o; $k_2$ )	Rule II
$\operatorname{full}(\mathbf{x};\mathbf{o};\mathbf{k}) \Leftrightarrow \operatorname{share}(x_1;\mathbf{o};k_1) \bigotimes \operatorname{pure}(x_2;\mathbf{o};k_2)$	Rule III
share(x;o;k) $\Leftrightarrow$ full( $x_1$ ;o; $k_1$ ) $\bigotimes$ pure( $x_2$ ;o; $k_2$ )	Rule IV
immutable(x;o;k) $\Leftrightarrow$ pure( $x_1$ ;o; $k_1$ ) $\bigotimes$ immutable( $x_2$ ;o; $k_2$ )	Rule V
unique(x;o;k) $\Leftrightarrow$ share( $x_1$ ;o; $k_1$ ) $\bigotimes$ share( $x_2$ ;o; $k_2$ )	Rule VI
immutable(x;o;k) $\Leftrightarrow$ immutable(x_1;o;k_1) $\bigotimes$ immutable(x_2;o;k_2)	Rule VII
share(x;o;k) $\Leftrightarrow$ share( $x_1$ ;o; $k_1$ ) $\bigotimes$ pure( $x_2$ ;o; $k_2$ )	Rule VIII
share(x;o;k) $\Leftrightarrow$ share( $x_1$ ;o; $k_1$ ) $\bigotimes$ share( $x_2$ ;o; $k_2$ )	Rule X
$\operatorname{full}(\mathbf{x};\mathbf{o};\mathbf{k}) \Leftrightarrow \operatorname{full}(x_1;\mathbf{o};k_1) \bigotimes \operatorname{pure}(x_2;\mathbf{o};k_2)$	Rule XI

Table 2: Access permissions splitting and joining rules (Bierhoff and Aldrich, 2007).

519

In Table 2, let x represent current reference, o represent the referenced object and k represent the fraction of permission assigned to a particular

reference, where at least one of  $x_1$  and  $x_2$  is x, and  $k_1 + k_2 = k$ . The opera-522 tor multiplicative conjunction (A  $\bigotimes$  B) denotes simultaneous occurrence of 523 permissions by multiple references, say  $x_1, x_2$ , on the same referenced object 524 o. The symbol  $\Leftrightarrow$  represents the two way operation of splitting and join-525 ing permissions. For example, a unique access permission (Rule-I) having 526 **k** fractions can be divided into  $k_1$  fraction of full and  $k_2$  fraction of pure 527 permission and then joined back accordingly. Likewise, a unique access per-528 mission (Rule VI) can be split into two **share** permissions but cannot be split 529 into a share and immutable permission as immutable cannot co-exist with 530 the share permission. Linearity of resources forces the unique permission to 531 be replaced by two **share** permissions which can be further split according 532 to splitting rules and then joined back. 533

# 534 4. Permission-based Verification of API Protocols

This section introduces the first dimension of the state of the art permissionbased verification approaches in detail. The focus is on the verification of API protocols for single- and multi-threaded programs. Table 3 provides a summary of the permission-based protocol verification approaches studied in this research.

Ref.	Prog	Lang	Tool	Analy	Perm	-Perm-	Perm-	Anno	Properties
					Kind	Specs	Infer		
Bierhoff and Aldrich (2007)	Seq	Plural (NSL)	Plural	(St,D)	Sym	(U,S,F,P,I	) N	Υ	VoUP
Beckman et al. (2008)	Con	Plural (NSL)	Plural	$\operatorname{St}$	Sym	(U,S,F,P,I)	) N	Υ	VoUP, RCwAB
Beckman (2009)	Con	Plural (NSL)	Sync-or -Swim	St	Sym	(U,I,S,F,P	) N	Υ	VoUP, RCwSB
Militao et al. (2010)	Seq	-	Plural	$\operatorname{St}$	Sym	(U,I,S,F,P	) N	Υ	RCs
Bierhoff (2011)	Seq	Java-lil (NSL)	<sup>ce</sup> JavaSyp	$\operatorname{St}$	Sym	(U,I)	Ν	Υ	CME
Aldrich et al. (2011) Aldrich et al. (2012)	Seq	$\frac{\text{Plaid}}{(\text{NSL})}$	Plaid	(St, D)	Sym	(U,S,I)	Ν	Υ	VoUP
Naden et al. (2012)	Seq& Con	Plaid (NSL)	Plaid	(St,D)	Sym	(U,S,I)	Ν	Υ	RCs

Table 3: Access permission-based protocol verification.

Keys to the table: Seq = sequential, Con = concurrent, St = static, D = dynamic, Sym = symbolic, Fract = fractional, U = unique, I = immutable, S = share, F = full, P = pure, NSL = new specification language, RCs = race conditions, VoUP = verification of usage protocls, RCwAB = race conditions with atomic blocks, RCwSB = race conditions with synchronized blocks, CME = concurrent modification exceptions,  $\mathbb{Z}$  set of Integers,  $\mathbb{R}$  set of Real numbers, N set of positive Integers.

# 540 4.1. Verification of API protocols in Single-threaded Programs

In object-oriented programs, objects define usage proto- cols. Usage pro-541 tocols are constraints on the order of method invocations that the client of 542 the protocol must follow for the underlying objects to work properly. Type-543 states have been used to specify usage protocols in many formal approaches 544 where state information is associated with a variable of a given type, which 545 is subsequently tracked to decide the valid operation sequence to be called 546 on an instance of that type. It is generally acknowledged that statically 547 tracking object's state is a challenging task in the presence of unrestricted 548 aliasing. This can be attributed to the improper and unexpected state tran-549 sition during program execution, and that can happen in situations such as 550 a) when a method uses a data structure having aliases deeply nested in the 551 hierarchy and causes side effects, or b) when aliased parameters are passed to 552 a method expecting non-aliased parameters. Access permissions have been 553 used, as part of formal specifications, to specify the intended design and 554 to verify the correctness of usage protocols in API protocols in many for-555 mal approaches such as Plural (Bierhoff and Aldrich, 2007, 2008; Bierhoff 556 and Kevin, 2009; Bierhoff et al., 2009b), JavaSyp (Bierhoff, 2011) and Plaid 557 (Aldrich et al., 2011, 2012) and other related techniques. 558

# 559 Typestate Verification using Linear Logic

Bierhoff and Kevin (2009) presented a formal specification language and a type system to soundly and modularly verify API protocols in sequential programs based on access permissions. The aim was twofold, firstly to verify protocol conformance with actual program implementation in the presence of aliasing, and secondly to check whether the client of the program obeys the specified protocol.

With this stream of work, Bierhoff presented a permission-based modu-566 lar protocol checking approach and a tool Bierhoff and Aldrich (2007) for a 567 Java-like object-oriented language. In this approach, programmers express 568 their design intents, as a valid sequence of events associated with a particular 569 object, and the aliasing information using permission-based typestate con-570 tracts, written in Linear logic-based specifications at the method level. The 571 system supports five kinds of symbolic permissions i.e., unique, immutable, 572 full, share and pure. The Linear logic operators are already explained in Sec-573 tion 3.2, except the additive disjunction  $(\oplus)$  that represents an alternative 574 occurrence of multiple tasks. 575

Listing 1 shows a sample permission-based typestate contract for a readonly iterator. The aim is to avoid the concurrent modification of Collection object when the iterator is in progress.

Listing 1: Permission-based typestate specifications of a read-only Iterator (Bierhoff and Aldrich, 2007).

```
1 interface Iterator<c : Collection, k : Fract> {
580
581
   2 states available. end alive
582
   3 boolean hasNext():
   4 pure(this) \neg (result = true \bigotimes pure(this) in available)
583
                \oplus (result = false \bigotimes pure(this) in end)
584
   5
   6 Object next(): full(this) in available - full(this)
585
   7 // other methods such as finalize() etc. can be written similarly
586
587
   8 interface Collection {
         methods such as add(), size(), remove(), contains() etc. comes here
588
   9 //
   10
      589
<del>3</del>99
   11 }
```

579

According to the usage protocol, an **iterator** can be in one of the two 592 states at any moment i.e. available or end. The state alive is a state 593 inherited from the root Object type (Line 2). In Line 11, an iterator 594 object is created with unique permission in Collection class. Importantly, 595 it can be observed that when an iterator is created, it stores a reference to 596 a collection object being iterated in one of its fields. This reference should 597 be associated with the appropriate permission i.e., immutable to guarantee 598 immutability of the Collection object while iteration is in progress. 599

The pre-condition (pure(this)) in Line 4 specifies that method hasNext() 600 requires **pure** permission on the referenced object as it just tests or reads 601 iterator's state. As method next() can change iterator's state, it needs 602 full (this) (Line 6). The method hasNext() determines whether another 603 object is already present in the Collection object with available state, or if 604 the iteration has reached its end. The post-condition (Line 4 and 5) specifies 605 if the result is true. It is legal to call the **next()** method on the same object 606 in the available state, Otherwise, it is illegal. The post-conditions of both 607 methods further show that they return the consumed permission back on 608 the referenced object when they exit. The system leverages this information 609 through method implementations to track state transition in the presence of 610 aliasing, to guarantee the absence of concurrent modifications, and to verify 611 whether program implementation follows the design intents. 612

This approach is realised in Plural (Bierhoff and Aldrich, 2008), a permissionbased automated protocol checking and conformance tool, implemented as a Java Eclipse plugin. In Plural, a program is specified with permission-based

pre and post-conditions at method (parameter) level using JSR-175<sup>1</sup> anno-616 tations to check whether the client of the APIs follows the specified protocol. 617 Plural performs intra-procedural static analysis, called DFA (Diagram 618 Flow Analysis) of the annotated code to identify and track specified pre-619 and post-permissions across method calls for every program variable (pa-620 rameter, receiver object, and local variable) and issues warning for protocol 621 violations in the program. As part of the analysis, it implements a Crystal 622 analyzer that performs dynamic state tests (branch-sensitive flow analysis) to 623 track and report exceptions to the underlying objects. It further checks the 624 structure of the provided specifications by implementing an Annotation ana-625 lyzer. The Effect checker in Plural identifies whether a method is immutable 626 or whether it produces side effects. The Fraction analyser tracks the flow of 627 permissions through the system to split and join permissions associated with 628 a referenced object. 620

Later, Bierhoff et al. (2009b) extended the modular protocol checking 630 approach to check the soundness and effectiveness of Plural in specifying large 631 case studies for real APIs and large third party open-source code bases, For 632 example, Database Connectivity (JDBC) API in Apache Beehive project<sup>2</sup> 633 and PMD, a static code analyzer from DaCapo benchmark<sup>3</sup> that implements 634 Java iterator API. The objective is to measure the precision in terms of false 635 positives, the computational cost and the annotation overhead associated in 636 manually specifying and verifying these APIs with Plural annotations. 637

The approach follows the Design by Contract Principle to explicitly specify state invariants at the method level. State invariants are permission-based typestate assertions with a valid typestate that should hold when an object is in a specific state. The approach uses the concepts of 'capture' and 'release' permissions to avoid inter-object dependencies at the method level. Listing 2 shows a sample code snippet in Plural for Connection class in Java JDBC API.

Listing 2: A Java JDBC connection interface (fragment) with Plural specifications (Bierhoff et al., 2009b).

```
645
646 1 @States({"open", "closed"})
647 2 public interface Connection {
648 3 @Capture(param = "conn")
649 4 @Perm(requires = "share(this, open)", ensures = "unique(result) in open")
```

<sup>&</sup>lt;sup>1</sup>https://jcp.org/en/jsr/detail?id=175
<sup>2</sup>http://beehive.apache.org/
<sup>3</sup>http://dacapobench.org/

```
650 | 5 Statement createStatement() throws SQLException;
651 | 6 @Full(ensures = "closed")
652 | 7 void close() throws SQLException;
853 | 8 }
```

In Listing 2, **@States** annotation (Line 1) specifies concurrent typestates 655 for a connection object. The method createStatement() (Line 5) creates 656 statements, in the Connection interface, with unique permission in 'open' 657 state. When the connection object is closed it invalidates all the statements 658 created with it, leading to runtime errors if a programmer uses invalidated 659 statements. To avoid this error, the approach captures the dependent conn 660 object, using **@Capture** annotation in Line 3. The annotation @Perm in Line 661 4 with requires, clause specifies share permission on the captured object as 662 pre-permission with open state guaranteed, and the ensures clause specifies 663 that the method returns a new statement object in open state with unique 664 permission on it. 665

The permissions on the conn object are explicitly released (using **@Release**) 666 when the statement object is no longer in use or when the connection object 667 is in closed state. The method close() closes the conn object by calling 668 method isClosed(), (not included in the sample program due to brevity), 660 that ensures the current state of the connection object to be closed (using 670 annotation @TrueIndicates ("closed")). The ensures clause in Line 6 speci-671 fies that method returns Full permission back to the connection object in the 672 closed state. The analysis tracks these specifications in the system to ver-673 ify protocol compliance with program implementation and to verify correct 674 usage of the protocol by the client program. 675

Table 4: Annotation overhead for sample APIs verified in (Bierhoff and Kevin, 2009; Bierhoff et al., 2009b).

	ad		
Program	#Annot.		
JDBC	9,866	440	838
Beehive	2,158	65	66
PMD	39,400	-	617

#### 676 Typestate Verification using Plaid

As discussed previously in Section 2, Plaid (Aldrich et al., 2011, 2012) is a new typestate-oriented programming language that verifies the correctness of programs based on access permissions.

Every type in Plaid is represented as tuples having a type structure and 680 associated permissions that express the aliasing and the mutability of the 681 corresponding object's typestate. Plaid borrows its grammar and lexical 682 structure from the Java Specification Language (JSL) and provides inter-683 operability with Java programs. Classes in Plaid are represented using the 684 keyword state and transitions between states are represented by the state 685 transition symbol '>' that distinguishes pre-state from post-state. Plaid sup-686 ports three types of symbolic access permissions i.e., unique, immutable and 687 share in the specifications. The keyword 'none' is used when no permissions 688 are required to access an object. 689

Listing 3 shows an annotated code fragment for a buffer implementa-690 tion in Plaid. A buffer can be in one of the two states i.e. EmptyBuffer 691 and FullBuffer (Line 1). The method put() is associated with an empty 692 buffer. The signature of the method put() (Line 4) specifies that the state 693 of the receiver object should change from EmptyBuffer to FullBuffer when 694 the buffer receives some element. The state FullBuffer requires a field el-695 ement elem that is passed as method parameter e. The permission-based 696 contract (Line 4) specifies that the passed element has unique permission in 697 the Element state and the method does not return any permission (none) 698 to the caller of the method. This is because a field reference with exclusive 699 rights (unique) has been created for that element in FullBuffer state. Oth-700 erwise returning permission back to a caller would cause a violation of the 701 uniqueness property. Likewise, the FullBuffer state has a single operation 702 get() (not included here due to brevity), that returns the current state of 703 the object in reference **elem** and ensures that the receiver object will go back 704 to an EmptyBuffer state. 705

Listing 3: A buffer implementation (fragment) in Plaid (Aldrich et al., 2011).

```
706
    1 state Buffer comprises EmptyBuffer, FullBuffer {}
707
708
    2
        state EmptyBuffer caseof Buffer {
709
    3
       method void put(unique Element \gg none e) [EmptyBuffer \gg FullBuffer] {
710
    4
711
    5
          this \leftarrow FullBuffer {elem = e};}
    6 }
<del>713</del>
```

<sup>714</sup> Plaid runtime leverages permission flow through the system along with asso-

ciated typestate information to ensure protocol compliance at runtime. The
type system allows permission splitting, joining and type casting automatically (when and where required) using the permission splitting and joining
rules given in Table 2.

Table 5: Summary of annotation overhead for sample programs in Plaid (Stork et al.,2014).

Program Statistics and Annotation Overhead						
Program	SLOC	#AnnoLOC	#Annot.			
webserver	227	47~(20.7%)	59			
dic/global	169	41~(24.2%)	65			
dic/fine	251	71 (28.3%)	109			

Naden et al. (2012) presented a type system and a flexible permission borrowing mechanism for unique, shared, and immutable permissions without using explicit fractions of permissions. Permission borrowing is an extraction of permission from a source field, temporarily using the borrowed permission, and returning part or all of it to the source field. The aim was to prevent the concurrent modifications of the shared objects based on symbolic permissions.

The type system is based on the Plaid language. Like Plaid, it supports three types of symbolic permissions i.e unique, immutable and share but unlike Plaid, where a field is reassigned with a new value to recover permission on the reference, in the Naden's type system the caller function itself returns the original permission consumed on the reference.

Unlike other techniques (Boyland, 2003; Bierhoff and Aldrich, 2007; Jacobs et al., 2011; Heule et al., 2011) that support permission borrowing, this approach provides a more intuitive and natural abstraction to model and to reason about the permission flow through the system, making permission tracking flexible and much easier for programmers. However, like Plaid, it wants programmers to explicitly specify permissions-based state information as a part of method specifications.

738 Typestate Verification using JML

<sup>739</sup> Bierhoff (2011) combines symbolic permissions (Bierhoff and Aldrich,

2007), with JML contracts (Leavens and Cheon, 2006) to reason about aliasing and to detect the absence of Concurrent Modification Exceptions (CMEs)
and other recurrent programming errors, such as IndexOutofBoundsExceptions
exceptions in realistic data structures such as Java ArrayList.

Although JML specifications have been used, in may formal approaches 744 (Rodríguez et al., 2005; Araujo et al., 2008; Kim et al., 2009; Cok, 2011), 745 to verify functional correctness and domain specific properties of sequential 746 and concurrent programs, the support for concurrency and aliasing is rather 747 limited. On the contrary, access permissions provide flexible aliasing control 748 mechanism to track all the references of a particular object and update state 749 changes to all such references. The presented approach defines permission-750 based class invariants as JML contracts. 751

The technique implements a permission tracking algorithm as a prototype 752 tool called JavaSyp<sup>4</sup>(Symbolic Permissions for efficient static program veri-753 fication). In JavaSvp. permission-based invariants are specified using Java 754 annotations and tracked as part of the type checking procedure, to ensure 755 that the specified invariants hold as long as a the client has permission to the 756 referenced object and to control aliasing. In this approach, permission track-757 ing is straightforward as tracking symbolic values is much easier than tracking 758 fractional permissions. However, it only supports two kinds of permissions, 759 i.e., unique and immutable using annotations @Excl and @Imm respectively 760 with the referenced object. 761

Listing 4 shows an annotated version of conventional Java ArrayList ob-762 ject a declared with unique permission in Line 2. The list object maintains 763 a list of elements in the order placed originally. The method getElem() 764 method returns the element on the given location (index) with immutable 765 permission on it (Line 6). The invariants (Line 4) for method getElem() 766 specifies that object **a** should be a non-null reference having at least one 767 element in it and the total number of elements in a should not exceed the 768 declared size. The annotation **@requires** in Line 5 specifies a pre-condition 769 for method getElem() that, before calling this method, the index parame-770 ter must be between 0 and size-1. JavaSyp performs static analysis of the 771 annotated code to generate verification conditions (VCs) based on the spec-772 ifications. The program is then verified against inferred conditions using the 773 SMT solver (C. Barrett A. Stump and Tinelli, 2010). 774

<sup>&</sup>lt;sup>4</sup>http://code.google.com/p/syper.

Listing 4: A Java array list (fragment) with permission-based JML contract in JavaSyp (Bierhoff, 2011).

```
775
    1 public class ArrayList <T> {
776
       @Excl private T[] a;
777
    2
      private int size ;
778
       //@invariant 0 <= size & a != null & size <= a.length;</pre>
779
    4
       //@requires 0 <= index & index < size ;</pre>
780
    5
781
    6
       @Imm public T getElem(int index) {
    7
         imm: return a[index];}
782
    8 }
783
```

#### 785 4.2. Verification of API Protocols in Multi-threaded Programs

Beckman (2010) extended the Bierhoff's modular automatic protocol checking approach to verify "if the object protocols work correctly even in the presence of concurrent modifications by multiple threads". In this stream of work, this section discusses the use of permission-based specifications to verify usage protocols in concurrent programs (Beckman et al., 2008; Beckman, 2009; Militao et al., 2010) using different mechanisms.

# 792 Typestate Verification with Atomic blocks

Beckman et al. (2008) extended the permission-based modular protocol 793 checking approach (Bierhoff and Aldrich, 2007) to verify the correctness of 794 usage protocols for a set of concurrent programs such as  $JChannel^5$  and 795 Reservation Manager that use atomic blocks as synchronization primitives. 796 The objective was to enforce the correct use of typestates at runtime and to 797 verify API protocol compliance with its specifications. The approach uses 798 five kinds of symbolic permissions to identify aliasing and to approximate 790 whether a referenced object can be thread-shared or not. 800

Listing 5 shows a sample method isConnected() in a Connection class with permission-based typestate specifications.

Listing 5: Permission-based typestate specifications for method isConnected() in a Connection class (Beckman et al., 2008).

```
1 class Connection {
804
    805
      (result == true \bigotimes share(this, CONNECTED)) \oplus
806
    3
       (result == false & share(this, IDLE)) {
807
    4
808
    5
      atomic:{
        return (this.socket != null);}
809
    6
    \overline{7}
      7
810
    8 }
8<del>1</del>2
```

803

<sup>&</sup>lt;sup>5</sup>http://www.cs.cmu.edu/~nbeckman/research/atomicver/

The pre-condition "share(this, ?)" in Line 2 asserts that the method needs share permission on the receiver object, which needs to be in the unknown (?) state. Likewise, the post-condition in Line 3 specifies that if the method returns true, the receiver object should get the original permission back while in the CONNECTED state Otherwise, it would get the same permission back but in the IDLE state in Line 4. Exclusive access to full, share and pure references are maintained using atomic blocks in Line 5.

The approach is realised as part of the Plural tool. The analysis identifies the abstract state of a referenced object before calling a method, and discovers the way it would be shared with other objects. If the permission on a particular reference (thread) indicates that the referenced object can be accessed simultaneously by other references (threads), as is the case with full, share and pure permissions, it assumes that the object is thread-shared.

The approach discards state information of the local variables having pure 826 and share permissions as the objects with these permissions can be modified 827 by other threads and it is difficult to track them statically through atomic 828 blocks. The limitation of this work is the use of atomic blocks as mutual ex-829 clusion primitives. Atomic blocks are generally associated with transactional 830 memory systems and have limited usage in today's applications. The cur-831 rent analysis generates false positives for programs having synchronization 832 primitives other than atomic blocks. 833

#### <sup>834</sup> Typestate Verification with Synchronized blocks

Beckman (2009) extended the previous type system to perform typestate verification of concurrent programs having synchronized blocks as mutual exclusion primitives.

In this approach, every program reference is associated with a permission 838 kind (having a permission type and an abstract state that is part of the ref-839 erence type). Like Beckman et al. (2008), the system distinguishes between 840 thread-local and thread-shared objects based on permission contracts. The 841 approach is implemented in a tool called Sync-or-Swim for Java that performs 842 static analysis of the program (within a method). It identifies the references 843 on which the current thread is known to have synchronized and tracks the 844 permissions associated with references as they flow with method's pre- and 845 post-conditions, to ensure that an object can be modified concurrently. The 846 analysis discards state information for thread-shared (modified by other ref-847 erences) objects unless it is statically known that the same references have 848 previously been synchronized. 849

Unlike other behavioral checkers for concurrent programs (Jacobs et al., 2005) that require lock-based specifications to identify the part of heap to be protected, the proposed approach can verify program behavior without requiring lock-based specifications. Like other single-threaded typestate verification approaches, it only requires aliasing (permission) and typestate information to verify the correctness of concurrent programs.

Program Statistics and Annotation Overhead									
Program         Classes         APIs         Methods         SLOC         #A									
JabRef	813	7	4,072	74,217	268				
JSpider	187	1	951	8,955	30				

Table 6: Annotation overhead for sample programs verified in (Beckman, 2010).

Table 6.2 summarizes the number and type of specications that were 856 required to be written in JabRef in order to check each API. These anno-857 tations are broken down into three categories. Invariant annotations are 858 used to specify state invariants. Method annotations are used to specify 859 pre- and post-conditions, and Poly. annotations are used by our polymorphic 860 extension. They are used to instantiate a polymorphic API with a specie 861 permission. Table 6.2 also summarizes the amount of time taken to specify 862 and verify the code under the column heading, Spec. Time. In this column, 863 h signies hours and m, minutes. But verifying these calls required 243 anno-864 tations, one quarter of which were state invariant annotations. In our case 865 much of the frustration could have been alleviated with good default permis-866 sions, a topic we will discuss further in Section 6.4. All in all, 357 annotations 867 were required to verify seven APIs used in a program of 74,217 lines, giving 868 a specication density of 1 annotation per 207 lines of code. 869

# 870 Typestate Verification with Views

Militao et al. (2010) expands on Bierhoff's permission type system for Plural (Bierhoff and Aldrich, 2007), by going beyond the five traditional types of permissions. The approach introduces a new abstraction called View that is a projection of an object with a small set of access permissions associated with individual components (fields and/or methods) of an object.

The type system combines view-based controlled aliasing with typestates and Boyland's fractional permissions to manage safe initialization of differ-

ent sections of an object reference, to track state information and to ensure 878 safe access of the referenced objects in a unique-writer and multiple-readers 879 scenario. An immutable view can be shared with an inbound number of 880 copies and a write request merges all the readers back to a single writer 881 using fractional permissions. However, it does not support aliasing of the 882 form where an object can be shared between multiple writers with a state 883 guarantee. In this approach, a view behaves as a state except that it can be 884 split, merged (recombined) using fractional permissions. Therefore, it resem-885 bles the permission accounting model (Bornat et al., 2005) where views are 886 treated as accountable parts of a typestate thereby, allowing local reasoning 887 of the shared-memory programs. 888

889

The analysis of all the Plural-based verification approaches in Section 4.1 890 and 4.2 shows that Plural can identify common challenges for specifying and 891 implementing usage protocols in real-world case studies. It helps program-892 mers to statically follow usage protocols without actually executing the pro-893 gram. However, Plural analysis is limited as it cannot identify errors in the 894 specifications and might use non-consistent specifications. Moreover, there 895 is no reachability analysis support in Plural, which means that a programmer 896 may write inconsistent specifications at the method level and consequently, 897 methods with these specifications will never be called by any client code, 898 resulting in unreachable code. 899

Complementing the Plural tool, research (Siminiceanu et al., 2012), has 900 been done to verify the correctness of manually added Plural specifications 901 as well as verifying the program behavior. Further, it provides limited sup-902 port to the verification of typestate invariants. It can only check invariants 903 on boolean properties, e.g. checking for non-nullness, However, it cannot 904 verify invariants that involve arithmetic predicates, e.g. x > 0 (where x is an 905 integer field). Moreover, it requires programmers to explicitly specify the de-906 sign intents of the API protocols as permission-based typestate specifications 907 in the program which results in annotation overhead for the programmers. 908

In the same fashion, Plaid and related approaches forces a client program to follow the desired sequence of method calls to verify the correctness of typestate-based programs, but at the cost of adding permission-based annotations as a part of type declarations at the code level. Moreover, Plaid does not support full and pure permission as for its analysis. In Plaid, permission-based typestate specifications are added as a part of the program to verify proper state transition between multiple objects during execution 916 of a method.

# <sup>917</sup> 5. Verification of Race conditions and Deadlocks in Multi-threaded <sup>918</sup> Programs

In imperative and object-oriented programming languages, the biggest 919 challenge has been the correctness of concurrent multi-threaded programs in 920 the presence of aliasing and to avoid domain specific problems such as dead-921 locks and race conditions. Access permissions have been used to characterize 922 the way a shared resource can be accessed by multiple threads and to han-923 dle aliasing in many verification approaches. The general idea is to assign 924 permission to program references to access memory locations and track the 925 permission flow through the system to enforce mutual exclusion mechanisms 926 in shared-memory concurrent programs. 927

This section discusses the second dimension of the survey i.e., the use of permission-based specifications for verification of domain specific problems in concurrent programs.

### 931 5.1. Permission Sharing and Accounting Models

Permission sharing and accounting models (Boyland, 2003; Bornat et al.,
2005; Parkinson and Bierman, 2005; Appel and Blazy, 2007; Hobor et al.,
2008; Dockins et al., 2009) facilitate thread-local reasoning for shared-memory
concurrent programs and to ensure race-free sharing of heap locations.

As discussed previously in Section 3, the Boyland's permission sharing 936 model (Boyland, 2003) also called fractional permission, defines shared own-937 ership of resources in a concurrent environment. The sharing policy maps a 938 permission fraction (share) as a rational number  $\mathbb{R}$ , in the range (0, 1], to al-930 low read or write operation on a particular memory location. The fractional 940 model has been used to handle problems that follow concurrent divide-and-941 conquer algorithms where a shared (read) permission can be divided into 942 multiple shared permissions, to an unbounded depth, for any possible pat-943 tern of divide-and-conquer. 944

Although fractional permissions are infinitely splittable, this permission sharing model does not satisfy the disjointness property because rational numbers are not ideal for resource sharing, as shown by Parkinson and Bierman (2005), who proposed a permission sharing model, that allows both read sharing and disjointness of resources, to formalize and verify a subset of single-threaded Java programs with Separation Logic. In this model, resource <sup>951</sup> invariants are defined using permission-based abstract predicates defined at
<sup>952</sup> the Object class level with an empty footprint, (permissions associated with
<sup>953</sup> a memory location), that each subclass extends to hold additional fields.

Bornat et al. (2005) proposed a permission accounting model and a light-954 weight verification approach to handle the accounting problem associated 955 with reader-writer locks in concurrent programs. The approach extends sep-956 aration relation  $\mapsto$  in classic Separation Logic (Reynolds, 2002) and asso-957 ciates fractional permissions with each heap location to allow read sharing 958 of heap locations. In this approach, each heap location x is treated as a map 950 having addresses E with a permission value z, where z represents the level 960 of permission carried by a heap location, as shown in Formula 1). 961

$$x \underset{z}{\mapsto} E \Rightarrow 0 \le z \le 1 \tag{1}$$

The idea is to count the number of shared tokens using an integer counter say s. It is incremented or decremented when a reader locks (receives a read token) or unlocks (returns the share token) respectively, s > 0 means there are outstanding read tokens but s = 0 implies the absence of outstanding readers, which means that a writer may acquire, hence ensuring a race-free sharing of heap locations.

Appel and Blazy (2007) presented the operational semantics and de-968 veloped a Sequential Separation Logic to extend the C Minor language, 969 a mid-level imperative programming language as a machine independent-970 intermediate language. The approach evaluates expressions as functions in 971  $\operatorname{Coq}^6$ , a formal language that combines mathematical functions, axioms and 972 theorems together with a semi-interactive environment to develop machine-973 checked proof assistants. The approach provides an end-to-end machine-974 checked correctness proof of the proposed logic in Coq. Unlike the classical 975 Separation Logic (O'Hearn et al., 2001) where expressions are evaluated in-976 dependent of heaps, the approach associates each expression evaluation with 977 a footprint. "A footprint is a mapping from memory addresses ( $\nu$ ) to per-978 missions" (Appel and Blazy, 2007). 970

In the proposed semantics, footprint  $(\phi)$  is considered as a set of fractional permissions (Bornat et al., 2005) to verify non-interference of load and store operations in memory. A memory store yields result only if reading

<sup>&</sup>lt;sup>6</sup>http://coq.inria.fr

or writing a chunk of memory type, say ch at location  $\nu$  is legal according to its footprint. For example, the semantics  $\phi \vdash load_{ch}\nu$  (or  $\phi \vdash store_{ch}\nu$ ) depicts that all the addresses from location  $\nu$  to  $\nu + |ch| - 1$  can be read (or write). Loading memory outside the footprint yields exceptions and causes expression evaluation to stop. The disjoint sum of two footprints  $\phi_0 \oplus \phi_1 = \phi$ ensures the exclusive read/write or read-only ownership of the underlying memory.

Hobor et al. (2008) extended the Appel and Blazy's machine-checked soundness proof and Leroy's compiler-correctness proof in a concurrent setting, and developed a concurrent C Minor language having shared-memory and first-class locks and thread. He proposed a modular concurrent operation semantics as a generalization of Concurrent Separation Logic (CSL) (OHearn, 2007) but it goes beyond CSL as it allows dynamic lock and thread creation.

In the semantics, a world w corresponds to a footprint  $\phi$  as in the work of Appel and Blazy (Appel and Blazy, 2007). A world specifies permissions for the current thread but this semantic deals with load (store) operations for multiple threads. The need was to evaluate an expression with a guarantee that footprints ( $\phi$ ) of different threads are disjoint. For this purpose, the approach defines permission-based lock invariants to grant or restrict ownership for the accessed memory by extending the classic separation relation  $\mapsto$ as follows:

$$e \underset{\pi}{\mapsto} R$$
 (2)

The relation shows an expression e maps to a memory address with resource 997 invariant R. Every expression acquires a lock before its evaluation. Each lock 998 is associated with a resource invariant R, where each invariant is supported 999 by a unique set of memory addresses and worlds that inform the lock owner-1000 ship  $\pi$  acquired or lost by each thread. The approach implements Parkinson's 1001 (Parkinson and Bierman, 2005) permission sharing model to define owner-1002 ship. A 100% share represents full ownership and a non-empty ownership 1003  $(0 < \pi < 100\%)$  represents read-only access. Any access without ownership 1004 means the program has no semantics and the evaluation stops. 1005

Dockins et al. (2009) proposed a tree share permission accounting model that is more powerful than Bornat's token accounting model. It rectifies the problems with Parkinson's permission model. Tree share is a booleanlabelled binary tree that supports both splitting and token accounting for the shared reading of resources in concurrent settings. Although Boyland's fractional share model is infinitely splittable, it does not satisfy the disjointness property and may pose read/write and write/write race conditions. Similarly, in Bornat's token accounting model, a central authority lends out the total permission into shares in the form of permission tokens when and where required. It counts the outstanding tokens to verify permission accounting. These models satisfy positivity of resources but not the disjointedness.

Unlike previous share accounting models, Dockins et al. (2009) defines heaps as partial functions from memory locations (L) to values (V) with pairs of non-unit shares (S). The token factories are represented using nonnegative integers and the tokens themselves are represented using negative integers. When a token is pushed back into the factory, the integers are added. The token factory's share becomes zero when it gets all its tokens back. The extended points-to operator relation is given below:

$$l \underset{s, n}{\mapsto} v \tag{3}$$

The relation specifies that memory location l contains a value v with a non-unit share s that is indexed by an integer n. If n is zero, the share sis full. If n is positive, it means that n tokens are missing over s in the token factory, and the negative value for n depicts that token factory has size - n shares. The extended model supports and satisfies all the required properties such as disjointedness, cross, and infinite splitting in permission sharing models.

#### 1024 5.2. Permission-based Verification Techniques and Tools

This section describes and discusses the permission-based verification ap-1025 proaches and tools such as Fluid (Zhao, 2007), Chalice (Leino et al., 2009; 1026 Leino and Müller, 2009), Verifast (Jacobs et al., 2010, 2011), Pulse (Siminiceanu 1027 et al., 2012; Cataño et al., 2014; Ahmed and Cataño, 2018), Heap-Hop (Vil-1028 lard et al., 2010), HIP/SLEEK (Hobor and Gherghina, 2012; Jacobs and 1029 Piessens, 2011), HJp (Westbrook et al., 2012), Vercors (Amighi et al., 2012, 1030 2014; Huisman and Mostowski, 2015; Amighi et al., 2015) and Viper (Juhasz 1031 et al., 2014; Müller et al., 2017) that have been developed to resolve concur-1032 rency problems in concurrent programs. 1033

Table 7 provides a summary of the existing permission-based verifications tools and related approaches to verify common concurrency problems.

Table 7: Permission-based verification of race conditions and deadlocks.									
Reference	Prog	Lang	Tool	Analy	Perm- Kind	Specs	Perm- Infer	Anno	Properties
Boyland (2003)	Con	PSM	-	-	Frac	$[0, 1] \cap \mathbb{R}$	Ν	-	RCs
Bornat et al. $(2005)$	Con	PSAM	-	-	Count	$[0, 1] \cap \mathbb{Z}$	Ν	-	RCs
Parkinson and Bierman (200	)5)Seq	Java-like (NSL)	-	-	Counting	$ ext{total}  ext{read}^*$	Ν	-	RCs, VoADT
Appel and Blazy (2007)	Seq	CMinor (PSAM)	Coq	-	Counting	$[0,1]\cap \mathbb{N}$	Ν	-	RCs, MCCP
Zhao (2007) Zhao et al. (2008)	Con	Java	$\mathrm{Fluid}^\dagger$	$\operatorname{St}$	Frac	read write <sup>*</sup>	Ν	Υ	RCs
Hobor et al. $(2008)$	Con	CMinor (PSAM)	$\operatorname{Coq}$	-	Counting	[0,  100]%	Ν	-	RCs, MCCP.
Dockins et al. (2009)	Con	CMinor	Coq	Frac	-	(s, n), $n \in \mathbb{Z}$	-	Y	MCP
Leino and Müller (2009)	Con	Chalice	Chalice	D	Frac	full some no	Ν	Υ	RCs, DLcks
Leino et al. (2009)	Con	Chalice	Chalice	D	Frac	$\operatorname{acc}(x)$ $\operatorname{rd}(x)^{\ddagger}$	access pure	Υ	RCs, DLcks
Villard et al. (2010)	Con	NSL	Heap- Hop	SFEA	-	$\mathbf{x} \mapsto \mathbf{C}$	Ν	Υ	RCs,DLcks MLeaks,VoUP
Jacobs et al. (2010, 2011)	Seq& Con	C and Java (NSL)	VeriFast	D	Frac Count	read write <sup>*</sup>	Ν	Υ	RCs,VoNullP, AIOBEx
Jacobs and Piessens (2011)	Con	NSL	VeriFast	D	Frac Count	$[0, 1] \cap \mathbb{R}$	Ν	Υ	FGSM, RCs
Hobor and Gherghina (2012)	Con	NSL	HIP/SLE	E <b>K</b>	Frac	$(0, 1] \cap \mathbb{R}$	Ν	Y	RCs
Westbrook et al. (2012)	Con	HJ	HJp	St	Frac	Shared read ( $\omega$ R) private write ( $\omega$ W) $\omega \in \{0, 1, \epsilon\}$	Ν	Y	RCs
Siminiceanu et al. (2012) Cataño et al. (2014)	Seq& Con	Plural	Pulse	St RGA	Sym	(U,I,S,F,P)	Ν	Y	RCs, DLcks VoAPs, VoNullI
Blom et al. (2014)	Con	OpenCL	VerCors	D	Frac	$\operatorname{Perm}(\mathbf{x},\pi)$ $\pi \in \{\mathrm{rd, rw}\}$	Ν	Υ	RCs, VoGPGPU
Amighi et al. (2014)	Con	Java	VerCors	D	Frac	$\operatorname{Perm}(\mathbf{x}, \pi) \\ \pi \in (0, 1] \cap \mathbb{R}$	Ν	Y	RCs, FCoJP
Huisman and Mostowski (2015)	Con	Java (PSM)	KeY and PVS	D	-	readPerm $(x, Perm)$ writePerm $(x, Perm)$	$_{\alpha}$ N	Υ	RCs, AVROfrac
Amighi et al. (2015)	Con	Java	Vercors	St	Frac Count	$(0,1]\cap\mathbb{R}$	Ν	Y	RCs.
Müller et al. (2017)	Seq	Java,Chal OpenCL,	lice Viper Scala	St	Frac	$\operatorname{acc}(x)$ $\operatorname{acc}(x, \operatorname{rd})$	N	Y	RCs
Ahmed and Cataño (2018)	Seq& Con	JML	Pulse	$\operatorname{St}$	Sym	(U,I,S,F,P)	Ν	Υ	VoJMLC, RCs

1036 Keys to the table: Seq = sequential, Con = concurrent, St = static, D = dynamic, Sym = sym-1037 bolic, Fract = fractional, U = unique, I = immutable, S = share, F = full, P = pure,  $\mathbb{Z}$  = set of Integers, 1038  $\mathbb{R}$  set of Real numbers,  $\mathbb{N}$  set of positive Integers, NSL = new specification language, PSM = permis-

sion sharing model, PSAM = permission sharing and accounting model, RCs = race conditions, DLcks 1039 = deadlocks, VoNullP = verification of null pointers, VoUP = verification of usage protcols, VoAPs= 1040 verification of access permission-based specifications, VoJMLC = verification of JML contracts, AROfrac 1041 1042 = Avoiding reasoning overhead associated with fractional permissions, FCoJP = functional correctness of Java programs, VoGPGPU = verification of GPGPU programs, FGSM = fine-grained synchroniza-1043 tion mechanism, AIOBEx, = ArrayIndexOutofBoundsExceptions, MLeaks = memory leaks, MCCP = 1044 machine-checked correctness proofs, VOADT = verification of abstract data types, x heap location, rd1045 read access, C = session type contract,  $\alpha$  a permission slice. \* permissions are read and write accesses, † 1046 implemented in the Fluid project, ‡ shows accessibility predicates. 1047

#### 1048 Fluid

1062

<sup>1049</sup> Zhao (2007) developed a permission-based language and a type system to <sup>1050</sup> enforce fixed locking order mechanism in Java-style multi-threaded programs <sup>1051</sup> having unstructured parallelism and synchronization.

The technique was realised as a prototype tool in the Fluid project 1052 (Greenhouse and Scherlis, 2002; Greenhouse, 2003). In their approach, a 1053 program is explicitly annotated with a method's effects and lock invariants. 1054 The method's effect specifies the read or write operation on the current 1055 object this or any of its field e.g., the annotations reads(this.x) and 1056 writes(this.x) in Listing 6, in Line 2 and 4. The lock invariant specifies 1057 the synchronized access of a referenced object, inside the method body, us-1058 ing the requires (this) clause in Line 7, that a method call for deposit() 1059 should be inside a synchronized block to acquire a lock on the receiver object 1060 this. 1061

Listing 6: Code segments showing read, write and lock usage annotations in Fluid.

```
1 class Account{
1063
        read (this.x)
1064
     2
              getBalance(){ x; }
1065
        void
1066
        writes (this.x)
        void setBalance(int newX) { x = newX; }
1067
1068
        void deposit(int x){
     6
        requires (this) { balance = balance + x; }
     7
1069
     8 }
<del>1</del>079
```

The system then translates high level access annotations into low-level (fractional) permissions to distinguish the read and write effects of a method on the referenced objects. The system assigns unique permission with a referenced object if it is being written in the method body and a value less than 1 is assigned for read operation. The type system uses this information to ensure that a given expression can be executed with assigned permissions <sup>1078</sup> but it does not verify program behavior based on input specifications.

<sup>1079</sup> Further, Zhao et al. (2008) proposed a synchronization policy to avoid the <sup>1080</sup> unnecessary synchronization effects in the previous approach. The system <sup>1081</sup> uses "permission nesting" to interpret the safe and correct usage of lock-<sup>1082</sup> based specifications associated with a field.

# 1083 Chalice

Leino and Müller (2009) presented a permission-based verification method 1084 to prevent problems such as deadlocks and race conditions, that arise due 1085 to dynamic locking orders in multi-threaded programs. The approach en-1086 sures concurrent sharing and un-sharing of objects among multiple threads 1087 based on Boyland's fractional permissions (Boyland, 2003). The system uses 1088 permission percentages (between 0 and 100) instead of permission fractions. 1089 A permission percentages is a fractional permission with a definite size that 1090 splits a field permission among several monitors or threads. A thread can 1091 access a shared object, (heap location), if it has permission to do so. The ap-1092 proach defines three types of permissions based on their percentages: 'Full', 1093 'Some' and 'No'. 1094

The technique was realized in Chalice (Leino et al., 2009), a concurrent 1095 program verifier that supports programs with fork/join parallelism, monitors 1096 invariants and automatically verifies the absence of deadlocks and data races. 1097 In Chalice, programmers annotate programs with permission-based contracts 1098 using access predicates for each heap location. The annotation acc(o.f) rep-1099 resents 'Full' permission (100%) on a field of object o that shows that a thread 1100 has exclusive access on o.f. A fractional permission having n percentage of 1101 the actual permission is represented as acc(o.f, n). A non-zero ('Some') 1102 permission depicts read-only access to location o.f, denoted as rd(o.f). 1103

Listing 7 shows a sample method specifications in Chalice. The pre-1104 condition of the method Clone() in Line 4 specifies that the caller of the 1105 method must possess non-zero (read) permission on location this.val be-1106 fore calling this method. Following the Design by Contract Principle, the 1107 post-condition in Line 5 specifies that the callee should generate Full per-1108 mission on result.val field and return the input (read) permission to lo-1109 cation this.val. Otherwise, the system will not be able to recover Full 1110 permission on it and in turn the location would remain immutable forever. 1111

Listing 7: A sample program with accessibility predicates in Chalice (Leino and Müller, 2009).

1112

<sup>1112 1</sup> class Cell {

```
2 int val ;
1114
1115
     3 Cell Clone()
     4 requires rd(this.val);
1116
       ensures acc(result.val) ^ rd(this.val);{
1117
1118
        Cell tmp := new Cell ;
        tmp.val := this.val ;
1119
     7
     8
        return tmp ;}
1120
     9 }
\frac{1121}{122}
```

The annotated program is analyzed to verify whether the code respects the permission contract for every thread schedule, as permissions flow between threads and monitors or between multiple threads. The analysis verifies that the sum of permissions for all threads remains less than or equal to 100% to ensure thread non-interference.

1128 VeriFast

Jacobs et al. (2010, 2011) developed a sound, modular automatic pro-1129 gram verification tool VeriFast to verify single- and multi-threaded programs 1130 written in C and Java. To enable verification, the programmer defines lemma 1131 functions in the program. Lemma functions are like ordinary C functions, 1132 except that lemma functions and calls of lemma functions are written within 1133 annotations. In VeriFast, lemma functions are interactively specified in the 1134 program following the Separation Logic style of specifications. They serve 1135 as proofs to ensure that a method terminates without producing any side 1136 effects in the system. 1137

The approach simulates shared variables as heap locations and associates 1138 a permission coefficient using Boyland's fractional permission to each heap 1139 location to represent its access rights. The coefficient lies within (0, 1] where 1140 1 represents exclusive rights to manipulate a particular heap location and any 1141 value smaller than 1 represents a shared (read) access by multiple threads. 1142 The analysis works in a way that each method is symbolically executed based 1143 on other methods' contracts to verify its calls. The logic-based specifications 1144 are tracked through the system to detect exceptions such as NullPointer 1145 and ArrayIndexOutOfBoundsExceptions exceptions in the program, and to 1146 verify the domain specific problems in a program such as race conditions. 1147 The system does not support permissions for the program's local variables. 1148

1149 Heap-Hop

1150

Villard et al. (2010) developed a program prover Heap-Hop<sup>1</sup> based on

<sup>&</sup>lt;sup>1</sup>http://www.lsv.fr/Software/heap-hop/
Program Stati	stics and Annotation O	verhead
Program	SLOC	#AnnoLOC
chat server	242	114
linked list and iterator	332	194
composite	345	263
JavaCard applet	340	95
GameServer	383	148

Table 8: Summary of annotation overhead for sample programs in VeriFast (Jacobs et al.,2010).

Hoare's monitors and copyless message-passing mechanism (Villard et al., 2009), an alternative to lock-based parallelism where only pointers to a mes-sage content in memory are transferred. The objective was to verify dead-locks, data races and to ensure the absence of memory leaks in heap manip-ulating concurrent programs, particularly those that involve communication protocols with list and tree structures. The approach uses channels as syn-chronization mechanisms where each channel consists of two endpoints, say e and f, dynamically allocated on the heap. 

Heap-Hop requires programmers to specify pre- and post-conditions, as ownership information for the heap locations, and loop invariants defined as Separation Logic (Reynolds, 2002) formulae. The communications between endpoints are governed using a contract C, a form of session types (Takeuchi et al., 1994), which specifies a valid sequence of message m passing on a channel. In Heap-Hop, ownership of cell to a heap location is represented using the notation  $x \mapsto$  and the point-to relation  $e \mapsto C\{a\}$  specifies a contract C in the state a with respect to a particular endpoint e. 

Listing 8: A sample method with permission-based contracts in Heap-Hop

```
1167

1168

11 contract C {//session type contract

1169

1170

1171

1171

1171

1171

1172

1172

1172

1173

1174

1174

1175

1175

1175

1176

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1177

1
```

The approach generates verification conditions based on the input specifications. It performs symbolic (forward) execution analysis of the generated conditions to determine what input (conditions) will cause each part of a program to execute and verifies the intended behavior of a program. The approach ensures that message sending never fails, and message reception should be blocked until the right message is received.

## 1181 HIP/SLEEK

Hobor and Gherghina (2012) developed a Hoare-style concurrent Sepa-1182 ration Logic that verifies Pthreads-style synchronization mechanism called 1183 barriers. Pthreads (POSIX Threads) is an API for threaded programming 1184 that manages various procedure calls for thread creation, destruction and 1185 synchronization (Butenhof, 1997). A common use of barrier calls is to man-1186 age a pool of threads in a pipeline. In Pthreads, barriers are used to re-1187 distribute ownership (as read and write access) of resources (memory cells) 1188 simultaneously between multiple threads. At barrier calls, every thread gives 1189 up its write access to the portion of memory allocated to it, and gets back 1190 the read-only access to the entire memory. 1191

The approach extends the concept of permission shares in DSA (Dockins at al., 2009) and assigns positive share to each thread to access a particular location. A full share is required to modify a particular location. A full share can be split into multiple partial shares that are merged back to get back the full share. The idea is to ensure that if a thread has a partial share for a particular location, no other thread has full share (permission) for that location.

Unlike previous Concurrent Separation Logics (OHearn, 2007; Hobor et al., 2008) that focuses on programs with critical sections, locks, and channels respectively, the approach uses barriers to model resource redistribution, and verifies if barriers are accessed safely in a concurrent environment. The idea is to associate some positive (fractional) share of the barrier itself as a precondition and to ensure that the sum of all preconditions entails full share of the barrier.

For example, the assertion  $\operatorname{barrier}(bn, \pi, cs)$  defines a pre-condition that specifies a barrier bn with a positive share  $\pi$  having a state cs that holds before entering a barrier. The state of barrier changes as threads are released from the barrier, and the next stage will follow based on the post-condition barrier( $bn, \pi, ns$ ), when the state transitions to a new state ns. A full permission ensures that no thread has a 'stale' view of the barrier state to ensure thread non-interference. The approach extends HIP/SLEEK tool set (Gherghina et al., 2011; Nguyen and Chin, 2008) to verify concurrent programs with barrier calls. SLEEK is based on Separation Logic and HIP applies Hoare's rules to program verification.

Later, Jacobs and Piessens (2011) in his work on fine-grained concur-1216 rency verified some of the program examples using the VeriFast tool. The 1217 sample programs were taken from the HIP/SLEEK project and experiments 1218 The experiwere performed by implementing the barrier calls as locks. 1219 ments revealed that VeriFast poses more annotation overhead compared to 1220 HIP/SLEEK tool. For example, authors reported that in the HIP/SLEEK 1221 project, programmers need to add approximately 30 lines of annotation for 1222 a program with 30 lines of source code where as in VeriFast, more than 600 1223 annotated lines were required, as user input, to verify the same program. 1224

1225 Pulse

Pulse (Siminiceanu et al., 2012) is an automatic formal verification ap-1226 proach and a tool that verifies the correctness of Plural (permission-based 1227 typestate contracts) specification itself, rather than the program implemen-1228 tation and its behavior. The goal was to write semantically correct specifi-1229 cations to verify program behavior based on these specifications. Like Plural. 1230 Pulse follows the Design by Contract Principle to specify permission-based 1231 contracts at method level. It supports five kinds of symbolic permissions: 1232 unique, full, share, pure and immutable in method specifications. 1233

In Pulse, programmers specify design intents as permission-based typestate invariants and lock-based specifications at the code level to avoid deadlocks and data races. State invariants are used to enforce the properties that should hold during program execution and to handle the design level inconsistencies in a program such as Null pointer references.

In an extended work, Cataño et al. (2014) evaluated the efficacy and expressiveness of Plural specifications on a multi-threaded application called Multi-threaded Task Server (MTTS), to evaluate its design and verify its behavior using Pulse.

Listing 9 shows lock-based typestate contracts in MTTS using Plural specifications.

Listing 9: Lock-based typestate contracts using access permissions in MTTS (Cataño et al., 2014).

<sup>1245</sup> 1246 1 @Perm(requires="Full(this) in NotAcq", ensures="Full(this) in Acq")

<sup>1247 2</sup> public abstract void acquire( ) { }

```
1248 | 3 @Perm( requires="Full(this) in Acq", ensures="Full(this) in NotAcq"
1249 | 4 public abstract void release() { }
1259 | 5 }
```

Pulse defines lock-based specifications to ensure mutual exclusion to a 1252 critical section. The annotations 'Acq' and 'NotAcq' are used to represent 1253 the state of a lock i.e., to be acquired or not-acquired respectively. The 1254 locks are acquired using method acquire() in Line 2 and released using 1255 method release() in Line 3. The permission contract in Line 1, dictates 1256 that the acquire() method needs Full permissions, as pre-permission on the 1257 mutex (lock) object while acquiring a lock, and transitions it from 'NotAcq' 1258 into 'Acq' typestate. Similarly, the specification in Line 3 shows that the 1259 release() method, before releasing the lock, needs Full permissions on the 1260 lock object that is in 'Acq' state and transitions it from 'Acq' into 'NotAcq' 1261 typestate. 1262

The code of the critical section is then enclosed between a call to method 1263 acquire() and a call to method release() to ensure mutual exclusion. 1264 The typestate transition in the given specification ensures that non-nested 1265 calls to method acquire() will always happen after a call to release() 1266 method. The permission contract ensures that if a thread has acquired a 1267 lock, it needs to be released before being used by other threads. However, 1268 as discussed previously, Plural does not support the reachability analysis of 1269 input specifications and cannot verify the absence of deadlocks caused by the 1270 input specifications. Pulse avoids deadlocks by using try-catch-finally 1271 statement in the code and enclosing call to method release() in a finally 1272 block to ensure that method release() is always called regardless of the 1273 termination status of the method. 1274

Additionally detects violations of intended semantics using the model 1275 checking power of evmdd-smc (Roux and Siminiceanu, 2010) symbolic model 1276 checker. It helps programmers write semantically correct specifications and 1277 find possible concurrency at the method level, but to exploit the full potential 1278 of the Pulse tool, the programmers need to manually add permission-based 1279 typestate specifications in the source program, resulting in annotation over-1280 head for the programmers. The authors reported that "it took six months 1281 on manually adding the specifications for 49 Java classes with 14,451 lines of 1282 source lines of Java code and 546 annotated lines of permission-based speci-1283 fications in MTTS". Moreover, in Pulse, the use of model checker can create 1284 state-space explosion problems even for a program of average size. 1285

Ahmed and Cataño (2018) proposed an automatic translation technique

that encodes JML-encoded Finite State Machine (FSM) specification of a Java program into Plural specifications (permission-based typestate contracts). The encoded specification was fed into Pulse to find problems such as unreachable states, unreachable methods and sink states (deadlocks) in the input specifications, and to reason about the correctness of the underlying program before it is implemented.

1293 HJp

Westbrook et al. (2012) proposed a permission-based type system that 1294 supports task parallelism, array parallelism, and object isolation. The system 1295 called 'Habanero Java with permissions (HJp)' is an extension of their pre-1296 vious work on the Habanero Java (HJ) language. "HJ itself is a task-parallel 1297 extension of Java language" (Cav et al., 2011). The main idea of HJp is that 1298 each object can be in any of the two permission modes, i.e., shared read 1299 or private read-write, at any moment of time. The shared read model 1300 specifies that any task (thread) can read from the object but none of them is 1301 allowed to write on it and private read-write mode shows that only one 1302 task is permitted to read from or write to the object. The system provides a 1303 practical solution to prevent data races for non-trivial parallel programs im-1304 plementing multiple synchronization primitives, and parallel patters instead 1305 of just one. 1306

The type system extend Boyland's fractional permissions with two new 1307 permission types: aliased write and storable permission. Unlike previ-1308 ous approaches (Bierhoff and Aldrich, 2007) where write (unique) permission 1309 is only supported for non-aliased objects at any one time, the aliased write 1310 permission supports write operations on aliased objects. In the system, mul-1311 tiple threads can write on multiple objects without actually having unique 1312 permissions on them, as long as the permissions are not passed to other 1313 threads. The storable permission provides a new and simple way for ex-1314 pressing transitive permission in complex objects such as linked list. Storable 1315 permission associates permission to the 'whole tree of objects' instead of as-1316 sociating it to a single object. The permissions transitivity between objects 1317 is managed by defining **exclusive fields**, using the exclusive keyword, at 1318 the class level. This feature makes the technique different from existing ap-1319 proaches that require more technical machinery, and sound approximations, 1320 to manage permission transitivity in complex objects. Further, the language 1321 adds keywords such as reading, shared reading, writing and exclusive, to 1322 indicate permissions held by the method arguments on entry to and exit 1323

from the method. Similarly, the keywords acquire and release are used to
acquire and release aliased-write and storable permissions on the referenced objects or fields.

Table 9: Summary of the annotation overhead for the sample programs in HJp (Westbrook et al., 2012).

Program Statistics and Annotation Overhead						
Program	SLOC(Methods)	#LOC-MK	#LOC-SP	# AnnLOC		
NPB.CG	1070(61)	25	7	32		
JGF.Series	225 (15)	6	3	9		
JGF.LUFact	467 (20)	16	11	27		
GF.SOR	175 (12)	6	4	10		
JGF.Moldyn	741 (57)	9	29	38		
JGF.RayTracer	810 (67)	57	22	79		
BOTS.NQueens	95~(3)	3	0	3		
BOTS.Fibonacci	70(3)	0)	0	0		
BOTS.FFT	4480 (46)	33	0	0		
PDFS	537 (26)	10	8	0		
DPJ.BarnesHut	682~(56)	18	10	28		
DPJ.MonteCarlo	2877 (287)	151	22	173		
DPJ.IDEA	228 (18)	9	8	17		
DPJ.CollisonTree	1032~(69)	108	24	132		
DPJ.K-Means	501 (38)	25)	6	31		

# 1327 VerCors

Blom et al. (2014) proposed a simplified version of Kernel Programming Language (Betts et al., 2012) and a permission-based Separation Logic to reason about the correctness of the GPU kernel written in OpenCl<sup>2</sup>. As the GPU kernel extensively uses threads to support parallelism, the objective was to verify the functional correctness of GPU programs and to ensure data race freedom in the underlying architecture.

The work follows an earlier work of Haack and Hurlin (2008) on verifi-1334 cation of the muti-threaded programs in which a thread can only access or 1335 update a particular memory location if it has permission to read or write. The 1336 permission in a program are provided, in the form of barrier specifications, 1337 using the read-write (rw) or read-only (rd) annotations. Multiple threads 1338 with read permissions can access the same location but only one thread can 1339 hold write permission at a time to change its content. The specifications com-1340 bine first-order logic formulas with permissions-based accessibility predicates 1341 and the separating conjunction operator (\*). The same idea was applied to 1342 delegate permissions across work groups and then to distribute permissions 1343 over threads. 1344

The approach was validated using the VerCors project<sup>3</sup> (Amighi et al., 1345 2012). VerCors is a program verifier for concurrent data structures. It sup-1346 ports both OpenCL and Java. The specification language and program logic 1347 in VerCors is based on earlier work on permission-based Concurrent Separa-1348 tion Logic (Haack et al., 2008) that supports Java. It uses Silicon (Juhasz 1349 et al., 2014) as a back-end verification tool which natively supports an ex-1350 pressive permission model. Moreover, VerCors exploits verification power of 1351 Chalice (Leino et al., 2009) and Boogie (Barnett et al., 2006) as program 1352 verifier. It encodes an annotated input program into series of intermedi-1353 ate representations and generates verification problems, understandable by 1354 Chalice or Boogie, to be verified by the SMT solvers based on the input 1355 specifications. 1356

However, the approach creates annotation overhead for programmers to add permission-based specifications in the program, as authors reported themselves "writing annotations can be very tedious. Not only is it necessary to write the contract for every method, it is also necessary to include many hints to the prover inside the code".

<sup>1362</sup> Further, VerCors in (Amighi et al., 2014) extends JML specifications with

<sup>&</sup>lt;sup>2</sup>Khronos OpenCL Working Group, The OpenCL specification, http://www.khronos.org/opencl/

<sup>&</sup>lt;sup>3</sup>https://fmttools.ewi.utwente.nl/redmine//projects/vercors-verifier/ wiki/Puptol/

fractional permissions to reason about the functional correctness of Java pro-1363 grams. The approach supports multiple synchronization primitives in a Java 1364 program. The permission-based contracts and class invariants are defined in 1365 the input program, using conjunction operator \* in the Separation Logic, as 1366 JML comments. Access permissions are specified using propositional formula 1367 of the form Perm(e.f,  $\pi$ ) where  $\pi$  represents fractional permission in the 1368 range (0, 1] assigned to an individual field f of object e. The permissions 1369 are then transferred between threads at synchronization points and analyzed 1370 with the execution of program. Although, VerCors tool can generate many 1371 of the specifications itself, the annotation overhead in specifying permission-1372 based contract in the input program is still a concern in this work, also 1373 evident from the authors perspective "our specification method in principle 1374 is very verbose, specifications at many different levels are required". 1375

Listing 10 shows a sample Java class point in VerCors with permission-1376 based access predicates. In Line 2, a state predicate state(frac p) speci-1377 fies that p permission is required on disjoint locations, this.x and this.y 1378 in memory. These predicates are then used to specify permission contract 1379 for the same locations at the method level. For example, the pre-condition 1380 state(1) of method set() in Line 5 specifies that the method requires full 1381 (write) permission on locations x and y, and the post-condition ensures 1382 state(1) ensures that the method returns the same permission on the cor-1383 responding locations when it exits. The invariant clause, in Line 4 specifies 1384 a functional property that both points should be in the first or third quarter 1385 of its cartesian space. 1386

Listing 10: A Java class Point example in (Amighi et al., 2014).

```
1 public class Point{
1388
    2 //@ resource state(frac p) = Perm(this.x, p) * Perm(this.y, p);
1389
         private int x, y;
1390
    3
            invariant (x \ge 0 \&\& y \ge 0) (x \le 0 \&\& y \le 0);
    4 //@
1391
1392
    5 //0
            requires state(1); ensures state(1);
       public void set(int xv, int yv){ this.x = xv; this.y = yv; }
1393
      //@ given frac p; requires state(p); ensures state(p);
1394
    8 public void plot(){}
1395
    9 //@ given frac p; requires state(p); ensures state(p);
1396
1397
    10 public int getQuarter(){}
    11 }
1388
```

1387

The contracts of methods plot() and getQuarter() in Line 8 and Line 10 respectively, specify that both methods require read permission p on locations x and y, which means that they can be executed simultaneously by multiple threads without the fear of data races. This is because the the pre-conditions <sup>1404</sup> of both are disjoint with respect to memory. This is not true for the method <sup>1405</sup> set() as it requires full permission on the same locations.

Instead of simply defining the amount in fractions of permission trans-1406 ferred, Huisman and Mostowski (2015) extended the previous fractional per-1407 mission model in VerCors by having symbolic expressions which include the 1408 kind of *transfer* applied to permission, and the owner of the transferred per-1409 mission. The approach facilitates high-precision, complex synchronization 1410 scenarios in concurrent data structures, and supports permission tracking 1411 at a high level of abstraction as compared to the previously mentioned ap-1412 proaches such as Veri-Fast (Jacobs et al., 2011) and Chalice (Leino et al., 1413 2009). 1414

In this approach, the program is annotated with symbolic (permission) 1415 expressions using JML annotations in the functional style. The analysis then 1416 tracks the owners using permission expressions and checks their permission 1417 return paths to reason about their behavior. The system identifies permission 1418 owners using object references and manages a list of owners. Whenever 1419 permissions are assigned to some owner, it is being added in the list and 1420 when an owner returns permissions, it is removed from the list. Each owner 1421 is considered as a permission slice. If all slices refer to the same owner, it 1422 means that the owner would have full permission. Otherwise, access is 1423 partial (read). 1424

Listing 11 shows a sample read and write resource locking mechanism in Java in fractional permissions style. The Line 2 specifies that acquiring a lock (c1) transfers full (1) permission for location o.x to the locking thread and a read permission (1/2) for location o.y to access these locations. When the lock is released, in Line 4, the current thread transfers the same permission back to the lock object.

Listing 11: A simple lock and its fractional permission style specifications (Huisman and Mostowski, 2015).

```
1 class Client {
1432
     2 cl.lock(); // produces Perm(o.x, 1) and Perm(o.y, 1/2)
1433
1434
     3 \text{ o.x} = \text{ o.y}; // \text{ write o.x, read o.y}
     4 cl.unlock(); // consumes Perm(o.x, 1) and Perm(o.y, 1/2)
1435
     5...}
1436
     6 class Lock {
1437
     7 //@ requires !locked; ensures locked;
1438
     8 //@ ensures Perm(o.x,1) ** Perm(o.y,1/2);
1439
1440
     0
       void lock():
1441
    10 //@ requires Perm(o.x,1) ** Perm(o.y,12);
1442
    11 //@ requires locked; ensures !locked;
1443
    12
       void unlock();
1444
    13 }
```

1445

1465

Listing 12 shows a code segment of a simple resource locking mechanism. 1446 Before acquiring the lock, permissions for location o.x and o.y are mapped 1447 to list cl (Line 2), as both of them belong to cl. In Line 3, acquiring a 1448 lock assigns full permission on o.x to the locking thread represented as ct 1449 while it gets partial permission (one slice only) for location o.y and the list 1450 becomes [ct, cl] in Line 3. which means that lock still owns the remaining 1451 1 slice on y that can be made available when required. The post-condition 1452 of method lock() specifies how permission to object o.x and o.y changes 1453 when function lock and unlock are called. It shows that the transfer function 1454 transPerm() in Line 8, forces its owner thread ct to give up all the rights 1455 completely on o.x while function transPermSplit(), in Line 9 transfers the 1456 old permission in slices to the specified thread ct. 1457

When function unlock() is called, (Line 5) permissions are returned to lock by replacing current thread ct with cl permission, on o.x and [cl, cl] on o.y that can be merged again into object lock permission [cl]. The method unlock is not specified here due to brevity.

This information is then used to manage permissions at synchronization points while threads are being forked or joined and to reason about their behavior

Listing 12: A simple lock specification in (Huisman and Mostowski, 2015).

```
1 class Client {
1466
     2 // Perm(o.x), Perm(o.y) are [cl]
1467
     3 cl.lock(); // Perm(o.x) becomes [ct], Perm(o.y) becomes [ct, cl]
1468
     4 o.x = o.y; // [ct] \rightarrow write access, [ct, cl] \rightarrow read access
1469
      cl.unlock(); // Perm(o.x) becomes [cl], Perm(o.y) becomes [cl, cl]
1470
1471
         . . }
      class Lock {
1472
1473
     8 //@ ensures Perm(o.x) == transPerm(this, ct, \old(Perm(o.x)));
     9 //@ ensures Perm(o.y) == transPermSplit(this, ct, \old(Perm(o.y)));
1474
    10 void lock(); . . .}
<del>1</del>478
```

The permission theory was formalized in the KeY tool (Beckert et al., 2007), an interactive verifier for Java, that is based on dynamic logic. The system extends KeY tool with permission accounting to verify program properties that are based on purely first-order reasoning. The general program properties that require structural induction proofs are validated using the PVS tool (Owre et al., 1992), because of its automated deduction and theorem proving capability.

Amighi et al. in (Amighi et al., 2015) proposed a variant of OHearn's

<sup>1485</sup> Concurrent Separation Logic (OHearn, 2007) to perform practical reasoning
<sup>1486</sup> of Java-like concurrent programs having main concurrency primitives such
<sup>1487</sup> as dynamic thread creation, thread joining, wait-notify scenarios and lock
<sup>1488</sup> reentrant mechanism.

The system combines Parkinson's share model with Boyland's fractional permissions to support inheritance of resource invariants and class parameters and to avoid data races in realistic applications.

In Parkinson share model, resource invariants are defined using abstract 1492 predicates at class Object level, with an empty footprint (permissions asso-1493 ciated with a memory location) that each subclass extends to hold additional 1494 fields. Like Parkinson's share model, access for a particular heap location is 1495 maintained using a resource's invariant property, where 1 represents full (ex-1496 clusive) permission to a heap location, and a fractional value in the interval 1497 (0, 1) defines the concurrent read access of a particular location. The idea 1498 is that a thread having partial permission is not allowed to write on a heap 1499 location and the total permission to access a heap location cannot exceed 1. 1500

Like the OHearn's approach, when a thread acquires a lock, it gets access 1501 to part of a heap location specified as a resource's invariant property. Upon 1502 unlocking, it transfers access of the same resource back to lock, to re-establish 1503 the resource's invariant property. The permissions are transferred between 1504 threads at the time of thread creation, thread joining and at lock entrances 1505 and reentrance points. However, the verification is performed at the cost 1506 of manually writing specification as a part of input program that creates 1507 annotation overhead for programmers. 1508

#### 1509 Viper

Müller et al. (2017) developed a verification infrastructure called Viper. 1510 It targets a sequential, object-based intermediate language Silver that en-1511 codes a flexible permission model and supports user-defined predicates and 1512 functions. The infrastructure includes two back-end verifiers and four front-1513 end tools for Chalice, Java, Scala, and OpenCL that was developed as a 1514 part of VerCors project (Blom and Huisman, 2014). A Viper program does 1515 not have classes and an object can access every field declared in a program. 1516 Moreover, there is no implicit receiver object for methods and functions. 1517

In a Viper program, a programmer defines accessibility predicates (Parkinson and Bierman, 2005), as permission-based pre- and post-conditions and loop invariants for heap structures to verify its behavior. For example, the predicate acc(e1.f, e2) represents the permission defined for the field f that belongs to a reference e1. The optional expression e2 represents the amount of permission, that is full (write) permission by default unless otherwise specified explicitly. A method can access a particular heap location if the appropriate permissions are held by that location. The permissions are then transferred between method execution and the loop body to verify program behavior based on the input specifications rather than using its implementation.

Listing 13 shows a sample sorted integer list data with access predicates in a Viper program. Line 1 declares an integer list as a data field of sequence data type. The macro sorted(s) in Line 2 sorts input list s in ascending order. The insert method adds a new element elem in the Ref list and returns the index idx where the new element was inserted.

Listing 13: A sorted integer list and its specifications in Viper (Müller et al., 2017).

```
1 field data: Seq[Int]
1535
      define sorted(s) forall i: Int, j: Int :: 0 <= i && i < j && j < s
     2
1536
                                                    ==> s[i] <= s[j]
1537
     3
        method insert(this: Ref, elem: Int) returns (idx: Int)
1538
     4
        requires acc(this.data) && sorted(this.data)
1539
     \mathbf{5}
        ensures acc(this.data) && sorted(this.data)
1540
     6
1541
     7
        ensures 0 <= idx && idx <= old(this.data)</pre>
        ensures this.data == old(this.data)[0..idx] ++
1542
     8
                     Seq(elem) ++ old(this.data)[idx..]{
1543
     9
         idx := 0
1544
    10
1545
    11
         while(idx < this.data && this.data[idx] < elem)</pre>
1546
    12
         invariant acc(this.data, 1/2)
1547
    13
1548
    14
         { idx := idx + 1 }
1549
    15
         . . .
       }
    16
1559
```

1534

The pre-condition of method insert() in Line 5 specifies that the method requires full permissions on the object list this.data and it should be sorted. The post-condition in Line 6 guarantees that when the method exits it returns the sorted list to the caller with the consumed permission. The second post-condition in Line 7 constrains and thus validates the index, while the third post-condition in Line 8 relates the current state of the list with the method's pre-state, using an old expression.

The insert() method iterates over data list to determine where to insert the new element elem in Line 11. The loop invariant (Line 12) specifies that loop body needs a half (read) permission on the list, while the second half permission would be held by method execution to ensure that the loop body does not modify the list. The Viper's front-end tools then encode the annotated program into an intermediate language acceptable by the by the <sup>1565</sup> back-ends tools, to verify its behavior.

## 1566 6. Automatic Inference of Access Permissions

Permission-based access notations have been generated as means for pro-1567 gram verification in many approaches (Bierhoff et al., 2009a; Leino et al., 1568 2009; Ferrara and Müller, 2012; Le et al., 2012; Heule et al., 2011, 2013; Sadiq 1569 et al., 2016; Dohrau et al., 2018). The generated specifications are either in 1570 the form of read/write accesses, fractional or symbolic permissions. The 1571 overall goal of these approaches was to relieve programmers from specifica-1572 tion overhead resulting from manually adding permission-based annotations 1573 in a source program for verification purpose. Table 10 shows a summary 1574 of the work done to infer permission-based specification in sequential and 1575 concurrent programs. 1576

Reference	Prog	Lang Tool	Analy	Perm- Kind	Perm- Specs	Perm- Infer	Anno	Properties
Bierhoff et al. (2009a)	Seq	Plural (NSL) Plural	(St,D)	Sym	(U,I,S,F,P)	Frac	Υ	VoUP
Leino et al. (2009)	Con	NSL Chalice	D	Frac	$\operatorname{acc}(x) \ \operatorname{rd}(x)^{\ddagger}$	access pure	Υ	RCs DLcks
Heule et al. (2013) Heule et al. (2013)	$_{3)}^{(1)}$ Con	- Chalice	$\operatorname{St}$	Frac	acc(x, 1) acc(x, rd)	full read	Υ	RCs
Le et al. (2012)	Con	NSL VPerm	D	Frac	$ \text{@full}[ u] \\ \text{@value}[ u]^{\phi} $	full zero	Υ	RCs
Ferrara and Müller (2012)	Con	Scala Sample	$\operatorname{St}$	-	$\begin{array}{l} \operatorname{acc}(\mathbf{x},\mathbf{p})\\ \mathbf{p}\in(0,1]\cap\mathbb{R}\\ \mathbf{p}\in(0,1]\cap\mathbb{Z},\\ \mathbf{p}\in(0,100]\% \end{array}$	Frac Count Chalice	Y	RCs
Dohrau et al. (2018)	Con	Viper Scala	$\operatorname{St}$	Frac	$[0,1]\cap\mathbb{R}$	read & write	Y	RCs
Sadiq et al. (2016, 2019)	Seq	Java $Sip4J^4$	$\operatorname{St}$	Sym	-	(U,I,S,F,I	P)N	EIC,CRA, INullP

Table 10: Access permission inference for sequential and concurrent programs.

1577 Keys to the table: Seq = sequential, Con = concurrent, St = static, D = dynamic, Sym = symbolic, Fract = 1578 fractional, U = unique, I = immutable, S = share, F = full, P = pure,  $\mathbb{Z}$  = set of Integers,  $\mathbb{R}$  set of Real numbers, NSL 1579 new specification language, RCs = race conditions, EIC = Enabling implicit concurrency, DLcks = deadlocks, INullP = 1580 identifying null pointers, CRA = code reachability analysis, VoUP = verification of usage protocols, x heap location,  $\nu$ 1581 represents a non-heap location. rd for the read access.  $\ddagger$  accessibility predicates.

<sup>4</sup>https://github.com/Sip4J/Sip4J

### 1582 6.1. Inference of Read & Write Accesses

As discussed previously, in Section 5.2, Chalice (Leino et al., 2009) is a verification framework that verifies correctness of multi-threaded programs written in the Chalice language. Chalice uses autoMagic, a command-line option, to infer the read and write accesses for the heap locations specified with accessibility predicates. The inferred notations are in the form of pure and access notations that represent read-only and full permission respectively for the specified heap locations.

Le et al. (2012) proposed a new permission system to avoid data races in multi-threaded applications having fork/join parallelism. The objective was to ensure the absence of data races for program variables that are not actually heap variables but can be accessed by multiple threads.

The scheme infers variable permissions at the method level using **procedure** 1594 specifications. However, in the procedure specification, a programmer ex-1595 plicitly specifies state changes (if any) for the referenced variable accessed by 1596 the current thread, using permission-based state invariants without actual 1597 variable permission. The generated permissions are in the form of notations 1598 such as full or zero where full represents exclusive rights on the refer-1599 enced variable and zero represents the absence of permission. The proposed 1600 technique then tracks permission flow between threads to ensure safe access 1601 to the shared variables. 1602

Listing 14 shows an example procedure specification example in a sample 1603 fork-join program. The method creator() takes two variables x and y as ref-1604 erence parameters. The requires clause in Line 2 specifies that the method 1605 needs full permission on the referenced variables x and y as pre-permissions 1606 when the method is called. The **ensures** clause specifies that the method 1607 should generate the same permissions as post-permissions on the same ref-1608 erenced variables when it exits (Line 3). The state changes are represented 1609 using prime  $\prime$  notation. For example, the specification " $y\prime = y + 2$ " in Line 3 1610 specifies state changes for the referenced variable y that should hold after the 1611 method completes its execution. These specifications are then tracked in the 1612 system to generate actual variable permissions for the referenced variables. 1613

Listing 14: A fork/join program fragment with procedure specifications (Le et al., 2012).

```
1614
     1 int creator(ref int x, ref int y)
1615
     2 requires Ofull [x, y]
1616
     3 ensures @full [y] \land y'
1617
                                    = y + 2 \land res = tid and @full[x] \land x/ = x + 1 \land
            thread = tid;{
1618
     4
        int tid = fork(inc, x, 1);
1619
1620
     \mathbf{5}
        inc(y, 2);
```

1621 | 6 return tid; 1623 | 7 }

The proposed scheme was realized in a concurrent program verifier called Vperm<sup>5</sup> that verifies the correctness of concurrent applications written in C/C++ language. The approach does not handle phased access to a shared variable by multiple threads, in which case a translation algorithm is used to simulate the affected variables as pseudo-heap locations.

# 1629 6.2. Inference of Fractional Permissions

Bierhoff et al. (2009a) proposed a deterministic algorithm to infer permission flow through the program while verifying usage protocols. The objective was to avoid the permission tracking overhead associated with splitting and joining the fractional permission during verification.

The algorithm is implemented in the Plural tool (Bierhoff and Aldrich, 1634 2008) that performs dataflow analysis of the program with in and out per-1635 missions as developer-provided annotations. The system collects linear con-1636 straints over fractional variables by tracking the flow of permissions in the 1637 program. The analysis then ensures the satisfiability of constraints in a mod-1638 ular fashion. The approach supports polymorphism over fractions that not 1630 only facilitates modular reasoning of the program, but also avoids impreci-1640 sion in loops by allowing permission consumption inside loops. Furthermore, 1641 the technique automatically infers loop invariants in a program. 1642

Ferrara and Müller (2012) proposed a permission inference technique to 1643 infer fractional and counting permissions for heap locations in a class-based 1644 language having threads and monitors. The technique performs static anal-1645 ysis of the source program and inference is based on abstract interpretations 1646 (Cousot and Cousot, 1977), a theory for defining and soundly approximating 1647 the semantics of a program. The approach firstly computes symbolic values 1648 (approximations) for each heap location using the pre- and post-conditions 1649 and lock invariants defined at the method and class level respectively. It 1650 then infers constraints over these symbolic values to reflect permission-based 1651 intermediate representation for the heap locations. Finally, it generates spec-1652 ifications in the form of fractional (value between 0 and 1) and counting (value 1653 between 0 and Integer:MAX\_VALUE) permissions for each heap location in 1654 the program. 1655

<sup>&</sup>lt;sup>5</sup>http://loris-7.ddns.comp.nus.edu.sg/project/vperm/.

The symbolic permissions  $(\overline{AV})$  for each heap location are calculated as "the summation of symbolic values  $s_i$  multiplied by integer coefficients  $a_i$  (to represent how many times the permission is consumed or returned) and an integer constant c" (see Formula 4). The integer constant c represents full permission that is inhaled when an object is created.

$$\overline{AV} = \sum_{i} a_i * s_i + c, \text{ where } a_i, c \in \mathbb{R}, s_i \in SV$$
(4)

For example, the expression 1 \* Pre(C,m,c:f) + 1 \* MI(C,c:f) + 01656 represents symbolic permissions computed for each heap location (c:f) in 1657 method m() of class C where Pre(C,m,c:f) represents the symbolic value 1658  $(s_i)$  assigned to location (c:f), as pre-condition before acquiring a lock, and 1659 the notation MI(C;c:f) represents monitor (MI) acquired on location c:f. 1660 Further, the notation Post(C,m,c:f) represents a symbolic value assigned 1661 to a heap location, as post-condition, to get the original permission back on it 1662 when the monitor is released. The technique then infers constraints over these 1663 symbolic values to generate actual permissions as fractional permissions. 1664

The inference technique is implemented in Sample (Static Analyzer of Multiple Programming LanguagEs)<sup>6</sup> that supports programs written in Scala (Odersky et al., 2004). Listing 15 shows the OwickiGries (Owicki and Gries, 1976) program fragment as an input program. In the example, all expressions are self explanatory except the old expression old (c.c1), in Line 4, that allows post-conditions to refer back to the pre-state of a referenced variable and its associated predicates.

Listing 15: The OwickiGries program (fragment) with method level specifications in (Ferrara and Müller, 2012)

```
1 class W1 {
1673
1674
     2 var c : Cell;
1675
     3 method Inc()
        ensures c.c1 = old(c.c1) + 1{
1676
1677
        acquire c;
        c.c1 := c.c1 + 1;
1678
     6
1679
        release c;
     7
1680
     8
       }
     9 }
1681
```

1672

In method Inc(), between acquire c and release c clauses (Line 5 and 7), the current thread is assigned with a symbolic permission 1 \* Pre(W1, Inc,

<sup>&</sup>lt;sup>6</sup>http://www.pm.inf.ethz.ch/research/sample.html

<sup>1685</sup> c:c1) + 1 \* MI(Cell, c:c1) for location c : c1. When method exits (Line 9), is <sup>1686</sup> gets the symbolic permission 1 \* Pre(W1, Inc, c:c1) = 1 \* Post(W1, Inc, c:c1) <sup>1687</sup> as post-condition since the monitor of c is released. Solving constraints over <sup>1688</sup> these symbolic values, the system generates full (1) as fractional permission <sup>1689</sup> for location c:c1 when method completes its execution and control is passed <sup>1690</sup> to Line 8.

The system works very well for Chalice lattice domain. However, the 1691 analysis based on fractional permissions is challenging and sometimes, the 1692 system converges the generated permissions back to zero to explicitly ter-1693 minate analysis. Moreover, it provides limited support to infer permission 1694 contracts for programs having recursive data structures. The rate of infer-1695 ring permission contracts for such programs is at minimum 36% and 68% at 1696 maximum. Moreover, like Chalice, manually annotating code with pre and 1697 post-conditions and monitor invariants creates significant annotation over-1698 head for programmers. 1699

In an extended work of Ferrara and Müller's permission inference, Dohrau 1700 et al. (2018) proposed a static analysis to infer permission-based contracts for 1701 array manipulating concurrent programs. The technique is based on Separa-1702 tion and related logics (Reynolds, 2002; Smans et al., 2009). The idea is to 1703 explicitly associate a separate (fractional) permission for each array element 1704 to specify its accessibility by parts of the program. The value 1 represents 1705 full access while rd, a positive fraction of permission, represents the con-1706 current (read) access to a memory location. The analysis then infers read 1707 and write accesses for the specified memory locations to generate permission 1708 contracts at the method-level and within loop. 1709

For example, the approach associates each array element say  $q_a[q_i]$  with a 1710 fraction of permission using a conditional expression of the form  $q_a = array \wedge$ 1711  $q_i = index ? 1 : 0$  that specifies full permission (1) for element array[index]1712 and no permission for all other elements. The permission required for each 1713 loop iteration is computed using a maximum expression that calculates the 1714 maximum of permission required by each referenced variable changed in a 1715 particular loop iteration. The whole (complete) loop execution depends on 1716 the maximum of all the fractions over all loop iterations. It is used to infer 1717 read and write specifications for all indices of an array, for the whole loop. 1718

However, it is generally acknowledged that tracking concrete fractional values is a cumbersome task for programmers especially when fractions continue to decrease indefinitely for a particular scenario (Heule et al., 2013). Moreover, the use of fractional permissions makes the specifications too lowlevel which can be tedious to add manually and harder to reuse and adaptfor programmers.

## 1725 6.3. Inference of Symbolic Permissions

Heule et al. (2011, 2013) proposed a technique to automatically convert fractional permissions into abstract read/write permissions for sharedmemory concurrent programs. The objective was to specify concurrent constructs such as fork/join threads, locks/monitors with abstract permissions to avoid the complex reasoning overhead associated with fractional shares.

The abstract read permissions allow programmers to reason at a high-1731 level of abstraction than using the fractional values for reading. The objec-1732 tive was to avoid the complex reasoning overhead associated with handling 1733 concrete values in fractional permissions during verification. The proposed 1734 methodology is implemented in Chalice. The system generates two kinds 1735 of permissions i.e., full and read. Like Chalice, it takes a program anno-1736 tated with accessibility predicates such as acc(x.f,1) and acc(x.f, rd)1737 at method level. The value 1 is mapped to represent the full (read and 1738 write) permission and rd represents the shared read permission (a part of 1739 permission that is not full) for the referenced object x.f. Moreover, the sys-1740 tem automatically computes read (rd) permission instead of programmers 1741 having to compute this value explicitly. 1742

Sadiq et al. (2016, 2019) developed a permission inference framework for sequential Java programs. The aim was to free programmers from specification overhead to manually annotate program with permission-based contracts to help enable implicit concurrency present in the system.

The permission inference approach performs inter-procedural static anal-1747 ysis (data flow and alias flow analysis) of the source code. It automatically 1748 extracts the implicit dependencies present between the code (methods) and 1749 shared states in the source program and maps them in the form of five types 1750 of symbolic permissions such as unique, full, immutable, pure, and share, 1751 without using any method level specifications. Further, it automatically 1752 generates Plural specifications i.e., access permission contracts, using a sin-1753 gle typestate 'alive', for the objects accessed at the method level, to verify 1754 correctness of the inferred specifications by the existing model-checking tool 1755 i.e., Pulse and to reason about their concurrent behavior. 1756

Listing 16 shows sample methods getColl() and initColl() with the inferred access permission contracts in Plural format. The proposed technique generates pure permission as pre-permission (Line 3) on the receiver object (this) as method getColl() read the referenced field (coll). Further, it generates full permission on this in method initColl() (Line 5) as method writes on coll. Following the Design by Contract Principle, the consumed permission are generated as post-permissions for the referenced objects when a method completes its execution.

Listing 16: The inferred permission contracts in Sip4J (Sadiq et al., 2019)

```
1 class ArrayColl{
1766
1767
     \mathbf{2}
       Integer[] coll;
     3 @Perm(requires="pure(this) in alive", ensures="pure(this) in alive")
1768
        public static void getColl(){ return this.coll;}
1769
     4
        @Perm(requires="full(this) in alive", ensures="full(this) in alive")
1770
     \mathbf{5}
     6 public static void initColl(){
1771
1772
     \overline{7}
         for (int i = 0; i < 10; i++){this.coll[i] = i*2;}</pre>
     8 }
1773
```

1765

	Program Statistics and Annotation Overhead				
Benchmark	Program	SLOC	Methods	#AnnLOC	#Annot.
Plural	Crystal	$17,\!512$	2,188	2,500	6,691
-	Pulse	$7,\!671$	461	513	4,850
	montecarlo	1,370	196	204	1,975
	euler	1,080	51	52	1,073
	search	666	50	60	691
	moldyn	608	43	52	901
ICD	lufact	549	42	50	437
JGB	$\operatorname{crypt}$	488	40	46	385
	series	359	37	43	207
	sor	354	34	42	267
	sparsemat	327	33	36	316
	blacksholes	437	21	56	694
Æminium	gaknapsack	232	50	38	250
	health	232	18	25	334
Plaid	webserver	143	12	11	11
	$\mathrm{fft}$	91	11	16	44
	quicksort	66	9	13	17
	shellsort	58	7	10	44
	integral	40	5	7	17
	fibonacci	22	4	8	9
-	Example	71	12	15	78
-	Total	32,376	3,111	$3,\!485$	18,647

Table 11: Annotation Overhead computed for the benchmark programs in (Sadiq et al., 2019).

# 1775 7. An Insight into Research Challenges and Future Directions

Access permissions (Boyland, 2003; Bornat et al., 2005; Bierhoff and Aldrich, 2007) are a novel abstraction that can model read and write effects of a referenced object as well as its aliasing information. Access permission sharing and accounting models attained considerable attention in the research community in the last decade because of their rich expressiveness and sound
reasoning capabilities to specify and verify the correctness of shared-memory
programs.

This survey organizes the permission-based verification of single- and multi-threaded programs into three dimensions according to the use of access permissions and their aims: verifying the correctness of API protocols (Section 4), avoiding synchronization problems in multi-threaded programs (Section 5) and inferring automatically permission-based specifications (Section 6). Table 3, 7 and 10 summarizes and contrasts the surveyed works based on the analysis criteria defined in Section 1.

Our analysis shows that although access permissions have been used to provide a sound reasoning mechanism to verifying program behavior, the existing permission-based verification approaches are limited in their support for verification due to the following reasons.

**Permission Annotation Overhead.** The common problem with all the 1794 permission-based verification approaches is the annotation overhead 1795 associated with the need to manually add permission-based dependen-1796 cies (invariants, contracts, assertions, etc.) or other access notations, to 1797 explicitly specify state changes and grant or restrict access to multiple 1798 references (threads), on the shared memory locations. It is generally 1799 acknowledged that manually annotating programs is laborious, chal-1800 lenging and time-consuming. 1801

Permission Verification Overhead. Given the intricacies in creating permissionbased specifications, it is very likely for a programmer to omit important dependencies or extract the wrong dependencies (read instead of
write). Moreover, there is no guarantee that the manually added specifications are spelled correctly (free of typo errors) and follows their
intended semantic (defined in Section 3) that in turn pose overhead for
verifying the correctness of the input specifications itself.

Although some of the existing approaches (Siminiceanu et al., 2012) present solution to this problem, by identifying the missing specifications and by verifying that, the specifications follow their intended semantics, the techniques themselves are limited in ensuring whether the program implementation complies with the input specifications and vice versa. This is because the analysis is based on input specifications rather than program implementation. Furthermore, although, access permissions support modular reasoning of a program behavior without
analyzing the entire program analysis, certain program properties such
as global invariants and liveness, are hard to verify.

- Permission Tracking Overhead. Existing verification and parallelization approaches use different forms of permission types such as fractional permissions, based on their expressiveness and to facilitate the ease of analysis. The runtime systems then analyzes the permission flow through the system to verify program behavior against the specification and computes the data dependencies in the system based on the specification.
- In particular, fractional permissions use fractional style to express the 1826 access rights for a reference in the range (0, 1] but its analysis is chal-1827 lenging for programmers, due to the overhead associated with tracking 1828 the splitting and the joining of concrete values. It is generally acknowl-1829 edged that tracking fractional values in a program is a cumbersome 1830 task for programmers. The situation becomes more serious when frac-1831 tions are split into multiple levels, which may create concerns relating 1832 to the proper termination of the analysis and affects the soundness of 1833 the technique itself. 1834
- Counting permissions are complementary to fractional permissions but 1835 they do not support all types of synchronization primitives. Symbolic 1836 permissions combine the access rights and aliasing information of a 1837 reference and have been used to allow programmers to reason about 1838 the program correctness, against specification at a higher level of ab-1839 straction than fractional or counting permissions. Therefore, automatic 1840 inference of permission-based specifications in the form of symbolic val-1841 ues is desirable to free the programmers from the low-level analysis 1842 overhead associated with adding and tracking concrete values in the 1843 program. 1844
- 1845

In addition to the specification overheads and related problems discussed
above, the following factors may also hinder the wider adoption of existing
permission-based verification approaches.

Languages and Tools. Existing permission-based verification and paral lelization approaches are mostly based on formal specification lan guages and type systems to support access permission as part of the

language. It can be challenging for most programmers who may not
be adept at the new syntax and semantics in order to exploit their
functionalities. Furthermore, most of these approaches are either research prototypes or developed in languages that are not commonly
used for general-purpose software development. These factors could
limit the adoption of the existing approaches to verify programs written in mainstream programming languages such as Java.

Program Constructs. Existing verification approaches have limited sup port for synchronization constructs such as fork-join parallelism, atomic
 blocks or semaphores. There is also limited support for the recursive
 data structures. Most of the approaches support verification of heap
 locations, while investigation on non-heap locations has not been as
 prevalent. These limitations restrict their ability to verify real-world
 applications.

Program analysis. Existing approaches either perform static and dynamic
program analysis or employ model-checking techniques to verify program properties. All have their own pros and cons that affect the
scalability of these approaches when analyzing program constructs and
verifying the program behavior. For example, the techniques employing model checkers may face state-space explosion problems even for a
program of average size.

Properties. Most verification tools focus on certain aspects of the program
behavior, such as verifying the correctness of API protocols or avoiding
synchronization issues such as data races are just two aspects. Some
tools can address a combination of issues, but none of them cover all
aspects of a program behavior.

Although, existing permission-based verification approaches are quite promis-1878 ing in verifying program behavior. The common problem in all the existing 1879 approaches is the annotation overhead associated with the need to manually 1880 add the permission-based specifications in the source program. Therefore, 1881 the automatic inference of permission-based specifications from the source 1882 program can be the first step to exploit the verification power of these ap-1883 proaches, without posing any extra burden on programmers, to enhance their 1884 applicability in the IT industry. 1885

Overall, given the number of different permission-based formal type theories and programming models that received remarkable attention over the last decade in the research community, there is an impending need to make these approaches adoptable and adaptable for general-purpose program development, and verification for mainstream programming languages such as Java.

The first step toward this goal can be an integration of the most widely used permission-based verification tools that at least support a common programming models such as Java. For example, Plural, Pulse, Verifast, VerCors and Sip4J fall in this category.

The analysis of existing approaches reveals that the inference of access 1896 permission contracts can be used to automatically compute the dependencies 1897 between methods while making the side effects explicit. As access permis-1898 sions pose their own ordering constraints, the computed dependencies can 1890 be used to automatically infer the synchronization primitives (locking and 1900 ordering constraints) from the source code of a sequential program. The 1901 generated specifications can then be used to enforce the locking policy to 1902 different program parts at different levels of granularity (method, instruction 1903 or task). This information can be used to automatically parallelize pro-1904 gram execution for the mainstream programming languages such as Java. 1905 to the extent permitted by the computed dependencies, without using any 1906 new programming language and runtime system to support access permis-1907 sions. The permission-based parallelization can free the programmers from 1908 the low-level synchronization and reasoning overhead associated with han-1909 dling multi-threading in sequential programming paradigms. 1910

In addition to supporting the race-free sharing of the heap or non-heap locations in sequential programs, the inference of permission-based synchronization constructs (such as acquire and release locks with permission invariants) can be used to verify the behavior of concurrent programs, that have already been written using multi-threading, without imposing extra work on the programmer side.

The ideal solution to all the above challenges can be the integration of the commonly used abstractions such as typestates and access permissions, as first-class language constructs in the mainstream programming models, to develop a complete, sound, modular, automated and economically feasible framework for everyday program development and verification. The new technique could be realized to provide modelling, verification and parallelization support without posing any extra burden on programmers.

#### 1924 Acknowledgment

This research was conducted under Endeavour Leadership Award, for P.hD., funded by the Department of Education and Training, Government of Australia and the Faculty of Information Technology, Monash University, under Postgraduate Publication Award.

- Abadi, M., Flanagan, C., Freund, S. N., 3 2006. Types for Safe Locking: Static Race Detection for Java.
   ACM Trans. Program. Lang. Syst. 28 (2), 207–255.
- Ahmed, I., Cataño, N., 2 2018. Checking JML-encoded finite state machine properties. In: 2018 Interna tional Conference on Advancements in Computational Sciences (ICACS). pp. 1–9.
- Aldrich, J., Beckman, N. E., Bocchino, R., Naden, K., Saini, D., Stork, S., Sunshine, J., 2012. The Plaid
   language: Typed core specification. Tech. rep., DTIC Document.
- Aldrich, J., Bocchino, R., Garcia, R., Hahnenberg, M., Mohr, M., Naden, K., Saini, D., Stork, S., Sunshine,
   J., Tanter, r., others, 2011. Plaid: a permission-based programming language. In: Proceedings of the
   ACM international conference companion on Object oriented programming systems languages and
   applications companion. ACM, pp. 183–184.
- Aldrich, J., Sunshine, J., Saini, D., Sparks, Z., 2009. Typestate-oriented Programming. In: Proceedings of
   the 24th ACM SIGPLAN Conference Companion on Object Oriented Programming Systems Languages
   and Applications. OOPSLA '09. ACM, pp. 1015–1022.
- Amighi, A., Blom, S., Darabi, S., Huisman, M., Mostowski, W., Zaharieva-Stojanovski, M., 2014. Verification of Concurrent Systems with VerCors. In: Formal Methods for Executable Software Models: 14th
  International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2014, Bertinoro, Italy, June 16-20, 2014, Advanced Lectures. Springer International Publishing, pp. 172–216.
- Amighi, A., Blom, S., Huisman, M., Zaharieva-Stojanovski, M., 2012. The {VerCors} Project: Setting Up
   Basecamp. In: Programming Languages meets Program Verification (PLPV 2012).
- Amighi, A., Haack, C., Huisman, M., Hurlin, C., 2015. Permission-based separation logic for multithreaded
   java programs. Logical Methods in Computer Science.
- Ancona, D., Bono, V., Bravetti, M., 2016. Behavioral Types in Programming Languages. Now Publishers
   Inc.
- Appel, A. W., Blazy, S., 2007. Separation Logic for Small-Step cminor. In: Theorem Proving in Higher
   Order Logics. Springer Berlin Heidelberg, pp. 5–21.
- Araujo, W., Briand, L., Labiche, Y., 2008. Concurrent contracts for Java in JML. In: Software Reliability
   Engineering, 2008. ISSRE 2008. 19th International Symposium on. IEEE, pp. 37–46.
- Artho, C., Havelund, K., Biere, A., 2003. High-level data races. In: Software Testing Verification and
   Reliability.
- Barnett, M., Chang, B.-Y. E., DeLine, R., Jacobs, B., Leino, K. R. M., 2006. Boogie: A modular reusable
  verifier for object-oriented programs. In: Formal Methods for Components and Objects. Springer Berlin
  Heidelberg, pp. 364–387.
- Beckert, B., Hhnle, R., Schmitt, P. H. P. H., 2007. Verification of object-oriented software : the KeY
   approach. Springer.

- Beckman, N. E., 2009. Modular typestate checking in concurrent Java programs. In: Proceedings of the
   24th ACM SIGPLAN conference companion on Object oriented programming systems languages and
   applications. ACM, pp. 737–738.
- Beckman, N. E., 12 2010. Types for Correct Concurrent API Usage, PhD thesis, technical report CMU ISR-10-131. Ph.D. thesis.
- Beckman, N. E., Bierhoff, K., Aldrich, J., 2008. Verifying Correct Usage of Atomic Blocks and Typestate.
  In: Proceedings of the 23rd ACM SIGPLAN Conference on Object-oriented Programming Systems
  Languages and Applications. OOPSLA '08. ACM, pp. 227–244.
- Beckman, N. E., Kim, D., Aldrich, J., 2011. An Empirical Study of Object Protocols in the Wild. In:
  Proceedings of the 25th European Conference on Object-oriented Programming. ECOOP'11. SpringerVerlag, pp. 2–26.
- Betts, A., Chong, N., Donaldson, A., Qadeer, S., Thomson, P., Betts, A., Chong, N., Donaldson, A.,
  Qadeer, S., Thomson, P., 2012. GPUVerify. In: Proceedings of the ACM international conference on
  Object oriented programming systems languages and applications OOPSLA '12. Vol. 47. ACM Press,
  p. 113.
- 1979 Bierhoff, Kevin, 2009. Api protocol compliance in object-oriented software.
- Bierhoff, K., 2011. Automated program verification made SYMPLAR: symbolic permissions for lightweight
   automated reasoning. In: Proceedings of the 10th SIGPLAN symposium on New ideas, new paradigms,
   and reflections on programming and software. ACM, pp. 19–32.
- Bierhoff, K., Aldrich, J., 2007. Modular typestate checking of aliased objects. Vol. 42 of OOPSLA '07.
   ACM.
- Bierhoff, K., Aldrich, J., 2008. PLURAL: Checking Protocol Compliance Under Aliasing. In: Companion of the 30th International Conference on Software Engineering. ICSE Companion '08. ACM, pp. 971–972.
- Bierhoff, K., Beckman, N. E., Aldrich, J., 2009a. Polymorphic fractional permission inference. In: Proceedings of the 18th ACM SIGSOFT International Symposium on Software Testing and Analysis. ISSTA
  09.
- Bierhoff, K., Beckman, N. E., Aldrich, J., 2009b. Practical API Protocol Checking with Access Permissions.
   In: ECOOP. pp. 195–219.
- Blom, S., Huisman, M., 2014. The vercors tool for verification of concurrent programs. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics).
- Blom, S., Huisman, M., Miheli, M., 2014. Specification and verification of GPGPU programs. Science of
   Computer Programming.
- Bornat, R., Calcagno, C., O'Hearn, P., Parkinson, M., 1 2005. Permission Accounting in Separation Logic.
   SIGPLAN Not. 40 (1), 259–270.
- Boyapati, C., Lee, R., Rinard, M., 2002. Ownership Types for Safe Programming: Preventing Data
  Races and Deadlocks. In: Proceedings of the 17th ACM SIGPLAN Conference on Object-oriented
  Programming, Systems, Languages, and Applications. OOPSLA '02. ACM, pp. 211–230.
- Boyapati, C., Rinard, M., 2001. A Parameterized Type System for Race-free Java Programs. In: Proceed ings of the 16th ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages,
   and Applications. OOPSLA '01. ACM, pp. 56–69.

- Boyland, J., 2003. Checking Interference with Fractional Permissions. In: Proceedings of the 10th Inter national Conference on Static Analysis. SAS'03. Springer-Verlag, pp. 55–72.
- Boyland, J., 2006. Why we should not add readonly to Java (yet). The Journal of Object Technology 5 (5), 5.
- Boyland, J. T., 8 2010. Semantics of Fractional Permissions with Nesting. ACM Trans. Program. Lang.
   Syst. 32 (6), 22:1–22:33.
- 2012 Brookes, S., 2004. A Semantics for Concurrent Separation Logic. Springer, Berlin, Heidelberg, pp. 16–34.
- 2013 Butenhof, D. R., 1997. Programming with POSIX threads. Addison-Wesley Professional.
- 2014 C. Barrett A. Stump, Tinelli, C., 2010. The SMTLIB Standard, Version 2.0.
- Caires, L., 7 2008. Spatial-behavioral Types for Concurrency and Resource Control in Distributed Systems.
   Theor. Comput. Sci. 402 (2-3), 120–141.
- Caires, L., Cardelli, L., 2002. A Spatial Logic for Concurrency (Part II). In: Proceedings of the 13th
   International Conference on Concurrency Theory. CONCUR '02. Springer-Verlag, pp. 209–225.
- 2019 Caires, L., Cardelli, L., 11 2003. A Spatial Logic for Concurrency (Part I). Inf. Comput. 186 (2), 194–235.
- Caires, L., Pfenning, F., 2010. Session types as intuitionistic linear propositions. In: Lecture Notes in
   Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in
   Bioinformatics).
- Caires, L., Seco, J. C., 2013. The Type Discipline of Behavioral Separation. In: Proceedings of the 40th
   Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '13.
   ACM, pp. 275–286.
- Capriccioli, A., Servetto, M., Zucca, E., 4 2016. An Imperative Pure Calculus. Electron. Notes Theor.
   Comput. Sci. 322 (C), 87–102.
- Carr, S. A., Logozzo, F., Payer, M., Aug 2017. Automatic contract insertion with ccbot. IEEE Transactions
   on Software Engineering 43 (8), 701–714.
- Cataño, N., Ahmed, I., Siminiceanu, R. I., Aldrich, J., 2014. A case study on the lightweight verification
   of a multi-threaded task server. Sci. Comput. Program. 80, 169–187.
- Cav, V., Zhao, J., Shirako, J., Sarkar, V., 2011. Habanero-Java: The New Adventures of Old X10. In:
   Proceedings of the 9th International Conference on Principles and Practice of Programming in Java.
   PPPJ '11. ACM, pp. 51–61.
- Chaki, S., Clarke, E., Groce, A., Jha, S., Veith, H., 6 2004. Modular verification of software components
   in C. IEEE Transactions on Software Engineering 30 (6), 388–402.
- Chaki, S., Gurfinkel, A., 2018. BDD-Based Symbolic Model Checking. In: Handbook of Model Checking.
   Springer International Publishing, pp. 219–245.
- Chaki, S., Rajamani, S. K., Rehof, J., 2002. Types As Models: Model Checking Message-passing Programs.
  In: Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming
  Languages. POPL '02. ACM, pp. 45–57.
- Clarke, D., stlund, J., Sergey, I., Wrigstad, T., 2013. Ownership Types: A Survey. Springer, Berlin,
   Heidelberg, pp. 15–58.

- Clarke, D., Wrigstad, T., 2003. External Uniqueness Is Unique Enough. Springer, Berlin, Heidelberg, pp.
   176–200.
- Clebsch, S., Drossopoulou, S., Blessing, S., Mcneil, A., 2015. Deny Capabilities for Safe, Fast Actors. In:
   Proceedings of the 5th ....
- Cohen, E., Dahlweid, M., Hillebrand, M., Leinenbach, D., Moskal, M., Santen, T., Schulte, W., Tobies, S.,
   2009. VCC: A Practical System for Verifying Concurrent C. Springer, Berlin, Heidelberg, pp. 23–42.
- Cok, D. R., 2011. OpenJML: JML for Java 7 by extending OpenJDK. In: Lecture Notes in Computer Sci ence (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics).
- Cousot, P., Cousot, R., 1977. Abstract interpretation: a unified lattice model for static analysis of programs
  by construction or approximation of fixpoints. In: Proceedings of the 4th ACM SIGACT-SIGPLAN
  symposium on Principles of programming languages. ACM, pp. 238–252.
- DeLine, R., Fähndrich, M., 2002. Adoption and focus: Practical linear types for imperative programming.
   Programming language design and implementation, ACM SIGPLAN.
- DeLine, R., Fähndrich, M., 2004. Typestates for Objects. In: Odersky, M. (Ed.), ECOOP 2004 Object Oriented Programming. Springer Berlin Heidelberg, pp. 465–490.
- Dezani-Ciancaglini, M., Yoshida, N., Ahern, A., Drossopoulou, S., 2005. A distributed object-oriented
   language with session types. In: Lecture Notes in Computer Science (including subseries Lecture Notes
   in Artificial Intelligence and Lecture Notes in Bioinformatics).
- Dias, R. J., Pessanha, V., Loureno, J. M., 2013. Precise Detection of Atomicity Violations. Springer,
   Berlin, Heidelberg, pp. 8–23.
- Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M. J., Vafeiadis, V., 2010. Concurrent Ab stract Predicates. In: Proceedings of the 24th European Conference on Object-oriented Programming.
   ECOOP'10. Springer-Verlag, pp. 504–528.
- Dockins, R., Hobor, A., Appel, A. W., 2009. A fresh look at separation algebras and share accounting.
   In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics).
- Dohrau, J., Summers, A. J., Urban, C., Münger, S., Müller, P., 2018. Permission inference for array programs. In: Chockler, H., Weissenbacher, G. (Eds.), Computer Aided Verification. Springer International
  Publishing, Cham, pp. 55–74.
- Dolby, J., Hammer, C., Marino, D., Tip, F., Vaziri, M., Vitek, J., 5 2012. A Data-centric Approach to
   Synchronization. ACM Trans. Program. Lang. Syst. 34 (1), 4:1–4:48.
- Engler, D., Ashcraft, K., 2003. RacerX: Effective, Static Detection of Race Conditions and Deadlocks. In:
  Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles. SOSP '03. ACM,
  pp. 237–252.
- Fähndrich, M., Logozzo, F., 2011. Static contract checking with abstract interpretation. In: Proceedings
  of the 2010 International Conference on Formal Verification of Object-oriented Software. FoVeOOS'10.
  Springer-Verlag, pp. 10–30.
- Ferrara, P., Müller, P., 2012. Automatic Inference of Access Permissions. In: Verification, Model Checking, and Abstract Interpretation: 13th International Conference, VMCAI 2012, Philadelphia, PA, USA, January 22-24, 2012. Proceedings. Springer Berlin Heidelberg, pp. 202–218.

- Fillitre, J. C., Paskevich, A., 2013. Why3 Where programs meet provers. In: Lecture Notes in Computer
   Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformat ics).
- Flanagan, C., Freund, S. N., Lifshin, M., Qadeer, S., 8 2008. Types for Atomicity: Static Checking and
   Inference for Java. ACM Trans. Program. Lang. Syst. 30 (4), 20:1–20:53.
- Flanagan, C., Qadeer, S., Flanagan, C., Qadeer, S., 2003. A type and effect system for atomicity. In: Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation
   PLDI '03. Vol. 38. ACM Press, pp. 338–349.
- Floyd, R. W., 1967. Assigning meanings to programs. Proceedings of Symposium on Applied Mathematics
   19, 19–32.

#### 2094 URL http://laser.cs.umass.edu/courses/cs521-621.Spr06/papers/Floyd.pdf

- Garcia, R., Tanter, r., Wolff, R., Aldrich, J., 10 2014. Foundations of Typestate-Oriented Programming.
   ACM Trans. Program. Lang. Syst. 36 (4), 12:1–12:44.
- Gay, S. J., Gesbert, N., Ravara, A., Vasconcelos, V. T., 2015a. Modular session types for objects. Logical
   Methods in Computer Science.
- Gay, S. J., Gesbert, N., Ravara, A., Vasconcelos, V. T., 2015b. Modular session types for objects. Logical
   Methods in Computer Science 11 (4).
- 2101 URL https://doi.org/10.2168/LMCS-11(4:12)2015
- 2102 Gay, S. J., Vasconcelos, V. T., 2010. Linear type theory for asynchronous session types. Journal of Func-2103 tional Programming.
- Gherghina, C., David, C., Qin, S., Chin, W.-N., 2011. Structured specifications for better verification of
   heap-manipulating programs. In: FM 2011: Formal Methods. Springer Berlin Heidelberg, pp. 386–401.
- Giannini, P., Richter, T., Servetto, M., Zucca, E., 2018a. Tracing sharing in an imperative pure calculus.
   CoRR abs/1803.0.
- Giannini, P., Servetto, M., Zucca, E., 2018b. A Type and Effect System for Uniqueness and Immutability.
   In: Proceedings of the 33rd Annual ACM Symposium on Applied Computing. SAC '18. ACM, pp. 1038–1045.
- 2111 Girard, J.-Y., 1987. Linear logic. Theoretical computer science 50 (1), 1–101.
- Gordon, C. S., Parkinson, M. J., Parsons, J., Bromfield, A., Duffy, J., 2012. Uniqueness and reference
  immutability for safe parallelism. In: Proceedings of the ACM international conference on Object
  oriented programming systems languages and applications OOPSLA '12.
- Gotsman, A., Berdine, J., Cook, B., Sagiv, M., 6 2007. Thread-modular Shape Analysis. SIGPLAN Not.
   42 (6), 266–277.
- 2117 Greenhouse, A., 2003. A Programmer-Oriented Approach to Safe Concurrency. Tech. rep.
- 2118 URL http://reports-archive.adm.cs.cmu.edu/anon/2003/CMU-CS-03-135.pdf
- Greenhouse, A., Scherlis, W. L., 2002. Assuring and evolving concurrent programs: annotations and
   policy. In: Proceedings of the 24th International Conference on Software Engineering.
- Haack, C., Huisman, M., Hurlin, C., 2008. Reasoning about java's reentrant locks. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics).

- Haack, C., Hurlin, C., 2008. Separation Logic Contracts for a Java-Like Language with Fork/Join. In:
   Algebraic Methodology and Software Technology. Springer Berlin Heidelberg, pp. 199–215.
- Hammer, C., Dolby, J., Vaziri, M., Tip, F., 2008. Dynamic Detection of Atomic-Set-Serializability Viola tions. In: Proceedigns of the 30th International Conference on Software Engineering (ICSE' 08).
- Heule, S., Leino, K. R. M., Müller, P., Summers, A. J., 2011. Fractional permissions without the fractions.
  In: Proceedings of the 13th Workshop on Formal Techniques for Java-Like Programs. ACM, p. 1.
- Heule, S., Leino, K. R. M., Müller, P., Summers, A. J., 2013. Abstract read permissions: Fractional permissions without the fractions. In: International Workshop on Verification, Model Checking, and Abstract Interpretation. Springer, pp. 315–334.
- Hoare, C. A. R., 1969. An axiomatic basis for computer programming. Communications of the ACM
   12 (10), 576–580.
- Hoare, C. A. R., 1972. Towards a Theory of Parallel Programming. In: The Origin of Concurrent Pro gramming. Springer New York, pp. 231–244.
- 2137 Hoare, C. A. R., 1974. Monitors: an operating system structuring concept. Communications of the ACM.
- Hobor, A., Appel, A. W., Nardelli, F. Z., 2008. Oracle semantics for concurrent separation logic. In:
  Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and
  Lecture Notes in Bioinformatics).
- Hobor, A., Gherghina, C., 2012. Barriers in concurrent separation logic: Now with tool support! Logical
   Methods in Computer Science.
- 2143 Honda, K., 1993. Types for dyadic interaction. Springer, Berlin, Heidelberg, pp. 509–523.
- Honda, K., Vasconcelos, V. T., Kubo, M., 1998. Language Primitives and Type Discipline for Structured
  Communication-Based Programming. In: Proceedings of the 7th European Symposium on Programming: Programming Languages and Systems. ESOP '98. Springer-Verlag, pp. 122–138.
- Honda, K., Yoshida, N., Carbone, M., 2008. Multiparty Asynchronous Session Types. In: Proceedings
  of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages.
  POPL '08. ACM, pp. 273–284.
- Hu, R., Kouzapas, D., Pernet, O., Yoshida, N., Honda, K., 2010. Type-safe eventful sessions in Java. In:
  Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and
  Lecture Notes in Bioinformatics).
- Hu, R., Yoshida, N., Honda, K., 2008. Session-based distributed programming in Java. In: Lecture Notes
  in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in
  Bioinformatics).
- Huisman, M., 2001. Reasoning about Java programs in higher order logic using PVS and Isabelle. PhD
   Thesis. Ph.D. thesis, Computing Science Institute, University of Nijmegen.
- Huisman, M., Mostowski, W., 2015. A symbolic approach to permission accounting for concurrent reasoning. In: Proceedings IEEE 14th International Symposium on Parallel and Distributed Computing,
  ISPDC 2015.
- Hüttel, H., Lanese, I., Vasconcelos, V. T., Caires, L., Carbone, M., Denilou, P.-M., Mostrous, D., Padovani,
  L., Ravara, A., Tuosto, E., Vieira, H. T., Zavattaro, G., 4 2016. Foundations of Session Types and
- 2163 Behavioural Contracts. ACM Comput. Surv. 49 (1), 3:1–3:36.

- 2164 Igarashi, A., Kobayashi, N., 2001. Resource usage analysis. In: POPL.
- Igarashi, A., Kobayashi, N., 3 2005. Resource Usage Analysis. ACM Trans. Program. Lang. Syst. 27 (2),
   264–313.
- Jacobs, B., Leino, K., Piessens, F., Schulte, W., 2005. Safe concurrency for aggregate objects with invariants. In: Third IEEE International Conference on Software Engineering and Formal Methods
  (SEFM'05). IEEE, pp. 137–146.
- Jacobs, B., Piessens, F., 2011. Expressive modular fine-grained concurrency specification. In: Proceedings
  of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages.
  POPL '11. ACM, pp. 271–282.
- Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F., 2011. VeriFast: A powerful,
  sound, predictable, fast verifier for C and java. In: NASA Formal Methods Symposium. Springer, pp.
  41–55.
- Jacobs, B., Smans, J., Piessens, F., 11 2010. A Quick Tour of the VeriFast Program Verifier. Springer,
   Berlin, Heidelberg, pp. 304–311.
- 2178 URL http://link.springer.com/10.1007/978-3-642-17164-2\_21
- Jensen, J. B., Birkedal, L., 2012. Fictional Separation Logic. In: Seidl, H. (Ed.), Programming Languages
   and Systems. Springer Berlin Heidelberg, pp. 377–396.
- 2181 Jones, C. B., 1983. Specification and Design of (Parallel) Programs. In: IFIP Congress.
- Juhasz, U., Kassios, I. T., Müller, P., Novacek, M., Schwerhoff, M., Summers, A. J., 2014. Viper. Tech.
   rep.
- Kidd, N., Reps, T., Dolby, J., Vaziri, M., 2011. Finding concurrency-related bugs using random isolation.
   International Journal on Software Tools for Technology Transfer.
- Kim, T., Bierhoff, K., Aldrich, J., Kang, S., 2009. Typestate protocol specification in JML. In: Proceedings
   of the 8th international workshop on Specification and verification of component-based systems. ACM,
   pp. 11–18.
- Kobayashi, N., Sangiorgi, D., 5 2010. A hybrid type system for lock-freedom of mobile processes. ACM
   Transactions on Programming Languages and Systems 32 (5), 1–49.
- Lai, Z., Cheung, S. C., Chan, W. K., 2010. Detecting atomic-set serializability violations in multithreaded
   programs through active randomized testing. In: Proceedings of the 32nd ACM/IEEE International
   Conference on Software Engineering ICSE '10.
- Le, D.-K., Chin, W.-N., Teo, Y.-M., 2012. Variable Permissions for Concurrency Verification. In: Formal Methods and Software Engineering: 14th International Conference on Formal Engineering Methods, ICFEM 2012, Kyoto, Japan, November 12-16, 2012. Proceedings. Springer Berlin Heidelberg, pp. 5–21.
- Leavens, G. T., Baker, A. L., Ruby, C., 2006. Preliminary design of JML: A behavioral interface specification language for Java. ACM SIGSOFT Software Engineering Notes 31 (3), 1–38.
- 2199 Leavens, G. T., Cheon, Y., 2006. Design by Contract with JML.
- Leino, K. R. M., Müller, P., 2009. A basis for verifying multi-threaded programs. In: European Symposium
   on Programming. Springer, pp. 378–393.

- Leino, K. R. M., Müller, P., Smans, J., 2009. Verification of Concurrent Programs with Chalice. In:
  Foundations of Security Analysis and Design V: FOSAD 2007/2008/2009 Tutorial Lectures. Springer
  Berlin Heidelberg, pp. 195–222.
- Lu, S., Park, S., Seo, E., Zhou, Y., Lu, S., Park, S., Seo, E., Zhou, Y., Lu, S., Park, S., Seo, E., Zhou, Y.,
  Lu, S., Park, S., Seo, E., Zhou, Y., 3 2008. Learning from mistakes. ACM SIGOPS Operating Systems
  Review 42 (2), 329.
- Marino, D., Hammer, C., Dolby, J., Vaziri, M., Tip, F., Vitek, J., 5 2013. Detecting deadlock in programs
  with data-centric synchronization. In: 2013 35th International Conference on Software Engineering
  (ICSE). pp. 322–331.
- Matsakis, N. D., Klock, F. S., Matsakis, N. D., Klock II, F. S., 2014. The rust language. In: Proceedings of the 2014 ACM SIGAda annual conference on High integrity language technology - HILT '14. Vol. 34.
  ACM Press, pp. 103–104.
- 2214 Meyer, B., 1988. Object-Oriented Software Construction, 1st Edition. Prentice-Hall, Inc.
- 2215 Meyer, B., 10 1992. Applying 'design by contract'. Computer 25 (10), 40–51.
- Militao, F., Aldrich, J., Caires, L., 2010. Aliasing Control with View-based Typestate. In: Proceedings of
   the 12th Workshop on Formal Techniques for Java-Like Programs. FTFJP '10. ACM, pp. 7:1–7:7.
- 2218 Militão, F., Aldrich, J., Caires, L., 2014a. Rely-Guarantee Protocols. In: ECOOP 2014 Object-Oriented 2219 Programming. Springer Berlin Heidelberg, pp. 334–359.
- Militão, F., Aldrich, J., Caires, L., 2014b. Substructural Typestates. In: Proceedings of the ACM SIG PLAN 2014 Workshop on Programming Languages Meets Program Verification. PLPV '14. ACM, pp.
   15–26.
- Militão, F., Aldrich, J., Caires, L., 2016. Composing interfering abstract: Protocols. In: 30th European
   Conference on Object-Oriented Programming, ECOOP 2016. Vol. 56. pp. 161–1626.
- Milito, F., Caires, L., 2009. An Exception Aware Behavioral Type System for Object-Oriented Programs.
   In: INFORUM 2009 Simpsio de Informitica. Faculdade de Cincias Universidade de Lisboa, pp. 1–12.
- Morrisett, G., Ahmed, A., Fluet, M., 2005. L3: A Linear Language with Locations. In: Typed Lambda
   Calculi and Applications. Springer Berlin Heidelberg, pp. 293–307.
- Müller, P., Rudich, A., 2007. Ownership transfer in universe types. In: Proceedings of the 22nd annual
   ACM SIGPLAN conference on Object oriented programming systems and applications OOPSLA '07.
- Müller, P., Schwerhoff, M., Summers, A. J., 2017. Viper: A verification infrastructure for permission-based reasoning. In: Dependable Software Systems Engineering.
- Naden, K., Bocchino, R., Aldrich, J., Bierhoff, K., 2012. A Type System for Borrowing Permissions. In:
   Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming
   Languages. POPL '12. ACM, pp. 557–570.
- Naik, M., Park, C. S., Sen, K., Gay, D., 2009. Effective static deadlock detection. In: Proceedings International Conference on Software Engineering.
- Nguyen, H. H., Chin, W.-N., 2008. Enhancing program verification with lemmas. In: Proceedings of the
   20th International Conference on Computer Aided Verification. CAV '08. Springer-Verlag, pp. 355–369.
- 2240 Nielson, F., Nielson, H. R., 1993. From CML to process algebras. Springer, Berlin, Heidelberg, pp. 493–508.

- 2241 Nielson, F., Nielson, H. R., 1996. From CML to its process algebra. Theoretical Computer Science.
- Noble, J., Vitek, J., Potter, J., 1998. Flexible alias protection. In: Lecture Notes in Computer Science(including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics).
- Odersky, M., Altherr, P., Cremet, V., Emir, B., Maneth, S., Micheloud, S., Mihaylov, N., Schinz, M.,
  Stenman, E., Zenger, M., 2004. An overview of the Scala programming language. Tech. rep.
- O'Hearn, P., Reynolds, J., Yang, H., 2001. Local Reasoning about Programs that Alter Data Structures.
   In: Computer Science Logic. Springer Berlin Heidelberg, pp. 1–19.
- 2248 OHearn, P. W., 4 2007. Resources, Concurrency, and Local Reasoning. Theor. Comput. Sci. 375 (1-3),
   2249 271–307.
- Owicki, S., Gries, D., 5 1976. Verifying Properties of Parallel Programs: An Axiomatic Approach. Com mun. ACM 19 (5), 279–285.
- Owre, S., Rushby, J. M., Shankar, N., 1992. PVS: a prototype verification system. 11th International
   Conference on Automated Deduction.
- Parkinson, M., Bierman, G., 2005. Separation logic and abstraction. ACM SIGPLAN Notices 40 (1), 247–258.
- Paulino, H., Parreira, D., Delgado, N., Ravara, A., Matos, A., 2016. From Atomic Variables to Data-centric
   Concurrency Control. In: Proceedings of the 31st Annual ACM Symposium on Applied Computing.
   SAC '16. ACM, pp. 1806–1811.
- Pradel, M., Jaspan, C., Aldrich, J., Gross, T. R., 2012. Statically checking API protocol conformance with
   mined multi-object specifications. In: Proceedings International Conference on Software Engineering.
- Reynolds, J. C., 1978. Syntactic Control of Interference. In: Proceedings of the 5th ACM SIGACT SIGPLAN Symposium on Principles of Programming Languages. POPL '78. ACM, pp. 39–46.
- Reynolds, J. C., 2002. Separation logic: A logic for shared mutable data structures. In: Logic in Computer
   Science, 2002. Proceedings. 17th Annual IEEE Symposium on. IEEE, pp. 55–74.
- Rodríguez, E., Dwyer, M., Flanagan, C., Hatcliff, J., Leavens, G., Robby, 2005. Extending JML for Modular Specification and Verification of Multi-threaded Programs. In: ECOOP 2005 Object-Oriented
  Programming.
- Roux, P., Siminiceanu, R., 2010. Model Checking with Edge-valued Decision Diagrams. In: Second {NASA} Formal Methods Symposium {NFM} 2010, Washington D.C., USA, April 13-15, 2010.
  Proceedings. pp. 222–226.
- Sadiq, A., Li, Y., Ling, S., Li, L., Ahmed, I., 2019. Sip4j: Statically inferring permission-based specifications for sequential java programs. CoRR abs/1902.05311.
   URL http://ownin.org/obc/1902.05311.
- 2273 URL http://arxiv.org/abs/1902.05311
- Sadiq, A., Li, Y.-F., Ling, S., Ahmed, I., 2016. Extracting Permission-Based Specifications from a Sequential Java Program. In: 21st International Conference on Engineering of Complex Computer Systems,
  United Arab Emirates, November 6-8, 2016. pp. 215–218.
- 2277 Sangiorgi, D., 1999. The name discipline of uniform receptiveness. Theoretical Computer Science.
- Schwinghammer, J., Birkedal, L., Reus, B., Yang, H., 2011. Nested Hoare Triples and Frame Rules for
   Higher-order Store. Logical Methods in Computer Science 7 (3).

- Siek, J., Taha, W., 2007. Gradual typing for objects. In: Ernst, E. (Ed.), ECOOP 2007 Object-Oriented
   Programming. Springer Berlin Heidelberg, pp. 2–27.
- Siminiceanu, R. I., Ahmed, I., Cataño, N., 2012. Automated Verification of Specifications with Typestates
   and Access Permissions. ECEASST 53.
- Smans, J., Jacobs, B., Piessens, F., 2009. Implicit Dynamic Frames: Combining Dynamic Frames and
   Separation Logic. Springer, Berlin, Heidelberg, pp. 148–172.
- Stork, S., Naden, K., Sunshine, J., Mohr, M., Fonseca, A., Marques, P., Aldrich, J., 2014. aem: A
  Permission-Based Concurrent-by-Default Programming Language Approach. ACM Trans. Program.
  Lang. Syst. 36 (1), 1–42.
- 2289 Strom, R. E., Yemini, S., 1986. Typestate: A programming language concept for enhancing software 2290 reliability. IEEE Transactions on Software Engineering (1), 157–171.
- Sunshine, J., Naden, K., Stork, S., Aldrich, J., Tanter, r., 2011. First-class state change in plaid. ACM
   SIGPLAN Notices.
- Takeuchi, K., Honda, K., Kubo, M., 1994. An interaction-based language and its typing system. In: In
   PARLE94, volume 817 of LNCS. pp. 398–413.
- Vafeiadis, V., Parkinson, M., 2007. A Marriage of Rely/Guarantee and Separation Logic. In: Caires, L.,
  Vasconcelos, V. T. (Eds.), CONCUR 2007 Concurrency Theory. Springer Berlin Heidelberg, pp.
  256–271.
- Vaziri, M., Tip, F., Dolby, J., 2006. Associating Synchronization Constraints with Data in an Object oriented Language. In: Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '06. ACM, pp. 334–345.
- Vaziri, M., Tip, F., Dolby, J., Hammer, C., Vitek, J., 2010. A Type System for Data-Centric Synchronization. In: ECOOP 2010 – Object-Oriented Programming. Springer Berlin Heidelberg, pp. 304–328.
- Villard, J., Lozes, É., Calcagno, C., 2009. Proving copyless message passing. In: Programming Languages
   and Systems. Springer Berlin Heidelberg, pp. 194–209.
- Villard, J., Lozes, t., Calcagno, C., 2010. Tracking Heaps That Hop with Heap-Hop. Springer, Berlin,
   Heidelberg, pp. 275–279.
- Visser, W., Havelund, K., Brat, G., Park, S., Lerda, F., 2003. Model Checking Programs. Automated
   Software Engineering 10 (2), 203–232.
- Voung, J. W., Jhala, R., Lerner, S., 2007. RELAY. In: Proceedings of the the 6th joint meeting of the
  European software engineering conference and the ACM SIGSOFT symposium on The foundations of
  software engineering ESEC-FSE '07. ACM Press, p. 205.
- 2312 Wadler, P., 2014. Propositions as sessions. In: Journal of Functional Programming.
- Westbrook, E., Zhao, J., Budimli, Z., Sarkar, V., 2012. Practical Permissions for Race-free Parallelism. In:
   Proceedings of the 26th European Conference on Object-Oriented Programming. ECOOP'12. Springer Verlag, pp. 614–639.
- Xu, M., Bodk, R., Hill, M. D., Xu, M., Bodk, R., Hill, M. D., 6 2005. A serializability violation detector
   for shared-memory server programs. ACM SIGPLAN Notices 40 (6), 1.
- Zhao, Y., 2007. Concurrency analysis based on fractional permissions. Ph.D. thesis, University of
   Wisconsin-Milwaukee.

Zhao, Y., Yu, L., Bei, J., 12 2008. The Permission Approach to Comprehend Lock-Based Synchronization
Policy. In: 2008 International Conference on Advanced Computer Theory and Engineering. IEEE, pp. 709–713.

# 2323 Appendix A. Vitae

- Miss. Ayesha Sadiq is a final year PhD student in the Faculty of Information Technology at Monash University. Her research interest primarily lies in software engineering with a focus on programming languages, program analysis and verifications and formal modelling of real time applications.
- 2329
- Dr. Yuan-Fang Li is a Senior Lecturer at Faculty of Information Technology, Monash University, Australia. He received his PhD in computer science from National University of Singapore in 2006. His research interests include knowledge graphs, knowledge representation and reasoning, ontology languages, and software engineering.
- 2335

• Dr. Sea Ling works in the Faculty of Information Technology at Monash University as a senior lecturer. His fundamental research interest lies in software engineering techniques with specific focus on formal methods and programming languages. Currently, his research is on extending and applying these techniques to the areas of ubiquitous and pervasive computing.