

Université de Paris

ED386 Sciences Mathématiques de Paris Centre — Laboratoire IRIF

Thèse de doctorat, spécialité informatique, sous la direction de Paul-André Melliès

Asynchronous and Relational Soundness Theorems for Concurrent Separation Logic

Léo Stefanesco

Jury

Adrien Guatto	MdC, Université de Paris	Examineur
Chung-Kil Hur	Assoc Prof, Seoul National University	Rapporteur
Damiano Mazza	DR, CNRS	Rapporteur
Daniela Petrisan	MdC, Université de Paris	Examinatrice
François Pottier	DR, INRIA	Examineur
Viktor Vafeiadis	Faculty, MPI-SWS	Examineur
Nobuko Yoshida	Prof, Imperial College	Examinatrice
Paul-André Melliès	DR, CNRS	Directeur
Lars Birkedal	Prof, Aarhus University	Invité

Présentée et soutenue publiquement le 12 novembre 2021.

Abstract – Asynchronous and Relational Soundness Theorems for Concurrent Separation Logic

The subject of this thesis is concurrent separation logic, a program logic for concurrent shared-memory languages. The relation between the proof of a program in a separation logic and the semantics of this program is expressed by the soundness theorem of this logic. This thesis introduces two soundness theorems. The first, the asynchronous soundness theorem, expresses the absence of data race in well specified programs in the language of template games in asynchronous graphs. The second part of this thesis extends the Iris concurrent separation logic with a relational soundness theorem which allows to establish simulations between a concrete program and an abstract model expressed as a state transition system. An application of this theorem is the proof of termination of concurrent programs under the assumption of a fair scheduler.

Keywords: concurrent separation logic, program logics, semantics, Iris, Coq

Résumé – Théorèmes de correction asynchrone et relationnelle de la logique de séparation concurrente

L'objet de cette thèse est la logique de séparation asynchrone, une logique de programme pour les langages de programmation concurrents et à mémoire partagée. Le lien entre une preuve d'un programme dans une logique de séparation concurrente et la sémantique de ce programme est exprimée par le théorème de correction de cette logique. Cette thèse introduit deux théorèmes de corrections. Le premier, le théorème de correction asynchrone, exprime dans le langage des jeux de gabarits sur des graphes asynchrones l'absence de courses des programmes bien spécifiés. L'autre étend la logique de séparation concurrente Iris avec un théorème de correction relationnel qui permet d'établir des simulations entre un programme concurrent concret et un modèle abstrait, formalisé comme un système de transition. Une application de ce théorème est la preuve de terminaison de programmes concurrent sous l'hypothèse d'un ordonnanceur équitable.

Mot-clés: logique de séparation concurrente, logique de programme, sémantique, Iris, Coq

Contents

Introduction	7
0.1 Program logics	7
0.1.1 Why use program logics?	7
0.1.2 Hoare logic	9
0.1.3 Program logics for concurrent programs	14
0.1.4 Limitations of Hoare logics	16
0.2 Separation logic	17
0.2.1 Ownership	18
0.2.2 Inference rules	19
0.2.3 Bunched implications	19
0.2.4 Towards concurrent separation logic	20
0.3 Concurrent separation logic	21
0.3.1 Resource invariants	21
0.3.2 Soundness theorem	22
0.3.3 Developments to concurrent separation logic	23
0.4 Organization of the thesis	26
I Asynchronous models of CSL	28
1 Proofs of soundness of CSL	29
1.1 Trace semantics of the language	29
1.1.1 Other proofs based on trace semantics	33
1.2 Operational semantics of the language	33
1.2.1 Step-indexed models of CSL	36
1.3 Other proofs soundness of CSL	36
1.3.1 Syntactic proofs	36
1.3.2 Other	37
1.4 Conclusion	38
2 Asynchronous soundness for CSL	39
2.1 Hoare logic as refinement systems	39
2.2 State transition systems	40

2.3	An imperative shared-memory concurrent language	41
2.4	Concurrent transition systems	43
2.4.1	Asynchronous graphs	43
2.4.2	Asynchronous machine models	45
2.4.3	Transition systems	50
2.4.4	Data-races	50
2.4.5	Polarized asynchronous transition systems	52
2.5	Concurrent separation logic	54
2.5.1	The logic	54
2.5.2	Semantic interpretation of CSL	57
2.5.3	The asynchronous soundness theorem	61
3	An asynchronous template game model of CSL	63
3.1	The double category $\mathbf{Cob}(\uplus)$ of games and cobordisms	63
3.1.1	Internal categories	63
3.1.2	Double categories	66
3.1.3	Polyads and internal J -opcategories	68
3.1.4	The double category $\mathbf{Cob}(\uplus)$ of games and cobordisms	70
3.2	Three internal J -opcategories: $\uplus_L, \uplus_S, \uplus_{\text{Sep}}$	72
3.2.1	The internal opcategories \uplus_L and \uplus_S for the code	72
3.2.2	The internal J -opcategory \uplus_{Sep} for the proofs	74
3.3	Parallel product	75
3.3.1	Plain internal functors	75
3.3.2	Acute spans of internal functors	76
3.3.3	Span monoidal internal J -opcategories	80
3.3.4	Parallel products of code and of proofs	82
3.4	Generalized sequential composition	84
3.5	Change of locks	87
3.5.1	Hiding	89
3.5.2	Critical sections	90
3.6	Sum of cobordisms	91
3.7	Interpretation of codes and proofs	91
3.7.1	Stateful and stateless interpretations of the code	92
3.7.2	Interactive and separated interpretations of the proofs	94
3.8	The asynchronous soundness theorem	95
3.8.1	Comparing the three interpretations	96
3.8.2	The asynchronous soundness theorem	98
3.9	Proof of the asynchronous soundness theorem	99
3.9.1	Well-formed cobordisms	99
3.9.2	Adhesivity of AsyncGraph	102
3.9.3	Preservation of well-formedness	103

3.9.4	Strict maps of cobordisms	108
3.9.5	Proof of 2-dimensional correctness	112
3.9.6	Proof of 1-dimensional correctness	114
II Relational soundness in Iris		117
4 Background on Iris		118
4.1	Overview of the approach	118
4.1.1	Iris predicates	119
4.1.2	The Iris standard weakest precondition	120
4.2	The Iris model	122
4.2.1	Ordered families of equivalences	123
4.2.2	RAs and Cameras	124
4.2.3	The Iris base logic	125
4.2.4	Interpretation of Iris predicates	126
4.2.5	Logical rules	127
4.3	High level logic	128
4.3.1	Combining cameras	128
4.3.2	Invariants	129
4.3.3	The standard Iris weakest precondition	130
5 A program logic to relate traces		131
5.1	Introduction	131
5.1.1	Relating STSs to programs in Iris	132
5.1.2	Simulations	133
5.1.3	Motivation	134
5.2	The trace weakest-precondition	135
5.2.1	Soundness of trwp	136
5.2.2	Proof of the soundness theorem	137
5.2.3	Infinite traces	139
5.3	Related works	140
6 Fair termination in Iris		141
6.1	Termination and fairness preserving simulation	142
6.1.1	Fairness models	142
6.1.2	Running example	143
6.1.3	A local criterion for the fair termination of models	144
6.1.4	Termination and fairness preserving refinements	146
6.1.5	The $\mathcal{L}\text{ive}$ construction	147

Contents

6.2	A logic for proving fairness and termination preserving refinements . .	151
6.2.1	Changes to the logic	151
6.2.2	Logical resources	153
6.2.3	Inference rules	154
6.2.4	Soundness theorem	156
6.2.5	Back to the example	157
6.3	Related works	158
	Conclusion	160
	Bibliography	164

Résumé en français

Les logiques de programme, dont la logique de séparation concurrente qui est l'objet de cette thèse, sont des *logiques axiomatiques*, en ce que les théorèmes qu'on prouve dans ces logiques ne sont que des moyens de prouver des propriétés *dans la logique ambiante* sur des programmes. Les propriétés des programmes sont déduites de la prouvabilité de certains théorèmes par le truchement des *théorèmes de correction* —aussi appelés théorèmes d'adéquation— de la logique. Cette thèse s'intéresse principalement à ces théorèmes de correction, pour deux variantes de la logique de séparation concurrente.

Après avoir introduit la logique de séparation concurrente, ce résumé décrit les deux parties largement indépendantes de ce manuscrit qui présentent chacun un théorème de correction : un théorème asynchrone pour la première partie, et un théorème relationnel pour la seconde.

La logique de séparation concurrente

La première logique de programme moderne, dans le sens où les logiques de programme contemporaines utilisent le même formalisme, est la logique de HOARE (1969), qui fut introduite par ce dernier dans son article fondateur *An axiomatic basis for computer programming*. Les prédicats principaux de la logique sont les *triplet de Hoare*¹

$$\{P\} C \{Q\}$$

où C est un programme écrit dans un langage de programmation impératif séquentiel et P et Q sont des prédicats sur l'état de la mémoire du programme, qu'on appelle respectivement la *pré-condition* et la *post-condition* du triplet de Hoare. Intuitivement, la signification du triplet de Hoare ci-dessus est que, si l'état initial de la mémoire avant que le programme C ne s'exécute satisfait la pré-condition P , alors d'une part le programme de recontera pas d'erreur qui stoppe son exécution, et d'autre part, si le programme termine, alors l'état final de la mémoire satisfait nécessairement la post-condition Q .

¹L'article original utilisait la notation duale $P \{C\} Q$ pour ce triplet, mais nous utilisons la notation devenue standard par la suite.

Les triplet de Hoare sont prouvés *dans la logique de Hoare* à l'aide de règles d'inférence, avec lesquelles sont construits des *arbres de dérivations*, aussi appelés *preuves*. Par exemple, la règle d'inférence qui correspond à la composition séquentielle de deux programmes est la suivante :

$$\frac{\{P\} C_1 \{Q\} \quad \{Q\} C_2 \{R\}}{\{P\} C_1; C_2 \{Q\}}$$

Comme nous l'avons expliqué précédemment, le lien entre l'existence d'un tel arbre de dérivation dont la conclusion est un triplet de Hoare $\{P\} C \{Q\}$ et la sémantique du programme C est exprimé par le théorème de correction de la logique de programme. Celui de la logique de Hoare (faible, car elle ne garantie pas la terminaison des programmes) est le suivant :

Theorème 1. *Soit un triplet de Hoare $\{P\} C \{Q\}$ qui soit prouvable à l'aide des règles d'inférence de la logique de hoare, et soit s un état de la mémoire qui satisfait la précondition P . Alors C exécuté à partir de l'état s ne rencontre pas d'erreur, et pour tout état mémoire s' tel que C exécuté à partir de s termine en s' , cet état s' satisfait la postcondition Q .*

Bien sûr, pour que le théorème ci-dessus soit complètement formel, il faut donner une sémantique formelle au langage de programmation, ce qui est aisé dans le cas d'un langage aussi simple que celui considéré par HOARE (1969).

La logique de Hoare a connu de nombreux développements ultérieurs qui l'ont étendu dans différentes directions telles que l'automatisation des preuves, l'augmentation de l'expressivité de la logique ou la gestion de langages de programmation avec des fonctionnalités plus avancées telles que la récursion.

L'une des limites de la logique de Hoare est qu'elle n'est pas adaptée à la gestion des structures de données contenant des pointeurs. En effet, il est difficile d'écrire une formule qui décrit, par exemple, un arbre binaire de recherche *sans partage*.

C'est pour résoudre ces problèmes que la *logique de séparation* fut inventée au tournant des années 2000. La différence fondamentale entre la logique de séparation et la logique de Hoare est que ses prédicates P, Q ne sont plus des formules de la logique du premier ordre avec des formules atomiques décrivant la mémoire, mais une logique sous-structurelle qui étent la logique du premier ordre avec une nouvelle forme de conjonction, la *conjonction séparante* $P * Q$. Un état mémoire s satisfait ce prédicat s'il est possible de découper cet état en deux parties disjointe, l'une satisfaisant P et l'autre satisfaisant Q .

Quelques années plus tard, O’Hearn découvrit que la logique de séparation était aussi adaptée à résoudre une autre limitation de la logique de Hoare, les programmes concurrents, et introduit la *logique de séparation concurrente*. En effet, les extensions de la logique de Hoare aux programmes concurrents, telles que la logique d’Owicki-Gries ou les logiques de Rely-Guarantee, doivent explicitement poser les hypothèses que l’environnement du programme qu’on prouve doit respecter.

Au contraire, la localité des prédicats de la logique de séparation permettent de mettre en place une discipline de *propriété* (en anglais : *ownership*) au niveau logique : si un prédicat mentionne une certaine ressource —par exemple de la mémoire— alors le programme peut utiliser cette ressource sans risque d’interférence de la part de l’environnement. En d’autres termes, la seule hypothèse sur l’environnement qui est faite est qu’il a aussi été prouvé avec la logique de séparation concurrente.

Un théorème de correction asynchrone pour la logique de séparation concurrente

La première partie de ce manuscrit étudie une version de la logique de séparation concurrente proche de la version originale proposée par O’Hearn. Le langage considéré étend celui considéré par Hoare avec une opération de composition parallèle $C_1 \parallel C_2$, ainsi que des verrous qui sont alloués statiquement avec la construction `resource r do C` et qui sont utilisés avec la construction `with r do C`. Ces constructions sont reflétées dans la forme des triplets de la logique de séparation concurrente : ils sont de la forme

$$\Gamma \vdash \{P\} C \{Q\}$$

où Γ est un contexte associant à chaque verrou disponible un invariant, qui est simplement un prédicat de la logique, du même type que les pré-conditions et que les post-conditions.

Dans un cadre concurrent, le théorème de correction de la logique doit aussi garantir qu’un programme qui a été prouvé ne fera pas d’*accès concurrent* (*data race* en anglais). Un accès concurrent est la situation où un deux threads d’un programme accèdent en même temps à la même case mémoire, et qu’au moins l’un de ces accès est une écriture.

Les deux parties du théorème de correction —le programme ne crash pas, le programme de fait pas d’accès concurrent— sont de nature largement différentes. En effet, la première propriété parle de chaque instruction exécutée par le programme, alors que la seconde parle des *paires* d’instructions exécutées par le programme. Dans notre formalisme, cette première condition est exprimée comme une propriété 1-dimensionnelle de la sémantique, alors que la seconde condition correspond à une propriété 2-dimensionnelle.

La grande majorité des preuves de correction de la logique de séparation concurrente, au contraire de notre approche, ramènent la seconde condition à la première en instrumentant la sémantique du langage de programmation pour provoquer un crash en cas d'accès concurrent.

Notre approche générale est la suivante. Étant donné un arbre de dérivation π d'un triplet de la logique de séparation concurrente

$$\Gamma \vdash \{P\} C \{Q\} \quad \begin{array}{c} \vdots \\ \pi \\ \vdots \end{array}$$

on interprète la preuve ainsi que le programme (de deux manières différentes) dans le même domaine sémantique de telle manière à ce ces trois intepretations soient liées de manière canoniques par des morphismes correspondant à des simulations :

$$\llbracket \pi \rrbracket_{\text{Sep}} \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_{\text{S}} \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_{\text{L}}.$$

Le domaine sémantique est basé sur la notion de *graphe asynchrone*, une notion de graphe dotés d'informations 2-dimensionnelles, les *tuiles* qui témoignent du fait que deux opérations sont "indépendantes".

Le théorème de correction peut alors s'exprimer à l'aide de notions de 1-fibration et de 2-fibration qui expriment que les arcs et les tuiles de la sémantique du programme C peuvent être relevées en des arcs et des tuiles de la sémantique du la preuve π , qui vivent dans un univers où, par construction, les erreurs et les accès concurrents sont impossibles.

Theorème 2 (Théorème de correction). *Étant donné une preuve π d'un triplet $\Gamma \vdash \{P\} C \{Q\}$, et avec les notations ci-dessus :*

1. \mathcal{S} est une 1-fibration du Code;
2. $\mathcal{L} \circ \mathcal{S}$ est une 2-fibration.

Un point important de cette partie est la méthodologie suivie pour définir les sémantiques des programmes et des preuves de la logique. Nous avons utilisés une version dualisée des *jeux de gabarits* de Melliès, basée sur des cospans au lieu de spans comme la version usuelle. L'idée la plus importante est que les interpretations vivents dans des doubles catégories induites par des *gabarits* (dans notre cas, des opcatégories internes à la catégorie des graphes asynchrones), et toutes les *opérations* sur ces interprétations sont induites par des *structures* sur ces gabarits.

Un théorème de correction relationnelle pour Iris

La seconde partie de cette thèse propose une nouvelle variante de la logique de séparation concurrente qui permet de construire des simulations entre un programme et un modèle abstrait qui peut soit être vu comme une spécification, soit comme une abstraction qui permet de prouver certaines propriétés du programme.

Cette logique est basée sur Iris, une logique de séparation concurrente moderne qui a plusieurs propriétés importantes pour ce travail. D’abord, contrairement aux logiques dont nous avons parlé jusqu’à présent, il n’y a pas de stratification entre les prédicats P, Q qui apparaissent en pré-condition par exemple, et les triplets de Hoare. Au contraire, les prédicats logiques de Iris sont suffisamment expressifs pour pouvoir définir les triplets de Hoare. Cela signifie que, pour définir une nouvelle logique de programme, il suffit de définir une nouvelle notion de triplet dans la “logique de base” de Iris, et c’est ce que nous avons fait ici.

Une autre différence importante est qu’Iris est une logique modale avec un opérateur de *point fixe gardé*. En particulier, Iris est doté d’une modalité \triangleright (“later”) et son modèle sémantique est basé sur le comptage de pas (*step-indexing*). Cela signifie en particulier que la logique n’est capable que d’exprimer des propriétés de sûreté, par opposition à des propriétés de vivacité telles que la terminaison.

Concrètement, nous proposons une logique de programme paramétrisée par un certain modèle et qui admet une unique nouvelle règle de raisonnement par rapport à la logique de programme standard d’Iris qui permet de “faire un pas” dans le modèle. Le théorème de correction de la logique permet de prouver qu’une certaine relation ξ est contenue dans une simulation entre le programme et le modèle abstrait si l’on a prouvé un certain triplet pour le programme en question.

Une application directe de ce théorème est de montrer qu’un programme distribué implémente l’algorithme de consensus distribué Paxos de Lamport, au sens où le programme et l’algorithme envoient les mêmes messages, modulo encodage. Un autre exemple que nous avons mis en œuvre est la preuve qu’un algorithme distribué converge nécessairement (*eventually consistent*), sous certaines hypothèses de vivacité du réseau et du programme.

Enfin, la dernière application de ce théorème de correction relationnelle est une logique permettant de prouver qu’un programme concurrent termine si l’ordonnanceur est équitable (*fair scheduler*). La limitation due au comptage de pas ne pose pas de problème car la technique utilisée est de construire une relation entre le programme et un modèle abstrait qui admet un ordre bien fondé satisfaisant quelques propriétés de décroissance. La relation préserve l’équité du programme vers le modèle, et la terminaison du modèle vers le programme, ce qui permet de déduire la terminaison des traces équitables du

programme de la terminaison des traces équitables du modèle abstrait, qui découle elle-même de l'existence d'un ordre bien fondé sur ses états.

Introduction

Proving the correctness of programs has been the object of scientific interest well before the use of computer software became widespread and the social need for proving their correctness arose, especially in critical contexts, such as cyber-physical systems where computer programs control physical actuators. Formally, the objective is to prove properties about the semantics of some program: most typically that it does not crash and that if it terminates, its final state satisfies a certain property which expresses that it did what was expected.

0.1 Program logics

The first paper on this topic is the now famous paper *Checking a large routine* by Turing (1949). Most current program logics are instances of what is now called *Hoare logic*, which was first developed by Hoare (1969), building on previous works by Floyd (1967). In every case, the method follows the structure of the program, be it in the form of a flowchart for Turing and Floyd, or of its syntax tree for Hoare.

Logics should provide reasoning rules which are as close as possible to the methods programmers use to reason about the correctness of their programs, such as using loop invariants to reason about the correctness of loops, and variants to prove their termination.

Related to the idea that proofs of programs carried out in a program logic should follow the structure of the program, these proofs should be *modular*: proofs of correctness of several part of a program should be carried out in isolation, and then combined into a proof of the whole program, mirroring the modularity mechanisms of the programming language.

0.1.1 Why use program logics?

Some researchers, such as L. Lamport (1994) have argued that instead of using specialized logics, one should prove programs using “usual” mathematics, or at least in a setting

as closed to it as possible. For example, the logic of his TLA+ tool is a set theory augmented with temporal modalities from LTL to express properties of execution steps of programs.

However, it seems to us that TLA+'s purpose is not to prove programs, instead it is to prove algorithms, or something intermediate between an algorithm and a program implementing it.

Proving programs and proving algorithms are widely different tasks: programming languages have intricate semantics, and proving that a piece of code is safe can be a subtle task, even if it corresponds to a simple action in the algorithm the program implements. For example, to prove that dereferencing a raw pointer in Rust is safe, which is an unsafe operation in general, the programmer needs to consider which other part of the program is using this pointer, whether the target of the pointer is valid, and what are the other pointers or references which may alias with it. Logics such as Rustbelt (Jung, Jourdan, et al., 2017) provide facilities to carry such arguments in a specialized logic.

Nowadays, realistic languages are formalized using an operational semantics (for example CompCert for C (Leroy, 2009), or JSCert for Javascript (Bodin et al., 2014)), because state transition systems are simple first order objects which are easy to define and manipulate, especially in a proof assistant. Operational semantics, however, are not well suited in practice to reason directly about programs, even programs of moderate size, because the number of cases to consider quickly explodes. This is specially true when, as it is common, the semantics leaves evaluation orders unspecified, leading to non-determinism in the semantics. This is even more true for concurrent programs, where one needs to consider all interleavings. In a way, the logic is required to recover the compositionality which operational semantics lack compared to denotational semantics.

Finally, most programs are composed of many algorithms: for instance, a simple concurrent program would contain a lock implementing the ticket-lock algorithm (Mellor-Crummey and Scott, 1991), which would prevent concurrent access to a skip-list for example. In the same way that the lock module can be used to protect any data structure, it should be given a specification and a proof of correctness that can be used in many contexts. Conversely, the skip list should be proved and specified without mention of locking or concurrency, as it is a sequential data structure.

One avenue we will explore in Part II is to cleanly separate the proof of the algorithm from the proof of the program module implementing it. For example, we consider an abstract model of the Paxos consensus algorithm taken from a TLA+ specification, and prove in the Iris logic that it simulates a concrete implementation. Then, the

consensus property of the program is deduced from the consensus property of the abstract algorithm.

0.1.2 Hoare logic

We begin by describing Hoare logic (Hoare, 1969), which is the core of most program logics for imperative programming languages. Its purpose is to prove the safety and the functional correctness of programs written in a simple WHILE language with the following grammar: It has Boolean expressions:

$$B ::= \text{true} \mid \text{false} \mid B \wedge B' \mid B \vee B' \mid \neg B \mid E = E'$$

arithmetic expressions:

$$E ::= 0 \mid 1 \mid \dots \mid x \mid E + E' \mid E * E'$$

and commands:

$$C ::= x := E \mid C; C' \mid \text{if } B \text{ then } C_1 \text{ else } C_2 \mid \text{while } B \text{ do } C \mid \text{skip}$$

This language is extremely simple: it does not have functions nor local variables, but it captures the essence of sequential imperative programming. For example, the following program `ediv` computes the Euclidian division of two integers. More precisely, it computes the division of the number stored in the global variable `x` by the one in global variable `y`, and stores the quotient in `q` and the remainder in `r`.

```

r := x;
q := 0;
while r ≥ y do
  (r := r - y;
   q := q + 1)

```

Of course, this piece of code only satisfies the informal specification above when `x` and `y` are positive.

Hoare logic specifies programs C using *Hoare triples*, which are predicates of the form:

$$\{P\} C \{Q\}$$

where P and Q are formulas which represent logical predicates over the state of the memory, written in a language which is part of Hoare logic. Intuitively, this states that any execution of the program C which begins with an initial memory state ξ_0 which

satisfies the *precondition* P , does not crash; moreover, if this execution of the program C terminates, then its final state s_f satisfies the *postcondition* Q .

In the version of Hoare logic we consider, P and Q are first-order formulas with atomic propositions $x = v$ asserting that, in the current memory state, the program variable x is well defined and contains the value v . The meaning of such a predicate is given by the satisfiability relation $s \models P$, where the memory state s is a finite partial map from variables to values; it is defined as follows

$$\begin{aligned} s \models x = v &\iff s(x) = v, \\ s \models P \wedge Q &\iff s \models P \wedge s \models Q, \end{aligned}$$

and similarly for every first order connectives.

In this thesis, we consider Hoare logics for *weak correctness*, which does not guarantee that programs terminate for initial states satisfying the precondition P . Such logics—called *strong*—do exist and, for sequential programs, are fairly similar to weak logics.

For instance, the `ediv` program above, which implements Euclidian division, would be specified using the following Hoare triple:

$$\{x \geq 0 \wedge y > 0\} \text{ ediv } \{x = y \cdot q + r \wedge 0 \leq r < y\}$$

To be accepted, a Hoare triple such as the one above needs to be the conclusion of a proof tree build using the inference rules of the logic. The usual inference rules for Hoare logic are presented in Figure 0.1. Each syntactic construction of the programming language has an associated rule. In this version of Hoare logic, the only rule which does not correspond to the syntax of the program is the *rule of consequence* `CONSEQ`. This rule, which allows strengthening the precondition and weakening the postcondition, is essential: it allows the prover to “massage” the pre and postconditions so that the syntax directed rules can be applied; for instance, the premise (and the conclusion) of the `WHILE` rule must be of a particular shape to be used. The premises of `CONSEQ` refer to the entailment relation on predicates, which is defined as follow:

$$P \Rightarrow Q \iff \forall s, s \models P \text{ implies } s \models Q,$$

or, equivalently, it holds iff the Hoare logic predicate $P \Rightarrow Q$ is a tautology. Importantly, entailment should be seen as a *semantic* construct, unlike the rest of Hoare logic which is a syntactic system.

Let us focus on the rule for `while`. To use it, the prover needs to find a Hoare logic predicate I which is called the *loop invariant*. The premise states that if, before the execution of an iteration, the invariant I and the guard B of the loop both hold, then the invariant is still true after that iteration.

$$\begin{array}{c}
 \frac{}{\{P\} \text{ skip } \{P\}} \text{SKIP} \qquad \frac{}{\{Q[x := a]\} x := E \{Q\}} \text{AFF} \\
 \\
 \frac{\{P\} C_1 \{Q\} \quad \{Q\} C_2 \{R\}}{\{P\} C_1; C_2 \{R\}} \text{SEQ} \qquad \frac{\{P \wedge B\} C_1 \{Q\} \quad \{P \wedge \neg B\} C_2 \{Q\}}{\{P\} \text{ if } B \text{ then } C_1 \text{ else } C_2 \{Q\}} \text{IF} \\
 \\
 \frac{\{I \wedge B\} C \{I\}}{\{I\} \text{ while } B \text{ do } C \{I \wedge \neg B\}} \text{WHILE} \qquad \frac{P \Rightarrow P' \quad \{P'\} C \{Q'\} \quad Q' \Rightarrow Q}{\{P\} C \{Q\}} \text{CONSEQ}
 \end{array}$$

Figure 0.1: Inference rules for Hoare logic

To see how this works out in practice, we prove that the program `ediv` does satisfy the specification we have given above. Instead of displaying the proof as a tree, because of the amount of redundancy which is contained in Hoare logic proof trees, we annotate the program with predicates enclosed in curly braces, and we denote the use of the `CONSEQ` rule with \Rightarrow .

```

{ x ≥ 0 ∧ y > 0 } ⇒
{ x = y · 0 + x ∧ x ≥ 0 ∧ y > 0 }
r := x;
{ x = y · 0 + r ∧ r ≥ 0 ∧ y > 0 }
q := 0;
{ x = y · q + r ∧ r ≥ 0 ∧ y > 0 }
while r ≥ y do
(
  { x = y · q + r ∧ r ≥ 0 ∧ y > 0 ∧ r ≥ y } ⇒
  { x = y · (q+1) + (r-y) ∧ r-y ≥ 0 ∧ y > 0 ∧ r-y ≥ 0 }
  r := r - y;
  { x = y · (q+1) + r ∧ r ≥ 0 ∧ y > 0 ∧ r ≥ 0 }
  q := q + 1
  { x = y · q + r ∧ r ≥ 0 ∧ y > 0 ∧ r ≥ 0 }
)
{ x = y · q + r ∧ r ≥ 0 ∧ y > 0 ∧ r < y } ⇒
{ x = y · q + r ∧ r < y }

```

This example demonstrates that proving the entailment relation used in the consequence rule requires general mathematical reasoning, in this case arithmetic reasoning. This is why considering it as a semantic notion instead of a syntactic one simplifies a lot the definition of the logic.

Weakest precondition

The rule `WHILE` has a special status in Hoare logic because it is the only rule for which the pre and postconditions in the premise cannot be computed from those in the conclusion. To see this, it is useful to consider another presentation of Hoare logic using *weakest preconditions*: the weakest precondition $\text{wp } C \{Q\}$ of a program C and a postcondition Q is the weakest —according to the entailment order— precondition such that the Hoare triple $\{\text{wp } C \{Q\}\} C \{Q\}$ holds. We can recover Hoare triples from weakest preconditions with:

$$\{P\} C \{Q\} \quad := \quad P \Rightarrow \text{wp } C \{Q\}$$

There are logics, such as the Iris logic, which we will present in Part II, which takes weakest preconditions as their primitive notion, from which Hoare triples are deduced.

One advantage of weakest preconditions is that, except for while loops as alluded to above, they can be computed by induction on the syntax of the program:

$$\begin{aligned} \text{wp skip } \{Q\} &:= Q \\ \text{wp } x := E \{Q\} &:= Q[x := E] \\ \text{wp } C_1; C_2 \{Q\} &:= \text{wp } C_1 \{\text{wp } C_2 \{Q\}\} \\ \text{wp if } B \text{ then } C_1 \text{ else } C_2 \{Q\} &:= (B \wedge \text{wp } C_1 \{Q\}) \vee (\neg B \wedge \text{wp } C_2 \{Q\}) \end{aligned}$$

Therefore, to check that a program (without loop) satisfies a Hoare triple $\{P\} C \{Q\}$, it suffices to prove the entailment

$$P \Rightarrow \text{wp } C \{Q\}$$

which, in practice, can be fairly well automated using automatic theorem provers.

A common strategy for loops is to ask the programmer to give a loop invariant, which can then be used to define the weakest precondition:

$$\text{wp while } B\{I\} \text{ do } C \{Q\} \quad := \quad I$$

assuming the side conditions $B \wedge I \Rightarrow \text{wp } C \{I\}$ and $\neg B \wedge I \Rightarrow Q$ hold. Examples (among many) of tools which use refinements of the method just outlined are Why3 (Filliâtre and Paskevich, 2013) for ML or Dafny (Leino, 2010) for C.

Soundness theorem

We have explained informally the meaning of a (weak) Hoare triple, in terms of executions of the programs. The corresponding formal statement is called the *soundness*

theorem of the logic. To be able to state it, we need to give the language a formal semantics. For example, we can consider its big step semantics $C, s \Downarrow s'$ which states that the program C started in the initial state s terminates, and its final state is s' , and $C, s \Downarrow \perp$ if the program crashes.

We can now state the soundness theorem:

Theorem 0.1.1. *Given a derivation tree π of a Hoare triple $\{P\}C\{Q\}$, and given a machine state s such that $s \models P$, the program is safe: $C, s \not\Downarrow \perp$ and if $C, s \Downarrow s'$ then $s' \models Q$*

Proof. Straightforward induction on the structure of the derivation tree π . □

The reason why the soundness theorem is important is that Hoare logic is an “axiomatic logic”, in the terminology of Girard, in that it is only useful as a tool to prove properties of the program expressed in the ambient logic; and this connection is given by the soundness theorem. This point will more salient in Part II because the distance between the Iris logic and the ambient logic is greater, as Iris is a step-indexed modal logic.

Semantic triples and completeness

We can rephrase the soundness theorem above by defining the *semantic validity* of a Hoare triple as:

$$\models \{P\}C\{Q\} \quad := \quad \forall s, s \models P \Rightarrow C, s \not\Downarrow \perp \quad \wedge \quad \forall s', C, s \Downarrow s' \Rightarrow s' \models Q$$

Then, writing $\vdash \{P\}C\{Q\}$ for the existence of a derivation tree for that Hoare triple, the soundness theorem can be expressed as

Theorem 0.1.2. $\vdash \{P\}C\{Q\}$ *implies* $\models \{P\}C\{Q\}$.

The usual way to prove this theorem amounts to proving that each inference rule of the logic has a corresponding lemma where we replace syntactic triples with semantic triples. Some logics such as Iris completely forgo syntactic triples and only have semantics triples with lemmas about them which correspond to the inference lemmas.

Despite its simplicity, Hoare logic is *complete*, in the sense that the converse of the second formulation of the soundness theorem holds:

Theorem 0.1.3. $\models \{P\}C\{Q\}$ *implies* $\vdash \{P\}C\{Q\}$.

In the sequel we will only consider the soundness of program logics.

0.1.3 Program logics for concurrent programs

As we have discussed, Hoare logic is a program logic well suited for a simple sequential programming language such as WHILE. Intuitively, the logic ensures that at each step of an execution of the program, it is in a state which guarantees that the execution is safe.

In a concurrent setting this is not sufficient anymore, because a non-trivial thread cannot be safe regardless of how the other threads executing concurrently are behaving. Consider indeed a thread which modifies randomly the values of every variables, which are shared between the threads. In these conditions, during the proof of any other thread, it does not make sense to have, say, a loop invariant which depends on the value of a variable since that other thread can invalidate it at any moment.

The fundamental difference between logics for sequential and concurrent languages is that, if we are to prove each thread separately, the proof needs to pose constraints on the behavior of its environment, which is the abstraction of all threads possibly executing concurrently.

A simple imperative shared memory language

To make things more concrete, let us consider a simple concurrent language with atomic sections, which adds the following commands to the WHILE language we considered in the previous section:

$$C ::= \dots \mid C_1 \parallel C_2 \mid \text{atomic } C$$

The command $C_1 \parallel C_2$ executes the two commands C_1 and C_2 in parallel and waits for both of them to finish, and the atomic section `atomic C` creates a critical section, in that at most one such critical section is executing in the complete program at any one time. Here is a very simple program taking advantage of these two new constructs:

```
atomic (y:=x; x:=y+1) || atomic (z:=x; x:=z+1)
```

If the initial values of x is 0, then the the final value will be 2. Without the critical sections, both threads could have read the value 0 from x , and then they both would have written 1 into x . (There would have been a data race as well, but we will talk about this later.)

Rely-guarantee

There have been several generalizations of Hoare logic to concurrent settings. An early milestone was proposed by Owicki and Gries (1978), who introduced an extension of Hoare logic to shared-memory concurrent programs, which now bears their name, which is compositional, in that each thread is proved separately. The central rule is the following

$$\frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\} \quad \text{the two proofs are not interfering}}{\{P_1 \wedge P_2\} C_1 \parallel C_2 \{Q_1 \wedge Q_2\}}$$

The issue however is that the non-interference side-condition is quite complicated and unpractical.

The logic which we consider more carefully is the *Rely-Guarantee* logic (RG), which was introduced by Jones (1983). The two key ideas are, first, to make explicit the assumption about the interferences of the environment on which the proof of the program *relies*, and, second, that conversely, the program should provide *guarantees* to its environment. Formally, a specification in RG is a 5-tuple which we write

$$R; G \vdash \{P\} C \{Q\}$$

where, as before, P and Q are predicates, and C is the program; the two new items R and G are *relations* on memory states (or a syntactic representation thereof). Intuitively, G contains the set of atomic modifications of the memory which the program is allowed to perform; dually, R represents the set of atomic modifications which the environment is allowed to perform while C is executing.

Since the program needs to be able to rely on its precondition P to hold at the beginning of its execution, the precondition P needs to be stable under the rely relations, in that $s \models P \wedge (s, s') \in R \Rightarrow s' \models P$. Otherwise, the proof of the program would have to take into account that the environment can invalidate P at any time.

The inference rule for the parallel product ensures that the two specifications are compatible, in that the *rely* relation of each thread follows from the *guarantee* of the other:

$$\frac{R_1; G_1 \vdash \{P_1\} C_1 \{Q_1\} \quad R_2; G_2 \vdash \{P_2\} C_2 \{Q_2\} \quad G_1 \Rightarrow R_2 \quad G_2 \Rightarrow R_1}{R_1 \cap R_2; G_1 \cup G_2 \vdash \{P_1 \wedge P_2\} C_1 \parallel C_2 \{Q\}}$$

for some Q computed from the Q_i 's and the R_i 's. The improvement over the Owicki-Gries logic is that the side conditions no longer refer to the proofs of C_1 and of C_2 , but only the specifications of C_1 and of C_2 .

The other rules ensure that each atomic step taken by a program are part of its guarantee relation, and that all preconditions are stable under the rely relation.

In conclusion, the fundamental difference between sequential Hoare logics and concurrent Hoare logics is that, for the logic to be compositional, the predicates it manipulates must be secure against interferences from any *well-specified* environment.

0.1.4 Limitations of Hoare logics

Despite the power of Hoare logic and its variant —they are complete— they have difficulties with compositionality. One instance of this is related to the *frame rule*, an important rule for compositionality of proofs:

$$\frac{\{P\} C \{Q\} \quad R \text{ does not contain variables written by } C}{\{P \wedge R\} C \{Q \wedge R\}}$$

It allows for proving a piece of code C using only the needed pre/post-condition P and Q , and the using it in any context where the precondition implies $P \wedge R$. A typical example where this is useful is in a program with many variables which contains a loop writing to a single variable x . Without such a rule, the loop invariant needs to state that every property about the other variables is preserved, which can be tiresome. More fundamentally, this allow a single piece of code to be proved once and used in several different contexts.

This rule does hold in “classical” Hoare logic, but tends to fail in extensions. In the context of concurrent programming languages, the Hoare triples of RG do not admit a suitable frame rule, making the rely and guarantee relations global.

The second programming language feature which invalidates Hoare logic’s frame rule, and which motivated separation logic originally, is the addition of pointers. The core of the problem is that, unlike for program variables, interference at the level of pointers (addresses in the memory heap) is a semantic property, which cannot be captured using a syntactic side condition like the frame rule presented above.

We add pointers to our simple (sequential) WHILE language. The memory states are now pairs $\mathfrak{s} = (s, h)$ of a stack s as before, and a heap h a finite partial mapping from addresses to values. We identify addresses with natural numbers, and we assume that addresses are values, allowing programs to store pointers. We write $[E]$ to dereference the address represented by the expression E , giving the following commands:

$$C ::= \dots \mid [E] := E' \mid x := [E]$$

It is natural to add an atomic predicate to be able to specify the value contained at some memory location, which mirrors the predicates for the stack:

$$\ell \mapsto v$$

specifies that the address ℓ is allocated and that it contains the value v . However, it is easy to see that the frame rule given above does not hold, as the following could be derived:

$$\frac{\{\ell \mapsto 1\} [\ell] := 3 \{\ell \mapsto 3\}}{\{\ell \mapsto 1 \wedge \ell \mapsto 1\} [\ell] := 3 \{\ell \mapsto 3 \wedge \ell \mapsto 1\}}$$

Of course this example can be ruled out using a version of the side condition mentioned above, but the address ℓ could be stored in a variable, or in the heap.

The fact that a program C admits a Hoare triple $\{P\} C \{Q\}$ implies, informally, that the memory it manipulates is described by the predicate P . Therefore, for a frame rule to exist, it would be sufficient to be able to state that the frame predicate R describes a part of the memory which is distinct from the memory described by P . The new tools provided by separation logic will allow us to do just that.

Another way this disjointness property is useful is in specifying programs which manipulate data structures which contain pointers. For example, the obvious program which mutably concatenates two linked lists produces a linked list only if the two lists passed as arguments do not *alias*, that is, only if they span two disjoint areas of the memory.

0.2 Separation logic

Separation logic was introduced around 2000 by Reynolds, O’Hearn, Ishtiaq and Yang to improve reasoning about pointers in Hoare logics. It acquired its name in a survey paper by Reynolds (2002), following works by Ishtiaq and P. W. O’Hearn (2001), Reynolds (2000), P. W. O’Hearn, Reynolds, and Yang (2001).

As we explained above, the main improvement over Hoare logic lies in how predicates can control aliasing. This is achieved using a new logical connector, *the separating conjunction*, written $P * Q$. It states that the predicates P and Q hold in two disjoint regions of the memory:

$$s \models P * Q \quad \Leftrightarrow \quad \exists s_1, s_2, s = s_1 \uplus s_2 \wedge s_1 \models P \wedge s_2 \models Q$$

This allows to state *at a semantic level* that two predicates are “independent”. More fundamentally, predicates now behave like resources: they are not duplicable in general, in that $P \Rightarrow P * P$ does not hold.

There are two flavors of separation logics, depending on whether the logic allows weakening: If it does, the logic is called *affine* (or intuitionistic), otherwise it is *linear* (or classical). Semantically, in logics of the first kind, the set of states which satisfy any predicate P is upward closed:

$$\forall s, s', s \models P \wedge s \subseteq s' \implies s' \models P$$

For example, the interpretation of the “points-to” predicates which specify the value at some address ℓ of the heap in an affine logic is defined in the same way as for the Hoare logic above:

$$s \models \ell \mapsto v \iff (\ell, v) \in s$$

In such a logic, the unit element of the separating conjunction $*$ is the predicate \top which holds for any memory state. In a separation logic which is linear, however, $\ell \mapsto v$ only holds for the corresponding singleton heap:

$$(s, h) \models \ell \mapsto v \iff h = [\ell \mapsto v]$$

In such a logic, the predicate \top is not the unit element of the separating conjunction. Actually, the mapping $P \mapsto P * \top$ defines a translation from a linear to an affine separation logic. The unit element is the atomic predicate **emp**, which holds exactly when the heap is empty. In the remainder of this introduction and in Part I we consider a linear separation logic. Part II uses the Iris separation logic, which is affine.

0.2.1 Ownership

A predicate P can be seen as representing the *ownership* of an area of the memory which satisfies it. As we will see later, this (purely logical) ownership can be transferred among threads when acquiring a lock. Another example of ownership transfer would be a program which sends a pointer in a channel, giving away its ownership of the memory it points to; the program which receives from the other end of the channel would be the new owner of this memory.

The separating conjunction $P * Q$ then represents the ownership of P and the ownership of Q . The usual conjunction $P \wedge Q$, on the other hand, states the ownership of memory which satisfies both P and Q . This is intuitively similar to intersection types, where a term has several types, and can be used as either.

The ownership interpretation of a separation logic Hoare triple $\{P\} C \{Q\}$ is that the program C needs to be given ownership of the resource P to run, and after it has finished, it gives back to its caller the ownership of the resource Q .

0.2.2 Inference rules

Separation logic allows us to state a frame rule which holds in the presence of pointers:

$$\frac{\{P\} C \{Q\} \quad R \text{ does not contain variables written by } C}{\{P * R\} C \{Q * R\}}$$

The reason why there still is a side condition about written program variables is that in the presentation we have given, predicates do not have a notion of ownership of variables. This means that non-interference at the level of program variables needs to be handled in side conditions.

This is why most modern separation logic handle variables like ML: a variable is immutable, but can contain a location whose content can change, as we will see in Part II. In other words, instead of memory states being a stack and a heap, they are a heap and an environment. In the remainder of this introductory chapter, we omit those side conditions, since they can be dealt with.

The remainder of the rules are similar to the ones for Hoare logic. For example, the rule for reading from the heap is the following:

$$\frac{}{\{x = v' * \ell \mapsto v\} x := [\ell] \{x = v * \ell \mapsto v\}} \quad (0.1)$$

The resource $\ell \mapsto v$ is given back in the postcondition because the command does not consume this resource. In contrast, the `dispose` operation which frees the memory has the following specification:

$$\frac{}{\{\ell \mapsto v\} \text{dispose}(\ell) \{\mathbf{emp}\}}$$

The linearity of the logic allows the user of the logic to ensure that all the memory has been freed using the `dispose` operation at the end of the execution of the program by choosing `emp` as the postcondition of the program.

0.2.3 Bunched implications

The predicates of separation logic are an instance of *the logic of bunched implications* (BI), a substructural logic introduced by P. W. O’Hearn and D. J. Pym (1999). It features two types of conjunctions, additive conjunction \wedge and multiplicative conjunction $*$, each with a right adjoint \rightarrow and \multimap (read “magic wand”), respectively.

A sequent in BI is of the form $\Gamma \vdash P$, where the context Γ is a tree of the form:

$$\Gamma ::= P \mid \Gamma_1, \Gamma_2 \mid \Gamma_1; \Gamma_2$$

The arborescent nature of the contexts gives its name to the logic. The internal nodes of the tree denoted with semicolons correspond to \wedge , and the nodes denoted with commas correspond to $*$, as demonstrated by the introduction rules of the implications:

$$\frac{\Gamma, P \vdash Q}{\Gamma \vdash P \multimap Q} \quad \frac{\Gamma; P \vdash Q}{\Gamma \vdash P \Rightarrow Q}$$

The main difference is that only $;$ admits structural rules. To express those rules, we denote by $\Gamma(\Delta)$ a context where Δ appears as a subtree, and where Γ is the “rest” of the tree (more formally, Γ is a context in the other meaning of the word).

$$\frac{\Gamma(\Delta) \vdash P}{\Gamma(\Delta; \Delta) \vdash P} \quad \frac{\Gamma(\Delta; \Delta) \vdash P}{\Gamma(\Delta) \vdash P}$$

As expected, there is a multiplicative unit I and an additive unit \top , which correspond to **emp** and \top in the separation logic we have presented.

Models of BI

A natural question to ask is how BI relates to linear logic, another substructural logic. One answer is to look at models of BI. A model of BI is a *doubly closed category*: a Cartesian closed category $(\mathcal{C}, \times, \Rightarrow, \top)$ together with a monoidal closed structure on the *same* category $(\mathcal{C}, *, \multimap, I)$. The semantics of a context Γ is given by interpreting $,$ using the monoidal product $*$, and interpreting $;$ using the Cartesian product \times . In particular, we have:

$$\mathcal{C}(I, P \multimap Q) \cong \mathcal{C}(P, Q) \cong \mathcal{C}(\top, P \Rightarrow Q)$$

Many separation logics have a magic wand; it is the basic notion of implication which is used in Iris for example.

One class of examples of models of BI are functor categories over monoidal categories, where the monoidal product is the Day tensor product. There is also a game model by McCusker and D. Pym (2007).

0.2.4 Towards concurrent separation logic

The great aptitude of separation logic for controlling interferences make it well-suited for the concurrent setting. Intuitively, the Frame rule implies the following rule for the parallel product:

$$\frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}}$$

Indeed, according to the frame rule, each program C_i leaves untouched the memory described by P_{3-i} , so it is reasonable to expect that C_1 and C_2 do not to interfere at all, which means that this rule is semantically valid.

Separation logic can be seen as an implicit instance of the Rely-Guarantee method: a program which can be specified with a separation logic Hoare triple $\{P\}C\{Q\}$ guarantees to only access memory it owns, which, at the start of its execution, is described by its precondition P , and it relies on the environment not to access the memory owned by C .

The purpose of concurrent separation logic is to provide means to prove concurrent programs which interfere in a controlled way.

0.3 Concurrent separation logic

Concurrent separation logic (CSL) was introduced by P. W. O’Hearn (2004) and Brookes (2004) as an extension of separation logic to a concurrent shared memory language with locks.

The language which was considered in the early CSL papers was slightly different from the ones we used in Section 0.1.3: Instead of atomic sections, programs can declare locks, and use critical sections associated to each of these locks. Its grammar replaces the command `atomic C` with:

$$C ::= \dots \mid \text{resource } r \text{ do } C \mid \text{with } r \text{ do } C$$

The command `resource r do C` declares a new lock r , which is available inside C ; `with r do C` acquires the lock r , runs C , and then unlocks r . Dynamically, two critical sections for the same lock are guaranteed not to run at the same time.

0.3.1 Resource invariants

The idea of CSL is to associate each lock with a predicate of CSL, which describes the resource protected by the lock while it is available (unlocked). In CSL, Hoare triple are parameterized with a context Γ , making them formally 4-tuples:

$$\Gamma \vdash \{P\} C \{Q\}$$

where P and Q are predicate of the same type as (sequential) separation logic, and Γ is a context

$$\Gamma ::= \emptyset \mid \Gamma, r : I$$

associating predicates I with locks r . The predicate I is called the *lock invariant* of the lock r . We consider contexts up to permutations and we assume the lock names are distinct.

The general idea is that the only way for threads to communicate is through locks: for example, suppose the lock r is associated with the resource $\exists v, \ell \mapsto v$, which expresses the ownership of the memory location ℓ . Suppose ℓ contains initially 0, and that one thread wants to signal another thread. It can write the value 1 to the memory location ℓ in a critical section:

$$\text{with } r \text{ do } [\ell] := 1$$

The other thread could be reading ℓ in a loop, waiting to read 1 before exiting the loop:

$$\begin{aligned} &x := 0; \\ &\text{while } x = 0 \text{ do with } r \text{ do } x := [\ell] \end{aligned}$$

At the logical level, the program acquires the resources protected by the lock when it enters the critical section, and it must release the resource at the end of the critical section.

$$\frac{\Gamma \vdash \{P * I\} C \{Q * I\}}{\Gamma, r : I \vdash \{P\} \text{with } r \text{ do } C \{Q\}} \quad \frac{\Gamma, r : I \vdash \{P\} C \{Q\}}{\Gamma \vdash \{P * I\} \text{resource } r \text{ do } C \{Q * I\}}$$

To declare a new lock r with invariant I , the program needs to let go of the resource I and put it “in the lock”; after the end of the command C the program gets back the resource I . Finally, the rule for the parallel product is the same as the one sketched at the end of the previous section, with the addition of the context which is shared between the two threads.

$$\frac{\Gamma \vdash \{P_1\} C_1 \{Q_1\} \quad \Gamma \vdash \{P_2\} C_2 \{Q_2\}}{\Gamma \vdash \{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}}$$

0.3.2 Soundness theorem

The soundness theorem of the CSL states, as for Hoare logic, that every execution of a well-specified program is safe, and that if it terminates, the final state satisfies the postcondition. If the logic uses linear predicates, one can allow the initial state of said execution to contain a sub-memory state which satisfies the precondition.

Theorem 0.3.1. *Suppose $\emptyset \vdash \{P\} C \{Q\}$, and suppose that $\varepsilon \vDash P * \top$, then any execution of C which starts from ε is safe, and its final state, if it exists, satisfies $Q * \top$.*

There is another property which is important to require for a concurrent shared-memory program to be considered safe: it must be data race free. Informally, an execution of a program contains a *data race* if there are two instructions which are executed at the same time and such that one writes to the location the other one is accessing. Therefore, the soundness theorem of CSL contains a second half:

Theorem 0.3.2. *Suppose $\emptyset \vdash \{P\} C \{Q\}$, and suppose that $\varepsilon \models P * \top$, then any execution of C which starts from ε is data-race free.*

This second half of the theorem is different in nature to the first because it requires reasoning about *pairs* of instructions of the same execution.

The proof of soundness of CSL is significantly more difficult than its sequential counterpart, as the meaning of a CSL Hoare triple needs to state formally what is the “rely” of the environment.

Most proofs of soundness of CSL add a rule in the semantics of the programming language which makes the program crash when a data race occurs, this allows one to deduce the second part of the theorem from the first.

Part I of this thesis, on the contrary, develops a proof of soundness of CSL which handles directly this aspect by giving an asynchronous semantics to both the program and the derivation trees of CSL, allowing us to talk about concurrent executions of instructions.

0.3.3 Developments to concurrent separation logic

There have been many advances in the field of concurrent separation logics since the seminal papers of O’Hearn and Brookes. The language we considered above is indeed very simple, and CSL was extended to support richer languages with higher-order functions, higher order store (mutable references which can contain closures), dynamic creation of threads, etc.

Another area of improvement is the concurrency primitives that the programming languages support. The concurrency primitive of the WHILE language above is statically declared locks. Logics for dynamically created locks were developed by Gotsman et al. (2007b), and later for atomic operations such as Compare-and-swap allowing lockless data structures to be proved (Parkinson, Bornat, and P. O’Hearn, 2007).

All the examples above consider programs with sequentially consistent (SC) semantics (Lamport, 1979): The execution of concurrent programs are interleaved together. Parallel computers exhibit non-SC behaviors, making it particularly difficult to reason

about concurrent programs. Logics based on CSL were developed to reason about such programs, such as GPS (Turon, Vafeiadis, and Dreyer, 2014).

Other directions include distributed systems —eg. Diesel (Sergey, Wilcox, and Tatlock, 2017), or Aneris (Krogh-Jespersen et al., 2020; Gondelman et al., 2021), and language based security (Georges et al., 2021), ...

The main technical means which contributed to this increase in expressive power of the logic is the decoupling of *logical state* from *machine state*. Here, we call machine state the notion of state which is used to define the semantics of the programming language, and we call logical state the type of object the predicates of the logic are predicates on. In the instance of CSL above, the meaning of a predicate P is defined by the satisfiability relation $s \models P$, where s is a machine state. In other words, the two notions are identified in this early example.

Permissions

An early instance of this decoupling is the introduction of *permissions* which were introduced by Bornat, Calcagno, P. O’Hearn, et al. (2005) in separation logic, following their invention by Boyland (2003), who used them in a type system.

To avoid data races, the situation to forbid is one thread writing to a location while another is reading or writing to that location. The variant of CSL we have presented above is too strict in that it also disallows concurrent reads. Indeed, to be able to read a location ℓ , using the rule (0.1), a command needs ownership of a corresponding resource $\ell \mapsto v$. Since this resource is not duplicable, either it is protected by a lock or only one thread can own it.

The solution is to consider a more refined notion of ownership: one can own a fraction p ($0 < p \leq 1$) of a location. Fractional ownership can be combined using the following law:

$$\ell \stackrel{p_1}{\mapsto} v * \ell \stackrel{p_2}{\mapsto} v \iff \ell \stackrel{p_1+p_2}{\mapsto} v$$

if $p_1 + p_2 \leq 1$. Any fractional ownership is sufficient to read from it, but full ownership ($p = 1$) is required to write to it. The invariant enforced by the logic is that if one thread owns a full permission over some location ℓ , then no other thread is able to own any fractional permission to that location, which means that no other thread can read from this location, avoiding data races on location ℓ .

Because predicates now contain permissions, they are predicates over a refined version of the memory which contains permissions. This logical state is related to the machine state by erasing permissions.

Ghost variables and logical state

In Hoare logics, it is very useful to have *ghost code*, code which is added to the program to help with reasoning but does not affect the behavior of the program. A classical example is proving that two threads which each add k to a shared reference end up adding $2k$ to it. Here is the program:

```
resource r do
  with r do (x := [ℓ]; [ℓ] := x + k) || with r do (y := [ℓ]; [ℓ] := y + k)
```

and we want to prove that, if ℓ contains initially 0, in the final state it contains $2k$. Since both threads write to ℓ , we need to put total ownership of it in the lock r , but the invariant cannot express that, depending on which instruction has been executed, the value at location ℓ is either 0, k or $2k$: it has to be of the form $\exists v, \ell \stackrel{1}{\mapsto} v$ or similar.

The solution is to use ghost code, and add two ghost variables λ and ρ such that λ contains the contribution of the left thread, and ρ that of the right. Because each is written to by a single thread and both are mentioned in the invariant, we split their ownership in two parts: one half is given to the relevant thread, and the other is in the invariant:

$$\exists v_1, v_2, \lambda \stackrel{1/2}{\mapsto} v_1 * \rho \stackrel{1/2}{\mapsto} v_2 * \ell \stackrel{1}{\mapsto} v_1 + v_2$$

We modify the code so that each critical section also increments the corresponding ghost variable by k . At the end of the execution, we can easily prove that λ and ρ both contain k and therefore ℓ contains $2k$.

Usually there are syntactic rules to ensure that non-ghost variables do not depend on ghost variables, and one can prove an erasure theorem stating that the program with the ghost code erased and the original program are equivalent in some sense. This approach has drawback however, as the source code of the program has to be modified, which is detrimental to modularity of proofs.

In recent concurrent separation logics such as Iris (Jung, Swasey, et al., 2015), ghost variables are subsumed by logical variables, which live in the logical state instead of the machine state. This increases further the distance between the machine and the logical states compared to permissions, as they no longer have the same “shape”.

Iris

In Part II of this thesis, we will use the Iris (Jung, Swasey, et al., 2015; Jung, Robbert Krebbers, Lars Birkedal, et al., 2016; Robbert Krebbers, Jung, et al., 2017) concurrent separation logic, which unifies under an extensible framework many of the previous

concurrent separation logics which had been developed in the 2010s. One powerful aspect of Iris is that Hoare triples are defined in its language of predicate. This has two consequences: First, a Hoare triple can appear in the precondition of some Hoare triple, allowing for specifying higher-order programs. Second, this lowers the effort required to develop new program logics, since one only needs to define a new version of Hoare triples in the base logic of Iris. An instance of this appears in Part II.

The semantic of Iris is defined a step-indexed Kripke semantics: an Iris predicate is a step-indexed predicate over a notion of world, which are defined as the solution to a guarded domain equation. Iris's facilities to solve domain equation have also been used to define step-indexed logical relations (Robbert Krebbers, Timany, and Lars Birkedal, 2017): for example to prove properties about the ST monad (Timany, Léo Stefanescu, et al., 2017), or the soundness of the DOT calculus (Giarrusso et al., 2020), but also to develop a model of the Rust language (Jung, Jourdan, et al., 2017).

The soundness theorem for the usual notion of Hoare triple of Iris is proved using the adequacy theorem of the base logic. The soundness theorem states, as usual, that a well-specified program does not get stuck (the way runtime errors are represented in Iris). Though it does not provide data race freedom, it is possible to define other program logics which do guarantee data race freedom, such as the RustBelt semantic model of the Rust language (Jung, Jourdan, et al., 2017).

In Part II of this thesis, we develop a program logic in the Iris framework with a more intensional soundness theorem. The logic is parameterized by a user-chosen state transition system (STS) which models the behavior they want their program to satisfy. For example, they can chose the TLA+ (L. Lamport, 1994) model of the single-decree Paxos algorithm. If they prove their implementation of Paxos in this logic, the soundness theorem will certify that their Paxos implementation is simulated by the abstract model.

Another use of this logic is to prove that a concurrent program terminates assuming the thread-scheduler is *fair*. This is done by relating the program with an STS whose states are ordered by a well-founded order, and which satisfy a property expressing that enough transitions of the STS make the state decrease.

0.4 Organization of the thesis

This thesis is made of two independent parts. The first part, which reports on joint works with Paul-André Melliès (Melliès and Léo Stefanescu, 2017; Melliès and Léo Stefanescu, 2018; Melliès and Léo Stefanescu, 2020), takes a new approach to the soundness of concurrent separation logic. Both the programs and the derivation trees of CSL are

given a truly concurrent interpretation using asynchronous graphs. These semantics are related by a forgetful map between them, and our version of the soundness theorem is expressed as topological properties of this map.

- Chapter 1 is a brief survey of previous proofs of soundness for CSL,
- Chapter 2 introduces our approach to the problem and defines the basic notions which we will use,
- Chapter 3 presents our asynchronous model of CSL and our proof of soundness.

Part II presents a joint work with Lars Birkedal, Léon Gondelman, Abel Nieto, Simon Oddershede Gregersen and Amin Timany which we described briefly just above (Timany, Gregersen, et al., 2021).

- Chapter 4 introduces Iris,
- Chapter 5 defines the new Hoare triples, its soundness theorem and examples relating programs with STSs describing their correctness,
- Chapter 6 develops a logic to prove the fair termination of shared memory concurrent programs based on the logic described in the previous chapter. This is the personal work of the author.

Links to definitions

This thesis is using the excellent *knowledge* package developed by **Thomas Colcombet**. Definitions are written in bold, as just above, and uses of these notions are shown like this: **tiny**. Clicking on it in a PDF viewer goes to the definition, and hovering it may display the definition.

Part I

Asynchronous models of CSL

1 Proofs of soundness of CSL

We can distinguish two kinds of proofs of soundness of concurrent separation logic, depending on how the semantics of the programming language is formalized: semantically using traces or operationally using a small-step semantics.

1.1 Trace semantics of the language

When O’Hearn developed the rules of CSL, he asked Brookes, a specialist in the semantics of concurrent shared-memory programming languages, to help him prove the soundness of the logic. We describe here Brookes seminal proof (Brookes, 2004; Brookes, 2011).

He considers a simple WHILE language with parallel composition, locks and critical sections. Its semantics is defined using the notion of *action*, which are of the form

$$\lambda := \text{nop} \mid i = v \mid i := v \mid [\ell] = v \mid [\ell] := v \mid \text{acq}(r) \mid \text{rel}(r) \mid \text{try}(r) \mid \text{abort} \dots$$

An action describes an event in the history of a program, for example, $i = v$ means that the program read variable i and its value was v , and $i := v$ means that the program wrote v into the variable i . The notation $[\ell]$ denotes the dereferencing of a memory address ℓ . The last three operations on a lock r respectively mean that the program successfully acquired the lock r , released it, or unsuccessfully tried to acquire it.

Interpretation of the language Programs are interpreted as sets of *traces*, finite or infinite sequences of such actions. But first, we can give a semantics to actions λ by giving them *effects* $\xRightarrow{\lambda} \subseteq \mathbf{State} \times \mathbf{State}$. The predicate $s \xRightarrow{\lambda} s'$ between states s and s' holds when the action λ is compatible with the initial state s and final state s' . For example, the action $i = v$ of reading the value v from the variable i has the following effect:

$$\begin{aligned} s &\xRightarrow{i=v} s && \text{if the state } s \text{ maps } i \text{ to } v \\ s &\xRightarrow{i=v} \perp && \text{if the variable } i \text{ is not in the state } s \end{aligned}$$

where $\perp \in \mathbf{State}$ is a special state denoting an error; here reading from an undefined variable. Note that, if s maps i to another value than v , then there is no s' such that $s \xrightarrow{i=v} s'$. We say that the action λ is *enabled* when there exists a state s' such that $s \xrightarrow{\lambda} s'$.

A trace $\vec{\lambda}$ is called *sequential* (or sequentially consistent) when it is enabled, in that there exist $s_0, \dots, s_{|\vec{\lambda}|}$ such that

$$s_0 \xrightarrow{\lambda_1} s_1 \xrightarrow{\lambda_2} s_2 \xrightarrow{\lambda_3} \dots$$

Such a trace describes an execution which makes sense without interference from the environment.

As expected for a denotational semantics, the operators used to compose programs are reflected as operations in the semantic domain. The interpretation of the sequential fragment of the language is standard. For the parallel product, given two sets T_1 and T_2 of traces, $T_1 \parallel T_2$ is defined as the union of the interleavings of $t_1 \in T_1$ and of $t_2 \in T_2$. The set of interleavings of two traces is defined coinductively on their structure, taking into account the set of locks which is held by the program and its environment. The definition is setup in such a way that $t_1 \parallel t_2$ contains the singleton trace *abort* if it contains a data race.

With these operations, we interpret the commands as sets of traces and the expressions as sets of pairs (t, v) of a trace t and a resulting value v in the standard way. Of course, the traces in the semantics of a program are not necessarily sequentially consistent; for example the expression $i + i$ which reads the variable i twice and adds the two values is interpreted as the following set of pairs of a trace and a resulting value:

$$\{ ((i = v) (i = w), v + w) \mid v \text{ and } w \text{ values} \}$$

Interpretation of the logic The Hoare triples $\Gamma \vdash \{P\} C \{Q\}$ which are considered by Brookes (2004) have contexts of the form

$$r_1(X_1) : J_1, \dots, r_n(X_n) : J_n$$

where the X_i are sets of program variables which are said to be owned by the invariant.

The soundness theorem is formalized by defining the notion of valid Hoare triple. Informally, if a Hoare triple $\Gamma \vdash \{P\} C \{Q\}$ is valid, then all the traces of the program C which start from a state which satisfy the precondition P and the invariants in Γ are safe and their final states satisfy the postcondition Q and the invariants still hold.

This definition is too weak to be an inductive invariant because it does not take into account the interference from the environment, which is necessary to prove that the parallel product preserves validity.

Brookes' solution is to define a local notion of reduction, which he calls the *local enabling relation*

$$(s, h, A) \xrightarrow[\Gamma]{\lambda} (s', h', A')$$

where s is the stack, h is the heap, and A is the set of locks held by the “current” thread. For actions which do not involve locks, local enabling essentially coincides with the global enabling we have talked about above, except that the use of variables which appear in invariants is forbidden: For instance, the read of a variable i

$$(s, h, A) \xrightarrow[\Gamma]{i:=v} \not\downarrow$$

is invalid if $i \notin \text{dom}(s)$, as before, or if i is free in an invariant J_k which is not owned, in that $r_k \notin A$.

The logical nature of the local enabling relation is clearer when looking, for example, at the rules for the action $acq(r)$:

$$(s, h, A) \xrightarrow[\Gamma, r(X):J]{acq(r)} (s \uplus s_r, h \uplus h_r, A \uplus \{r\})$$

for any stack s_r and heap h_r such that $\text{dom}(s_r) = X$ and $(s \uplus s_r, h_r) \models J$. When the program acquires a lock, its memory state can be adjoined with any memory state satisfying the invariant associated with the lock r .

Remark 1.1.1. This proof technique puts a requirement of the semantics of the programming language we consider: they must be insensitive to any extension of the memory:

$$s \uplus s_F \xrightarrow{\lambda} s' \quad \Rightarrow \quad \exists s'', s' = s'' \uplus s_F \quad \wedge \quad s \xrightarrow{\lambda} s''$$

A common counter-example is semantics with a deterministic allocation strategy, such as `malloc` allocating the least free address. A naive notion of validity would make the following Hoare triple valid

$$\Gamma \vdash \{\mathbf{emp}\} \text{malloc}(E) \{0 \mapsto E\} \tag{1.1}$$

because we assert that in the precondition that the initial state is empty, and therefore that address 0 is the least free address. This would, however, contradict the Frame-rule of separation logic, since it would let us prove the following rule:

$$\Gamma \vdash \{R\} \text{malloc}(E) \{R * 0 \mapsto E\}$$

which is clearly unsound (take $R := 0 \mapsto 88$ for example). The technique used by Brookes (2004) ensures the validity of the Frame rule by putting conditions on the semantics of the programming language. Another more flexible proof, for example the proof by Vafeiadis (2011) which we will describe in Section 1.2, bakes in the frame into the definition of validity, making rule (1.1) invalid.

The core of the proof are lemmas relating local traces of a command C with local traces of its subcommands C_i . For example, for the parallel product:

Theorem 1.1.2 (Thm 21 in (Brookes, 2004)). *Let C_1 and C_2 be two commands and Γ a CSL context, and suppose that every variable which is free in one and written to in the other is owned by some invariant in Γ . Consider a trace $\alpha_1 \in \llbracket C_1 \rrbracket$ and a trace $\alpha_2 \in \llbracket C_2 \rrbracket$. Given a trace $\alpha \in \alpha_1 \parallel \alpha_2$ in their parallel product and $h = h_1 \uplus h_2$:*

1. *If $(s, h, \emptyset) \xrightarrow{\alpha}_{\Gamma} \downarrow$ then $(s \setminus \text{wr}(\alpha_2), h_1, \emptyset) \xrightarrow{\alpha_1}_{\Gamma} \downarrow$ or $(s \setminus \text{wr}(\alpha_1), h_2, \emptyset) \xrightarrow{\alpha_2}_{\Gamma} \downarrow$.*
2. *If $(s, h, \emptyset) \xrightarrow{\alpha}_{\Gamma} (s', h', \emptyset)$ then either:*
 - *$(s \setminus \text{wr}(\alpha_2), h_1, \emptyset) \xrightarrow{\alpha_1}_{\Gamma} \downarrow$, or*
 - *$(s \setminus \text{wr}(\alpha_1), h_2, \emptyset) \xrightarrow{\alpha_2}_{\Gamma} \downarrow$, or*
 - *there exists $h'_1 \uplus h'_2 = h'$ such that*
 - *$(s \setminus \text{wr}(\alpha_2), h_1, \emptyset) \xrightarrow{\alpha_1}_{\Gamma} (s' \setminus \text{wr}(\alpha_2), h'_1, \emptyset)$, and*
 - *$(s \setminus \text{wr}(\alpha_1), h_2, \emptyset) \xrightarrow{\alpha_2}_{\Gamma} (s' \setminus \text{wr}(\alpha_1), h'_2, \emptyset)$.*

If we assume that α_1 and α_2 do not crash, then this imply that neither do the interleavings. This result is proved using a difficult induction on the length of the traces α_1 and α_2 , where the statement is generalized to an arbitrary set A of locked resources instead of \emptyset .

Notice also that, because the separation of which program variables are used in which part of the program is handled using side rules, the statements of the lemmas as the one above need to pose subtle restrictions. This was solved in subsequent works with the variables as resources discipline of Bornat, Calcagno, and Yang (2006) (see the next chapter) and later by considering languages with immutable variables which can contain mutable memory-cells, in the tradition of ML.

We can now state the definition of validity of a Hoare triple:

Definition 1.1.3. A Hoare triple $\Gamma \vdash \{P\} C \{Q\}$ is *valid* if, for any trace $\alpha \in \llbracket C \rrbracket$, for any stacks s whose domain contains all program variables which appear in C and in Γ but are not owned by an invariant in Γ , and for any heap h , if $(s, h, \emptyset) \models P$ and $(s, h, \emptyset) \xrightarrow{\alpha} s'$, then $s' \models Q$. In particular, s' is not the error state \perp since it satisfies some formula.

As expected, the soundness theorem states:

Theorem 1.1.4 (Soundness). *Every provable Hoare triple $\Gamma \vdash \{P\} C \{Q\}$ is valid.*

It is proved by induction on the derivation tree of the Hoare triple $\Gamma \vdash \{P\} C \{Q\}$, using decomposition results such as Theorem 1.1.2.

Finally, all that remains to be done is to link validity of a Hoare triple, which is a statement about the *local* enabling relation of the program, to its global semantics, as mentioned in Remark 1.1.1.

1.1.1 Other proofs based on trace semantics

In their work on Subjective Concurrent Separation Logic, Ley-Wild and Nanevski (2013) interpret programs as sets of *action trees*, which one can see as more interactive variants of traces. An instead of having one trace for, say, each value return by a load to memory, the tree contains a branching with one subtree for each value. Unlike in Brookes's proof, action trees are finite approximations of the program, and programs are interpreted as sets of such trees, which form a complete lattice to interpret recursion. A more recent research program which interprets program as (possibly infinite) interactive trees in the Coq proof assistant is the work on *Interaction trees* by Xia et al. (2020) and Zakowski et al. (2020).

1.2 Operational semantics of the language

In this section, we explain how to prove the soundness of concurrent separation logic which is based on a small step operational semantics of the language, instead of a denotational semantics as in Brookes original proof which we described in the previous section. This section is based on the proof of soundness by Vafeiadis (2011). There are other proofs based on operational semantics, such as the Views framework by Dinsdale-Young et al. (2013), the TaDa logic by da Rocha Pinto, Dinsdale-Young, and Gardner (2014) or, more recently, the Iris framework, which we will discuss in the second part of this thesis.

$$\frac{\varepsilon \models Q}{\text{safe}(\text{skip}, \varepsilon, Q)} \qquad \frac{(C, \varepsilon) \not\rightarrow \perp \quad \forall C', \varepsilon', (C, \varepsilon) \rightarrow (C', \varepsilon') \Rightarrow \text{safe}(C', \varepsilon', Q)}{\text{safe}(C, \varepsilon, Q)}$$

Figure 1.1: Semantic validity for Hoare logic

Vafeiadis expresses the semantics of the programming language under consideration, a simple WHILE language with a parallel product and locks as small-step semantics: a state transition system whose states are pairs (C, ε) of a program C and a machine state ε which contains the state of the memory and of the locks. A transition

$$(C, \varepsilon) \longrightarrow (C', \varepsilon')$$

expresses that, in machine state ε , the program C reduces to C' and updates the machine state to ε' . The definition is the expected one for a concurrent shared memory programming language, with the following two rules for the parallel product:

$$\frac{(C_1, \varepsilon) \longrightarrow (C'_1, \varepsilon')}{(C_1 \parallel C_2, \varepsilon) \longrightarrow (C'_1 \parallel C_2, \varepsilon')} \qquad \frac{(C_2, \varepsilon) \longrightarrow (C'_2, \varepsilon')}{(C_1 \parallel C_2, \varepsilon) \longrightarrow (C_1 \parallel C'_2, \varepsilon')}$$

Note that this is a slight simplification of the setting in the proof of Vafeiadis (2011) in the way we deal with locks.

Safety of a Hoare triple Before we explain the validity of a concurrent separation logic Hoare triple in the setting of Vafeiadis (2011), we go back to the simpler case of Hoare logic for (weak correctness of) sequential programs. In that setting, one can express the semantic validity of a Hoare triple $\vdash \{P\} C \{Q\}$ at machine state ε as $\text{safe}(C, \varepsilon, Q)$, where the predicate safe is defined in Figure 1.1. The predicate $\text{safe}(C, \varepsilon, Q)$ explores the tree of all computation which start in state (C, ε) and check that no path leads to the error state \perp , and that all leaves satisfy the postcondition Q . We can then define the validity of a Hoare triple as follows:

$$\vdash \{P\} C \{Q\} := \forall \varepsilon, \varepsilon \models P \Rightarrow \text{safe}(C, \varepsilon, Q).$$

The predicate of Figure 1.1 needs to be adapted to the setting of concurrent separation logic: First, separation logic predicates P (in linear variants such as the ones we consider here) hold for “small” memory states; therefore we would have, for example, to replace $\varepsilon \models P$ by $\exists \varepsilon', \varepsilon' \subseteq \varepsilon \wedge \varepsilon' \models P$. More fundamentally, the predicate safe cannot handle an environment modifying the state, and therefore does not validate the inference rule for the parallel product.

$$\begin{array}{c}
 \frac{\mathfrak{s} \models Q}{\text{safe}(\text{skip}, \mathfrak{s}, \Gamma, Q)} \\
 \\
 \forall \mathfrak{s} \uplus \mathfrak{s}_F, (C, \mathfrak{s} \uplus \mathfrak{s}') \not\rightarrow \frac{1}{2} \\
 \\
 \forall \mathfrak{s} \uplus \mathfrak{s}_\Gamma \uplus \mathfrak{s}_F, \mathfrak{s}_\Gamma \models \bigcirc \Gamma(r) \Rightarrow \\
 \quad r \in \text{unlocked}(\mathfrak{s}) \\
 \forall C', \mathfrak{s}', (C, \mathfrak{s} \uplus \mathfrak{s}_\Gamma \uplus \mathfrak{s}_F) \rightarrow (C', \mathfrak{s}') \Rightarrow \\
 \quad \exists \mathfrak{s}'' \uplus \mathfrak{s}'_\Gamma \uplus \mathfrak{s}_F = \mathfrak{s}' \wedge \mathfrak{s}'_\Gamma \models \bigcirc \Gamma(r) \wedge \text{safe}(C', \mathfrak{s}'', Q) \\
 \quad \quad \quad r \in \text{unlocked}(\mathfrak{s}') \\
 \hline
 \text{safe}(C, \mathfrak{s}, \Gamma, Q)
 \end{array}$$

Figure 1.2: Semantic validity for concurrent separation logic

Vafeiadis' solution is to use the predicate defined coinductively¹ in Figure 1.2. The predicate $\text{safe}(C, \mathfrak{s}, \Gamma, Q)$ is defined as a predicate over the the executions of C . Unlike its counterpart for Hoare logic, at each step, we consider what happens if C is executed in a state which *contains* \mathfrak{s} and which contains memory which satisfies the invariants of the unlocked resources. One important thing to note and to contrast with the notion of local enabling in the proof of Brookes is that we always consider the reduction of the program in the global state of the memory, which is decomposed as $\mathfrak{s} \uplus \mathfrak{s}_\Gamma \uplus \mathfrak{s}_F$: the memory \mathfrak{s} which is owned by the program, \mathfrak{s}_Γ which is being protected by the locks, and \mathfrak{s}_F which is owned by the environment, which we also call the frame. The reason this definition is adequate for the concurrent setting is that, at each step, the pieces of state \mathfrak{s}_Γ and \mathfrak{s}_F are universally quantified. As long as the environment does not modify \mathfrak{s} and preserves the invariants, the program C will be safe, and will preserve the memory of the environment and the invariants. Therefore, one can show that, assuming both $\text{safe}(C_1, \mathfrak{s}_1, \Gamma, Q_1)$ and $\text{safe}(C_2, \mathfrak{s}_2, \Gamma, Q_2)$ hold, with $\mathfrak{s}_1 \uplus \mathfrak{s}_2$, then so does $\text{safe}(C_1 \parallel C_2, \mathfrak{s}_1 \uplus \mathfrak{s}_2, \Gamma, Q_1 * Q_2)$.

Finally, Vafeiadis proves the soundness theorem:

Theorem 1.2.1. *If $\Gamma \vdash \{P\} C \{Q\}$ is provable, and if $\mathfrak{s} \models P$, then $\text{safe}(C, \mathfrak{s}, \Gamma, Q)$.*

To prove that the absence of data-races in well-specified programs, similarly to Brookes'

¹The paper and its associated mechanizations use the corresponding step-indexed predicate safe_n ; and we can define $\text{safe} := \bigwedge_n \text{safe}_n$ to recover the coinductive predicate.

proof before, the following rule RACEDetect which crashes the program when there is a data-race:

$$\frac{(\text{accesses}(C_1, \varsigma) \cap \text{writes}(C_2, \varsigma)) \cup (\text{accesses}(C_2, \varsigma) \cap \text{writes}(C_1, \varsigma)) = \emptyset}{(C_1 \parallel C_2, \varsigma) \longrightarrow \downarrow}$$

which states that if C_1 is about to access a variable or a location which C_2 is about to write, or vice-versa, then the program crashes. Since the safe predicate implies safety of executions of closed programs without interference from the environment, this implies the absence of data-race.

1.2.1 Step-indexed models of CSL

There is a line of research (Svendsen and Lars Birkedal, 2014; Dodds et al., 2016) which uses a step-indexed type theory to define the validity of a Hoare triple. This line of research developed into the Iris framework (Jung, Swasey, et al., 2015; Jung, Robbert Krebbers, Lars Birkedal, et al., 2016; Robbert Krebbers, Jung, et al., 2017; Jung, Robbert Krebbers, Jourdan, et al., 2018).

The proof of soundness of the Iris separation logic proceeds by defining a weakest precondition similar to safe expressed in the Iris base logic, which is a modal logic with guarded recursion and a separating conjunction. Instead of syntactic inference rules, the user of the logic is directly manipulating the semantic predicate, with lemmas which roughly correspond to the usual inference rules of the logic. We give a more complete description of Iris in Chapter 4.

One thing to note here is that the “generic” notion of Hoare triple of Iris does not guarantee data-race freedom, or rather, every memory access is considered atomic. One can recover data-race freedom by instrumenting the semantics: For example, in the semantics of λ_{rust} , in the RustBelt development by Jung, Jourdan, et al. (2017), each memory cell is associated with what amounts to a reader-writer lock which crashes the program if two threads access and write the same memory location concurrently.

1.3 Other proofs soundness of CSL

1.3.1 Syntactic proofs

There are many proofs of soundness based on an operational semantics of the programming language and which are modeled after the familiar technique of progress

and preservation proofs of safety of type systems. The idea is to see a (derivation tree of) a Hoare triple as a typing derivation and to prove two lemmas: First, a version of *progress* stating that a well-specified program does not immediately crash in a state which satisfies the precondition of the Hoare triple.

Lemma 1.3.1. *Given a Hoare triple $\Gamma \vdash \{P\} C \{Q\}$, and a machine state $s \models P$, the next step of reduction of C are safe:*

$$(C, s) \not\rightarrow \downarrow.$$

Second, one also proves a version of *preservation*, which states that a well-specified program necessarily reduces to a well-specified program.

Lemma 1.3.2. *Given a Hoare triple $\Gamma \vdash \{P\} C \{Q\}$, and a machine state $s \models P$, for any possible reduction*

$$(C, s) \longrightarrow (C', s')$$

there exists a predicate P' of CSL such that $s' \models P'$ and a derivable Hoare triple $\Gamma \vdash \{P'\} C' \{Q\}$.

The safety of well specified programs and the validity of the postcondition at final states are easy consequences of the two preceding lemmas.

This proof techniques have been used by many authors, for example François Pottier (2013) and Balabonski, François Pottier, and Protzenko (2014) in works which prove the type soundness of type systems which use features borrowed from separation logic. Gotsman et al. (2007a) use a predicate transformer semantics to prove the soundness of their variant of concurrent separation logic.

1.3.2 Other

There are many proofs of soundness of concurrent separation logic. We cite a few notable ones here, but we do not claim any exhaustiveness. Hayman and Winskel (2008) have given a proof of soundness of CSL using a truly concurrent semantics of the underlying programming language using Petri nets. Unlike the other proofs we have cited here, data-races are not detecting by crashing or clocking the program, but instead, data races are detected using the notion of independence of Petri nets, which state that two transitions have disjoint neighborhoods.

1.4 Conclusion

To our knowledge, most proofs deduce data race-freedom from the safety of programs where the semantics crashes in case of data race. Moreover, in most proofs of correctness of concurrent separation logics, the meaning of a Hoare triple $\{P\} C \{Q\}$ is a property of the program.

Instead, our proof of soundness interprets the derivation tree π of $\{P\} C \{Q\}$ itself as an interactive semantics of a program operating with a refined notion of state which tracks the owner of each memory cell.

The asynchronous soundness theorem which we will describe in the next chapter explains the absence of data race in valid executions by the fact that any concurrent access in the semantics of the program C corresponds to an independent access in the semantics of the derivation tree π .

2 Asynchronous soundness for CSL

2.1 Hoare logic as refinement systems

The starting point of our semantics of concurrent separation logic (CSL) is to see such a program logic as a refinement type system: the program C , which we see as being of type $\mathbf{State} \rightarrow \mathbf{State}$, is *refined* by the Hoare triples $\Gamma \vdash \{P\} C \{Q\}$ it satisfies, which we could write $\Gamma \vdash C : \{P\} \rightarrow \{Q\}$. Indeed, if we consider the input and the output of the program to be the initial and the final states of its executions, the predicates P and Q describe subsets of the input and output respectively. This is characteristic of type refinement systems.

Melliès and Zeilberger (2015) have argued that the semantics of a type refinement system should be given as a functor

$$\mathcal{L} : \mathcal{D} \rightarrow \mathcal{C}$$

from a category \mathcal{D} where the refined programs are interpreted to a category \mathcal{C} where the programs are interpreted. In the case of a Hoare logic for example, the functor \mathcal{L} corresponds to the syntactic “forgetful” map

$$\begin{array}{c} \vdots \pi \\ \vdots \\ \Gamma \vdash \{P\} C \{Q\} \end{array} \mapsto C$$

projecting the program C out of the proof tree π of the Hoare triple $\Gamma \vdash \{P\} C \{Q\}$.

Our model will follow this guiding idea. In particular, the proofs π of the Hoare triples $\Gamma \vdash \{P\} C \{Q\}$ will be given their own semantic interpretations $\llbracket \pi \rrbracket$. The programs C will also be given a semantics $\llbracket C \rrbracket$, and the two will be related through a semantic map

$$\mathcal{L} : \llbracket \pi \rrbracket \rightarrow \llbracket C \rrbracket.$$

The existence of the map \mathcal{L} suggests that the interpretations of the proofs and of the programs should be similar in nature so that they can easily be compared.

This chapter explains the rationale behind the choices we have made in the construction of our model of concurrent separation logic, and in particular how we deal with data races.

2.2 State transition systems

Perhaps the simplest way to give a semantics to a sequential imperative program is as a *state transformer*: a program C is seen as a function

$$\llbracket C \rrbracket : \mathbf{State} \rightarrow \mathbf{State} + \perp$$

which takes as input an initial memory state and outputs either the final state of the program, when the program terminates on that input, or \perp if the program does not halt.

In the semantics of state transformers, the following two programs have the exact same semantics:

$$\begin{array}{ll} X := 0; & X := 4; \\ Y := 0; & Y := 0; \\ & X := Y; \end{array} \quad \text{In}$$

some context, too much information is lost. When he gave his object based semantics of Idealized Algol, a higher-order sequential language, Reddy (1996) argued the information missing in this semantics is the *changing* of the state, seen as a first-class phenomenon and not simply a piece of information which can (partially) be recovered as the difference between the initial and the final state the program.

Another, more obvious, obstacle appears when adding concurrency to the language: the environment is able to observe the transient state of the program on the right above where the variable X is equal to the value 4. For example, the context $(- \mid \mid Z := X)$ is able to distinguish between the two programs by reading X during the time in which it holds the value 4.

To capture the changes induced by a program C , we interpret it as a [state transition systems](#) $\llbracket C \rrbracket$: a graph where each intermediate state the program goes through is represented by a node, and each action of the program is represented as an edge.

It is important, if not so common, to distinguish between the *internal* and the *external state* of the program. The external state is characterized by the fact that it is observable from the context, and therefore it can be used for synchronization. In the case of the simple imperative and concurrent language we consider, it consists essentially on the state of the memory, since it is what is shared between the program and its environment. In a typed language, the type of the program would also be part of its external state; to a certain extent, this can be seen in how we interpret proofs of CSL, which we can see as programs typed using a refinement type system. The internal state, on the other hand, corresponds in our case to the position of the program the control flow, which would correspond to the program counter in a concrete machine.

To enforce this discipline, the state transition system $\llbracket C \rrbracket$ is labeled explicitly using a map $\lambda_{\llbracket C \rrbracket} : \llbracket C \rrbracket \rightarrow \mathfrak{A}$ which associate an external state in \mathfrak{A} to each node of the state transition system $\llbracket C \rrbracket$. We explain the nature of \mathfrak{A} in the sequel. The nodes of $\llbracket C \rrbracket$ are left “opaque”: all we know about them is their image under $\lambda_{\llbracket C \rrbracket}$.

2.3 An imperative shared-memory concurrent language

The language we consider is the same as the one considered in the early concurrent separation logic papers, for example in the proofs of soundness of Brookes (2004) in Vafeiadis (2011).

The grammar of this language is the following. It has Boolean expressions

$$B ::= \text{true} \mid \text{false} \mid B \wedge B' \mid B \vee B' \mid \neg B \mid E = E'$$

arithmetic expressions

$$E ::= 0 \mid 1 \mid \dots \mid x \mid E + E' \mid E * E'$$

and commands

$$\begin{aligned} C ::= & x := E \mid x := [E] \mid [E] := E' \mid \text{skip} \mid \text{dispose}(E) \mid x := \text{malloc}(E) \\ & \mid C; C' \mid C_1 \parallel C_2 \mid \text{resource } r \text{ do } C \mid \text{with } r \text{ do } C \\ & \mid \text{if } B \text{ then } C_1 \text{ else } C_2 \mid \text{while } B \text{ do } C \end{aligned}$$

where $x \in \mathbf{Var}$ are program (stack) variables, and the notation $[E]$ dereferences the address denoted by the arithmetic expression E . The action of the construction $\text{with } r \text{ do } C$ is to acquire the lock r , execute the command C , and then release the lock r . Locks are introduced with the construction $\text{resource } r \text{ do } C$: the lock r is only available inside of the command C .

Let us define the external states of our state transition systems. We decompose the memory in two parts: the *stack* $s : \mathbf{Var} \rightarrow_{\text{fin}} \mathbf{Val}$ which is a partial finite map from the set \mathbf{Var} of variables to the set \mathbf{Val} of values, and the *heap* $h : \mathbf{Loc} \rightarrow_{\text{fin}} \mathbf{Val}$, which associates a value to each allocated location $\ell \in \mathbf{Loc}$. We identify the set \mathbf{Loc} of locations with the natural numbers \mathbb{N} , and we assume that $\mathbf{Loc} \subseteq \mathbf{Val}$ so that values can serve as addresses in the heap. A *memory state* μ is a pair (s, h) of a stack s and of a heap h .

There is one more piece of state which is needed: the locks. Given a set $\mathcal{L} \subseteq \mathbf{Locks}$ of locks valid in the current context, the *lock state* keeps track of the set $L \subseteq \mathcal{L}$ of locks

$$\begin{array}{c}
 \frac{E(\mu) = v}{(\mu, L) \xrightarrow{x:=E} (\mu[x \mapsto v], L)} \qquad \frac{E(\mu) \text{ not defined}}{(\mu, L) \xrightarrow{x:=E} \perp} \\
 \\
 \frac{E(\mu) = \ell \quad \ell \in \text{dom}(\mu) \quad E'(\mu) = v}{(\mu, L) \xrightarrow{[E]:=E'} (\mu[\ell := v], L)} \qquad \frac{r \notin L}{(\mu, L) \xrightarrow{P(r)} (\mu, L \uplus \{r\})} \\
 \\
 \frac{}{(\mu, L \uplus \{r\}) \xrightarrow{V(r)} (\mu, L)} \qquad \frac{E(\mu) = v \quad \ell \notin \text{dom}(\mu)}{(\mu, L \uplus \{r\}) \xrightarrow{\text{alloc}(E)} (\mu[x := \ell, \ell := v], L)}
 \end{array}$$

Figure 2.1: Semantics of machine instructions

which are being held (locked). Finally, a **machine state** $\mathfrak{s} \in \mathbf{State}(\mathfrak{Q})$ is a pair (μ, L) of a memory state μ and of a set of a lock state L .

We label each transition in the state transition system $\llbracket C \rrbracket$ with a label describing the external action the transition corresponds to. There is a strong correspondance between the external actions of the code and the “leaf” commands. The transitions of $\llbracket C \rrbracket$ are labeled with the following set $\mathbf{Instr}(\mathfrak{Q})$ of actions, which we call **machine instructions** m :

$$\begin{aligned}
 & x := E \mid x := [E] \mid [E] := E' \mid \text{nop} \\
 & x := \text{alloc}(E, \ell) \mid \text{dispose}(E) \mid P(r) \mid V(r)
 \end{aligned}$$

The first three correspond to reading and writing the memory, while `nop` does nothing, $x := \text{alloc}(E, \ell)$ corresponds to allocating a new memory cell at address ℓ initializing it with E , $\text{dispose}(E)$ frees the address E . The last two machine instructions $P(r)$ and $V(r)$ respectively acquire and release the lock $r \in \mathfrak{Q}$. The semantics of these instructions are expressed by the transition system $\mathfrak{s}_{\mathfrak{S}}$ between **machine states** described in Figure 2.1. The set of nodes of $\mathfrak{s}_{\mathfrak{S}}$ is the set $\mathbf{State} + \perp$ of machine states adjoined with an error state \perp . Its edges $\cdot \xrightarrow{m} \cdot$ are labeled with machine instructions m .

Remark 2.3.1. The symbol \mathfrak{s} was chosen by Melliès (2019) for its similarity with boat anchors. It was introduced by Jeremy Gibbons as a symbol for tree constructors, as it is inspired by \mathfrak{M} , the Chinese ideogram for tree. It can be pronounced “anchor”, “template” or “moo”.

A state transition system for interpreting programs can be defined as a pair $(G, \lambda : G \rightarrow \mathfrak{s})$ of a graph G and a graph homomorphism λ . Explicitly, each node x of the graph G

is labeled with a machine state $\lambda(x) \in \mathbf{State}$ or the error state \downarrow . Each edge in G , seen as a transition,

$$x \longrightarrow y$$

is labeled with a transition of the state transition system \mathfrak{s}_S defined in Figure 2.1 of one of the two forms

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}' \qquad \mathfrak{s} \xrightarrow{m} \downarrow$$

depending on whether executing the **machine instruction** m in **machine state** \mathfrak{s} is safe (left) or produces an error (right). Since λ is a graph homomorphism, $\lambda(x) = \mathfrak{s}$ and, respectively, $\lambda(y) = \mathfrak{s}'$ or $\lambda(y) = \downarrow$. A quick inspection of the definition of \mathfrak{s}_S shows that the error state is terminal, in that there are no edges outgoing from \downarrow ; hence the description of the image under λ of the edge $x \rightarrow y$ above is exhaustive. We sometimes write the situations above as

$$\mathfrak{s} \xrightarrow{m} \mathfrak{s}' \qquad \mathfrak{s} \xrightarrow{m} \downarrow$$

always keeping in mind that \mathfrak{s} , \mathfrak{s}' and \downarrow are merely labels attached to the (unnamed) nodes of the graph G .

2.4 Concurrent transition systems

The transition systems we have described above are inadequate to interpret shared memory programs for two reasons: First, the treatment of data-races is impossible because the semantics of a program which we have sketched so far does not track which pairs of instructions are executed in parallel. Second, the semantics of programs are not stable under an environment which is executing concurrently with the program, this precludes a compositional definition of the parallel product $C_1 \parallel C_2$ of two programs. To handle the first issue, we switch from using graphs to using **asynchronous graphs** for our transition systems and our machine models. For the second, we consider polarized transition systems which contain both transitions of the program and transitions of the environments.


2.4.1 Asynchronous graphs

To solve these problems without abandoning the idea of using state transition systems based on graphs, we use asynchronous graphs, which are a notion of graphs with 2-dimensional information.


Definition 2.4.1. A graph $G = (V, E, \partial_-, \partial_+)$ consists in a set V of vertices, or nodes, a set E of edges, and two maps $\partial_-, \partial_+ : E \rightarrow V$ assigning a source and a target vertex to each edge, respectively.

Given a graph G , a **square** in G is a pair of paths of length 2 in G with a common source vertex and a common target vertex.

An **asynchronous graph** $\mathcal{G} = (G, \diamond)$ is a graph G equipped with a tuple $\diamond = (T, \partial_\diamond, \sigma)$ of a set T , the elements of which we call **permutation tiles**, a function ∂_\diamond from the set T to the set of squares of the graph G , and an endofunction σ on T such that, if $\partial_\diamond(t) = (f, g)$, then $\partial_\diamond(\sigma(t)) = (g, f)$. A permutation tile mapped to a square $(u \cdot v', v \cdot u')$, where u, u', v, v' are edges of the graph G which form a square



is depicted as a 2-dimensional tile sitting on that square:



A square on which there are no tiles is called a **hole**. We write $T : p \sim q$ to express that T is a tile which sits between the two paths p and q of length 2.

The intuition is that a tile such as the one depicted above witnesses that the transitions u and v are *independent*. The transitions u' and v' should be seen as being the same, respectively, as u and v , but executed from another state. In particular, the fact that they form a square means that the action of u and of v commute.

A morphism \mathcal{F} between **asynchronous graphs** $\mathcal{G} = (G, \diamond_G)$ and $\mathcal{H} = (H, \diamond_H)$ is called an **asynchronous morphism**. It is a pair $\mathcal{F} = (\mathcal{F}_G, \mathcal{F}_\diamond)$ of a graph homomorphism $\mathcal{F}_G : G \rightarrow G'$ and of a map $\mathcal{F}_\diamond : T_G \rightarrow T_H$ between their tiles such that it preserves borders:

$$\mathcal{F}_G \circ \partial_\diamond = \partial_\diamond \circ \mathcal{F}_G, \quad \mathcal{F}_\diamond \circ \sigma = \sigma \circ \mathcal{F}_\diamond.$$

This definition of asynchronous graph, without uniqueness properties on tiles is motivated by the simplicity of the limits and colimits in the category **AsyncGraph** of asynchronous graphs and asynchronous morphisms. This is because the category **AsyncGraph** is the category of presheaves over the category presented by the following graph:

$$[0] \begin{array}{c} \xrightarrow{s} \\ \xleftarrow{t} \end{array} [1] \begin{array}{c} \xrightarrow{dl} \\ \xrightarrow{dr} \\ \xrightarrow{ur} \\ \xrightarrow{ul} \end{array} [2] \begin{array}{c} \xrightarrow{\sigma} \\ \xleftarrow{\sigma} \end{array}$$

and with the equations:

$$\begin{array}{llll} dl \circ s = ul \circ s & dl \circ t = dr \circ s & dr \circ t = ur \circ t & ul \circ t = ur \circ s \\ \sigma \circ dl = ul & \sigma \circ ul = dl & \sigma \circ dr = ur & \sigma \circ ur = dr. \end{array}$$

As such, both limits and colimits are computed object-wise. We detail below the case of pullbacks, and consequently of Cartesian products. The pullback of a cospan

$$A_1 \xrightarrow{f} B \xleftarrow{g} A_2$$

of asynchronous graphs is defined as the asynchronous graph $A_1 \times_B A_2$ below. Its nodes are the pairs (x_1, x_2) consisting of a node x_1 in A_1 and of a node x_2 in A_2 , such that $f(x_1) = g(x_2)$. Its edges $(u_1, u_2) : (x_1, x_2) \rightarrow (y_1, y_2)$ are the pairs consisting of an edge $u_1 : x_1 \rightarrow y_1$ in A_1 and of an edge $u_2 : x_2 \rightarrow y_2$ in A_2 , such that $f(u_1) = g(u_2)$. In the same way, a tile $(\alpha_1, \alpha_2) : (u_1, u_2) \cdot (v'_1, v'_2) \diamond (v_1, v_2) \cdot (u'_1, u'_2)$ is a pair consisting of a tile $\alpha_1 : u_1 \cdot v'_1 \diamond v_1 \cdot u'_1$ of the asynchronous graph A_1 and of a tile $\alpha_2 : u_2 \cdot v'_2 \diamond v_2 \cdot u'_2$ of the asynchronous graph A_2 , such that the two tiles $f(\alpha_1)$ and $g(\alpha_2)$ are equal in the asynchronous graph B . The Cartesian product $A_1 \times A_2$ of two asynchronous graphs is obtained by considering the special case when B is the terminal asynchronous graph, with one node, one edge, and one tile. The definition of $A_1 \times A_2$ thus amounts to forgetting the equality conditions in the definition of the pullback $A_1 \times_B A_2$ above.

2.4.2 Asynchronous machine models

We now add permutation tiles to the machine model \mathfrak{A}_S which we have defined above to express the behavior of instructions which are executed in parallel with respect to data-races. We will then define another machine model \mathfrak{A}_L which will contain the synchronization behavior of the instructions. First, we define the notion of **machine state footprint** which we will use to define data races and the permutation tiles in our machine models.

m	rd	wr	lk	mem
nop	\emptyset	\emptyset	\emptyset	\emptyset
$x := E$	$\text{fv}(E)$	$\{x\}$	\emptyset	\emptyset
$x := [E]$	$\text{fv}(E) \cup \{E\}$	$\{x\}$	\emptyset	\emptyset
$[E] := E'$	$\text{fv}(E')$	$\{E\}$	\emptyset	\emptyset
$x := \text{alloc}(E, \ell)$	$\text{fv}(E)$	$\{x\}$	\emptyset	$\{\ell\}$
dispose(E)	$\text{fv}(E)$	$\{E\}$	\emptyset	$\{E\}$
$P(r)$	\emptyset	\emptyset	$\{r\}$	\emptyset
$V(r)$	\emptyset	\emptyset	$\{r\}$	\emptyset

Table 2.1: Footprint of machine instructions

Data races

We define **data races** using the notion of footprint of a **machine instruction** which summarizes both the part of the memory and the set of locks which is accessed by the instruction.

Definition 2.4.2. A **machine state footprint** (with valid locks \mathcal{L})

$$p \in \wp(\mathbf{Var} + \mathbf{Loc}) \times \wp(\mathbf{Var} + \mathbf{Loc}) \times \wp(\mathcal{L}) \times \wp(\mathbf{Loc})$$

contains: 1. $\text{rd}(p)$ the set of locations that are read, 2. $\text{wr}(p)$ the set of locations that are written, 3. $\text{lk}(p)$ the set of locks that are touched (either acquired or released), and 4. $\text{mem}(p)$ the set of locations which are allocated or deallocated.

Given a **machine instruction** m and a **machine state** s , we define the footprint $p_s(m)$ of m executed at s in Table 2.1. The set of free program variables of an expression E is denoted by $\text{fv}(E)$, and we write E for its value seen as a location in \mathbf{Loc} when evaluated in the state s .

Informally, two instructions executing at the same time define a data race when their footprints are *not* independent:


Definition 2.4.3. Two **machine state footprints** p_1 and p_2 are defined to be **independent** when:

$$\begin{aligned} (\text{rd}(p_2) \cup \text{wr}(p_2)) \cap \text{wr}(p_1) &= \emptyset & \text{lk}(p_1) \cap \text{lk}(p_2) &= \emptyset \\ (\text{rd}(p_1) \cup \text{wr}(p_1)) \cap \text{wr}(p_2) &= \emptyset & \text{mem}(p_1) \cap \text{mem}(p_2) &= \emptyset \end{aligned}$$


It may not be obvious why the last component $\text{mem}(p)$ of the footprint is needed. If two instructions alloc and dispose over the same memory location ℓ are executed, then there must have been some synchronization to propagate the information that the location ℓ is free or no longer free. Hence, the two instructions do not run in parallel. A dedicated component is needed so that the map from the stateful to the stateless semantics, which is defined below, is well defined.

Stateful machine model

We can now complete the definition of the stateful machine model $\mathfrak{A}_S \in \mathbf{AsyncGraph}$ by defining its set of tiles. All the tiles will sit on **squares** of the form



which represent the two schedules of m_1 and m_2 being executed on two different threads. For every such square in \mathfrak{A}_S , there is a unique tile



when $p_{s_1}(m_1)$ and $p_{s_1}(m_2)$ are **independent**.

Remark 2.4.4. This condition above does not depend on the state at which are computed the footprints: When the footprints are independent for one of the four states in the square, they are independent at the other three.

The stateful machine model contains the information of when two instructions executed in two different threads create a data race. We use a second machine model to encode the information of when two instructions are unsynchronized when run in two different threads.

Stateless machine model

We define the *stateless machine model* \mathfrak{M}_L , which is also parameterized by the set \mathcal{L} of available locks. Because its purpose is to embody the synchronization behavior of instructions in the language of *asynchronous graphs*, it needs to contain “more squares” than the *stateful machine model* \mathfrak{M}_S . Indeed, the fact that the actions of two instructions on the *state of the machine* commute with each other does not depend on whether these two instructions synchronize or not.

The solution to this problem is to use a coarser notion of state for the machine model \mathfrak{M}_L : a *lock state*

$$L \in \wp(\mathcal{L}) + \downarrow$$

only keeps track of the set of locks which are being held in the current state.

The edges of the *asynchronous graph* \mathfrak{M}_L are of the form

$$L \xrightarrow{m} L'$$

where m is a *lock instruction*

$$m ::= P(r) \mid V(r) \mid \text{alloc}(\ell) \mid \text{dispose}(\ell) \mid \text{nop}.$$

The effect of a lock instruction on a lock state is simple: $\text{alloc}(\ell)$, $\text{dispose}(\ell)$ and nop act as the identity, and there are transitions

$$L \xrightarrow{P(r)} L \uplus \{r\} \qquad L \uplus \{r\} \xrightarrow{V(r)} L$$

Moreover, since lock states do not contain enough information about whether fallible instructions crash, we add pessimistically the following transitions:

$$\forall L, \forall m, \quad L \xrightarrow{m} \downarrow$$

In order to define the tiles of the asynchronous graph \mathfrak{M}_L , we define the notion of *stateless footprint*.


Definition 2.4.5. A *stateless footprint*

$$p \in \wp(\mathcal{L}) \times \wp(\mathbf{Loc})$$

is the data of a set $\text{lk}(p)$ of locks, and of a set $\text{mem}(p)$ of memory locations. Two such footprints p_1 and p_2 are *independent* when

$$\text{lk}(p_1) \cap \text{lk}(p_2) = \emptyset \qquad \text{mem}(p_1) \cap \text{mem}(p_2) = \emptyset.$$

A square



in the asynchronous graph \mathfrak{A}_L is then associated with a unique tile when the footprints of m_1 and of m_2 are independent.

Relationship between the models

The two models are constructed in such a way that there exists an **asynchronous morphism**

$$\mathcal{L}_{\pm} : \mathfrak{A}_S \rightarrow \mathfrak{A}_L$$

defined on nodes by projecting out the lock state out of the machine state

$$\mathcal{L}(\mu, L) := L$$

and on edges as follows.

$$\begin{array}{ll} (\mu, L) \xrightarrow{P(r)} (\mu, L') & \mapsto L \xrightarrow{P(r)} L' \\ (\mu, L) \xrightarrow{V(r)} (\mu, L') & \mapsto L \xrightarrow{V(r)} L' \\ (\mu, L) \xrightarrow{\text{alloc}(E, \ell)} (\mu, L') & \mapsto L \xrightarrow{\text{alloc}(\ell)} L' \\ (\mu, L) \xrightarrow{\text{dispose}(E)} (\mu, L') & \mapsto L \xrightarrow{\text{dispose}(\llbracket E \rrbracket)(\mu)} L' \\ (\mu, L) \xrightarrow{m} (\mu', L') & \mapsto L \xrightarrow{\text{nop}} L' \end{array}$$

where, in the last case, m is any instruction which is not mentioned above. A rapid inspection of the definitions of the stateful machine model \mathfrak{A}_S and the stateless machine model \mathfrak{A}_L shows that the map \mathcal{L} maps tiles to tiles, and hence defines an **asynchronous morphism**.

2.4.3 Transition systems

We can now refine the way we interpret programs which we sketched in Section 2.3. A state transition system is a pair $(G, \lambda : G \rightarrow \mathfrak{t})$ of an **asynchronous graph** and of an **asynchronous morphism** $\lambda : G \rightarrow \mathfrak{t}$ from G to a machine model \mathfrak{t} . The machine model \mathfrak{t} contains all the information which is intrinsic to the behavior of instructions: their effect on state, and how pairs of instructions behave when they run in parallel. The asynchronous graph G on the other hand, expresses how and when instructions are dispatched, adding the constraints of the program to their inherent constraints in the machine model.

Another way to look at it is that the program adds *program-order* dependencies to instructions to the dependencies such as *synchronizes-with* which are part of the specification of instructions. The property that the program can only *add* dependencies corresponds to the requirement that an asynchronous morphism sends tiles in the program (intuitively, non-dependent transitions) to tiles in the machine model.

To interpret programs with fine-grained information about their concurrent behavior using asynchronous graphs, we interpret each program with two different machine models. As we will see in the next section, the stateful interpretation $\llbracket C \rrbracket_S$ and the stateless interpretation $\llbracket C \rrbracket_L$ will be related by a map \mathcal{L} , resulting in a commutative diagram in the category **AsyncGraph** of the form:

$$\begin{array}{ccc} \llbracket C \rrbracket_S & G_S \xrightarrow{\lambda_S} & \mathfrak{t}_S \\ \mathcal{L} \downarrow & = & \downarrow \mathcal{L}_{\mathfrak{t}} \\ \llbracket C \rrbracket_L & G_L \xrightarrow{\lambda_L} & \mathfrak{t}_L \end{array} \quad (2.1)$$

To demonstrate why this double interpretation is useful, we describe data-races in this semantics.

2.4.4 Data-races

Recall that a data-race happens when two instructions are executed in parallel and touch the same area of memory, with at least an instruction writing. In particular, we need to know when two instructions are executed in parallel without synchronization. This corresponds in our semantics to a tile in the asynchronous graph G_L , with the notations of diagram (2.1) above.

Now, we need to characterize the second half of the definition of a data race, which says that the two instructions are in conflict. Since this is defined with respect to the





Figure 2.2: A data race in the interpretation of a program

memory, this is a property of the stateful interpretation $\lambda_S : G_S \rightarrow \mathfrak{A}_S$. There are two cases. If the instructions involved in the data-race do not commute, for example $x := 4$ and $x := 5$, then there will be a “fork” mapped under \mathcal{L} to the tile in G_L as depicted on the left of Figure 2.2. The second case occurs when the two instructions happen to commute, but still produce a data-race, for example $x := 1$ in parallel with $x := 1$. In that case there will be a **hole**, since there a tile in \mathfrak{A}_S only when there is no data-race, which is mapped to the tile in G_L ; this situation is depicted on the right of Figure 2.2.

Because our goal is to prove the *absence* of data races, we need to prove the negation of the above: for every tile in G_L which is part of an execution, there exists a tile in G_S which is mapped to that tile under the map $\mathcal{L} : G_S \rightarrow G_L$. The condition that the tile is part of an execution is necessary because the stateless semantics $[[C]]_L$ contains nonsensical executions; the stateless semantics is a tool to for adding *local* information to transitions in the stateful semantics $[[C]]_S$.

More precisely, a set of executions of $[[C]]_S$ is data race-free if, whenever two consecutive transitions $m_1 \cdot m_2$ of some execution are mapped under \mathcal{L} to the edge of some tile (depicted below left), the tile can be lifted along \mathcal{L} to a tile whose edge is $m_1 \cdot m_2$ (below right):



2.4.5 Polarized asynchronous transition systems

We now explain how to solve the second problem outlined at the beginning of Section 2.4: the interpretation of programs should be closed under the action of the environment. For this, we explicitly add Environment transitions in the machine models and in the program interpretations so that any possible interference of the environment is accounted for.

First, we adapt our stateful and stateless machine models to have two copies of each edge, which we write

$$\mathfrak{s} \xrightarrow{m:\mathbf{C}} \mathfrak{s}' \qquad \mathfrak{s} \xrightarrow{m:\mathbf{F}} \mathfrak{s}'$$

for, respectively, a transition of the Code, and a transition of the Frame (or Environment). We write \mathfrak{s}° and \mathfrak{L}° for the version with two polarities. To make the notation more explicit about the number of polarities, we write the unpolarized machine models as \mathfrak{s}^{\bullet} and \mathfrak{L}^{\bullet} in the sequel instead of notation \mathfrak{s} and \mathfrak{L} which we have been using until now.

A transition system is therefore a pair (G, λ) of an asynchronous graph G and of an [asynchronous morphism](#)

$$\lambda : G \rightarrow \mathfrak{s}^{\circ}$$


Because the transitions with the Frame polarity F need to account for every possible transition by the Environment, there must be “enough” Frame transitions in G . From the point of view of the Frame, in any state x in G , the set of transitions available to the Frame must be the same as the set of transitions available from $\lambda(x) \in \mathfrak{s}^{\circ}$. This is made formal by stating that λ is an Environment 1-fibration.

Definition 2.4.6. Let $f : G \rightarrow G'$ be an [asynchronous morphism](#) between two [asynchronous graphs](#) G and G' . Let \mathcal{M} be a set of edges in G' .


The asynchronous morphism f is called an **\mathcal{M} -1-fibration** if, for all $x \in G$, and for all edges $e' : f(x) \rightarrow y'$ in \mathcal{M} , there exists an edge $e : x \rightarrow y$ in G such that $f(y) = y'$ and $f(e) = e'$, as depicted below. When there is a notion of Code or Environment edges, this defines the notions of **Code 1-fibration** and **Environment 1-fibration**.

$$\begin{array}{ccc} \begin{array}{c} x \\ \vdots \\ f \\ \vdots \\ f(x) \end{array} & \xrightarrow{\quad} & \begin{array}{c} x \\ \vdots \\ f \\ \vdots \\ f(x) \end{array} \\ \xrightarrow{\quad} & & \xrightarrow{\quad} \\ y' & & y' \end{array} \quad \Rightarrow \quad \begin{array}{ccc} \begin{array}{c} x \\ \vdots \\ f \\ \vdots \\ f(x) \end{array} & \xrightarrow{\quad} & \begin{array}{c} y \\ \vdots \\ f \\ \vdots \\ y' \end{array} \\ \xrightarrow{\quad} & & \xrightarrow{\quad} \\ y' & & y' \end{array}$$

The asynchronous morphism f is called a **2-fibration** if, for all paths $p : x \rightarrow z$ of length 2 in G , and for all tiles $T' : f(p) \sim q'$ in G' , there exists a tile $T : p \sim q$ in G such that $f(q) = q'$ and $f(T) = f(T')$.



Such a transition system, which interpret a program C can be depicted as follows:




At each state, there are many Environment transitions, depicted in red. The Code transitions, the Code transitions, make the execution of the program progress. A path from an initial state on the left to a final state on the right will begin with any number of Frame transitions, followed by one Code transitions, followed by any number of Frame transitions, and so on.

We finally have a suitable definition of a transition system.

Definition 2.4.7. A transition system $(G, \lambda, G_{in}, G_{out})$ over the machine model $\mathfrak{A} \in \{\mathfrak{A}_S, \mathfrak{A}_L\}$ is the data of:


1. an asynchronous morphism $\lambda : G \rightarrow \mathfrak{A}^\bullet$, which is an Environment 1-fibration;
2. two sub-asynchronous graphs $G_{in}, G_{out} \subseteq G$ containing only Frame transitions denoting the initial and final states of the program.

This definition will be reformulated in a more categorical and convenient language in the next chapter. The basic idea is that a transition system as above can be seen as a cospan



in the category of asynchronous graphs. These asynchronous graphs will be organized to be horizontal morphisms in a certain double category.

Moreover, following Mellies methodology of *template games* (Mellies, 2019), but in a dual setting, with cospans instead of spans, the spans we consider are “typed” using *templates* \varkappa :



The purpose of the templates is that every operation which we will define on transition systems will be induced by structures at the level of templates. This has the practical advantage that the templates are simple objects, and so are easier to manipulate.

2.5 Concurrent separation logic

In the previous section, we have explained how to interpret programs as transition systems, and we have sketched the formal statement of the data-race freedom property we wish to prove about programs which have been proved using CSL.

In this section, we describe the version of CSL which we consider, and we explain how to interpret proofs, that is, proof trees, of CSL as transition systems over an appropriate machine model \varkappa_{Sep} . We conclude the section and the chapter by stating the asynchronous soundness theorem of the logic.

2.5.1 The logic

The version of CSL which we consider here is the same as the original papers by P. O’Hearn (2007) and Brookes (2004) as well as the first operational soundness proof, by Vafeiadis (2011). The main differences is that the version of CSL we consider has **permissions** which were introduced for separation logic by Bornat, Calcagno, P. O’Hearn, et al. (2005), following their introduction by Boyland (2003). The purpose of permissions allow the logic to prove programs where several threads read memory without synchronization. We also use a variant of the “variables as resources” discipline introduced by Bornat, Calcagno, and Yang (2006) to remove side conditions to the rules of the logic such as the inference rule for the parallel product and the frame rule.

The judgments of CSL are Hoare triples of the form

$$\Gamma \vdash \{P\} C \{Q\}$$

where P and Q are formulas of CSL, and $\Gamma = r_1 : J_1, \dots, r_n : J_n$ is a context associating a formula J_i to each lock r_i .

Formulas of CSL

Formulas denote predicates over a logical version of the memory augmented with permissions. Their syntax is the following:

$$\begin{aligned} P, Q, R, J ::= & \mathbf{emp} \mid P * Q \mid \top \mid \perp \mid P \vee Q \mid P \wedge Q \mid \neg P \\ & \mid \forall a. P \mid \exists a. P \mid v \stackrel{p}{\mapsto} w \mid \mathbf{Own}_p(x) \mid E'_1 = E'_2 \end{aligned}$$

where $x \in \mathbf{Var}$, $p \in \mathbf{Perm}$ is a permission, $v, w \in \mathbf{Val}$, and the E'_i are arithmetic expressions which can contain logic variables a, b, c . The idea behind permissions is that a permission $p \in \mathbf{Perm}$ represents a degree of ownership of a memory cell either in the heap through the atomic formula $x \stackrel{p}{\mapsto} v$ or in the stack with $\mathbf{Own}_p(x)$. Any permission is enough to *read* the memory, but only the **total permission** \top allows for writing to the memory. This leads to the following definition.

Definition 2.5.1. A **permission monoid** \mathbf{Perm} is a partial cancellative commutative semi-group with a distinguished element $\top \in \mathbf{Perm}$ called the **total permission** of \mathbf{Perm} which allows no multiple:

$$\forall x \in \mathbf{Perm}, x + \top \text{ is not defined.}$$

This implies in particular that \mathbf{Perm} has no neutral element. Cancellativity means that for all elements $x, y, z \in \mathbf{Perm}$, if $x + y = x + z$ (and both sides of the equality are well defined), then $y = z$.

Of course, such permission monoids are partial semi-groups, but this name is traditional.

The canonical example of **permission monoid** are the fractional permissions $(0, 1]$ of the additive monoid \mathbb{Q} restricted to values greater than 0 and less than or equal to 1. It is easy to check that 1 is a **total permission** since, for any $x > 0$, $x + 1$ is greater than 1 and thus not defined in the partial monoid of fractional permissions.

Since the formulas can specify the level of permission over pieces of memory, they cannot be interpreted as predicates over **memory states**. Instead, they are predicates over **logical states**.

$$\begin{aligned}
 (s, h) \models v \overset{p}{\mapsto} w &\Leftrightarrow v \in \mathbf{Loc} \wedge s = \emptyset \wedge h = [v \mapsto (w, p)] \\
 (s, h) \models \mathbf{Own}_p(x) &\Leftrightarrow \exists v \in \mathbf{Val}, s = [x \mapsto (v, p)] \wedge h = \emptyset \\
 \sigma \models \mathbf{emp} &\Leftrightarrow \sigma = (\emptyset, \emptyset) \\
 \sigma \models P * Q &\Leftrightarrow \exists \sigma_1 \sigma_2, \sigma = \sigma_1 * \sigma_2, \sigma_1 \models P \wedge \sigma_2 \models Q \\
 \sigma \models P \wedge Q &\Leftrightarrow \sigma \models P \text{ and } \sigma \models Q \\
 \sigma \models E_1 = E_2 &\Leftrightarrow \llbracket E_1 \rrbracket = \llbracket E_2 \rrbracket \wedge \text{fv}(E_1 = E_2) \subseteq \text{vdom}(\sigma) \\
 \sigma \models \exists a. P &\Leftrightarrow \exists v \in \mathbf{Val}, \sigma \models P[v/a]
 \end{aligned}$$

Figure 2.3: Satisfaction of CSL formulas

Definition 2.5.2. A *logical state* $\sigma \in \mathbf{LogState}$ is a pair (s, h) of a logical stack $s \in \mathbf{Var} \rightarrow_{fin} (\mathbf{Val} \times \mathbf{Perm})$ and a logical heap $h \in \mathbf{Loc} \rightarrow_{fin} (\mathbf{Val} \times \mathbf{Perm})$ which associate to each variable and each heap location a value and a permission.

Permissions for a resource can be split and combined:

$$x \overset{p+q}{\mapsto} v \dashv\vdash x \overset{p}{\mapsto} v * x \overset{q}{\mapsto} v$$

therefore, to interpret the separating conjunction, we need an operator to play the role of the disjoint union of heaps in a way that takes permissions into account. Reading the equivalence above tells us that we can combine logical states which contain the same values, and in that case the permissions are added (if the sum is well-defined of course).

Definition 2.5.3. Given two *logical states* σ_1 and σ_2 , their product $\sigma_1 * \sigma_2$, if it is defined, is a logical state whose domain is the union of the domains of σ_1 and of σ_2 . For $a \in \text{dom}(\sigma_1) \cup \text{dom}(\sigma_2) \subseteq \mathbf{Var} \uplus \mathbf{Loc}$,

$$(\sigma_1 * \sigma_2)(a) := \begin{cases} \sigma_1(a) & \text{if } a \in \text{dom}(\sigma_1) \setminus \text{dom}(\sigma_2) \\ \sigma_2(a) & \text{if } a \in \text{dom}(\sigma_2) \setminus \text{dom}(\sigma_1) \\ (v, p + p') & \text{if } \sigma_1(a) = (v, p) \text{ and } \sigma_2(a) = (v, p') \end{cases}$$

If any the right hand-side is undefined for any a , then the product $\sigma_1 * \sigma_2$ as a whole is undefined.

We define the satisfaction relation $\sigma \models P$ in Figure 2.3, using the operation above for the case of the separating conjunction. We write $\text{dom}(\sigma) \subseteq \mathbf{Var} \uplus \mathbf{Loc}$ for the complete domain of the state σ , and $\text{vdom}(\sigma)$ for the domain of its stack component s . The

formula $v \overset{p}{\mapsto} w$ is only satisfied for logical states of the form $\sigma = (\emptyset, [v \mapsto (w, p)])$, similarly, the formula $\text{Own}_p(x)$ is only satisfied by singleton stacks whose domains are $\{x\}$. The value of variables is specified using the predicates $E_1 = E_2$, which can only be true if the heap contains all the variables mentioned in E_1 and E_2 . Therefore, one can use the formula $\text{Own}_1(x) \wedge x = 5$ to state that x must contain the value 5. As expected, $P * Q$ holds for σ if the logical state σ can be split using the operation of Definition 2.5.3 in two logical states σ_1 and σ_2 which satisfy P and Q respectively.

Given the satisfaction relation, we can define the entailment of predicates as

$$P \Rightarrow Q \quad \Leftrightarrow \quad \forall \sigma, \sigma \models P \Rightarrow \sigma \models Q.$$

Inference rules of CSL

Hoare triples $\Gamma \vdash \{P\} C \{Q\}$ of CSL are constructed by building derivation trees, or proof trees, using the inference rules presented in Figure 2.4. The first premise of the rule IF , $P \Rightarrow \text{def}(B)$ states that if P holds for some **logical state** σ , then all the program variables appearing in B are owned with some permission.

The rules AFF for writing into variables and ST for writing into the heap require full permission over the respective memory resources; while we only require partial permission for reading, implicitly for variables—otherwise equalities involving expressions would be false—and explicitly in rule LD for heap locations. The other rules are standard and were explained in Section 0.1.2.

2.5.2 Semantic interpretation of CSL

Our method for proving the soundness of the logic we describe above is to interpret proof trees of CSL as transition systems over an adequate machine model. The first task is to define the notion of state which will be used by the machine model.

The predicates P which appear in the Hoare triples of the logic describe the logical memory which is owned by the Code. Since we have chosen to describe explicitly the actions of the Environment, or Frame, in our transition systems, our states need to contain the (logical) memory which is owned by the Frame.

This dichotomy does not account for all the memory: shared resources are owned by the locks, and therefore are owned by neither the Code nor the Frame. These states will keep track, for each lock, of the logical state it owns or the identity of its holder (Code or Frame). Putting all of this together, we define the notion of state which we will use to interpret proofs of CSL.

$$\begin{array}{c}
 \frac{}{\Gamma \vdash \{(\text{Own}_{\top}(x) * P) \wedge E = v\} x := E \{(\text{Own}_{\top}(x) * P) \wedge x = v\}} \text{AFF} \\
 \\
 \frac{}{\Gamma \vdash \{(P \wedge E = w \wedge E' = v) * w \overset{\top}{\mapsto} -\} [E] := E' \{(P \wedge E = w \wedge E' = v) * E \overset{\top}{\mapsto} w\}} \text{ST} \\
 \\
 \frac{P \Rightarrow \text{def}(B) \quad \Gamma \vdash \{P \wedge B\} C_1 \{Q\} \quad \Gamma \vdash \{P \wedge \neg B\} C_2 \{Q\}}{\Gamma \vdash \{P\} \text{if } B \text{ then } C_1 \text{ else } C_2 \{Q\}} \text{IF} \\
 \\
 \frac{P \Rightarrow \text{def}(B) \quad \Gamma \vdash \{I \wedge B\} C \{I\}}{\Gamma \vdash \{I\} \text{while } B \text{ do } C \{I \wedge \neg B\}} \text{WHILE} \\
 \\
 \frac{}{\Gamma \vdash \{((\text{Own}_{\top}(x) * P) \wedge E = w) * w \overset{p}{\mapsto} v\} x := [E] \{((\text{Own}_{\top}(x) * P) \wedge x = v) * w \overset{p}{\mapsto} v\}} \text{LD} \\
 \\
 \frac{\Gamma \vdash \{P\} C_1 \{Q\} \quad \Gamma \vdash \{Q\} C_2 \{R\}}{\Gamma \vdash \{P\} C_1; C_2 \{R\}} \text{SEQ} \qquad \frac{\Gamma \vdash \{P_1\} C \{Q_1\} \quad \Gamma \vdash \{P_2\} C \{Q_2\}}{\Gamma \vdash \{P_1 \vee P_2\} C \{Q_1 \vee Q_2\}} \text{DISJ} \\
 \\
 \frac{\Gamma, r : J \vdash \{P\} C \{Q\}}{\Gamma \vdash \{P * J\} \text{resource } r \text{ do } C \{Q * J\}} \text{RES} \qquad \frac{\Gamma \vdash \{(P * J) \wedge B\} C \{Q * J\}}{\Gamma, r : J \vdash \{P\} \text{with } r \text{ do } C \{Q\}} \text{WITH} \\
 \\
 \frac{\Gamma \vdash \{P_1\} C_1 \{Q_1\} \quad \Gamma \vdash \{P_2\} C_2 \{Q_2\}}{\Gamma \vdash \{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}} \text{PAR} \qquad \frac{\Gamma \vdash \{P\} C \{Q\}}{\Gamma \vdash \{P * R\} C \{Q * R\}} \text{FRAME} \\
 \\
 \frac{P \Rightarrow P' \quad \Gamma \vdash \{P'\} C \{Q'\} \quad Q' \Rightarrow Q}{\Gamma \vdash \{P\} C \{Q\}} \text{CONSEQ}
 \end{array}$$

Figure 2.4: Inference rules of CSL

Definition 2.5.4. Given a CSL context Γ , a *separated state* is a triple

$$(\sigma_C, \vec{\sigma}, \sigma_F) \in \mathbf{SepState}(\Gamma)$$

with $\sigma_C, \sigma_F \in \mathbf{LogState}$, and $\vec{\sigma} : \text{dom}(\Gamma) \rightarrow \mathbf{LogState} + \{\mathbf{C}, \mathbf{F}\}$, and which satisfies two properties:

1. Locked resources satisfy their invariants: $\forall r \in \text{unlocked}(\vec{\sigma}), \vec{\sigma}(r) \models \Gamma(r)$;
2. The components are coherent: the logical state below is defined:

$$\sigma_C * \left\{ \bigotimes_{r \in \text{dom}(\vec{\sigma})} \vec{\sigma}(r) \right\} * \sigma_F \in \mathbf{LogState} \quad (2.2)$$

A separated state $(\sigma_C, \vec{\sigma}, \sigma_F)$ defines as *machine state* $\otimes(\sigma_C, \vec{\sigma}, \sigma_F) = (\mu, L)$ whose memory state μ is obtained by forgetting permissions in the *logical state* (2.2) above, and whose set L of available locks is $\text{unlocked}(\vec{\sigma})$.

Unlike the two machine models \mathfrak{M}_S and \mathfrak{M}_L which we use to interpret programs, the machine model for the proofs distinguishes between transitions of the Code and of the Frame. It is defined as follows.

Definition 2.5.5. The *machine model of separated states* $\mathfrak{M}_{\text{Sep}}^{\bullet\bullet}(\Gamma)$ parameterized by a CSL context Γ is an *asynchronous graph* whose nodes are the *separated states* over the context Γ and with Code transitions

$$(\sigma_C, \vec{\sigma}, \sigma_F) \xrightarrow{m:\mathbf{C}} (\sigma'_C, \vec{\sigma}', \sigma_F)$$

where $m \in \mathbf{Instr}$ is an instruction, such that

$$\otimes(\sigma_C, \vec{\sigma}, \sigma_F) \xrightarrow{m} \otimes(\sigma'_C, \vec{\sigma}', \sigma_F)$$

is a transition in \mathfrak{M}_S^{\bullet} , and such that the following conditions are satisfied:

$$\begin{array}{ll} \forall \ell \notin \text{wr}(m), \sigma_C(\ell) = \sigma'_C(\ell) & \text{wr}(m) \cup \text{rd}(m) \subseteq \text{dom}(\sigma_C) \\ \text{lk}(m) \subseteq \text{dom}(\vec{\sigma}) \cup \vec{\sigma}^{-1}(\mathbf{C}) & \forall r \notin \text{lk}(m), \vec{\sigma}(r) = \vec{\sigma}'(r). \end{array}$$


Frame transitions are of the form $(\sigma_C, \vec{\sigma}, \sigma_F) \xrightarrow{m:\mathbf{F}} (\sigma_C, \vec{\sigma}', \sigma'_F)$ with symmetric conditions:

$$\otimes(\sigma_C, \vec{\sigma}, \sigma_F) \xrightarrow{m} \otimes(\sigma_C, \vec{\sigma}', \sigma'_F)$$

is a transition in \mathfrak{M}_S^{\bullet} and

$$\begin{array}{ll} \forall \ell \notin \text{wr}(m), \sigma_F(\ell) = \sigma'_F(\ell) & \text{wr}(m) \cup \text{rd}(m) \subseteq \text{dom}(\sigma_F) \\ \text{lk}(m) \subseteq \text{dom}(\vec{\sigma}) \cup \vec{\sigma}^{-1}(\mathbf{F}) & \forall r \notin \text{lk}(m), \vec{\sigma}(r) = \vec{\sigma}'(r). \end{array}$$

The **tiles** are defined in the same way as in the **stateful machine model**, using the same notion of footprint: a square in $\mathfrak{A}_{\text{Sep}}$ of the form:



has a unique tile if the footprint of m_1 and of m_2 computed at the **machine state** $\textcircled{*}a$ are **independent**.

This new machine model of separated states induces a notion of asynchronous transition system which can be used to interpret proofs trees of CSL by induction of the inference rules. A proof $\pi : \Gamma \vdash \{P\} C \{Q\}$ is interpreted as a pair $(G_\pi, \lambda_\pi, G_{\text{in}}, G_{\text{out}})$ where G_π is an **asynchronous graph**, and λ_π is an **asynchronous morphism**

$$\lambda_\pi : G_\pi \longrightarrow \mathfrak{A}_{\text{Sep}}^{\circ\bullet}$$


which maps all initial states in $G_{\text{in}} \subseteq G_\pi$ to separated states $(\sigma_C, \vec{\sigma}, \sigma_F)$ which satisfy the precondition P (in that $\sigma_C \models P$) and, symmetrically, all final states in $G_{\text{out}} \subseteq G_\pi$ are mapped to separated states which satisfy the postcondition Q .

In the same way that there is a map from the stateful machine model to the stateless machine model $\mathcal{L}_\pm : \mathfrak{A}_S^{\circ\bullet} \rightarrow \mathfrak{A}_L^{\circ\bullet}$, the operation $(\sigma_C, \vec{\sigma}, \sigma_F) \mapsto \textcircled{*}(\sigma_C, \vec{\sigma}, \sigma_F)$ defines an asynchronous morphism $\mathcal{S}_\pm : \mathfrak{A}_{\text{Sep}}^{\circ\bullet} \rightarrow \mathfrak{A}_S^{\circ\bullet}$.

As it turns out, we will see in the next chapter that, if π is a proof of the Hoare triple $\Gamma \vdash \{P\} C \{Q\}$, the interpretation of π and the two interpretations of the program C are related as follows, completing the diagram (2.1) page 50:

$$\begin{array}{ccccc}
 \llbracket \pi \rrbracket_{\text{Sep}} & G_\pi & \xrightarrow{\lambda_\pi} & \mathfrak{A}_{\text{Sep}} & \\
 \mathcal{S} \downarrow & \mathcal{S} \downarrow & & \mathcal{S}_\pm \downarrow & \\
 \llbracket C \rrbracket_S & G_S & \xrightarrow{\lambda_S} & \mathfrak{A}_S & \\
 \mathcal{L} \downarrow & \mathcal{L} \downarrow & & \mathcal{L}_\pm \downarrow & \\
 \llbracket C \rrbracket_L & G_L & \xrightarrow{\lambda_L} & \mathfrak{A}_L &
 \end{array}$$

The map \mathcal{S} induces a sub-asynchronous graph of $\llbracket C \rrbracket_S$ of the executions which are supported by the CSL proof:



This sub-asynchronous graph cannot be characterized at the level of the stateful semantics, it is only makes sense seen as the paths of $\llbracket C \rrbracket_S$ which are refined by the proof π . This is the basis of the asynchronous soundness theorem of CSL.

2.5.3 The asynchronous soundness theorem

The soundness theorem of CSL states that programs which have been proved with the logic are safe, satisfy their postconditions if they terminate, and do not produce any data-race. In our setting, these properties are expressed as properties of the maps

$$\llbracket \pi \rrbracket_{\text{Sep}} \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_S \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_L$$

which relate the interpretation $\llbracket \pi \rrbracket_{\text{Sep}}$ of a proof π of the Hoare triple $\Gamma \vdash \{P\} C \{Q\}$ and the stateful $\llbracket C \rrbracket_S$ and stateless $\llbracket C \rrbracket_L$ interpretations of the program C . This property is expressed using the language of fibrations, which we defined in Definition 2.4.6. First, the safety of the program is expressed as a property about \mathcal{S} :

Theorem 2.5.6. *The asynchronous morphism \mathcal{S} is a Code 1-fibration.*

Unfolding the definitions, this means that for all Code transitions

$$\mathfrak{s}_1 \xrightarrow{m} \mathfrak{s}_2$$

in $\llbracket C \rrbracket_S$, if there is a separated state $(\sigma_C, \vec{\sigma}, \sigma_F)$ which corresponds to the machine state \mathfrak{s}_1 thought \mathcal{S} , then there exists a transition

$$(\sigma_C, \vec{\sigma}, \mathfrak{S}_F) \xrightarrow{m} (\sigma'_C, \vec{\sigma}', \mathfrak{S}_F)$$

with $\odot(\sigma'_C, \vec{\sigma}', \mathcal{S}_F) = \varepsilon_2$. In particular, since the image of the \odot operation is always a well-defined machine state and not the error state \downarrow , this means that the operation m did not fail.

There are two ways to understand why this implies the safety of the whole program. The first is that all Code transitions preserve the safety of the program, and that it is up to the Environment to make sure that its transitions follow the discipline of CSL and preserve the correctness of the program. The other point of view is that, given a closed and complete program, one can forbid actions from the Environment. Then the asynchronous morphism \mathcal{S} guarantees that any path which starts in a machine state which corresponds to a separated state which satisfies the precondition P does not contain the error state \downarrow .

Now, let us state the second part of the soundness theorem, which states that there are no data-races in well-specified programs.

Theorem 2.5.7. *The asynchronous morphism $\mathcal{L} \circ \mathcal{S}$ is a 2-fibration.*

This means that whenever there is a tile in the stateless semantics $\llbracket C \rrbracket_L$ of the program C , meaning that two instructions are executed in parallel and do not synchronize between each other, if one of the edges of the tile corresponds to a path in the semantics of the proof $\llbracket \pi \rrbracket_{\text{Sep}}$, which means that this path corresponds to a piece of execution which follows the memory management rules of CSL, then there is a tile completing the path in $\llbracket \pi \rrbracket_{\text{Sep}}$.

The asynchronous morphism \mathcal{S} maps this tile in $\llbracket \pi \rrbracket_{\text{Sep}}$ into a tile in $\llbracket C \rrbracket_S$ which, by definition of tiles in \oplus_S , means that there is no data-race between the two instructions.

3 An asynchronous template game model of CSL

Our semantic interpretation of CSL is based on the methodology of *templates* (Melliès, 2019). The general idea is that the interpretations live in a semantic domain which is parameterized by a template \pm . In the work of Melliès (2019) for example, a template \pm is defined to be an *internal category*, which induces a bicategory of games $\mathbf{Games}(\pm)$ based on spans.

The slogan of this approach is that every operation at the level of the interpretation, say tensor products, should be induced by a *structure* at the level of templates, for tensors: span-monoidal internal categories in (Melliès, 2019).

In this work, our notion of template is almost dual, as they are a generalization of *internal opcategories*, and programs and proofs of CSL are interpreted as *cobordisms*, which are based on cospans.

3.1 The double category $\mathbf{Cob}(\pm)$ of games and cobordisms


3.1.1 Internal categories

Internal categories are the result of spelling out the definition of a category in the language of categories, in the same way that one can define a monoid object or a group object in any category with enough structures. We will use this notion in two ways: First, double categories are by definition weakly internal categories in the 2-category \mathbf{Cat} of categories; and, second, they will be generalized and dualized in the sequel to define *internal J -opcategories*, the notion of *template* which our model will be based on.

To a first approximation, a category is a reflexive graph (or, more precisely, quiver): a set $\mathcal{C}[0]$ of objects, and a set $\mathcal{C}[1]$ of morphisms. Each morphism has a source and a target: this defines two maps $s, t : \mathcal{C}[1] \rightarrow \mathcal{C}[0]$. Similarly, each object $A \in \mathcal{C}[0]$ is

associated with an identity morphism $\text{Id}_A : A \rightarrow A$, this defines a map $e : \mathcal{C}[0] \rightarrow \mathcal{C}[1]$. One then needs to express that, for instance, the source and target of the identity on A are equal to A , and to formalize the composition operator. This leads to the following definition, discovered by Ehresmann (1963).

Definition 3.1.1. We consider a category \mathbb{S} with finite limits. An *internal category* \mathcal{C} in \mathbb{S} contains the following data. Two objects $\mathcal{C}[0]$ and $\mathcal{C}[1]$ of \mathbb{S} which we call the space of objects and the space of morphisms of \mathcal{C} respectively. Morphisms $s, t : \mathcal{C}[1] \rightarrow \mathcal{C}[0]$ and $e : \mathcal{C}[0] \rightarrow \mathcal{C}[1]$ as explained above. A composition morphism $m : \mathcal{C}[2] \rightarrow \mathcal{C}[1]$ which associates morphisms to composable pairs of morphisms defined as the pullback



In the category of sets, this corresponds to pairs of arrows such that the target of the first is equal to the source of the second.

This structure $\mathcal{C} = (\mathcal{C}[0], \mathcal{C}[1], s, t, e, m)$ must satisfy certain properties:

1. The identities have the right endpoints: $s \circ e = t \circ e = \text{Id}_{\mathcal{C}[0]}$;
2. Morphism composition preserves endpoints: $s \circ l = s \circ m$ and $t \circ r = t \circ m$;
3. Identities are identities for composition: the following diagram commutes

$$\begin{array}{ccccc}
 \mathcal{C}[0] \times_{\mathcal{C}[0]} \mathcal{C}[1] & \xrightarrow{e \times_{\mathcal{C}[0]} \text{Id}} & \mathcal{C}[2] & \xleftarrow{\text{Id} \times_{\mathcal{C}[0]} e} & \mathcal{C}[1] \times_{\mathcal{C}[0]} \mathcal{C}[0] \\
 & \searrow \pi_2 & \downarrow m & & \swarrow \pi_1 \\
 & & \mathcal{C}[1] & &
 \end{array} \tag{3.1}$$

4. and composition must be associative: writing $\mathcal{C}[3] = \mathcal{C}[1] \times_{\mathcal{C}[0]} \mathcal{C}[1] \times_{\mathcal{C}[0]} \mathcal{C}[1]$ the space of triples of composable arrows, the following diagram must commute:

$$\begin{array}{ccc}
 \mathcal{C}[3] & \xrightarrow{m \times_{\mathcal{C}[0]} \text{Id}} & \mathcal{C}[2] \\
 \text{Id} \times_{\mathcal{C}[0]} m \downarrow & & \downarrow m \\
 \mathcal{C}[2] & \xrightarrow{m} & \mathcal{C}[1]
 \end{array} \tag{3.2}$$

As expected, a (small) category is an internal category in the category **Set** of sets and functions. We ask that the ambient category \mathbb{S} has finite limits in the definition above as this is general enough for our purposes, in general only the pullbacks of s along t needs to exist. The instance we are most interested in are **strict double categories**, which are categories internal to **Cat**. More precisely, we need the weaker notion of (*pseudo*) double category, where the composition operator $m : \mathcal{C}[2] \rightarrow \mathcal{C}[1]$ is only associative up to an isomorphism. Accordingly, we define a **weakly internal category** in a 2-category \mathbb{S} with finite limits to be the same as an internal category except that diagrams (3.1) and (3.2) only needs to commute up to isos which satisfy the pentagon and triangle identities of monoidal categories. See Appendix A.2 of the PhD thesis of Courser (2020) for the full definitions of double categories and related objects.

A nice characterization of internal categories, which is going to be useful to generalize internal categories to **polyads**, is the following:

Lemma 3.1.2. *An internal category in \mathbb{S} is the same as a monad in the bicategory $\mathbf{Span}(\mathbb{S})$ of spans in \mathbb{S} . Equivalently, this is also the same as a lax pseudo-functor $\mathbf{1} \rightarrow \mathbf{Span}(\mathbb{S})$, where $\mathbf{1}$ is the terminal bicategory.*

In the same way the definition of a category was expressed in the language of categories, there is an internal notion of functor between internal categories:

Definition 3.1.3. An **internal functor of categories** F from \mathcal{C} to \mathcal{D} , two internal categories in a category \mathbb{S} with finite limits, is the data of two morphisms in \mathbb{S} between the two spaces of objects and the two spaces to morphisms:

$$F[0] : \mathcal{C}[0] \rightarrow \mathcal{D}[0] \quad \text{and} \quad F[1] : \mathcal{C}[1] \rightarrow \mathcal{D}[1]$$

satisfying the expected conditions:

1. F acts coherently with respect to the source and target maps, in that the following diagram commutes:

$$\begin{array}{ccccc} \mathcal{C}[0] & \xleftarrow{s} & \mathcal{C}[1] & \xrightarrow{t} & \mathcal{C}[0] \\ F[0] \downarrow & & \downarrow F[1] & & \downarrow F[0] \\ \mathcal{D}[0] & \xleftarrow{s} & \mathcal{D}[1] & \xrightarrow{t} & \mathcal{D}[0] \end{array}$$

2. The commutation of the diagram above induces a map $F[2] : \mathcal{C}[2] \rightarrow \mathcal{D}[2]$ by the universality of the pullback defining $\mathcal{D}[2]$. The following diagram must commute,

expressing the fact that the composite of the images is equal to the image of the composite and that F preserves identity maps:

$$\begin{array}{ccccc}
 \mathcal{C}[2] & \xrightarrow{m} & \mathcal{C}[1] & \xleftarrow{e} & \mathcal{C}[0] \\
 F[2] \downarrow & & \downarrow F[1] & & \downarrow F[0] \\
 \mathcal{D}[0] & \xrightarrow{m} & \mathcal{D}[1] & \xleftarrow{e} & \mathcal{D}[0]
 \end{array}$$

3.1.2 Double categories

Double categories are a notion of two-dimensional categories introduced by Ehresmann (1963). Unlike 2-categories and bicategories, double categories have two distinct classes of 1-cells, called vertical and horizontal. This notion is central in this work, as the programs and the proof tree of CSL will be interpreted as horizontal morphisms in a double category.

Definition 3.1.4. A *double category* \mathcal{D} is a *weakly internal category* in \mathbf{Cat} , the 2-category of small categories. Let us spell it out. The double category \mathcal{D} is a tuple $\mathcal{D} = (\mathcal{D}[0], \mathcal{D}[1], s, t, e, m)$ which contains two categories $\mathcal{D}[0]$ and $\mathcal{D}[1]$. The objects of $\mathcal{D}[0]$ are called the objects of the double category \mathcal{D} , and its maps are called the *vertical morphisms* of \mathcal{D} and are denoted with plain arrows $A \rightarrow B$. Each object of $\mathcal{D}[1]$ is associated with a source and a target object in $\mathcal{D}[0]$ through the two functors s and t , and is called a *horizontal morphism*. They are denoted with crossed arrows $A \dashrightarrow B$. Finally, a morphism α in $\mathcal{D}[1]$ is called a *2-cell* of the double category \mathcal{D} . The 2-cell α fills a square of two horizontal and two vertical maps of the form:


$$\begin{array}{ccc}
 A_1 & \dashrightarrow^F & B_1 \\
 f \downarrow & \Downarrow \alpha & \downarrow g \\
 A_2 & \dashrightarrow^G & B_2
 \end{array}$$

A 2-cell as above is called *special* if the two vertical morphisms f and g are identities. Composition of horizontal morphisms is associative up to coherent special two-cells.


Most examples of double categories have a similar flavor: *vertical morphisms* are the usual notion of map for the objects under consideration, and the horizontal morphisms correspond to relation-like mappings between objects. For example, there is a double category where the horizontal morphisms between rings A and B are the AB -bimodules,

while vertical morphisms are simply ring homomorphisms. In other words, the underlying vertical category ($\mathcal{D}[0]$ in the notation of the definition) is the usual category of rings. This intuition is captured by the notion of *framed bicategory* introduced by Shulman (2008).


Another important family of example of double categories which follow the same principle are the double categories $\mathbf{Cospan}(\mathbb{S})$, parameterized by a category \mathbb{S} with pushouts. The object and vertical morphisms are the same as \mathbb{S} , and horizontal morphisms between objects A and B of \mathbb{S} are cospans:



The maps l and r are called its legs, and S is its support. We sometimes write $S : A \twoheadrightarrow B$ when the legs are clear from context. The composite of two cospans $S : A \twoheadrightarrow B$ and $T : B \twoheadrightarrow C$ is defined using the pushout of the right leg of S and the left leg of T as depicted in this diagram:



The 2-cells between horizontal maps $S : A \twoheadrightarrow B$ and $T : C \twoheadrightarrow D$ and vertical maps $f : A \rightarrow C$ and $g : B \rightarrow D$ are the maps $\alpha : S \rightarrow T$ which make the following diagram commute:



Of course, there is a similar double category $\mathbf{Span}(\mathbb{S})$ of spans when the category \mathbb{S} has all pullbacks. We can recover the more common *bicategory* of cospans if we only consider *special 2-cells*; this comes from a general fact:

Lemma 3.1.5. *Any double category \mathcal{D} induces a bicategory \mathbf{HD} whose objects are the same as \mathcal{D} 's, whose 1-cells are the horizontal morphisms of \mathcal{D} , and whose 2-cells are the special 2-cells of \mathcal{D} .*

There is a natural notion of map between **double categories**:

Definition 3.1.6. A **double functor** \mathcal{F} between two **double categories** \mathcal{C} and \mathcal{D} is an **internal functor** between them. Explicitly, \mathcal{F} maps objects, vertical 1-cells and horizontal 1-cells and 2-cells of the double category \mathcal{C} to the same type of cells of the double category \mathcal{D} while preserving source and targets, as expected.

Moreover, for two composable vertical morphisms f, g in \mathcal{C} , $\mathcal{F}(g \circ f) = \mathcal{F}(g) \circ \mathcal{F}(f)$ in \mathcal{D} , and, for two composable horizontal morphisms F, G in \mathcal{C} , $\mathcal{F}(G \circ F)$ and $\mathcal{F}(G) \circ \mathcal{F}(F)$ are related by an invertible **special 2-cell**. Finally, \mathcal{F} preserves vertical identities on the nose and vertical identities up to invertible special 2-cells.

A weaker notion which will be useful is the following

Definition 3.1.7. A **lax double functor** \mathcal{F} between two **double categories** \mathcal{C} and \mathcal{D} is the same as a **double functor** except that the **special 2-cells** which relates $\mathcal{F}(G \circ F)$ and $\mathcal{F}(G) \circ \mathcal{F}(F)$

$$\begin{array}{ccccc} \mathcal{F}(A) & \xrightarrow{\mathcal{F}F} & \mathcal{F}(B) & \xrightarrow{\mathcal{F}G} & \mathcal{F}(A) \\ \parallel & & \Downarrow \alpha & & \parallel \\ \mathcal{F}(A) & \xrightarrow{\quad\quad\quad} & \mathcal{F}(C) & \xrightarrow{\quad\quad\quad} & \mathcal{F}(C) \end{array}$$

is not necessarily invertible. Again, we refer the reader to Appendix A.2 of (Courser, 2020) for the full definition.

3.1.3 Polyads and internal J -opcategories

Polyads are a generalization of monads with several objects which were introduced by Bénabou (1967). They are used to define the notion of **internal J -opcategory**, which will be the basic notion of **template** in this work.

The forgetful functor $| - | : \mathbf{Cat} \rightarrow \mathbf{Set}$ which transports every small category to its set of objects has a right adjoint **chaos** : $\mathbf{Set} \rightarrow \mathbf{Cat}$ which transports every set J to its chaotic category defined as the category **chaos**(J) whose objects are the elements i, j, k of the set J , with a unique morphism between each pair of objects.

Definition 3.1.8. Given a set J , a **J -polyad** in a double category \mathcal{D} is a lax double functor

$$\mathbf{chaos}(J) \rightarrow \mathcal{D}$$


from the chaotic category over the set J , seen as a double category whose horizontal morphisms are the morphisms of **chaos**(J), to \mathcal{D} . The elements of J are called the colors of the polyad. A polyad is a J -polyad for some set J .

Polyads generalize monads as a monad is simply a $\{*\}$ -polyad. This leads us, inspired by Lemma 3.1.2, to define an **internal J -opcategory** to be a polyad in the double category $\mathbf{Cospan}(\mathbb{S})$. They are, in a sense, dual to internal categories in that they are based on cospans instead of spans; they also are a generalization to multiple objects.

The definition can be expounded as follows. An internal J -opcategory \multimap consists of an object $\multimap[0, i]$ of \mathbb{S} for each element $i \in J$, and of a cospan

$$\multimap[1, ij] : \multimap[0, i] \multimap \multimap[0, j] = \multimap[0, i] \xrightarrow{\text{in}_{ij}} \multimap[1, ij] \xleftarrow{\text{out}_{ij}} \multimap[0, j]$$


for each pair $i, j \in J$, together with two coherent families η and μ of **special 2-cells** between cospans:



In particular, the left condition implies that the cospans $\multimap[1, ii]$ are of the form

$$\multimap[0, i] \xrightarrow{\eta_i} \multimap[1, ii] \xleftarrow{\eta_i} \multimap[0, i]$$

More explicitly, μ_{ijk} is given by the map g_{ijk} in the following commutative diagram in the category \mathbb{S} :



Finally, there are some conditions about the associativity of the composition, and about the identity; they can be found in (Bénabou, 1967, def. 5.5.1). An **internal opcategory** \multimap is defined as an internal $\{*\}$ -opcategory. We write in that case $\multimap[0]$ and $\multimap[1]$ for the objects $\multimap[0, *]$ and $\multimap[1, **]$ of the ambient category \mathbb{S} , respectively.

The notion of internal J -opcategory comes with a natural notion of morphism between them: internal functors of J -opcategories:

Definition 3.1.9. An *internal functor* $F = (f, F[0, \cdot], F[1, \cdot])$ from \mathfrak{A} to \mathfrak{A}' is a triple consisting of a function $f : J \rightarrow J'$ between the sets of colors, and:

1. for each $i \in J$, a map $F[0, i]$ in the ambient category \mathbb{S} :

$$F[0, i] : \mathfrak{A}[0, i] \rightarrow \mathfrak{A}'[0, f(i)],$$

2. for each $i, j \in J$, a map $F[1, ij]$ in the ambient category \mathbb{S} :

$$F[1, ij] : \mathfrak{A}[1, ij] \rightarrow \mathfrak{A}'[1, f(ij)]$$

where we use the lighter notation $f(ij)$ for $f(i)f(j)$. One asks moreover that the diagram below commutes

$$\begin{array}{ccccc} \mathfrak{A}[0, i] & \xrightarrow{\text{in}_{ij}} & \mathfrak{A}[1, ij] & \xleftarrow{\text{out}_{ij}} & \mathfrak{A}[0, j] \\ F[0, i] \downarrow & & \downarrow F[1, ij] & & \downarrow F[0, j] \\ \mathfrak{A}'[0, f(i)] & \xrightarrow{\text{in}_{f(i)f(j)}} & \mathfrak{A}'[1, f(ij)] & \xleftarrow{\text{out}_{f(i)f(j)}} & \mathfrak{A}'[0, f(j)] \end{array}$$

and that the internal functor is compatible with the identities: the diagram on the right below commutes; and with composition: the maps $F[2, ijk] : \mathfrak{A}[2, ijk] \rightarrow \mathfrak{A}'[2, f(ijk)]$ induced by universality of the pushout must make the left diagram commute.

$$\begin{array}{ccc} \mathfrak{A}[2, ijk] & \longrightarrow & \mathfrak{A}[1, ik] \\ F[2, ijk] \downarrow & & \downarrow F[1, ik] \\ \mathfrak{A}'[2, f(ijk)] & \longrightarrow & \mathfrak{A}'[1, f(ik)] \end{array} \qquad \begin{array}{ccc} \mathfrak{A}[0, i] & \xrightarrow{\eta_i} & \mathfrak{A}[1, ii] \\ F[0, i] \downarrow & & \downarrow F[1, ii] \\ \mathfrak{A}[0, f(i)] & \xrightarrow{\eta'_{f(i)}} & \mathfrak{A}'[1, f(ii)] \end{array}$$

The internal J -opcategories and internal functors form a category $\mathbf{opCat}(\mathbb{S})$. This category admits products: the index set of the product of two internal categories is the product of their index sets, and $(\mathfrak{A} \times \mathfrak{A}')[0, (i, j)] = \mathfrak{A}[0, i] \times \mathfrak{A}'[0, j]$. This fact will play an important role in the sequel.


3.1.4 The double category $\mathbf{Cob}(\mathfrak{A})$ of games and cobordisms

Suppose given an internal J -opcategory \mathfrak{A} in a category \mathbb{S} with pushouts. A *j -colored game* (A, λ_A) is defined as an object A of \mathbb{S} equipped with a morphism $\lambda_A : A \rightarrow \mathfrak{A}[0, j]$. An *ij -colored cobordism* from an i -colored game (A, λ_A) to a j -colored game (B, λ_B)

is defined as a cospan in \mathcal{C} together with colors $i, j \in J$ and a map λ_σ such that the following diagram commutes:

$$\begin{array}{ccccc}
 A & \xrightarrow{\text{in}} & S & \xleftarrow{\text{out}} & B \\
 \lambda_A \downarrow & & \downarrow \lambda_\sigma & & \downarrow \lambda_B \\
 \mathfrak{z}[0, i] & \xrightarrow{\text{in}_{ij}} & \mathfrak{z}[1, ij] & \xleftarrow{\text{out}_{ij}} & \mathfrak{z}[0, j]
 \end{array}$$

We think of a cobordism as a space S of all executions of a program, visualized as a topological space, with two “borders” which are interpreted as the initial (or input) states and the final (or output) states of the program. The cospan at the level of templates is a way of typing, or coloring this object. They are depicted like this:



Two such ij -colored cobordism $\sigma : A \rightarrow B$ and jk -colored cobordism $\tau : B \rightarrow C$ can be composed into an ik -colored cobordism $\sigma; \tau : A \rightarrow C$ using a pushout and a relabeling along the 2-cell μ_{ijk} provided by the internal J -opcategory, as shown below:

$$\begin{array}{ccccc}
 A & \xrightarrow{\quad C \quad} & B & \xrightarrow{\quad C' \quad} & C \\
 \downarrow & \Downarrow & \downarrow & \Downarrow & \downarrow \\
 \mathfrak{z}[0, i] & \xrightarrow{\quad \mathfrak{z}[1, ij] \quad} & \mathfrak{z}[0, j] & \xrightarrow{\quad \mathfrak{z}[1, jk] \quad} & \mathfrak{z}[0, i] \\
 & \searrow & \downarrow \mu_{ijk} & \nearrow & \\
 & & \mathfrak{z}[1, ik] & &
 \end{array}$$

In order to give cobordisms over an internal J -opcategory \mathfrak{z} a structure of double category, we need identities, which are given by:

$$\begin{array}{ccc}
 A & \xrightarrow{\text{Id}_A} & A \\
 \lambda_A \downarrow & \Downarrow \lambda_A & \downarrow \lambda_A \\
 \mathfrak{z}[0, i] & \xrightarrow{\text{Id}_{\mathfrak{z}[0, i]}} & \mathfrak{z}[0, i] \\
 & \searrow & \downarrow \eta_i \\
 & & \mathfrak{z}[1, ii]
 \end{array}$$

The vertical morphisms are the maps $f : A \rightarrow A'$ such that $\lambda_A = \lambda_B \circ f$, and the two-cells are triples of maps $f_{\text{in}} : A \rightarrow A'$, $f_{\text{out}} : B \rightarrow B'$, $f : S \rightarrow S'$ such that the following diagram commutes

$$\begin{array}{ccccc}
 A & \xrightarrow{\text{inc}} & S & \xleftarrow{\text{outc}} & B \\
 f_{\text{in}} \downarrow & & \downarrow f & & \downarrow f_{\text{out}} \\
 A' & \xrightarrow{\text{inc}'_C} & S' & \xleftarrow{\text{out}'_C} & B'
 \end{array}$$

such that moreover, f_{in} and f_{out} are vertical morphisms, and similarly $\lambda_{\sigma} = \lambda_{\sigma'} \circ f$. One obtains:

Theorem 3.1.10. *Every internal J -opcategory \mathfrak{A} induces a double category $\mathbf{Cob}(\mathfrak{A})$ whose objects are the j -colored games and whose horizontal maps are the cobordisms with composition defined as above.*

3.2 Three internal J -opcategories: \mathfrak{A}_L , \mathfrak{A}_S , $\mathfrak{A}_{\text{Sep}}$

We instantiate the framework described above to define the three J -opcategories in the category $\mathbb{S} = \mathbf{AsyncGraph}$ of asynchronous graphs which we will use to give a semantics to the concurrent shared memory language we describe in Section 2.3 and to the proofs of CSL, described in Section 2.5.

They are based on the machine models \mathfrak{A}_L^\bullet , \mathfrak{A}_S^\bullet , and $\mathfrak{A}_{\text{Sep}}^\bullet$ which we defined in Chapter 2. The machine models describe the basic operations supported by the underlying machine, and how they interact when they are executed in parallel. The cobordisms over these internal J -opcategories will contain the information of how programs cause these instructions to be executed and how they react to changes in the global memory state.

3.2.1 The internal opcategories \mathfrak{A}_L and \mathfrak{A}_S for the code

Recall that both the stateless machine model \mathfrak{A}_L^\bullet and the stateful machine model \mathfrak{A}_S^\bullet are parameterized by a set \mathcal{L} of allocated locks. In this section we assume given such a set and we do not write it explicitly.

The construction of the internal opcategories from the machine models is the same for both the *stateless internal opcategory* \mathfrak{A}_L and the *stateful internal opcategory* \mathfrak{A}_S . We write \mathfrak{A}^\bullet for either \mathfrak{A}_L^\bullet or \mathfrak{A}_S^\bullet and we write \mathfrak{A} for either \mathfrak{A}_L or \mathfrak{A}_S , respectively.

Recall that defining an internal opcategory \mathfrak{z} comes down to defining a cospan in the ambient category **AsyncGraph**. In our case, it will be of the form:

$$\mathfrak{z}[0] \begin{array}{c} \swarrow \iota \\ \searrow \iota \end{array} \mathfrak{z}[1]$$

That is, both its legs are equal and are monomorphisms.

We construct an asynchronous graph $\mathfrak{z}[1]$ with two players (for Code and Environment) from the asynchronous graph \mathfrak{z}^\bullet which we see as having a single player. This shift from one player to two players is performed in a very simple way. We consider the functor $\Omega : \mathbf{Set} \rightarrow \mathbf{AsyncGraph}$ which transports a given set \mathcal{L} of labels to the asynchronous graph $\Omega(\mathcal{L})$ with one single node, the elements of \mathcal{L} as edges, and one tile for each square. We are particularly interested in the case when the set of labels $\mathcal{L} = \{\mathbf{C}, \mathbf{F}\}$ contains the two polarities **C** and **F** associated with the Code and to the Frame (or the Environment), respectively. Note that the functor preserves limits: in particular, the asynchronous graph $\Omega(1)$ is the terminal object of **AsyncGraph**. The asynchronous graph $\Omega(\{\mathbf{C}, \mathbf{F}\})$ enables us to define the *two-player* machine models $\mathfrak{z}^{\circ\bullet}$ as the Cartesian product of asynchronous graphs

$$\mathfrak{z}^{\circ\bullet} = \mathfrak{z}^\bullet \times \Omega(\{\mathbf{C}, \mathbf{F}\}).$$

The resulting asynchronous graph $\mathfrak{z}^{\circ\bullet}$ has the same nodes as \mathfrak{z}^\bullet and contains two edges

$$x \xrightarrow{m:\mathbf{C}} y \qquad x \xrightarrow{m:\mathbf{F}} y$$

the first one labeled with a polarity **C** for Code, the second one labeled with a polarity **F** for Frame, for each edge

$$x \xrightarrow{m} y$$

in the original one-player, or apolar, asynchronous graph \mathfrak{z}^\bullet . The two circles \circ and \bullet in the notation $\mathfrak{z}^{\circ\bullet}$ are mnemonics designed to remind us that there are two players in the game: the Code playing the white side (\circ) and the Environment playing the black side (\bullet).

We have accumulated enough material at this stage to define the internal opcategories $\mathfrak{z} = \mathfrak{z}_S, \mathfrak{z}_L$ in the ambient category **AsyncGraph**. The asynchronous graph $\mathfrak{z}[1]$ is defined as the two-player machine model $\mathfrak{z}^{\circ\bullet}$ while the asynchronous graph $\mathfrak{z}[0]$ is defined as the one-player machine model, which we see here as played only by the Frame. In summary:

$$\mathfrak{z}[1] = \mathfrak{z}^{\circ\bullet} \qquad \mathfrak{z}[0] = \mathfrak{z}^\bullet.$$


The asynchronous graph $\mathfrak{z}[0]$ with one player is then embedded in the two-player asynchronous graph $\mathfrak{z}[1]$ by the **asynchronous morphism** $\iota = \mathfrak{z}[0] \rightarrow \mathfrak{z}[1]$ obtained

by transporting every node of $\mathfrak{A}[0]$ to the corresponding node in $\mathfrak{A}[1]$, and every edge of $\mathfrak{A}[0]$ to the corresponding edge in $\mathfrak{A}[1]$ with the polarity \mathbf{F} of the Frame.

3.2.2 The internal J -opcategory $\mathfrak{A}_{\text{Sep}}$ for the proofs

The case of the separated state internal J -opcategory differs from the two internal opcategories we have just defined in two ways: First, the underlying **machine model of separated states** already has two players because the notion of state contains distinguishes several the area owned by the Code from the area owned by the Environment.

Second, the set J of colors is not a singleton: it is the set of predicates of the logic. The reason is that a proof π of a Hoare triple $\Gamma \vdash \{P\} C \{Q\}$ has a meaning only when, in the initial state, the memory owned by the Code satisfies the precondition P , and dually it guarantees that the final state satisfies Q . This typing information is contained in the space of objects of the internal J -opcategory of a cobordism interpreting the proof π :



This expresses that the initial (internal) states of the proof, represented by the asynchronous graph I have (external) states which satisfy the predicate P , and similarly for final states and the postcondition Q .

Following this intuition, the nodes of the asynchronous graph $\mathfrak{A}_{\text{Sep}}[0, P]$ are all the separated states whose Code-components satisfy the predicate P :

$$(\sigma_C, \vec{\sigma}, \sigma_F) \in \mathfrak{A}_{\text{Sep}}[0, P] \quad \Leftrightarrow \quad \sigma_C \models P.$$

As in the previous section for the internal opcategories for the code, we omit the parameter Γ . The transitions in $\mathfrak{A}_{\text{Sep}}[0, P]$ between such states are those of $\mathfrak{A}_{\text{Sep}}^{\circ}$ played by Frame.

The asynchronous graph $\mathfrak{A}_{\text{Sep}}[1]$ is simply defined to be $\mathfrak{A}_{\text{Sep}}^{\circ}$, and the maps

$$\mathfrak{A}_{\text{Sep}}[0, P] \hookrightarrow \mathfrak{A}_{\text{Sep}}[1]$$

are the obvious inclusions.

3.3 Parallel product

The double categories $\mathbf{Cob}(\pm)$, where \pm is one of the three internal opcategories \pm_L , \pm_S or \pm_{sep} , need additional structures to interpret the constructions of the programming language we consider and the corresponding rules of CSL. We begin with the parallel product.

It would be natural to interpret it as a monoidal product \parallel in the model. However, given maps $C_i : A_i \rightarrow B_i$ and $D_i : B_i \rightarrow E_i$, for $i = 1, 2$, there would have to be an invertible **special 2-cell**:

$$(C_1 \parallel C_2); (D_1 \parallel D_2) \Rightarrow (C_1; D_1) \parallel (C_2; D_2) \quad (3.3)$$

between two **horizontal morphisms** $A_1 \parallel A_2 \rightarrow E_1 \parallel E_2$. This contradicts the well-known fact that in general the concatenation of interleavings are *included* in the interleaving of the concatenations of traces, expressed by the so-called Hoare inequality:

$$(C_1 \parallel C_2); (D_1 \parallel D_2) \subseteq (C_1; D_1) \parallel (C_2; D_2).$$


Hence, the right structure to interpret the parallel product is that of *lax monoidal products*, where the 2-cell (3.3) is not necessarily invertible.

As for the other constructions on cobordisms, this lax monoidal product is going to be induced by a structure at the level of templates: **span monoidal internal J -opcategories**. In order to define them, we introduce **plain internal functors**, which are **internal functors of J -opcategories** which enjoy the property that one can pull along them.

3.3.1 Plain internal functors

An important family of internal functors are **plain internal functors**, which are functors whose action on the colors is the identity. With the notations of Definition 3.1.9, they are the **internal functors of J -opcategories** such that $f = \text{Id}$. Plain internal functors $F : \pm_1 \rightarrow \pm_2$ are useful because they define a pull operation on **cobordisms** \mathbf{C} in $\mathbf{Cob}(\pm_2)$ operation when \mathbb{S} has pullbacks, given by the obvious three pullbacks along

the components of F in the following diagram:



This construction induces a **lax double functor** between $\mathbf{Cob}(\mathfrak{z}_2)$ and $\mathbf{Cob}(\mathfrak{z}_1)$. The converse operation of pushing exists for any internal functor.

Lemma 3.3.1. *A plain internal functor $F : \mathfrak{z}_1 \rightarrow \mathfrak{z}_2$ induces a **lax double functor** defined by taking the pointwise pullbacks along the components of F*

$$\text{pull}[F] : \mathbf{Cob}(\mathfrak{z}_2) \rightarrow \mathbf{Cob}(\mathfrak{z}_1).$$

*Conversely, any internal functor $F : \mathfrak{z}_1 \rightarrow \mathfrak{z}_2$ induces a **double functor** by post-composition*

$$\text{push}[F] : \mathbf{Cob}(\mathfrak{z}_1) \rightarrow \mathbf{Cob}(\mathfrak{z}_2).$$

Proof. The double functor $\text{pull}[F]$ is lax, in that there is a (not necessarily invertible) map $\text{pull}[F](C); \text{pull}[F](D) \rightarrow \text{pull}[F](C; D)$ induced by the universal properties of colimits and limits. \square

3.3.2 Acute spans of internal functors

We start by introducing the notion of acute span of internal functors, and then describe how transport of structure works along an acute span.

Acute spans An *acute span* between an internal J_1 -opcategory \mathfrak{A}_1 and an internal J_2 -opcategory \mathfrak{A}_2 is defined as a span of **internal functors**

$$\mathfrak{A}_1 \xleftarrow[=]{F} \mathfrak{A}_0 \xrightarrow{G} \mathfrak{A}_2$$

where \mathfrak{A}_0 is an internal J_0 -opcategory and the $=$ sign indicates that the internal functor F is plain. Here, we suppose that the ambient category \mathbb{S} has pushouts and pullbacks. In that case, the definition gives rise to a double category **AcuteSpan**, whose objects are the internal J -opcategories in the ambient category \mathbb{S} , whose horizontal 1-cells are acute spans, whose vertical 1-cells are **internal functors of J -opcategories** and whose 2-cells are **maps of acute spans**: commuting diagrams of the form:

$$\begin{array}{ccccc} \mathfrak{A}_1 & \xleftarrow{=} & \mathfrak{A}_2 & \xrightarrow{\quad} & \mathfrak{A}_3 \\ \downarrow & & \Downarrow \phi & & \downarrow \\ \mathfrak{A}'_1 & \xleftarrow{=} & \mathfrak{A}'_2 & \xrightarrow{\quad} & \mathfrak{A}'_3 \end{array}$$

where the internal functor $\phi : \mathfrak{A} \rightarrow \mathfrak{A}'$ is plain. Acute spans are composed using pullbacks of plain internal functors along internal functors, in the category of internal J -opcategories. The pullback is defined as follows: Assume that we are given two internal functors:

$$\mathfrak{A}_1 \xrightarrow{F} \mathfrak{A}_2 \xleftarrow[=]{G} \mathfrak{A}_3$$

between an internal J_1 -opcategory \mathfrak{A}_1 and a J_3 -opcategory \mathfrak{A}_3 . Their pullback is defined as the internal J_1 -opcategory $\mathfrak{A}_1 \times_{\mathfrak{A}_2} \mathfrak{A}_3$ described as follows. At the level of its components and objects of \mathbb{S} , for each $i \in J_1$, the two internal functors give the following diagram:

$$\mathfrak{A}_1[0, i] \xrightarrow{F[0, i]} \mathfrak{A}_2[0, f(i)] \xleftarrow{G[0, f(i)]} \mathfrak{A}_3[0, f(i)]$$

where f is the action of the internal functor F on colors; and we simply define $(\mathfrak{A}_1 \times_{\mathfrak{A}_2} \mathfrak{A}_3)[0, i]$ as the pullback of that diagram in the ambient category \mathbb{S} . We proceed similarly to define $(\mathfrak{A}_1 \times_{\mathfrak{A}_2} \mathfrak{A}_3)[1, ij]$, and the universality of the pullbacks in \mathbb{S} gives the structural maps between the two.

Transport along acute spans The *raison d'être* of acute spans is to induce an operation of transport by “pull-then-push” along the two legs of a span. This operation plays a fundamental role in Part I of this thesis, in particular because we derive our definition of parallel product, and later of **change of lock** operations, from it.

To be able to characterize this operation, we need the notion of **vertical transformation**, a generalization of natural transformations to double categories.

Definition 3.3.2. Given two double categories \mathcal{C} and \mathcal{D} , and two lax double functors $\mathcal{F}, \mathcal{G} : \mathcal{C} \rightarrow \mathcal{D}$ between them, a **vertical transformation** $\alpha : \mathcal{F} \Rightarrow \mathcal{G}$ is defined as the following.

- For all object $A \in \mathcal{C}$, a vertical 1-cell

$$\begin{array}{c} \mathcal{F}(A) \\ \downarrow \alpha_A \\ \mathcal{G}(A) \end{array}$$

in \mathcal{D} which is natural with respect to composition with vertical 1-cells in the double category \mathcal{C} .

- For every horizontal 1-cell $F : A \rightarrow B$ in \mathcal{C} , a 2-cell

$$\begin{array}{ccc} \mathcal{F}(A) & \xrightarrow{\mathcal{F}(F)} & \mathcal{F}(B) \\ \downarrow \alpha_A & \Downarrow \alpha_F & \downarrow \alpha_B \\ \mathcal{G}(A) & \xrightarrow{\mathcal{G}(F)} & \mathcal{G}(B) \end{array}$$

- For composable horizontal 1-cells $F : A \rightarrow B$ and $G : B \rightarrow C$, the horizontal composite of α_F and α_G is $\alpha_{G \circ F}$.

The operation of transport along an acute span can be formulated as a double functor

$$\mathbf{Cob} : \mathbf{AcuteSpan} \rightarrow \mathbf{DbI}_{vstrict, hlax}$$

between the double category **AcuteSpan** of acute spans defined just above, and the following double category:

Definition 3.3.3. The double category of double categories, double functors and vertical transformations $\mathbf{DbI}_{vstrict, hlax}$ is defined as follows:

- 0-cells are double categories;
- vertical 1-cells are double functors;
- horizontal 1-cells are lax double functors;
- 2-cells are vertical transformations between the two composite lax double functors.

Lemma 3.3.4. *The operation **Cob** defines a double functor*

$$\mathbf{Cob} : \mathbf{AcuteSpan} \rightarrow \mathbf{DbI}_{\text{vstrict,hlax}}$$

from the double category **AcuteSpan** just defined to the double category $\mathbf{DbI}_{\text{vstrict,hlax}}$ of double categories, double functors, lax double functors and vertical transformations. It is defined as follows:


$$\begin{aligned} \mathbf{Cob} & : \mathbf{AcuteSpan} \rightarrow \mathbf{DbI}_{\text{vstrict,hlax}} \\ \varepsilon & \mapsto \mathbf{Cob}(\varepsilon) \\ f & \mapsto \text{push}[f] \\ (F, G) & \mapsto \text{push}[G] \circ \text{pull}[F] \end{aligned}$$

and the 2-cells are given by the universal properties of the pullback in the definition of $\text{pull}[\cdot]$.

Proof. To prove that this induces a double functor, one needs in particular to check that, given two composable acute spans $\mathbf{F} : \varepsilon_1 \dashrightarrow \varepsilon_3$ with support ε_2 and $\mathbf{G} : \varepsilon_3 \dashrightarrow \varepsilon_5$ with support ε_4 , there is an invertible special 2-cell, that is, an invertible vertical transformation $\alpha : \mathbf{Cob}(\mathbf{G} \circ \mathbf{F}) \cong \mathbf{Cob}(\mathbf{G}) \circ \mathbf{Cob}(\mathbf{F})$. In particular, given an object $\lambda : A \rightarrow \varepsilon_1[0, i]$ of $\mathbf{Cob}(\varepsilon_1)$, there must be a vertical isomorphism

$$\alpha_A : \mathbf{Cob}(\mathbf{G} \circ \mathbf{F})(A) \xrightarrow{\sim} (\mathbf{Cob}(\mathbf{G}) \circ \mathbf{Cob}(\mathbf{F}))(A)$$

The following diagram shows how to construct this isomorphism.



Here, $\pm_6[1]$ is the pullback defining the composition of F and of G , and A_1 and A_2 are the pullbacks in the action of the acute spans F and G . Finally, B is the pullback of the action of $G \circ F$, along the map from $\pm_6[1]$ to $\pm_1[1]$.

Using the pasting lemma on pullbacks, the face of the cube with B, A_1 and $\pm_6[1]$ is a pullback. Using the same lemma with the other faces of the cube, we conclude that the top face is a pullback as well.

This implies that map from B to A_2 is an isomorphism, since it is the pullback of the identity on A_1 . This means that B and A_2 are related by an isomorphism which commutes with the labeling maps:

$$\begin{array}{ccc} B & \xrightarrow{\alpha_A} & A_2 \\ & \searrow & \swarrow \\ & \pm_5[1] & \end{array}$$

We can do the same with the $\pm[0, i]$'s to construct the desired isomorphism. □

3.3.3 Span monoidal internal J -opcategories

We use the span operation above to define the parallel product of two programs. Categorically, we construct a lax-monoidal structure on $\mathbf{Cob}(\pm)$, that is to say a lax double functor

$$\parallel : \mathbf{Cob}(\pm) \times \mathbf{Cob}(\pm) \rightarrow \mathbf{Cob}(\pm).$$

By lax double functor, we mean that there exists a natural and coherent family of maps:

$$(C_1 \parallel C_2); (D_1 \parallel D_2) \Rightarrow (C_1; D_1) \parallel (C_2; D_2).$$

In our setting, this follows from Lemma 3.3.1, and therefore from the fact that there is always a map from a colimit of limits to the corresponding limit of colimits.

The formal statement is that $\mathbf{Cob}(\pm)$ is a monoid object in the bicategory $\mathbf{DbI}_{\text{lax}}$ of double categories and lax double functors. This is why we will use the fact (Lemma 3.1.5) that \mathbf{Cob} induces a lax pseudo functor of bicategories \mathbf{HCob} between the horizontal bicategories $\mathbf{HAcuteSpan}$ and $\mathbf{HDbI}_{\text{vstrict,hlax}} = \mathbf{DbI}_{\text{lax}}$ of double categories and **lax double functors**. The objects of bicategory $\mathbf{AcuteSpan}$ are the same as the corresponding double category, its 1-cells are the **acute spans** and a two cell $\phi : (F', G') \Rightarrow (F', G')$ is

a **special 2-cell** in the double category, that is, a **plain internal functors** ϕ which make the following diagram commute:

$$\begin{array}{ccccc}
 & & \mathfrak{A}_0 & & \\
 & F & \swarrow & G & \\
 \mathfrak{A}_1 & \xleftarrow{=} & & & \mathfrak{A}_2 \\
 & F' & \swarrow & G' & \\
 & & \mathfrak{A}'_0 & & \\
 & & \Downarrow \phi & &
 \end{array}$$

In this section, **Cob** will denote this lax pseudo functor and **AcuteSpan** will denote the bicategory of acute spans.

This tensor product is built the same way as in the case of template games (Melliès, 2019), but using the notion of **span-monoidal internal J -opcategory**. We remarked below Definition 3.1.9 that the category **opCat**(\mathbb{S}) of internal opcategories in the ambient category \mathbb{S} is a Cartesian category.

As a consequence, the lax pseudo functor **Cob** is lax monoidal, where we equip both **AcuteSpan** and **Dbi_{lax}** with the monoidal structure induced by their Cartesian products. In particular, there exist coercions living in the Cartesian category **Dbi_{lax}**, in the form of lax double functors

$$\begin{array}{lcl}
 \mathbf{m}_{\mathfrak{A}_1, \mathfrak{A}_2} & : & \mathbf{Cob}(\mathfrak{A}_1) \times \mathbf{Cob}(\mathfrak{A}_2) \rightarrow \mathbf{Cob}(\mathfrak{A}_1 \times \mathfrak{A}_2) \\
 \mathbf{m}_1 & : & \mathbb{1} \rightarrow \mathbf{Cob}(\mathbb{1})
 \end{array}$$

The first coercion is obtained in the natural way by taking the “pointwise” Cartesian product of the two cobordisms, and the second is the trivial cobordism, given by the initial opcategory **1**.

Definition 3.3.5. A **span-monoidal internal J -opcategory** $(\mathfrak{A}, \parallel, \eta)$ is a symmetric pseudomonoid object in the symmetric monoidal bicategory **AcuteSpan**. In particular, \parallel and η are two acute spans:

$$\mathfrak{A} \times \mathfrak{A} \xleftarrow{\text{pick}} \mathfrak{A} \parallel \xrightarrow{\text{pince}} \mathfrak{A} \qquad \mathbb{1} \xleftarrow{\eta} \mathfrak{A} \xrightarrow{\eta} \mathfrak{A}.$$

together with invertible 2-cells that witness the associativity of \parallel , and the fact that η is a left and right identity of \parallel . See (Day and Street, 1997, §3) for the complete definition.

Now, combining the external operations $\mathbf{m}_{\mathfrak{A}, \mathfrak{A}}$ and \mathbf{m}_1 on the one hand, and the internal operations \parallel, η on the other, we get a lax-monoidal structure on **Cob**(\mathfrak{A}) whose multiplication and unit are respectively given by:

$$\begin{array}{ccccc}
 \mathbf{Cob}(\mathfrak{A}) \times \mathbf{Cob}(\mathfrak{A}) & \xrightarrow{\mathbf{m}_{\mathfrak{A}, \mathfrak{A}}} & \mathbf{Cob}(\mathfrak{A} \times \mathfrak{A}) & \xrightarrow{\mathbf{Cob}(\parallel)} & \mathbf{Cob}(\mathfrak{A}) \\
 \mathbb{1} & \xrightarrow{\mathbf{m}_1} & \mathbf{Cob}(\mathbb{1}) & \xrightarrow{\mathbf{Cob}(\eta)} & \mathbf{Cob}(\mathfrak{A})
 \end{array}$$

Theorem 3.3.6. *Every span-monoidal internal J -opcategory \mathfrak{A} induces a symmetric lax-monoidal double category **Cob**(\mathfrak{A}).*

3.3.4 Parallel products of code and of proofs

Now that we have a general method for equipping the double category $\mathbf{Cob}(\ddagger)$ with a lax-monoidal structure, we use it to define the parallel product for the stateful and the stateless semantics of the Code.

Parallel product of codes The basic idea is to think of the code $C_1 \parallel C_2$ as a situation where (1) there are *three* players involved: the Code of C_1 , the Code of C_2 and the overall Frame F , but where (2) we have forgotten the identities of the codes C_1 and C_2 by considering both of them to be the Code C . This idea leads us to define the three-player machine models $\ddagger_L^{\circ\circ}$ and $\ddagger_S^{\circ\circ}$, with polarities C_1, C_2 and F . In the same way the two-player versions of the machine models are deduced from their one-player versions \ddagger_L^\bullet and \ddagger_S^\bullet in Section 3.2.1 using a product, we use here the pullback:

$$\begin{array}{ccc}
 \ddagger^{\circ\circ} & \xrightarrow{\pi_{\text{pol}}^{(12)(F)}} & \Omega\{C_1, C_2, F\} \\
 \downarrow \pi_{\text{state}}^{(12)(F)} & & \downarrow \Omega(12)(F) \\
 \ddagger^{\circ} & \xrightarrow{\pi_{\text{pol}}} & \Omega\{C, F\}
 \end{array}$$

where $(12)(F)$ denotes the function $\{C_1, C_2, F\} \rightarrow \{C, F\}$ which maps the polarities C_1 and C_2 to C , and the polarity F to itself. More explicitly, $\ddagger^{\circ\circ}$ has the same states as \ddagger^\bullet and three copies for each transition of \ddagger^\bullet , one copy for each of the three polarities C_1, C_2, F . The span-monoidal structures on the internal opcategories $\ddagger = \ddagger_S, \ddagger_L$ are defined in a common way:

$$\ddagger^{\circ\circ} \times \ddagger^{\circ\circ} \xleftarrow{\langle \pi_{\text{state}}^{(1)(2F)}, \pi_{\text{state}}^{(2)(1F)} \rangle} \ddagger^{\circ\circ} \xrightarrow{\pi_{\text{state}}^{(12)(F)}} \ddagger^{\circ} \quad \mathbf{1} \xleftarrow{\iota_F} \ddagger^\bullet \xrightarrow{\iota_F} \ddagger^{\circ\circ}.$$

where, for instance, $\pi_{\text{state}}^{(1)(2F)}$ maps every transition with polarity C_1, C_2, F to the corresponding transition with polarity given by the map $C_1 \mapsto C$ and $C_2, F \mapsto F$; while the homomorphism ι_F embeds the asynchronous graph \ddagger^\bullet into $\ddagger^{\circ\circ}$ with every edge transported to the corresponding edge of polarity F .

The product synchronizes transitions of the Code in C_1 with transitions of the Environment in C_2 which are mapped to the same transition in \ddagger . The intuition here is that a transition of C_1 is seen by C_2 as a transition of its Environment. Note that the parallel product preserves tiles from C_1 and from C_2 , and adds Code/Code tiles when one transition comes from C_1 and the other transition from C_2 —meaning that the two instructions are executed on two different “threads”—and their image in $\ddagger^{\circ\circ} \times \ddagger^{\circ\circ}$ forms a tile—meaning that these two instructions are independent.

Parallel product of proofs In order to define the parallel product at the level of proofs, which we will use to interpret the CSL rule for the parallel product, we endow $\mathfrak{A}_{\text{Sep}}$ with a span-monoidal structure. The main piece of this structure is the asynchronous graph $\mathfrak{A}_{\text{Sep}}^{\circ\circ\bullet}$ of *three-player separated states*, which will be the basis for the support of the span

$$\mathfrak{A}_{\text{Sep}} \times \mathfrak{A}_{\text{Sep}} \xleftarrow{\text{pick}} \mathfrak{A}_{\text{Sep}}^{\parallel} \xrightarrow{\text{pince}} \mathfrak{A}_{\text{Sep}}$$

of Definition 3.3.5. It is the straightforward generalization of separated states to the case where there are three players, \mathbf{C}_1 , \mathbf{C}_2 and \mathbf{F} .

Definition 3.3.7. A *three-player separated state* $(\sigma_1, \sigma_2, \vec{\sigma}, \sigma_F)$ is a triple of three *logical states* σ_1 , σ_2 and σ_F and a map $\vec{\sigma} : \mathfrak{L} \rightarrow \mathbf{LogState} + \{\mathbf{C}_1, \mathbf{C}_2, \mathbf{F}\}$ which satisfies the following properties. For all $r \in \vec{\sigma}^{-1}(\mathbf{LogState})$, $\vec{\sigma}(r) \models \Gamma(r)$, and the following logical state is well defined:

$$\sigma_1 * \sigma_2 * \left\{ \bigotimes_{r \in \vec{\sigma}^{-1}(\mathbf{LogState})} \vec{\sigma}(r) \right\} * \sigma_F \in \mathbf{LogState}$$

As for two-player separated states, using the logical state above, we can define a map \bigotimes which maps a three-player separated state to its underlying machine state.

The asynchronous graph $\mathfrak{A}_{\text{Sep}}^{\circ\circ\bullet}$, the three-player analogue to $\mathfrak{A}_{\text{Sep}}^{\circ\bullet}$ which is the basis to the $\mathfrak{A}_{\text{Sep}}$ internal J -opcategory is defined as follows. Its nodes are the three-player separated states defined above. It has three kinds of transitions: The \mathbf{C}_1 transitions are of the form:

$$(\sigma_1, \sigma_2, \vec{\sigma}, \sigma_F) \xrightarrow{m:\mathbf{C}_1} (\sigma'_1, \sigma_2, \vec{\sigma}', \sigma_F)$$

where $m \in \mathbf{Instr}$ is an instruction, such that

$$\bigotimes(\sigma_1, \sigma_2, \vec{\sigma}, \sigma_F) \xrightarrow{m} \bigotimes(\sigma'_1, \sigma_2, \vec{\sigma}', \sigma_F)$$

is a transition in $\mathfrak{A}_{\text{S}}^{\bullet}$, and such that the following conditions are satisfied:

$$\begin{aligned} \forall \ell \notin \text{wr}(m), \sigma_1(\ell) = \sigma'_1(\ell) & \quad \text{wr}(m) \cup \text{rd}(m) \subseteq \text{dom}(\sigma_1) \\ \text{lk}(m) \subseteq \text{dom}(\vec{\sigma}) \cup \vec{\sigma}^{-1}(\mathbf{C}_1) & \quad \forall r \notin \text{lk}(m), \vec{\sigma}(r) = \vec{\sigma}'(r). \end{aligned}$$

The \mathbf{C}_2 transitions are of the form

$$(\sigma_1, \sigma_2, \vec{\sigma}, \sigma_F) \xrightarrow{m:\mathbf{C}_2} (\sigma_1, \sigma'_2, \vec{\sigma}', \sigma_F)$$

with symmetric conditions. Finally, the \mathbf{F} transitions are of the form

$$(\sigma_1, \sigma_2, \vec{\sigma}, \sigma_F) \xrightarrow{m:\mathbf{F}} (\sigma_1, \sigma_2, \vec{\sigma}', \sigma'_F)$$

such that they correspond to a transition in \mathfrak{A}_S as above, and such that

$$\begin{aligned} \forall \ell \notin \text{wr}(m), \sigma_F(\ell) = \sigma'_F(\ell) & & \text{wr}(m) \cup \text{rd}(m) \subseteq \text{dom}(\sigma_F) \\ \text{lk}(m) \subseteq \text{dom}(\vec{\sigma}) \cup \vec{\sigma}^{-1}(\mathbf{F}) & & \forall r \notin \text{lk}(m), \vec{\sigma}(r) = \vec{\sigma}'(r). \end{aligned}$$

We can now define the internal J -opcategory $\mathfrak{A}_{\text{Sep}}^{\parallel}$, where J is the set of ordered pairs of formulas of CSL. As for $\mathfrak{A}_{\text{Sep}}$, we define $\mathfrak{A}_{\text{Sep}}^{\parallel}[1]$ to be $\mathfrak{A}_{\text{Sep}}^{\circ\circ}$, and $\mathfrak{A}_{\text{Sep}}^{\parallel}[0, (P, Q)]$ is the sub-asynchronous graph of $\mathfrak{A}_{\text{Sep}}^{\circ\circ}$ induced by keeping all \mathbf{F} transitions (and tiles over them) between three-player separated states $(\sigma_1, \sigma_2, \vec{\sigma}, \sigma_F)$ such that $\sigma_1 \models P$ and $\sigma_2 \models Q$.


The internal functor pince has the following action on the supports of the internal J -categories: it maps a three-player state to the (two-player) separated state $(\sigma_1 * \sigma_2, \vec{\sigma}', \sigma_F)$, where $\vec{\sigma}'$ is the same as $\vec{\sigma}$ where C_1 and C_2 have been remapped to C . It maps $\mathfrak{A}_{\text{Sep}}^{\parallel}[0, (P, Q)]$ to $\mathfrak{A}_{\text{Sep}}^{\parallel}[0, P * Q]$ by restriction of the above asynchronous morphism. The morphism pick maps that three-player state to the pair $\langle (\sigma_1, \vec{\sigma}_1, \sigma_2 * \sigma_F), (\sigma_2, \vec{\sigma}_1, \sigma_1 * \sigma_F) \rangle$. The unit $\mathfrak{A}_{\text{Sep}}^{\parallel\eta}$ of the structure is the subgraph of $\mathfrak{A}_{\text{Sep}}^{\parallel}[1]$ with only Frame moves.

3.4 Generalized sequential composition

So far, to compose cobordisms horizontally, the target of the first cobordism must exactly match the source of the second. This is not a reasonable assumption, because the initial and the final states of a program are part of its internal state. For example, the two cobordisms corresponding to the two programs being sequentially composed in the following program

$$(\text{if } B \text{ then } C_1 \text{ else } C_2) ; D$$

could look like this:



In summary, the two cobordisms we wish to compose generally look like:

$$\begin{array}{ccc}
 A & \xrightarrow{\quad} & B \\
 \downarrow & & \downarrow \lambda_{\text{out}} \\
 \mathfrak{z}[0, i] & \xrightarrow{\mathfrak{z}[1, ij]} & \mathfrak{z}[0, j]
 \end{array}
 \qquad
 \begin{array}{ccc}
 B' & \xrightarrow{\quad} & C' \\
 \lambda_{\text{in}} \downarrow & & \downarrow \\
 \mathfrak{z}[0, i'] & \xrightarrow{\mathfrak{z}[1, i'j']} & \mathfrak{z}[0, j']
 \end{array}$$

In practice, in all the cases we consider, $\mathfrak{z}[0, j]$ and $\mathfrak{z}[0, i']$ will be equal, but B and B' will be different. To bridge the gap between $(B, \lambda_{\text{out}})$ and $(B', \lambda_{\text{in}})$, we will use a **filling system** over the internal J -opcategory \mathfrak{z} . With each pair of interfaces $\lambda_{\text{out}} : B \rightarrow \mathfrak{z}[0, j]$ and $\lambda_{\text{in}} : B' \rightarrow \mathfrak{z}[0, i']$, the filling system associates a cobordism

$$\text{fill}((B, \lambda_{\text{out}}), (B', \lambda_{\text{in}})) : (B, \lambda_{\text{out}}) \dashrightarrow (B', \lambda_{\text{in}})$$

from $(B, \lambda_{\text{out}})$ to $(B', \lambda_{\text{in}})$. Thanks to this mediating cobordism, it becomes possible to compose the two cobordisms using the usual composition of cobordism. Given such a filling system, we write $C_1 \mathbin{\text{\textcircled{;}}} C_2$ for this generalized form of composition.

Proposition 3.4.1. *Suppose that for all games $\lambda, \lambda', \mu, \mu'$ there exists a map $\text{fill}(\lambda \parallel \lambda', \mu \parallel \mu') \rightarrow \text{fill}(\lambda, \mu) \parallel \text{fill}(\lambda', \mu')$. In that case, the Hoare inequality holds: $(C_1 \parallel C'_1) \mathbin{\text{\textcircled{;}}} (C_2 \parallel C'_2) \rightarrow (C_1 \mathbin{\text{\textcircled{;}}} C_2) \parallel (C'_1 \mathbin{\text{\textcircled{;}}} C'_2)$.*

Proof. The Hoare inequality for the usual composition applied twice gives us the map below, from which we can conclude using the hypothesis on the filling system.

$$(C_1 \parallel C'_1); (\text{fill}(\lambda, \mu) \parallel \text{fill}(\lambda', \mu')); (C_2 \parallel C'_2) \rightarrow (C_1 \mathbin{\text{\textcircled{;}}} C_2) \parallel (C'_1 \mathbin{\text{\textcircled{;}}} C'_2) \quad \square$$

When the base category \mathcal{C} has pullbacks as well as pushouts, which is the case in the examples we are considering, a filling system always exists. It is defined by the following diagram:

$$\begin{array}{ccccc}
 & & B_j \times_{\mathfrak{z}[1, jj']} A_{j'} & & \\
 & \swarrow & & \searrow & \\
 B_j & & & & A_{j'} \\
 & \searrow & & \swarrow & \\
 & & B_j \cup_{\mathfrak{z}[1, jj']} A_{j'} & & \\
 \downarrow \lambda & & \downarrow & & \downarrow \lambda' \\
 \mathfrak{z}[0, j] & \xrightarrow{\text{in}_{jj'}} & \mathfrak{z}[1, jj'] & \xleftarrow{\text{in}_{jj'}} & \mathfrak{z}[0, j']
 \end{array}$$

where $B_j \times_{\mathfrak{z}[1, jj']} A_{j'}$ is a pullback, and where $B_j \cup_{\mathfrak{z}[1, jj']} A_{j'}$ is a pushout above that pullback. Intuitively, it identifies all nodes of B_j and of $A_{j'}$ that have the same underlying state in $\mathfrak{z}[1, jj']$. This is the filling system which will be used in our interpretation of sequential composition.

Lemma 3.4.2. *The hypothesis of Proposition 3.4.1 holds in the case of the stateful and stateless templates \pm_S and \pm_L .*

Proof. Let \pm be either \pm_L or \pm_S . Given two games $\lambda : A \rightarrow \pm[0]$ and $\mu : B \rightarrow \pm[0]$ formulated as asynchronous morphisms, let us describe $\text{fill}(\lambda, \mu)$. Every node x of A comes with a state $\text{in}_{\lambda(x)} \in \pm[1]$ which we call the underlying state of x . Similarly, every node y of B comes with an underlying state $\text{in}_{\mu(y)} \in \pm[1]$. The pullback $A \times_{\pm[1]} B$ contains all the pairs $(a, b) \in A \times B$ which share the same underlying state. Hence, when we perform the pushout, we identify all such nodes a and b ; in particular, in the case where there is another node a' of A with the same underlying states, the pullback will also contain (a', b) , and the nodes a and a' will be identified in the pushout. In summary, the support S of the filling $\text{fill}(\lambda, \mu)$ is made of three kinds of nodes:

1. the nodes x of A such that no node of B has the same underlying state;
2. the nodes y of B such that no node of A has the same underlying state;
3. the states $s \in \pm[1]$ such that there exists nodes in A and nodes in B whose underlying states is s ; we use the notation $[s]$ in order to denote these specific nodes.

Let us prove now that the filling system defined just above satisfies the property that there exists a map:

$$\text{fill}(\lambda \parallel \lambda', \mu \parallel \mu') \rightarrow \text{fill}(\lambda, \mu) \parallel \text{fill}(\lambda', \mu')$$

The map is constructed in the following way. Consider a node of the support of $\text{fill}(\lambda \parallel \lambda', \mu \parallel \mu')$. As we have just mentioned, there are three possibilities:

1. In the first case, the element is a node of $A \parallel A'$, and thus a pair $(x, x') \in A \times A'$ consisting of two elements $x \in A$ and $x' \in A'$ with the same underlying state s . Recall indeed that the asynchronous morphism $\text{pince}[1] : \pm^{\parallel}[1] \rightarrow \pm[1]$ is *injective on states*. Since we are in the first case, there exists no node of $B \parallel B'$ with underlying state s . This means that either:
 - a) neither B nor B' have nodes whose underlying state is s ; in that case the node x is in $\text{fill}(\lambda, \mu)$, and the node x' is in $\text{fill}(\lambda', \mu')$, and we map (x, x') to (x, x') .
 - b) B has a node whose underlying state is s , but not B' ; in the same way essentially as in the preceding case, we can map (x, x') to $([s], x')$,
 - c) B' has a node whose underlying state is s , but not B ; this case is symmetric to the previous one.

2. the second case is symmetric to the previous one.
3. last case: the node of the support of fill $(\lambda \parallel \lambda', \mu \parallel \mu')$ is of the form $[s]$. In that case, there are nodes in each of the four graphs A, A', B, B' whose underlying states is s . We are thus allowed to map $[s]$ to the pair $([s], [s])$.

□

Interestingly, this is not necessarily the case for the template $\mathfrak{A}_{\text{Sep}}$ of separated states regulating the interpretation of CSL proofs. The reason is that there are several ways to decompose a given separated state between two players \mathbf{C}, \mathbf{F} in $\mathfrak{A}_{\text{Sep}}^{\bullet}$ into a separated state between three players $\mathbf{C}_1, \mathbf{C}_2, \mathbf{F}$ in $\mathfrak{A}_{\text{Sep}}^{\bullet\bullet}$.

3.5 Change of locks

In the language we consider, locks names follow lexical scoping: A lock r is introduced with the resource r do C construction, and is only available inside of the command C . This would be formalized syntactically by a judgment $\mathfrak{L} \vdash C$ which states that all the locks in C are in the set \mathfrak{L} , with the expected introduction and elimination rules:

$$\frac{\mathfrak{L} \cup \{r\} \vdash C}{\mathfrak{L} \vdash \text{resource } r \text{ do } C} \qquad \frac{\mathfrak{L} \vdash C}{\mathfrak{L} \cup \{r\} \vdash \text{with } r \text{ do } C}$$

Correspondingly, the three machine models $\mathfrak{A} = \mathfrak{A}_{\text{Sep}}, \mathfrak{A}_{\text{S}}, \mathfrak{A}_{\text{L}}$ considered in this paper are parameterized by the set of locks that the programs can access. In the case of the two internal opcategories \mathfrak{A}_{S} and \mathfrak{A}_{L} , the free locks are simply described by a set \mathfrak{L} of lock names. Then, a program $\mathfrak{L} \vdash C$ is interpreted as a cobordism in $\mathbf{Cob}(\mathfrak{A}(\mathfrak{L}))$.

In the case of the *internal J-opcategories* $\mathfrak{A}_{\text{Sep}}$ of separated states, the free locks are described by a context $\Gamma = r_1 : I_1, \dots, r_n : I_n$ which associates with each free lock r_k the predicate I_k of a CSL invariant. A proof π of a Hoare triple $\Gamma \vdash \{P\} C \{Q\}$ is interpreted as cobordism in $\mathbf{Cob}(\mathfrak{A}_{\text{Sep}}(\Gamma))$. The rules RES and WITH, which manipulate the context Γ of locks and invariants refine the two rules defined above.

As we wish to consider both cases at once, we write Γ, r to denote a context in either of the two cases. The operations of *introducing a new lock* and of *creating a critical section* transport cobordisms across machine models parameterized with different free locks. This change-of-basis technique is similar in nature to the way categorical models of Algol such as (Reynolds, 1997) deal with lexical scoping of variables. Hence, we call the two operations defined in this section *change of locks* operations.

The change-of-basis operations are induced by the two acute spans below:

$$\begin{array}{ccc} & \text{when}[r] & \\ & \curvearrowright & \\ \ddagger(\Gamma) & & \ddagger(\Gamma, r) \\ & \curvearrowleft & \\ & \text{hide}[r] & \end{array}$$

where we write explicitly the dependence of the models $\ddagger(\Gamma)$ on the contexts (or lists) Γ of free locks. We formalize the situation by defining the graph **LockGraph** whose vertices are the lock contexts Γ , with edges defined as transitions *adding* or *removing* one specific lock in the context:

$$\begin{array}{ccc} & \iota_r^\Gamma & \\ & \curvearrowright & \\ \Gamma & & \Gamma, r \\ & \curvearrowleft & \\ & \pi_r^\Gamma & \end{array}$$

Consider **LockGraph**^{*} the locally posetal bicategory freely generated by this graph, with invertible 2-cells between paths that are equal up to the reorderings: $\pi_r \circ \iota_{r'} \sim \iota_{r'} \circ \pi_r$, $\pi_r \circ \pi_{r'} \sim \pi_{r'} \circ \pi_r$ and $\iota_r \circ \iota_{r'} \sim \iota_{r'} \circ \iota_r$, for $r \neq r'$, and leaving the Γ 's implicit. By locally posetal, we mean that there is at most one tile (or invertible 2-cell) between any pair of morphisms, witnessing that they are equal up to the reorderings above. We call a **LockGraph-template** a double functor from this bicategory see as a double category to **AcuteSpan**(\mathbb{S}), the double category of internal J -opcategories and *acute spans* we defined in Section 3.3.2. The invertible 2-cells in the domain reflect the fact that the order of the operations does not matter, up to isomorphism. The three families $\ddagger = \ddagger_{\text{sep}}, \ddagger_S, \ddagger_L$ of internal J -opcategories defined so far are **LockGraph**-templates; we explain how in the remainder of this section. Practically, this means that, by post-composing with the lax double functor **Cob**(\cdot), we are able to transport a cobordism defined on the internal category $\ddagger(\Gamma)$ to one defined on the internal category $\ddagger(\Gamma, r : I)$, to interpret critical sections, and back, for resource introduction. We abuse the notation and write these lax double functors as follows:

$$\begin{array}{ccc} & \text{when}[r] & \\ & \curvearrowright & \\ \mathbf{Cob}(\ddagger(\Gamma)) & & \mathbf{Cob}(\ddagger(\Gamma, r)) \\ & \curvearrowleft & \\ & \text{hide}[r] & \end{array}$$

The fact that it is possible to rename bound lock names at the syntactic level is reflected in the semantics by the fact that the target of $\text{hide}[r]$ does not contain the name r of the hidden lock.

3.5.1 Hiding

To hide a lock r , we proceed in two steps for all the templates that we consider $\pm = \pm_L, \pm_S, \pm_{\text{Sep}}$. First, we prevent the Environment from touching that lock, and then we remove this lock from the states and we transform all transitions $P(r), V(r)$ into nops. Formally, hiding is defined by a pull and push operation along the acute span $\text{hide}[r]$:

$$\pm(\Gamma, r) \xleftarrow{\text{inj}_C} \pm^{(r)}(\Gamma, r) \xrightarrow{\text{hide}_C} \pm(\Gamma).$$

The support $\pm^{(r)}(\Gamma, r)$ of the span describes programs where the lock r is hidden from the environment, which means that only the Code can use it. Hence, $\pm^{(r)}(\Gamma, r)[1]$ is defined to be the same as the template $\pm(\Gamma, r)$, except that all Environment transitions $P(r), V(r)$ are deleted, and, in the case of the template of separated states, we remove the states of the form $(\sigma_C, \vec{\sigma}, \sigma_F)$, where the lock r is held by the Environment (that is: $\vec{\sigma}(r) = \mathbf{F}$).

Moreover, we know that at the beginning and at the end of the execution of C in resource r do C , the resource r is unlocked. This is why we define $\pm^{(r)}(\Gamma, r)[0]$ to only contain states where the resource r is unlocked. Note that it means that $\pm^{(r)}(\Gamma, r)[0]$ is isomorphic to $\pm(\Gamma)[0]$; a fact which will prove useful later.

By definition of $\pm^{(r)}(\Gamma, r)$ as a restriction of $\pm(\Gamma, r)$, there is a canonical injection $\text{inj}_C : \pm^{(r)}(\Gamma, r) \rightarrow \pm(\Gamma, r)$. The map hide_C is defined differently depending on the kind of template we are considering. In the case of the templates used for the code, respectively \pm_L and \pm_S , we map the states $L \subseteq \mathcal{L}$ to $L \setminus \{r\}$, and $s = (\mu, L)$ to $(\mu, L \setminus \{r\})$ respectively. In the case of the template of separated states \pm_{Sep} used to interpret proofs, it is defined as follows on separated states:

$$(\sigma_C, \vec{\sigma}, \sigma_F) \mapsto \begin{cases} (\sigma_C * \vec{\sigma}(r), \vec{\sigma} \setminus r, \sigma_F) & \text{if } \vec{\sigma}(r) \in \mathbf{LogState} \\ (\sigma_C, \vec{\sigma} \setminus r, \sigma_F) & \text{if } \vec{\sigma}(r) = \mathbf{C} \end{cases}$$

Intuitively, this means that after the lock is hidden, the resource it protects belongs to the Code. The fact that the lock has been hidden from the Environment ensures that the mapping above defines an asynchronous morphism.

In all cases, hide_C maps $P(r)$ and $V(r)$ to nop, and otherwise preserves the edges. The intuition behind the definition of hide_C for the separated states is that when we forget about the lock r , even when nobody is holding it, we need to do something with the resources that is associated with the lock: we chose to give this resource to the Code. This means that outside of the binder for r , the resource associated with r is not shared, but belongs to the Code.

3.5.2 Critical sections

The case of critical sections is more delicate, as the definition differs between \pm_L , \pm_S on the one hand, and \pm_{Sep} on the other. In the case of the semantics of the code, we wish to let the Environment be wild and change the states of locks whenever it wants, even if they happen to be held by the Code. On the other hand, when it comes to the semantics of the proofs, we enforce the discipline that a lock can only be unlocked by the player that locked it (they are not semaphores). This difference of requirements is reflected by differences in the definition of $\text{when}[r]$.

Stateful and stateless semantics We consider the machine model $\pm_{\langle r \rangle}(\mathcal{Q} \uplus \{r\})$ which is the restriction of $\pm(\mathcal{Q} \uplus \{r\})$ where *only the Environment* is able to use the lock r . This corresponds to the fact that inside a critical section the Code loses syntactically access to the corresponding lock until the end of the critical section. There exists a map ∇ from $\pm_{\langle r \rangle}(\mathcal{Q} \uplus \{r\})$ to $\pm(\mathcal{Q})$ given by:

$$\begin{aligned} \text{Nodes: } & (\mu, L) \mapsto (\mu, L \cap \mathcal{Q}) \\ \text{Edges: } & \begin{cases} P(r):F & \mapsto \text{nop} \\ V(r):F & \mapsto \text{nop} \\ m & \mapsto m \end{cases} \end{aligned}$$

Then, the lifting of a cobordism for the critical sections is given by the following acute span:

$$\pm(\mathcal{Q}) \xleftarrow{\nabla} \pm_{\langle r \rangle}(\mathcal{Q} \uplus \{r\}) \xrightarrow{\text{incl}} \pm(\mathcal{Q} \uplus \{r\})$$

where the right leg is the obvious inclusion. Intuitively, the operation of pulling along ∇ , which defines the $\text{when}[r]$ operation, duplicates the whole transition system, with one version for each state of the lock r . It is the Environment which is in control of switching between the copies, and the Code is oblivious to the state of the lock r .

Separated state semantics The lifting operation, that we will use to deal with critical sections, is simply defined as the push-forward along the map

$$\pm_{\text{Sep}}(\Gamma) \rightarrow \pm_{\text{Sep}}(\Gamma, r : I)$$

which sends a separated state $(\sigma_C, \vec{\sigma}, \sigma_F)$ over Γ to the separated state over $\Gamma, r : I$ where the lock is held by the code $(\sigma_C, \vec{\sigma} \uplus [r \mapsto C], \sigma_F)$.

3.6 Sum of cobordisms

We define the disjoint sum of two cobordisms. This will be useful as a building block for interpreting conditionals, and for the CSL rule for disjunction.

First, we remark that the category $\mathbf{opCat}(\mathbb{S})$ of *internal J -opcategories* has binary sums: the sum of an internal J -opcategory \mathfrak{z} and an internal J' -opcategory \mathfrak{z}' is an internal $J \times J'$ -opcategory $\mathfrak{z} + \mathfrak{z}'$. For example, for $j \in J$ and $j' \in J'$, define

$$(\mathfrak{z} + \mathfrak{z}')[0, (j, j')] := \mathfrak{z}[0, j] + \mathfrak{z}'[0, j'].$$

Commutation of colimits ensures that this operation is well defined.

As usual, this operation on cobordisms is defined using a structure at the level of *internal J -opcategories*.

Definition 3.6.1. An *internal J -opcategory with sums* (\mathfrak{z}, \vee) is an internal J -opcategory \mathfrak{z} together with a symmetric injections $\vee : J \times J \rightarrow J$ in \mathbb{S} and an *internal functor of J -opcategories* $\nabla : \mathfrak{z} + \mathfrak{z} \rightarrow \mathfrak{z}$ whose action on colors is the map \vee .

In particular, every *internal opcategory* has a canonical structure of internal opcategory with sums.

Given two cobordisms \mathbf{C} and \mathbf{D} over an *internal J -opcategory with sums*, their sum $\mathbf{C} \oplus \mathbf{D}$ is defined as the following cobordism, with the obvious notations:

$$\begin{array}{ccccc}
 I + I' & \xrightarrow{\text{in+in}} & S + S' & \xleftarrow{\text{out+out}} & O + O' \\
 \lambda_I + \lambda_{I'} \downarrow & & \downarrow \lambda_S + \lambda_{S'} & & \downarrow \lambda_O + \lambda_{O'} \\
 \mathfrak{z}[0, i] + \mathfrak{z}[0, i'] & \xrightarrow{\text{in}_{ij} + \text{in}_{i'j'}} & \mathfrak{z}[1, ij] + \mathfrak{z}[1, i'j'] & \xleftarrow{\text{out}_{ij} + \text{out}_{i'j'}} & \mathfrak{z}[0, j] + \mathfrak{z}[0, j'] \\
 \downarrow \nabla & & \downarrow \nabla & & \downarrow \nabla \\
 \mathfrak{z}[0, i \vee j] & \xrightarrow{\text{in}_{i \vee i' j \vee j'}} & \mathfrak{z}[1, i \vee i' j \vee j'] & \xleftarrow{\text{out}_{i \vee i' j \vee j'}} & \mathfrak{z}[0, j \vee j']
 \end{array}$$

3.7 Interpretation of codes and proofs

We have studied in previous sections how to express the main operations of concurrent separation logic (sequential composition, parallel product and change of lock) in the language of template games and cobordisms. Now that each of these basic operations has been defined, the interpretation of the code and of the proofs of concurrent separation logic (CSL) is uniform and essentially straightforward.

3.7.1 Stateful and stateless interpretations of the code

We begin by describing how the stateful and stateless interpretations $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ of a given code C are computed in our template game model. Since the interpretation is uniform in \mathfrak{s} and \mathfrak{L} and only depends on the combinators of C , we find it convenient to write \mathfrak{s} alternatively for \mathfrak{s}_S or \mathfrak{s}_L . The code C with free locks \mathfrak{L} is interpreted by structural induction as a cobordism of the form

$$\llbracket C \rrbracket = \begin{array}{ccccc} \llbracket C \rrbracket_{in} & \longrightarrow & \llbracket C \rrbracket_{support} & \longleftarrow & \llbracket C \rrbracket_{out} \\ \lambda_{in} \downarrow & & \downarrow & & \downarrow \lambda_{out} \\ \mathfrak{s}(\mathfrak{L})[0] & \longrightarrow & \mathfrak{s}(\mathfrak{L})[1] & \longleftarrow & \mathfrak{s}(\mathfrak{L})[0]. \end{array} \quad (3.4)$$

The interpretation of every non-leaf command of the language corresponds to an operation on cobordisms already defined, and it is thus straightforward.

Instructions We explain first how to define the cobordism that interprets a single instruction m . The cobordism $\llbracket m \rrbracket$ is constructed in the following way: its input and output borders $\llbracket m \rrbracket_{in}$ and $\llbracket m \rrbracket_{out}$ are defined as the asynchronous graph $\mathfrak{s}(\mathfrak{L})[0]$, while its support $\llbracket m \rrbracket_{support}$ consists of the disjoint union of $\llbracket m \rrbracket_{in}$ and $\llbracket m \rrbracket_{out}$ augmented with an edge $s_1 \rightarrow s_2$ from $s_1 \in \llbracket m \rrbracket_{in}$ to $s_2 \in \llbracket m \rrbracket_{out}$ for every machine transition $m : s_1 \rightarrow s_2$ performed by the instruction m . Note that $\llbracket m \rrbracket_{in}$ and $\llbracket m \rrbracket_{out}$ contain only Frame transitions, and that all the “transverse” edges $s_1 \rightarrow s_2$ from $\llbracket m \rrbracket_{in}$ to $\llbracket m \rrbracket_{out}$ are Code transitions, with the state s_2 potentially equal to the error state.

This definition can be formulated using a well-chosen pullback: Consider the following asynchronous graph

$$A = \begin{array}{ccc} F_1 & & F_2 \\ \Downarrow & & \Downarrow \\ \bullet & \xrightarrow{C} & \bullet \end{array}$$

with a tile $F_1 \cdot C \sim C \cdot F_2$. Then, we can construct the pullback, where **Instr** is the set of instruction and $\Omega : \mathbf{Set} \rightarrow \mathbf{AsyncGraph}$ is the functor we used to define \mathfrak{s}^\bullet page 72.

$$\begin{array}{ccc} G(m) & \hookrightarrow & A \times \mathfrak{s}^\bullet \\ \downarrow & & \downarrow f \\ \Omega(\{F, m\}) & \hookrightarrow & \Omega(\{F\} \cup \mathbf{Instr}) \end{array}$$

where the map f sends edges of the form (F_1, \cdot) and (F_2, \cdot) to F and edges of the form (C, m') to the edge m' in $\Omega(\{F\} \cup \mathbf{Instr})$, for all instructions m' . It is then easy to deduce the maps from the borders of the cobordism to its support using the universal property of the pullback.

Leaf codes Except for `malloc` which is treated below, every leaf command $x := E \mid x := [E] \mid [E] := E' \mid \text{skip} \mid \text{dispose}(E)$ of the language corresponds to a specific machine instruction m which is interpreted as the cobordism $\llbracket m \rrbracket$ we have just defined. The case of `malloc` is different because it needs to choose non-deterministically the address which is allocated; as such, we take the disjoint sum of all the semantics of the basic instructions $\text{alloc}(E, \ell)$:

$$\llbracket \text{malloc}(E) \rrbracket := \text{fill}(\text{Id}_{\mathbb{Z}[0]}, \dots) ; \bigoplus_{\ell \in \text{Loc}} \llbracket \text{alloc}(E, \ell) \rrbracket$$

where we pre-compose with a **filling system** to normalize the initial states (the second component of `fill` is determined by the cobordism on its right.)

We now detail how to give a semantics to any code C as a cobordism $\llbracket C \rrbracket$, by induction on its structure. This lets us build $\llbracket C \rrbracket_S$ and $\llbracket C \rrbracket_L$ in the same way.

Conditionals Conditional branching is interpreted as

$$\llbracket \text{if } B \text{ then } C_1 \text{ else } C_2 \rrbracket$$

defined as

$$\llbracket \text{test}(B) \rrbracket ; \llbracket C_1 \rrbracket \oplus \llbracket \text{test}(\neg B) \rrbracket ; \llbracket C_2 \rrbracket \tag{3.5}$$

precomposed with

$$\text{fill}\left(\mathbb{Z}[0] \xrightarrow{id} \mathbb{Z}[0], \mathbb{Z}[0] + \mathbb{Z}[0] \xrightarrow{\nabla} \mathbb{Z}[0]\right)$$

as depicted at the beginning of Section 3.4, page 84.

Here, the purpose of precomposing with the filling is to identify the two copies of the input $\mathbb{Z}[0]$ appearing on each sides of the disjoint sum (3.5).

Sequential and parallel compositions We use the sequential and the parallel product of cobordisms to interpret their syntactic counterparts:

$$\llbracket C_1 \parallel C_2 \rrbracket = \llbracket C_1 \rrbracket \parallel \llbracket C_2 \rrbracket \quad \llbracket C_1 ; C_2 \rrbracket = \llbracket C_1 \rrbracket ; \llbracket C_2 \rrbracket.$$

Resource introduction and critical sections We use the change of locks operations. The interpretation of resource r do C is defined as

$$\llbracket \text{resource } r \text{ do } C \rrbracket = \text{hide}[r](\llbracket C \rrbracket)$$

The interpretation of $\llbracket \text{with } r \text{ do } C \rrbracket$ is defined as:

$$\llbracket P(r) \rrbracket ; \text{when}[r](\llbracket C \rrbracket) ; \llbracket V(r) \rrbracket.$$

Loops To interpret a loop `while B do C`, we build its infinite unfolding as the least fixpoint of the map:

$$F(X) = \llbracket \text{test}(B) \rrbracket ; \llbracket C \rrbracket ; X \oplus \llbracket \text{test}(\neg B) \rrbracket$$

seen as an endofunctor on the category of arrows of the double category $\mathbf{Cob}(\pm)$ seen as an internal category in \mathbf{Cat} . It exists because that category has all colimits of ω -chains, and F preserves such colimits because it is itself defined using colimits.

3.7.2 Interactive and separated interpretations of the proofs

Our purpose is to interpret by structural induction every CSL proof $\pi : \Gamma \vdash \{P\} C \{Q\}$ as a cobordism of the form

$$\llbracket \pi \rrbracket_{\text{Sep}} = \begin{array}{ccccc} \llbracket \pi \rrbracket_{\text{Sep},in} & \longrightarrow & \llbracket \pi \rrbracket_{\text{Sep},support} & \longleftarrow & \llbracket \pi \rrbracket_{\text{Sep},out} \\ \downarrow & & \downarrow & & \downarrow \\ \pm_{\text{Sep}}(\Gamma)[0, P] & \longrightarrow & \pm_{\text{Sep}}(\Gamma)[1] & \longleftarrow & \pm_{\text{Sep}}(\Gamma)[0, Q]. \end{array} \quad (3.6)$$

living in the double category $\mathbf{Cob}(\pm_{\text{Sep}}(\Gamma))$ associated with the template $\pm_{\text{Sep}}(\Gamma)$ of separated states, parameterized by the context Γ . As it stands, the interpretation is essentially straightforward, since most of the rules of the logic correspond to an operation on cobordisms already carefully defined. There is apparently one exception however: the `FRAME` rule does not seem, at least at first sight, to correspond to a basic operation on cobordisms. Given a cobordism $\llbracket \pi \rrbracket_{\text{Sep}}$ which interprets a proof π of the Hoare triple $\Gamma \vdash \{P\} C \{Q\}$, we need to define a new cobordism associated with the Hoare triple $\Gamma \vdash \{P * R\} C \{Q * R\}$. The solution is not difficult to find however: we define the new cobordism as the parallel product $\llbracket \pi \rrbracket_{\text{Sep}} \parallel \pm_{\text{Sep}}[0, R]$, where the asynchronous graph $\pm_{\text{Sep}}[0, R]$, which contains only states which satisfy the predicate R and Environment transitions, is lifted to the identity cobordism defined in the expected way in the double category $\mathbf{Cob}(\pm_{\text{Sep}})$.

The rules that correspond to machine instructions $m \in \mathbf{Instr}$ (such as `LD`) are interpreted in a way which is similar to the interpretation of the corresponding codes, always preserving the permission associated with affected locations. For instance, the axiom rule

$$\frac{\Gamma \vdash \{((\text{Own}_{\top}(x) * P) \wedge E=w) * w \xrightarrow{p} v\} x := [E] \{((\text{Own}_{\top}(x) * P) \wedge x=v) * w \xrightarrow{p} v\}}{\text{LD}}$$

is interpreted as a cobordism whose Code edges are all the edges from an initial state to an final state of the form

$$((s, h), \vec{\sigma}, \sigma_F) \xrightarrow{x := [E] : C} ((s[x \mapsto v], h), \vec{\sigma}, \sigma_F)$$

where h contains a mapping $[w \rightarrow v]$, and where $E = w$ holds in h .

$$\left[\left[\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P\} C_1 \{Q\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{Q\} C_2 \{R\} \end{array}}{\Gamma \vdash \{P\} C_1; C_2 \{R\}} \right] \right]_{\text{Sep}} = \llbracket \pi_1 \rrbracket_{\text{Sep}} \circledast \llbracket \pi_2 \rrbracket_{\text{Sep}}$$

For the parallel product rule PAR, we use the parallel product of cobordisms using the above notion of *compatibility*:

$$\left[\left[\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\} C_1 \{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\} C_2 \{Q_2\} \end{array}}{\Gamma \vdash \{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}} \right] \right]_{\text{Sep}} = \llbracket \pi_1 \rrbracket_{\text{Sep}} \parallel \llbracket \pi_2 \rrbracket_{\text{Sep}}$$

$$\left[\left[\frac{\begin{array}{c} \vdots \pi \\ \Gamma \vdash \{P\} C \{Q\} \end{array}}{\Gamma \vdash \{P * R\} C \{Q * R\}} \right] \right]_{\text{Sep}} = \llbracket \pi \rrbracket_{\text{Sep}} \parallel \text{Id}_{\text{Sep}[0,R]}$$

$$\left[\left[\frac{\begin{array}{c} \vdots \pi \\ \Gamma, r : J \vdash \{P\} C \{Q\} \end{array}}{\Gamma \vdash \{P * J\} \text{resource } r \text{ do } C \{Q * J\}} \right] \right]_{\text{Sep}} = \text{hide}[r](\llbracket \pi \rrbracket)$$

$$\left[\left[\frac{\begin{array}{c} \vdots \pi \\ \Gamma, r : J \vdash \{(P * J) \wedge B\} C \{Q * J\} \end{array}}{\Gamma, r : J \vdash \{P\} \text{with } r \text{ do } C \{Q\}} \right] \right]_{\text{Sep}} = \text{acquire}[r]; \text{when}[r](\llbracket \pi \rrbracket_{\text{Sep}}); \text{release}[r]$$

$$\left[\left[\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash \{P_1\} C \{Q_1\} \end{array} \quad \begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash \{P_2\} C \{Q_2\} \end{array}}{\Gamma \vdash \{P_1 \vee P_2\} C \{Q_1 \vee Q_2\}} \right] \right]_{\text{Sep}} = \llbracket \pi_1 \rrbracket_{\text{Sep}} \oplus \llbracket \pi_2 \rrbracket_{\text{Sep}}$$

In the definition of the interpretation of WITH, the cobordisms $\text{acquire}[r]$ and $\text{release}[r]$ are the counterparts of $\llbracket P(r) \rrbracket_S$ and $\llbracket V(r) \rrbracket_S$ at the level of separated states.

3.8 The asynchronous soundness theorem

We explain how to state the asynchronous soundness theorem of concurrent separation logic which we described in the previous chapter in Section 2.5.3 in the language of cobordisms.

3.8.1 Comparing the three interpretations

Given a code C and a proof π of $\Gamma \vdash \{P\}C\{Q\}$, the proof of the asynchronous soundness relies on the existence of a chain of *translations*

$$\llbracket \pi \rrbracket_{\text{Sep}} \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_{\text{S}} \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_{\text{L}} \quad (3.7)$$

between cobordisms living respectively in $\mathfrak{A}_{\text{Sep}}(\Gamma)$, $\mathfrak{A}_{\text{S}}(\mathfrak{Q})$ and $\mathfrak{A}_{\text{L}}(\mathfrak{Q})$. Here, we write \mathfrak{Q} for the domain of Γ . In order to clarify the functorial nature of these translations, one starts by observing the existence of internal functors between the **internal J -opcategories**:

$$\mathfrak{A}_{\text{Sep}}(\Gamma) \xrightarrow{u_{\mathcal{S}}} \mathfrak{A}_{\text{S}}(\mathfrak{Q}) \xrightarrow{u_{\mathcal{L}}} \mathfrak{A}_{\text{L}}(\mathfrak{Q}).$$

The first internal functor $u_{\mathcal{S}}$ transports every separated state $(\sigma_C, \vec{\sigma}, \sigma_F)$ into the machine state obtained by multiplying all its components as in Definition 2.5.4 and by forgetting all the permissions. The second internal functor $u_{\mathcal{L}}$ forgets the memory from a machine state in order to obtain the corresponding lock state. The three systems of tiles which equip the synchronization templates $\mathfrak{A}_{\text{Sep}}$, \mathfrak{A}_{S} and \mathfrak{A}_{L} were carefully designed in order to ensure that these internal functors do indeed exist. Since we can compare by a simulation two cobordisms defined over the same internal J -opcategory, we can also compare by “change-of-basis” two cobordisms over different internal J -opcategories. This leads to the following definition.

Definition 3.8.1. Given two **internal J -opcategories** \mathfrak{A} and \mathfrak{A}' , and two cobordisms

$$\mathbf{C} : A \rightarrow B \text{ in } \mathbf{Cob}(\mathfrak{A}) \quad \text{and} \quad \mathbf{C}' : A' \rightarrow B' \text{ in } \mathbf{Cob}(\mathfrak{A}')$$

where we write A to denote the 0-cell $A \rightarrow \mathfrak{A}[0, i]$ of $\mathbf{Cob}(\mathfrak{A})$, etc, a **map of cobordisms** $(\mathcal{F}, \mathcal{F}_A, \mathcal{F}_B, \mathcal{F}_{\mathfrak{A}})$ from \mathbf{C} to \mathbf{C}' is the data of an **internal functor of J -opcategories** $\mathcal{F}_{\mathfrak{A}} : \mathfrak{A} \rightarrow \mathfrak{A}'$ and of a 2-cell \mathcal{F} and two vertical 1-cells \mathcal{F}_A and \mathcal{F}_B in $\mathbf{Cob}(\mathfrak{A}')$ of the form:

$$\begin{array}{ccc} \text{push}[\mathcal{F}_{\mathfrak{A}}](A) & \xrightarrow{\text{push}[\mathcal{F}_{\mathfrak{A}}](\mathbf{C})} & \text{push}[\mathcal{F}_{\mathfrak{A}}](B) \\ \mathcal{F}_A \downarrow & \Downarrow \mathcal{F} & \downarrow \mathcal{F}_B \\ A' & \xrightarrow{\mathbf{C}'} & B' \end{array}$$

This means that the translations in (3.7) are simulations:

$$\text{push}[u_{\mathcal{S}}](\llbracket \pi \rrbracket_{\text{Sep}}) \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_{\text{S}} \quad \text{and} \quad \text{push}[u_{\mathcal{L}}](\llbracket C \rrbracket_{\text{S}}) \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_{\text{L}}.$$

Because the semantic interpretations of programs and proofs are defined by structural induction, we need the semantic combinators we use to “preserve” maps of cobordisms. The following lemmas provide conditions for their preservation.

The definition of *internal functor of J -opcategories* implies that we have a natural family of isomorphisms

$$\text{push}[u](C; D) \xrightarrow{\text{iso}} \text{push}[u](C); \text{push}[u](D)$$

in the ambient category $\mathbb{S} = \mathbf{AsyncGraph}$. Intuitively, the reason why this is an isomorphism is that, in both cases, we perform the same pushout for cospan composition $C, D \mapsto C; D$ at the level of the supports of games and cobordisms. The following lemma gives a condition for the sequential product to preserve maps of cobordisms.

Lemma 3.8.2. *Let \pm and \pm' be two *internal J -opcategories* related by an *internal functor of J -opcategories* $u : \pm \rightarrow \pm'$. Consider two composable cobordisms $\mathbf{C} : I \rightarrow M$ and $\mathbf{D} : M \rightarrow O$ in $\mathbf{Cob}(\pm)$, and two composable cobordisms $\mathbf{C}' : I' \rightarrow M'$ and $\mathbf{D}' : M' \rightarrow O'$ in $\mathbf{Cob}(\pm')$. Suppose finally that they are related by two *maps of cobordisms* $\mathcal{F} : \mathbf{C} \rightarrow \mathbf{C}'$ and $\mathcal{G} : \mathbf{D} \rightarrow \mathbf{D}'$ such that $u = \mathcal{F}_{\pm} = \mathcal{G}_{\pm}$ and $\mathcal{F}_M = \mathcal{G}_M$. Then there is a map of cobordisms*

$$\mathcal{F}; \mathcal{G} : \mathbf{C}; \mathbf{D} \rightarrow \mathbf{C}'; \mathbf{D}'$$

Proof. The desired 2-cell in $\mathbf{Cob}(\pm')$ is the following, where the top invertible 2-cell is the isomorphism given above.

$$\begin{array}{ccccc}
 \text{push}[u](I) & \xrightarrow{\text{push}[u](\mathbf{C}; \mathbf{D})} & & & \text{push}[u](O) \\
 \parallel & & \Downarrow \wr & & \parallel \\
 \text{push}[u](I) & \xrightarrow{\text{push}[u](\mathbf{C})} & \text{push}[u](M) & \xrightarrow{\text{push}[u](\mathbf{D})} & \text{push}[u](O) \\
 \mathcal{F}_I \downarrow & \mathcal{F} \Downarrow & \mathcal{F}_M \downarrow & \mathcal{G} \Downarrow & \mathcal{F}_O \downarrow \\
 I' & \xrightarrow{\mathbf{C}'} & M' & \xrightarrow{\mathbf{D}'} & O'
 \end{array}$$

□

We defer to Lemma 3.9.2 the explanation of why the condition $\mathcal{F}_M = \mathcal{G}_M$ is always true for the cobordisms which we compose sequentially.

The preservation of the maps of cobordisms by the action of the double functor \mathbf{Cob} of Lemma 3.3.4 is expressed in the following lemma.

Lemma 3.8.3. Let $\mathfrak{A}_1, \mathfrak{A}_2, \mathfrak{A}'_1$ and \mathfrak{A}'_2 be *internal J-opcategories*, $\mathbf{C} : A \dashrightarrow B$ be a cobordism in $\mathbf{Cob}(\mathfrak{A}_1)$ and $\mathbf{C}' : A' \dashrightarrow B'$ be one in $\mathbf{Cob}(\mathfrak{A}'_1)$. Suppose moreover given two acute spans (F, G) and (F', G') related as follows in the double category $\mathbf{AcuteSpan}$:

$$\begin{array}{ccc} \mathfrak{A}_1 & \xrightarrow{(F,G)} & \mathfrak{A}_2 \\ u \downarrow & \Downarrow \alpha & \downarrow v \\ \mathfrak{A}'_1 & \xrightarrow{(F',G')} & \mathfrak{A}'_2 \end{array}$$

Then a *map of cobordisms* $\mathcal{F} : \mathbf{C} \rightarrow \mathbf{C}'$ such that $\mathcal{F}_{\mathfrak{A}} = u$ induces a map of cobordisms from $\mathbf{Cob}((F, G))(\mathbf{C})$ to $\mathbf{Cob}((F', G'))(\mathbf{C}')$ which is above v .

Proof. The induced map of cobordism is given by the following vertical composition of 2-cells in $\mathbf{Cob}(\mathfrak{A}'_2)$, and the lax double functor structure of $\mathbf{Cob}(\cdot)$.

$$\begin{array}{ccccc} \text{push}[v](\mathbf{Cob}((F, G))(A)) & \xrightarrow{\text{push}[v](\mathbf{Cob}((F, G))(\mathbf{C}))} & \text{push}[v](\mathbf{Cob}((F, G))(B)) & & \\ \text{Cob}(\alpha)_A \downarrow & & \text{Cob}(\alpha)_B \downarrow & & \\ \mathbf{Cob}((F', G'))(\text{push}[u](A)) & \xrightarrow{\mathbf{Cob}((F', G'))(\text{push}[u](\mathbf{C}))} & \mathbf{Cob}((F', G'))(\text{push}[u](B)) & & \\ \text{Cob}((F', G'))(\mathcal{F}_A) \downarrow & & \text{Cob}((F', G'))(\mathcal{F}_B) \downarrow & & \\ \mathbf{Cob}((F', G'))(A') & \xrightarrow{\mathbf{Cob}((F', G'))(\mathbf{C}')} & \mathbf{Cob}((F', G'))(B') & & \end{array}$$

□

3.8.2 The asynchronous soundness theorem

As we have explained in the previous chapter, the asynchronous theorem of CSL uses the notion of fibration. Definition 2.4.6 introduces different notions of fibration for asynchronous morphisms. We can lift this notion to maps of cobordisms by considering the asynchronous morphism between the supports which is induced by such a map.

We can now state the asynchronous soundness theorem for CSL which relies on these two notions of fibrations in order to express the *safety* and *data-race freedom* of the code in a clean topological way. In order to prove these two properties, the theorem focuses on the nature of the asynchronous comparison maps between cobordisms

$$\llbracket \pi \rrbracket_{\text{Sep}} \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_{\text{S}} \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_{\text{L}}$$

discussed in §3.8.1. The theorem states:

Theorem 3.8.4 (Asynchronous soundness theorem). *For every CSL proof π of $\Gamma \vdash \{P\}C\{Q\}$, the comparison map $\mathcal{S} : \llbracket \pi \rrbracket_{\text{Sep}} \rightarrow \llbracket C \rrbracket_{\mathcal{S}}$ is a Code 1-fibration, and the comparison map $\mathcal{L} \circ \mathcal{S} : \llbracket \pi \rrbracket_{\text{Sep}} \rightarrow \llbracket C \rrbracket_{\mathcal{L}}$ is a 2-fibration.*

As we explained in Section 2.5.3, the first part of the theorem ensures that a program specified in CSL does not crash. Indeed, the cobordism $\llbracket \pi \rrbracket_{\text{Sep}}$ lives above $\varkappa_{\text{Sep}}[1]$ which does not include the error state \downarrow . Since every step performed by the Code in $\llbracket C \rrbracket_{\mathcal{S}}$ can be lifted to $\llbracket \pi \rrbracket_{\text{Sep}}$ by the 1-fibrational property, the specified Code cannot produce any error. The second part of the theorem ensures that a specified program does not produce nor encounter any data race. Indeed, every time two instructions are executed in parallel in the machine, they define a tile in the cobordism $\llbracket C \rrbracket_{\mathcal{L}}$, which can be lifted by the 2-fibrational property to a tile in the cobordism $\llbracket \pi \rrbracket_{\text{Sep}}$ of separated states. There, the very existence of the tile implies that these two instructions are independent and do not produce any data race.

3.9 Proof of the asynchronous soundness theorem

The general method for proving Theorem 3.8.4 is to prove by structural induction over the proof trees of CSL that the maps \mathcal{S} and \mathcal{L} exist, have the fibrational properties stated in the theorem for the axioms of the logic, and that the maps and these properties are preserved by the rules of the logic.

3.9.1 Well-formed cobordisms

To do so, we need to refine the kind of cobordisms we use to interpret programs and proofs: We require that they satisfy a number of properties.

Definition 3.9.1. A cobordism

$$\begin{array}{ccccc}
 I & \xrightarrow{\text{in}} & S & \xleftarrow{\text{out}} & O \\
 \lambda_I \downarrow & & \downarrow \lambda_\sigma & & \downarrow \lambda_O \\
 \varkappa[0, i] & \xrightarrow{\text{in}_{ij}} & \varkappa[1, ij] & \xleftarrow{\text{out}_{ij}} & \varkappa[0, j]
 \end{array}$$

in an ambient category \mathbb{S} is said to be *well-formed* if:

1. the map $\lambda_I : I \rightarrow \varkappa[0, i]$ is an isomorphism;

2. the map λ_σ is an **Environment 1-fibration**;
3. the map $\text{out} : O \rightarrow S$ is a mono;
4. the pullback of $\text{in} : I \rightarrow S$ along $\text{out} : O \rightarrow S$ is the initial object of \mathcal{S} .

The conditions 1. and 2. are related to receptivity in games semantics: the first means that any initial state in $\pm[0, i]$ must be “accepted” by the program or the proof, and the second means that the Frame can always play any valid transition, or move. Condition 3. and 4. are more technical and are needed to ensure that sequential composition behaves nicely with respect to fibrational properties.

The condition that λ_I be a monomorphism (since it is an isomorphism) is used to ensure that the maps of cobordisms \mathcal{S} and \mathcal{L} always exist.

Lemma 3.9.2. *Consider two composable cobordisms $\mathbf{C} : I \rightarrow M$ and $\mathbf{D} : M \rightarrow O$ in $\mathbf{Cob}(\pm)$, and two composable cobordisms $\mathbf{C}' : I' \rightarrow M'$ and $\mathbf{D}' : M' \rightarrow O'$ in $\mathbf{Cob}(\pm')$. Suppose finally that they are related by two *maps of cobordisms* $\mathcal{F} : \mathbf{C} \rightarrow \mathbf{C}'$ and $\mathcal{G} : \mathbf{D} \rightarrow \mathbf{D}'$ such that $\mathcal{F}_\pm = \mathcal{G}_\pm$. If \mathbf{D}' is *well-formed*, then $\mathcal{F}_M = \mathcal{G}_M$.*

Proof. This follows from $\lambda'_M \circ \mathcal{F}_M = u \circ \lambda_M = \lambda'_M \circ \mathcal{G}_M$, where $u = \mathcal{F}_\pm = \mathcal{G}_\pm$ and from the fact that λ'_M is a mono. \square


We also require that the **internal J -opcategories** we use to interpret have good fibrational properties:

Definition 3.9.3. An **internal J -opcategory** \pm is called **\mathcal{M} -fibrational** if the maps $\text{in}_{ij} : \pm[0, i] \rightarrow \pm[1, ij]$ and $\text{out}_{ij} : \pm[0, j] \rightarrow \pm[1, ij]$ are monos, **\mathcal{M} -1-fibrations** and **2-fibrations**, and if the map $\pm[2, ijk] \rightarrow \pm[1, ik]$ are an **\mathcal{M} -1-fibration**. This induces the notion of **Environment-fibrational** internal J -opcategories.

Every internal J -opcategories \pm_{Sep} , \pm_S and \pm_L we have seen so far is Environment-fibrational.

Before we proceed to proving that all operations we use on cobordisms preserve well-formedness, we characterize fibrations as maps which satisfy a right lifting property. Let us write $\cdot \rightarrow \cdot$ for the asynchronous graph with two nodes 0 and 1 and a unique edge $0 \rightarrow 1$. We write \rightarrow_C when this unique edge is considered a Code transition, and \rightarrow_F when it is considered as a Frame transition. We write \cdot for the asynchronous graph with a unique node, and no edge or tile. We denote by $s : \cdot \rightarrow \cdot \rightarrow \cdot$ the asynchronous morphism which maps that single node to the node 0 of $\cdot \rightarrow \cdot$.


Lemma 3.9.4. Let $f : G \rightarrow G'$ be an *asynchronous morphism* between asynchronous graphs with a notion of Code and Frame edge. Then f is a *Code 1-fibration* (respectively an *Environment 1-fibration*) iff, for all squares of polarity preserving morphisms as in the left-hand diagram below (respectively the right-hand diagram), there exists a map denoted by the dashed arrow which makes both triangles commute.



Proof. Follows easily from the definitions. □

We can characterize *2-fibrations* in the same way. Let Λ be the asynchronous graph with three nodes $0, 1, 2$ and two edges $0 \rightarrow 1 \rightarrow 2$, and let \blacklozenge be the asynchronous graph with 4 nodes $0, 1, 1', 2$ and four edges $0 \rightarrow 1 \rightarrow 2$ and $0 \rightarrow 1' \rightarrow 2$ and a unique tile which sits on the two paths of length two above. Write $u : \Lambda \rightarrow \blacklozenge$ for the inclusion, mapping the path of length two in Λ to the “upper path” of the tile in \blacklozenge . Similarly, we have:


Lemma 3.9.5. Let $f : G \rightarrow G'$ be an *asynchronous morphism* between two asynchronous graphs. Then f is a *2-fibration* iff, for all squares as in the diagram below, there exists a map denoted by the dashed arrow which makes both triangles commute.



Proof. Follows easily from the definitions. □

One immediate consequence of this diagrammatic characterization of fibrations is that they are stable under pullbacks, in the following sense.

Lemma 3.9.6. Given an morphism $f : G \rightarrow G'$ in \mathbb{S} , suppose g its pullback along some map $u : H' \rightarrow G'$:



then g is a fibration of some type (Code or Environment 1-fibrations, or 2-fibrations) if f is.

Proof. Assume f is a fibration characterized by a right lifting property with respect to a map $i : X \rightarrow Y$

$$\begin{array}{ccccc}
 X & \longrightarrow & H & \longrightarrow & G \\
 \downarrow i & & \downarrow g & \text{p.b.} & \downarrow f \\
 Y & \longrightarrow & H' & \xrightarrow{u} & G'
 \end{array}$$

then there exists a map $a : Y \rightarrow G$ making the two triangles commute. The universal property of the pullback implies that there exists a unique map $a' : Y \rightarrow H$ which makes the two triangles commute, and therefore g is a fibration as well. \square

Another important observation, especially when dealing with colimits, is that of all the “test” objects $(\cdot, \cdot \rightarrow \cdot, \wedge, \blacklozenge)$ only \wedge is not representable in **AsyncGraph**, and hence is not tiny, since, in presheaf categories, tiny objects are the retracts of representables.

Definition 3.9.7. An object A of a category \mathcal{C} is called *tiny* if the functor $\mathcal{C}(A, -)$ preserves colimits.

We now introduce another notion, *adhesivity*, which is going to be useful to reason about sequential composition of cobordisms.


3.9.2 Adhesivity of AsyncGraph

There are several equivalent definitions of *adhesivity*, the one we present here, which is not the tersest, has been chosen because it is used directly in the sequel.

Definition 3.9.8. A category \mathcal{C} is *adhesive* if all pushouts along monos exist, and if all such pushout squares are *van Kampen squares*. A commutative square

$$\begin{array}{ccc}
 A & \xrightarrow{g} & C \\
 \downarrow f & & \downarrow v \\
 B & \xrightarrow{u} & D
 \end{array}$$

is said to be a *van Kampen square* if, for every cubical commutative diagram of the form



whose two “back faces” (W, A, X, B and W, A, Y, C) are pullbacks, the top face is a pushout if and only if the two front faces (X, Z, B, D and Y, Z, C, D) are pullbacks. We call such a commutative cube a *van Kampen cube*.

The category **AsyncGraph** of asynchronous graphs is adhesive since it is a presheaf category, and therefore a topos.

The following consequences of adhesivity will be useful, see for example (Lack and Sobociński, 2004).

Lemma 3.9.9. *In an adhesive category, pushouts of monos are monos, and the resulting pushout squares are also pullback squares.*

Another useful fact is that a the coproducts are well-behaved as soon as an adhesive category has strict initial objects.

Lemma 3.9.10. *Any adhesive category \mathcal{S} which has a strict initial object is extensive. In particular, this means that its coproducts are disjoint, which means that coprojections are monic, and their pullback is an initial object in \mathcal{S} .*

3.9.3 Preservation of well-formedness

We prove that the operations that we use to interpret programs and proofs preserve the *well-formedness* of cobordisms. The cases of operations such as the parallel product and the two change of locks operations follows from the lemma below, which relies on the following property of the acute spans we use.

Definition 3.9.11. An acute span **S** of internal J -opcategories

$$\mathfrak{A}_1 \longleftarrow^F \mathfrak{A}_2 \longrightarrow^G \mathfrak{A}_3$$

is called **strict** if the following squares are pullbacks:

$$\begin{array}{ccccc}
 \mathfrak{A}_2[0, i] & \xrightarrow{\text{in}_i} & \mathfrak{A}_2[1] & \xleftarrow{\text{out}_j} & \mathfrak{A}_2[0, j] \\
 F[0, i] \downarrow & & \downarrow F[1] & & \downarrow F[0, j] \\
 \mathfrak{A}_1[0, i] & \xrightarrow{\text{in}_i} & \mathfrak{A}_1[1] & \xleftarrow{\text{out}_j} & \mathfrak{A}_1[0, j]
 \end{array}$$

Lemma 3.9.12. *Given a **strict acute span** \mathbf{S} of **internal J -opcategories***

$$\mathfrak{A}_1 \xleftarrow{F} \mathfrak{A}_2 \xrightarrow{G} \mathfrak{A}_3$$

*the operation on cobordisms $X \mapsto \mathbf{Cob}(\mathbf{S})(X)$ preserves well-formedness (except for condition 1 of Definition 3.9.1) if $G[1]$ is an **Environment 1-fibration** and a mono. If moreover the maps $G[0, i]$ are isos, then Condition 1 is also preserved.*

Proof. Let the following be a well-formed cobordism over the **internal J -opcategory** \mathfrak{A}_1 :

$$\begin{array}{ccccc}
 I & \xrightarrow{\text{in}} & S & \xleftarrow{\text{out}} & O \\
 \lambda_I \downarrow & & \downarrow \lambda_\sigma & & \downarrow \lambda_O \\
 \mathfrak{A}_1[0, i] & \xrightarrow{\text{in}_{ij}} & \mathfrak{A}_1[1, ij] & \xleftarrow{\text{out}_{ij}} & \mathfrak{A}_1[0, i]
 \end{array}$$

and let the following cobordism over \mathfrak{A}_2 be obtained by pulling back along the components of the **internal functor** F using the pull[F] **lax double functor** of Lemma 3.3.1:

$$\begin{array}{ccccc}
 I' & \xrightarrow{\text{in}'} & S' & \xleftarrow{\text{out}'} & O' \\
 \lambda_{I'} \downarrow & & \downarrow \lambda_{\sigma'} & & \downarrow \lambda_{O'} \\
 \mathfrak{A}_2[0, i] & \xrightarrow{\text{in}_{ij}} & \mathfrak{A}_2[1, ij] & \xleftarrow{\text{out}_{ij}} & \mathfrak{A}_2[0, i]
 \end{array}$$

We need to check the four conditions of Definition 3.9.1.

- (1) Isos are stable under pullback in **AsyncGraph**, so $\lambda_{I'}$ is an iso as well. If the condition that $G[0, i]$ is an iso holds, the operation $\text{push}[F]$ also preserves Condition 1.
- (2) Fibrations are stable under pullback, so $\lambda_{\sigma'}$ is also a 1-fibration. The hypothesis on $G[1]$ implies that Condition (2) is preserved by post-composing.
- (3) The map $\text{out}' : O' \rightarrow S'$ is a mono because, according to the pasting lemma on pullbacks and the fact that $G[1]$ is a mono, it is the pullback of $\text{out} : O \rightarrow S$ along the map $O' \rightarrow O$ which is part of the pullback defining O' .


- (4) The universal property of pullbacks implies that there is a map from the pullback X of I' and O' to the pullback of I and O , which is an initial object of **AsyncGraph** since it is a **well-formed** cobordism. Since **AsyncGraph** has strong initial objects, X is an initial object as well.

□

All the operations on cobordisms, the parallel product, hiding and lifting of locks (except for Condition 1, see below), are defined using acute spans which satisfy the conditions of the lemma above, and therefore they preserve well-formedness.

To prove that the other operations also preserve well-formedness, we need to prove that sequential composition preserves well-formedness. This fact relies on the following result.

Lemma 3.9.13. *Suppose given the following diagram:*




where λ_1, λ_2, i and j are Environment 1-fibrations, and i, j are monos. Then the map μ induced by the universality of the pushout is an Environment 1-fibration as well.

Proof. First, because **AsyncGraph** is **adhesive**, i' and j' are monos as well. Suppose we have a node x in Z which is mapped to the source node of an Environment transition t in D . Because the walking arrow is tiny, either B or C contains a transition t' which is mapped to the transition t . Without loss of generality, suppose that it is in B . Similarly, the node x has an antecedent x' in either X or Y . Suppose first that it is in X . Then, because j' is a monomorphism, it must be that this node is mapped by λ_1 to the source node of t' , and then we conclude using the fact that λ_1 is an Environment 1-fibration. Suppose now that this node x' is in Y . Then the node $\lambda_2(x')$ in C is mapped to the same node in D as the source node of t' , therefore there is a node y in A which is mapped under j to the source node of t' , and under i to $\lambda_2(x')$. We get back to the previous case by using the fact that j is an Environment 1-fibration. □

It is now easy to prove that sequential composition preserves well-formedness.

Lemma 3.9.14. *Let $\mathbf{C} : I \rightarrow I'$ and $\mathbf{D} : I' \rightarrow O'$ be two well-formed cobordisms over one of the *internal J -opcategories* we have defined above. Then $\mathbf{C}; \mathbf{D}$ is well-formed. This is also the case if \mathbf{D} does not satisfy Condition 1 of Definition 3.9.1.*

Proof. Condition 1 is obvious, Condition 2 is the lemma above, Condition 3 follows from the stability of monos under pullback. Condition 4 follows from the fact that pushouts along monos are pullbacks in the diagram




where X is the pullback of the map $I \rightarrow C;D$ along $O' \rightarrow C;D$, and where the four “small” squares are pullbacks for the reason above. □

The operation for which we do not yet have the tools to prove that it preserves well-formedness is the generalized sequential composition with *filling systems*, because they are not well-formed cobordisms, as they fail to satisfy Conditions 3 and 4.


Lemma 3.9.15. *Given the following commutative diagram, where the square is a pushout*

and o' is a mono:



if the pullback of i' and o' is the initial object, then the composite $o' \circ g$ is a mono. Moreover, the pullback of O' and C is the initial object.


Proof. First, we establish that the C' contains the disjoint union of the image of M under i' and the image of O' under o' . For that purpose, consider an epi-mono factorization $i'_2 \circ i'_1$ of i' and the following diagram, where the two squares are pullbacks:



As epimorphisms are preserved under pullbacks in **ASyncGraph**, the map $\emptyset \rightarrow X$ is an epi, which means that X is the initial object \emptyset . From that, we deduce that the map

$$I' + O' \longrightarrow C'_M$$

is a monomorphism, using the fact that colimits are universal. This is an instance of the general fact that the union of two subobjects is given by the pushout above the pullback. Now, consider the following diagram, where all squares are pushouts:



Because monos are preserved under pushouts, the composite $k \circ (h + \text{Id}) \circ \iota_2 = k_2 = o'$ is indeed a mono, where we write $k = [k_1, k_2]$.

Finally, $C \times_D O' \cong \emptyset$ because $C \times_{A+O'} O' \cong \emptyset$ (coproducts are disjoint, because adhesive categories are extensive if they have a strict initial object as we noted in Lemma 3.9.10) and $k : A + O' \rightarrow D$ is a mono. \square

Fibrational properties of the disjoint sum Suppose the ambient category \mathbb{S} has disjoint sum, which is the case of **AsyncGraph**. Then

Lemma 3.9.16. *The sum of two fibrations is a fibration.*

The three family of template which we use have a natural structure of **internal J -opcategory with sums**: for the two template \mathfrak{A}_S and \mathfrak{A}_L , this is automatic as they are **internal opcategories**, and for the separated states, the map \vee simply computes the disjunction of the formulas. We note that in all cases, the map $\nabla : \mathfrak{A}[1] + \mathfrak{A}[1] \rightarrow \mathfrak{A}[1]$ is an **Environment 1-fibration** and a **2-fibration**. This ensures in particular that the disjoint sum of cobordisms preserves **well-formedness**, except for Condition 1, which can be recovered by precomposing with a suitable filling system when adequate.

Putting all of this together, we prove:

Proposition 3.9.17. *All operations which are used to interpret programs preserve **well-formedness**.*

Proof. For sequential composition, while loops and conditionals, this follows from Lemma 3.9.14 and Lemma 3.9.15. Parallel composition and lock allocation follows from Lemma 3.9.12 about the action of acute spans, and critical sections follow from the aforementioned lemma and sequential composition (this is where the caveats about Condition 1 in the lemmas above are useful). \square

3.9.4 Strict maps of cobordisms

The proof of the asynchronous soundness theorem relies on the fact that the **maps of cobordisms** \mathcal{S} and \mathcal{L} are **strict** in the following sense:

Definition 3.9.18. A map of cobordisms $(\mathcal{F}, \mathcal{F}_\pm)$ from $\mathbf{C} \in \mathbf{Cob}(\pm_1)$ to $\mathbf{D} \in \mathbf{Cob}(\pm_2)$ is *strict* if the two squares below, with the same notations as Definition 3.8.1, are pullbacks.

$$\begin{array}{ccccc} I & \xrightarrow{\text{inc}} & C & \xleftarrow{\text{out}_C} & O \\ \mathcal{F}_I \downarrow & & \downarrow \mathcal{F}_C & & \downarrow \mathcal{F}_O \\ I' & \xrightarrow{\text{inc}} & D & \xleftarrow{\text{out}_D} & O' \end{array} \quad \begin{array}{c} pb \\ pb \end{array}$$

We prove in this section that the maps of cobordisms \mathcal{S} and \mathcal{L} are *strict*. We begin with operations based on the pull-and-push construction along *acute spans*

Lemma 3.9.19. Let $\mathbf{C} \in \mathbf{Cob}(\pm_1)$ and $\mathbf{C}' \in \mathbf{Cob}(\pm'_1)$ be cobordisms related by a *strict map of cobordisms* $\mathcal{F} : \mathbf{C} \rightarrow \mathbf{C}'$, and let

$$\pm_1 \xleftarrow{F} \pm_2 \xrightarrow{G} \pm_3 \quad \pm'_1 \xleftarrow{F'} \pm'_2 \xrightarrow{G'} \pm'_3$$

be two *strict acute spans* related by three *internal functors of J-opcategories* u_1, u_2 and u_3 as in the following commutative diagram

$$\begin{array}{ccccc} \pm_1 & \xleftarrow{F} & \pm_2 & \xrightarrow{G} & \pm_3 \\ u_1 \downarrow & & \downarrow u_2 & & \downarrow u_3 \\ \pm'_1 & \xleftarrow{F'} & \pm'_2 & \xrightarrow{G'} & \pm'_3 \end{array}$$

Then the induced map of cobordisms $\mathcal{F}' : \mathbf{Cob}((F, G))(\mathbf{C}) \rightarrow \mathbf{Cob}((F', G'))(\mathbf{C}')$ is a *strict map of cobordisms*.

Proof. Write the cobordism \mathbf{C} as the diagram

$$\begin{array}{ccccc} I & \xrightarrow{\text{in}} & S & \xleftarrow{\text{out}} & O \\ \lambda_I \downarrow & & \downarrow \lambda_\sigma & & \downarrow \lambda_o \\ \pm_1[0, i] & \xrightarrow{\text{in}_{ij}} & \pm_1[1, ij] & \xleftarrow{\text{out}_{ij}} & \pm_1[0, i] \end{array}$$

and write

$$\begin{array}{ccccc} J & \xrightarrow{\text{in}_{pb}} & T & \xleftarrow{\text{out}_{pb}} & P \\ \lambda_I \downarrow & & \downarrow \lambda_\tau & & \downarrow \lambda_o \\ \pm_2[0, i] & \xrightarrow{\text{in}_{ij}} & \pm_2[1, ij] & \xleftarrow{\text{out}_{ij}} & \pm_2[0, i] \end{array}$$

the result of pulling back the cobordism \mathbf{C} along the **plain internal functor** F , which is part of the definition of the action of the **acute span** (F, G) on \mathbf{C} .

First, we claim that the following squares induced by the pullbacks is a pullback square itself.

$$\begin{array}{ccccc}
 I & \xrightarrow{\text{in}} & S & \xleftarrow{\text{out}} & O \\
 \uparrow & & \uparrow & & \uparrow \\
 J & \xrightarrow{\text{in}_{pb}} & T & \xleftarrow{\text{out}_{pb}} & P
 \end{array}$$

Since the two cases are symmetric, we prove that the left square is a pullback. This follows from the pasting lemma of pullbacks in the following commutative cube

$$\begin{array}{ccccc}
 I & \xrightarrow{\quad} & S & & \\
 \downarrow & \swarrow & \downarrow & \swarrow & \\
 J & \xrightarrow{\quad} & T & & \\
 \downarrow & \swarrow & \downarrow & \swarrow & \\
 \mathfrak{t}_1[0] & \xrightarrow{\quad} & \mathfrak{t}_1[1] & & \\
 \downarrow & \swarrow & \downarrow & \swarrow & \\
 \mathfrak{t}_2[0] & \xrightarrow{\quad} & \mathfrak{t}_2[1] & &
 \end{array}$$

since the two side faces are pullbacks by definition, and the bottom face is one because (F, G) is a **strict acute span** by assumption. We have the same result for the pullback of \mathbf{C}' along F' .

Finally, we conclude by using again the pasting lemma for pullback in the cube

$$\begin{array}{ccccc}
 I & \xrightarrow{\quad} & S & & \\
 \downarrow \mathcal{F}_I & \swarrow & \downarrow \mathcal{F}_S & \swarrow & \\
 J & \xrightarrow{\quad} & T & & \\
 \downarrow & \swarrow & \downarrow & \swarrow & \\
 I' & \xrightarrow{\quad} & S' & & \\
 \downarrow & \swarrow & \downarrow & \swarrow & \\
 J' & \xrightarrow{\quad} & T' & &
 \end{array}$$


We have just proved that the top and the bottom faces are pullbacks, and the hypothesis that \mathcal{F} is a **strict map of cobordisms** means that the back face is a pullback as well. We conclude that the front face is one as well, which is what we needed to prove. \square

The second basic operation we need to prove preserves strictness of maps of cobordisms is sequential composition. This fact relies heavily on the fact that the ambient category **AsyncGraph** is **adhesive**.

Lemma 3.9.20. *Let $\mathbf{C}, \mathbf{D} \in \mathbf{Cob}(\oplus_1)$ and $\mathbf{C}', \mathbf{D}' \in \mathbf{Cob}(\oplus'_1)$ be cobordisms related by two *strict maps of cobordisms* $\mathcal{F} : \mathbf{C} \rightarrow \mathbf{C}'$ and $\mathcal{G} : \mathbf{D} \rightarrow \mathbf{D}'$. Suppose moreover that \mathbf{C}' and \mathbf{D}' satisfy Condition 3 of Definition 3.9.1.*

*Then the induced map of cobordisms $\mathcal{F}; \mathcal{G} : \mathbf{C}; \mathbf{D} \rightarrow \mathbf{C}'; \mathbf{D}'$ is *strict*.*

Proof. The situation is described by the following diagram:



where the non-dashed vertical arrows are given by \mathcal{F} and \mathcal{G} , and the map $M' \rightarrow C'_1$ is a mono, since the cobordism \mathbf{C}' satisfies Condition 3.

We focus on the central commutative cube. The bottom face is a pushout along a mono, and therefore it is a *van Kampen square* in **AsyncGraph**. The two back faces are pullbacks since we assume that \mathcal{F} and \mathcal{G} are *strict*. Since the top face is a pushout by the definition of sequential composition, *adhesivity* of **AsyncGraph** implies, by definition, that the two front faces are pullbacks. We conclude by the pasting lemma of pullbacks. \square

Strictness of maps of *well-formed* cobordisms is stable under generalized sequential composition with *filling systems* follows from the following lemma.

Lemma 3.9.21. *Let $\lambda : A \rightarrow \oplus[0, i]$ and $\mu : B \rightarrow \oplus[0, j]$ be two *colored-games* over an *internal J-opcategory* which is *Environment-fibrational*. If μ is a mono, then the cobordism $\text{fill}(\lambda, \mu)$ satisfies Condition 3 of Definition 3.9.1.*

Proof. Follows from the definition of $\text{fill}(\lambda, \mu)$, and from the fact that monos are stable under pullback and pushout in **AsyncGraph**. \square

Similarly to the previous section, the lemmas above imply the following.

Proposition 3.9.22. *Given a proof π of some CSL Hoare triple $\Gamma \vdash \{P\} C \{Q\}$, the two maps*

$$\llbracket \pi \rrbracket_{\text{Sep}} \xrightarrow{\mathcal{S}} \llbracket C \rrbracket_{\mathcal{S}} \xrightarrow{\mathcal{L}} \llbracket C \rrbracket_{\mathcal{L}}$$

are strict maps of cobordisms.

3.9.5 Proof of 2-dimensional correctness

We begin with the proof that the constructions on cobordisms which we use preserve **2-fibrations** because it is simpler than proving that they preserve **Code 1-fibrations**.

We prove by induction on the structure of the derivation trees of the logic that the map $\mathcal{L} \circ \mathcal{S} : \llbracket \pi \rrbracket_{\text{Sep}} \rightarrow \llbracket C \rrbracket_{\mathcal{L}}$ is a **2-fibration**. The fact that the rule PAR for the parallel product preserves 2-fibration in the sense above uses the following technical lemma.

Lemma 3.9.23. *Consider an asynchronous morphism $s : S_1 \rightarrow S_2$ which we see as defining a notion of fibration through a right-lifting property, and another asynchronous morphism $g : G \rightarrow H$ which satisfies the following property:*

$$\forall l, r : S_2 \rightarrow G, (rs = ls \wedge gr = gl) \implies l = r. \quad (3.8)$$

Then, for another map $f : K \rightarrow G$, if $g \circ f$ is a $S_1 \rightarrow S_2$ fibration, then so is f . Moreover, maps which satisfy Condition 3.8) are stable under pullback. Monos are a particular case of such maps.

Proof. The second statement follows from the uniqueness property of the universal property of pullbacks, similarly to preservation of monos under pullbacks. The third is easy.

Let us consider the first statement. This situation is depicted with the following diagram:


$$\begin{array}{ccc} K & \xleftarrow{w} & S_1 \\ \downarrow f & \nearrow \ell & \downarrow s \\ G & \xleftarrow{t} & S_2 \\ \downarrow g & & \\ H & & \end{array}$$

where the full arrow commute and where the map ℓ is such that $s\ell = w$ and $gf\ell = gt$. Since moreover $ts = fw$, we can use the hypothesis to conclude that $f\ell = t$, which means that ℓ is indeed a lifting and that f is a fibration. \square

Lemma 3.9.24. *Let $(F, G) : \mathfrak{t}_1 \dashrightarrow \mathfrak{t}_3$ and $(F, G') : \mathfrak{t}'_1 \dashrightarrow \mathfrak{t}'_3$ be two acute spans related by a map of acute spans $\alpha : F \rightarrow G$. Suppose that $F : \mathfrak{t}_2 \rightarrow \mathfrak{t}_1$ is a 2-fibration, and that $F' : \mathfrak{t}_2 \rightarrow \mathfrak{t}_3$ satisfies Condition (3.8).*

If $\mathcal{L} \circ \mathcal{S} : \llbracket \pi \rrbracket_{Sep} \rightarrow \llbracket C \rrbracket_L$ is a 2-fibrations, then so is the induced map of cobordisms $\alpha(\mathcal{L} \circ \mathcal{S}) : \mathbf{Cob}(F)(\llbracket \pi \rrbracket_{Sep}) \rightarrow \mathbf{Cob}(G)(\llbracket C \rrbracket_L)$.

Proof. Consider the following diagram, where f is the component of $\mathcal{L} \circ \mathcal{S}$ between the supports:



and where the two squares on the sides are pullbacks.

First, since the map $F : \mathfrak{t}_2 \rightarrow \mathfrak{t}_1$ is a 2-fibration, and because the fibrations are preserved under pullbacks, the map $T \rightarrow S$ is a 2-fibration as well. Similarly, the map $T' \rightarrow S'$ satisfies Condition (3.8), since it is the pullback of F' . We conclude using Lemma 3.9.23 above. \square


Since 2-fibrations are closed under product, and $\text{pick}[1] : \mathfrak{t}_L^{\parallel}[1] \rightarrow \mathfrak{t}_L[1] \times \mathfrak{t}_L[1]$ is a mono, the lemma above can be applied to the parallel product.

Corollary 3.9.25. *If $\mathcal{L}_1 \circ \mathcal{S}_1 : \llbracket \pi_1 \rrbracket_{Sep} \rightarrow \llbracket C_1 \rrbracket_L$ and $\mathcal{L}_2 \circ \mathcal{S}_2 : \llbracket \pi_2 \rrbracket_{Sep} \rightarrow \llbracket C_2 \rrbracket_L$ are 2-fibrations, then so is the induced map of cobordisms $(\mathcal{L}_1 \circ \mathcal{S}_1) \parallel (\mathcal{L}_2 \circ \mathcal{S}_2) : \llbracket \pi_1 \rrbracket_{Sep} \parallel \llbracket \pi_2 \rrbracket_{Sep} \rightarrow \llbracket C_1 \rrbracket_L \parallel \llbracket C_2 \rrbracket_L$.*

Similarly, this lemma can be applied for the two change of locks operations as well.

The proof that sequential composition preserves 2-fibrations follows from the fact that van Kampen cubes preserve 2-fibrations in the following sense:

Lemma 3.9.26. Consider the following diagram in **AsyncGraph**:



such that:

1. λ_1 and λ_2 are 2-fibrations,
2. the cube is van Kampen,
3. the two back faces are pullbacks.

Then, the asynchronous map μ is a 2-fibration.

Proof. Let us assume there is a path of length 2 in Z which is mapped by μ to the top of a tile in D . We will use the fact that representables are tiny, in particular, tiles are tiny. Assume without loss of generality that the tile in D has a preimage T_C in C . Then the upper path of the tile in D is the target of a path of length 2 in Z and the target of a path of length 2 in C (the upper path of T_C). Since the cube is van Kampen, Y is the pullback of i' along μ , which means that there is a path of length 2 in Y that is mapped to the one in C . We can conclude from there using the fact that λ_2 is a 2-fibration. \square

3.9.6 Proof of 1-dimensional correctness

Preservation of Code 1-fibrations by sequential composition follows directly from a variant of Lemma 3.9.13 where all occurrences of Environment 1-fibrations are replaced with Code 1-fibrations. The proof for change of lock operations follows from Lemma 3.9.24.

The proof that the parallel product preserves Code 1-fibrations is a bit more intricate than the case of 2-fibrations. The reason is that the map $\text{pick} : \mathfrak{t}_S^\otimes[1] \rightarrow \mathfrak{t}_S[1] \times \mathfrak{t}[1]$ is not a Code 1-fibration, which means we cannot use Lemma 3.9.24. Instead, we need to rely on the fact that the map $S \rightarrow \mathfrak{t}_S[1]$ from the support of the cobordism to the template is an Environment 1-fibration and on the fact that the morphism pick above never pairs two Code transitions.

Lemma 3.9.27. *Given two maps of cobordisms*

$$F_i : C_i \rightarrow C'_i$$

for $i = 1, 2$, and if F_1 and F_2 are 1-fibrations on Code transitions, then the induced morphism

$$F_1 \parallel F_2 : C_1 \parallel C_2 \rightarrow C'_1 \parallel C'_2$$


is a 1-fibration on Code transitions as well.

Proof. Let us focus on the maps between the supports. We call the supports S_1, S_2, S'_1 and S'_2 , and we call the maps between that are contained in F_1 and F_2 :

$$f_1 : S_1 \rightarrow S'_1 \quad f_2 : S_2 \rightarrow S'_2$$


Finally, we write pick both for the map pick in the span-monoidal structure, and for the map it induces by pullback at the level of the supports of the cobordisms.

Recall that the map $f_1 \parallel f_2$ is defined by the following diagram, where (u, u^\parallel) is the span-monoidal functor structure:



The polarity of the image of the unique edge of \rightarrow_c in \pm^\parallel is either C_1 or C_2 . Because the situation is symmetric, we can suppose without loss of generality that this polarity

is C_1 . In that case, the edge \rightarrow_C is mapped to a Code transition in S'_1 . By assumption, f_1 is a 1-fibration on Code transitions, so we can lift this arrow to S_1 . Moreover, because $\pi \circ \text{pick} : \mathfrak{z}^\parallel \rightarrow \mathfrak{z}$ is a 1-fibration on Code transitions as well, we can also lift \rightarrow_C to \mathfrak{z}^\parallel . Therefore, we have the following situation:



The map (1) exists because f_1 is a 1-fibration on Code transitions, as mentioned above, and the map (2) exists because λ_2 is a 1-fibration on Environment moves. Therefore, we can lift the arrow \rightarrow_{C_1} to $S_1 \times S_2$. We conclude by using the fact that 1-fibrations on Code transitions are stable by pushout, and that $\text{pick} : \mathfrak{z}^\parallel \rightarrow \mathfrak{z} \times \mathfrak{z}$ is a 1-fibration. \square

Part II

Relational soundness in Iris

4 Background on Iris

Iris is a state of the art concurrent separation logic. Its goal was to unify the many variants of concurrent separation logic which were being developed to handle increasingly complicated concurrent algorithms or language features. The first step (Jung, Swasey, et al., 2015) was to axiomatize the notion of logical resource: all the resources the predicates are describing, such as the heap, or the heap with permission are *cameras*, a structure similar to a partial commutative monoid. *Cameras* can be combined to enable reasoning about complex concurrent data structures and algorithms with subtle synchronization patterns.

The Iris program logic is also parameterized by a programming language, which is expressed as a small step operational semantics. This allows users to prove properties of programs written in their language of choice without having to redefine and reprove the soundness of the logic from scratch.

Early versions of Iris were program logics, in that concepts such as Hoare triples or invariants were primitive notions. In later versions (Robbert Krebbers, Jung, et al., 2017), such notions became defined in terms of a more basic modal logic, which is called the *base logic*.

Importantly, the soundness of the program logic follows from an adequacy theorem of the base logic. This allows one to develop new program logics which provide other guarantees than the usual Hoare triples of Iris in a much simpler way than if one had to define it and prove its soundness from scratch. In this part of the thesis, we define a new notion of Hoare triple, which allows its users to construct simulations between a program and a state transition system which we see as an abstract specification for the program. First, we present in this Chapter the basics of the Iris logic.

All the result and the background in this part are formalized in the Coq proof assistant.

4.1 Overview of the approach

One way to explain Iris is to start from the semantic definition of safety which can be used to define the meaning of a Hoare triple $\{P\} C \{Q\}$ from Chapter 1 and which

$$\frac{\varepsilon \vDash Q}{\text{safe}(\text{skip}, \varepsilon, Q)} \qquad \frac{(C, \varepsilon) \not\rightarrow \downarrow \quad \forall C', \varepsilon', (C, \varepsilon) \rightarrow (C', \varepsilon') \Rightarrow \text{safe}(C', \varepsilon', Q)}{\text{safe}(C, \varepsilon, Q)}$$

Figure 4.1: Semantic validity for Hoare logic

we reproduce here in Figure 4.1 (recall the double bar denotes a coinductive predicate) Recall that validity of a Hoare triple was defined as:

$$\vDash \{P\} C \{Q\} \iff \forall \varepsilon, \varepsilon \vdash P \Rightarrow \text{safe}(C, \varepsilon, Q)$$

It is easy to check that defining a new predicate $\text{wp } C \{Q\}$ as follows:

$$\varepsilon \vDash \text{wp } C \{Q\} \iff \text{safe}(C, \varepsilon, Q)$$

defines a semantic weakest precondition for the Hoare triple. Iris proceeds by defining a semantic weakest precondition which follows the same intuition as the one in Figure 4.1. One difference is that, in Iris, the machine state ε is not a parameter of the predicate, instead it is part of the logical state, or world.

4.1.1 Iris predicates

The Iris base logic propositions are predicates over *worlds*, and the Iris logic is parameterized over the set \mathcal{W} of words. Iris propositions are predicates over worlds, that is, functions of type $iProp := \mathcal{W} \rightarrow \mathbb{P}$, where \mathbb{P} is the two-element set $\mathbb{2}$ or the type Prop in Coq. It is conventional to call the world $w \in \mathcal{W}$ which is passed to a predicate the *current world*.

Predicates can access the current state (that is, the state contained in the current world) using the predicate $\text{state_interp } \varepsilon$ which states that the current memory state is equal to ε . In other words, it is the predicate:

$$\text{state_interp } \varepsilon := \lambda w. \text{“the memory state in } w \text{ is } \varepsilon\text{”}$$

Using this predicate, we can already express that an expression must not be stuck:

$$\text{not_stuck } e := \exists \varepsilon, \text{state_interp } \varepsilon \wedge \ulcorner \exists e', \varepsilon', (e, \varepsilon) \rightarrow (e', \varepsilon') \urcorner$$

where $\ulcorner \cdot \urcorner : \mathbb{P} \rightarrow iProp$ injects propositions from the ambient logic into Iris propositions. This proposition expresses that the term e is reducible in the current memory state.

Iris is a separation logic, which means that there is a notion of the “part of the world” which verifies some predicate P . This is achieved by requiring that the set \mathcal{W} of worlds have a partial commutative monoid structure (actually, as we will see in the sequel, a **camera** structure.) This partial commutative monoid structure induces an order relation on elements of \mathcal{W} :

$$w \preceq w' \quad \Leftrightarrow \quad \exists w_f, w' = w \cdot w_f$$

We model the partiality of \mathcal{W} by choosing a subset $\mathcal{V} \subseteq \mathcal{W}$ of valid elements. Intuitively, two subworlds w_1 and w_2 are compatible if $(w_1 \cdot w_2) \in \mathcal{V}$. This defines a separating conjunction of predicates: $w \models P * Q$ if w is the product of two worlds which satisfy P and Q respectively.

This generalizes our notion of logical state in Part I: The $*$ product which we defined on **LogState** gave a partial commutative monoid structure to logical states.

4.1.2 The Iris standard weakest precondition

Until the end of this thesis, we consider expression-based languages: there is no distinction between commands and expressions, which can also have side effects.

Definition 4.1.1 (Language). We consider languages with expressions denoted by $e \in \mathbf{Expr}$, values $v \in \mathbf{Val}$ with $\mathbf{Val} \subseteq \mathbf{Expr}$. A language has a notion of state $s \in \mathbf{State}$, and a (thread-local) small-step operational semantics

$$\rightarrow \subseteq (\mathbf{Expr} \times \mathbf{State}) \times (\mathbf{Expr} \times \mathbf{State} \times \mathbf{List}(\mathbf{Expr}))$$

A step $(e, s) \rightarrow (e', s', \vec{e}_f)$ means that the expression e reduces to e' while changing the state from s to s' , and creating the threads in the list \vec{e}_f . A value must not be reducible, and a non-value expression which cannot reduce is how Iris characterizes runtime errors.

This thread-local reduction relation is lifted to a threadpool reduction relation

$$\rightarrow_{\text{tp}} \subseteq (\mathbf{List}(\mathbf{Expr}) \times \mathbf{State}) \times (\mathbf{List}(\mathbf{Expr}) \times \mathbf{State})$$

defined using the following rule:

$$\frac{(tp[k], s) \rightarrow (e'_k, s', e_f)}{(tp, s) \xrightarrow{k}_{\text{tp}} (tp[k := e'_k] ++ e_f, s')}$$

where $tp[k]$ denotes the k th element of the list tp , and $tp[k := e'_k]$ denotes tp where the k th element has been replaced by e'_k . The threadpool is the list of all the threads of the

system which have been spawned. The thread ID of a thread, denoted in general with the variable tid is the index of the thread in the threadpool. This index is stable because irreducible programs are not removed from the threadpool. One reducible thread is chosen non-deterministically, and is reduced for one step. The threads it has spawned are added to the threadpool.

The next step to restate the definition of weakest precondition given by Figure 4.1 is to be able to *update* the memory state so that we can state that the weakest precondition holds in the new state s' after the program has reduced. To do so, we use the **viewshift** \Rightarrow connector of the logic. The proposition $P \Rightarrow Q$ asserts that we can replace any part of the world which satisfies P by a new subworld which satisfies Q in a way which “does not disturb” the rest of the world.

Remark 4.1.2. Readers familiar with Kripke semantics should keep in mind that the \Rightarrow connector is not related to the notion of “future world”. As we will see, however, the Iris model has a notion of future world for the step-indexing, and propositions are closed under that relation.

Given an Iris predicate $\Phi : \mathbf{Val} \rightarrow iProp$, we can define a weakest precondition as the following recursive predicate:

$$\begin{aligned} \text{wp } e \{Q\} &:= (\ulcorner e \in \text{Val} \urcorner \wedge Q e) \vee \\ &(\exists s, \text{state_interp } s \text{ } \multimap \ulcorner \exists e', s', (e, s) \rightarrow (e', s', \vec{e}_f) \urcorner \wedge \\ &\forall e', s', \vec{e}_f, \ulcorner (e, s) \rightarrow (e', s', \vec{e}_f) \urcorner \Rightarrow \text{state_interp } s' \multimap \\ &\triangleright \text{wp } e' \{Q\} * \bigotimes_{e_f \in \vec{e}_f} \text{wp } e_f \{\top\}) \end{aligned}$$

Note that, because e is an expression, which evaluates to a value, the postcondition Q takes a program value as an argument: it is of type $Q : \mathbf{Val} \rightarrow iProp$.

If e is already a value, the predicate is true if e satisfies the postcondition Q . Otherwise, if the current memory state is s , the program e must be reducible at that state, and, for all possible reductions $(e, s) \rightarrow (e', s', \vec{e}_f)$, where \vec{e}_f is the list of threads which are spawned during the reduction steps, we can update the world in such a way that the new memory state is s' , the weakest precondition holds for the reduct e' and all the spawned threads are safe.

The \triangleright modality in front of the recursive occurrence of wp is needed because, the Iris weakest precondition is defined as a guarded fixpoint instead of coninductively. A recursive definition is guarded if each recursive occurrence is under a \triangleright modality. This modality was introduced by Nakano (2000), and its relation with step-indexing was discovered by Appel et al. (2007).

Because Iris is a separation logic, the proposition

$$\text{wp } e_1 \{P_1\} * \text{wp } e_2 \{P_2\}$$

means that the resources needed to guarantee the safety of e_1 and those for e_2 must be “disjoint”, and e_1 and e_2 are safe to execute in parallel. This is why this is enough to require the separated product of the weakest preconditions of the reduct of e and of every forked thread in the definition of the weakest precondition.

4.2 The Iris model

One requirement of the Iris logic is that the world be able to contain Iris predicates. This is particularly useful to implement invariants, which are a generalization of lock invariants in the version of CSL we have studied in Part I. Let us make the simplifying assumption that we want our world to only contain one Iris predicate. Then, we have the following isomorphisms:

$$\begin{aligned} iProp &\cong \mathcal{W} \rightarrow \mathbb{P} \\ \mathcal{W} &\cong iProp \end{aligned}$$

The first isomorphism comes from the definition of Iris proposition as predicate over worlds. If we rewrite the second isomorphism into the first, we get:

$$iProp \cong iProp \rightarrow \mathbb{P}$$

This kind of recursive equation with a negative occurrence (on the left of an arrow) does not have a solution in most domains.

The solution Iris uses to solve this problem is to use step-indexing: intuitively, every set X is replaced by a sequence $(X_n)_{n \in \mathbb{N}}$ of sets. The unsolvable equation above is replaced with a system of equations which uses these indices break the cyclic nature of the previous equation:

$$\begin{aligned} iProp_0 &\cong \mathbb{P} \\ iProp_{n+1} &\cong iProp_n \rightarrow \mathbb{P} \end{aligned}$$

We can use operators on such indexed sets to simplify its presentation: Given a set X , we write $\mathbf{G}X$ the constant sequence $(X)_{n \in \mathbb{N}}$. We can also define a “next” operator \blacktriangleright which takes care of decreasing the step-indices:

$$\begin{aligned} (\blacktriangleright X)_0 &:= \{\star\} \\ (\blacktriangleright X)_{n+1} &:= X_n \end{aligned}$$

Using these two operators, we can rewrite the system of equations above as:

$$iProp \cong \blacktriangleright iProp \rightarrow \mathbf{GP}$$

This presentation of the model of Iris has so far been informal, as we have not defined the notion of maps, etc. One way to formalize this is to work in the category of presheaves over the ordinal ω , seen as a category; this category is also known as the topos of trees (L. Birkedal et al., 2012).

4.2.1 Ordered families of equivalences

In the actual model of Iris, sequences of sets are replaced by sets equipped with a sequence of equivalence relations:

Definition 4.2.1. An *ordered family of equivalences (OFE)* is the data $(X, (\stackrel{n}{=})_n)$ of a set X and a sequence of equivalence relations over X , such that:

1. $\stackrel{n}{=} \subseteq \stackrel{m}{=}$ for $n \geq m$;
2. $x = y$ iff $\forall n, x \stackrel{n}{=} y$.

The first condition states that the equivalence relation becomes more discriminating as n grows, and the second states that the limit of $(\stackrel{n}{=})$ is the equality relation.

An OFE X induces a sequence of sets as in the informal overview above by quotienting by the equivalence relation at each step-index:

$$X \mapsto (X / \stackrel{n}{=})_n$$

There are two useful notions of maps of OFEs:

Definition 4.2.2. Given two OFEs X and Y , a map $f : X \rightarrow Y$ between the underlying sets is said to be:

- *non-expansive* if, for all $x, y \in X$, and step-index n , $x \stackrel{n}{=} y \Rightarrow f(x) \stackrel{n}{=} f(y)$;
- and it said to be *contractive* if, for all $x, y \in X$, $x \stackrel{n}{=} y \Rightarrow f(x) \stackrel{n+1}{=} f(y)$.

OFEs and non-expansive maps form a Cartesian-closed category **OFE**. This vocabulary comes from the intuition that X is an ultrametric space, and that $x \stackrel{n}{=} y$ means that the points x and y are at a distance $\leq 2^{-n}$. An OFE which will be useful to define the Iris

model is the the OFE $SProp$ of **step-indexed propositions**: its underlying set $SProp$ is the set

$$\{A \subseteq \mathbb{N} \mid \forall n \leq m, m \in A \Rightarrow n \in A\}$$

of downward closed subsets of \mathbb{N} , and two step-indexed propositions are related by $\neq n$ if

$$\forall m \leq n, x \in A \Leftrightarrow x \in B$$

An intuitive reading of $n \in A$ is that it is not possible to disprove the proposition A in the first n steps of “execution”. Now that we have a step-indexed notion of proposition, it remains to define a step-indexed notion of partial commutative monoid.

4.2.2 RAs and Cameras

Iris is defined in terms of an extension of the notion of partial commutative monoid known as *resource algebras* (RA). Instead of having a neutral element, an RA M contains a partial map $| \cdot | : M \rightarrow M$ called the *core*, which associates to some element a its duplicable part, if it exists; in which case we have:

$$|a| \cdot a = a \quad \text{and} \quad ||a|| = |a|$$

which does imply that the core is duplicable:

$$|a| \cdot |a| = ||a|| \cdot |a| = |a|$$

The step-indexed version of RAs are Cameras:

Definition 4.2.3. A **camera** M is the data $(M, \mathcal{V}, |\cdot|, \cdot)$ of:

- an OFE M of elements;
- a “step-indexed subset” of valid elements $\mathcal{V} : M \rightarrow SProp$ expressed as a **non-expansive** map of OFEs;
- a partial map $|\cdot| : M \rightarrow M^?$, where $M^? := M \uplus \perp$ adjoins a new neutral element;
- and a commutative and associative **non-expansive** operation $\cdot : M \times M \rightarrow M$

which satisfy a number of axioms given in (Iris Team, 2021). This definition induces a step-indexed order on elements of M with $a \stackrel{n}{\leq} b \Leftrightarrow \exists a', b \stackrel{n}{=} a \cdot a'$.

4.2.3 The Iris base logic

The Iris logic is parameterized by a **camera** M ; we will explain later how the choice of M is done in a way that is modular.

The formulas of Iris base logic comprise higher-order BI

$$P ::= P \wedge Q \mid P \vee Q \mid P \Rightarrow Q \mid P * Q \mid P \multimap Q \mid \exists x : \tau. P \mid \forall x : \tau. P \mid x$$

where the type τ of the quantified variable x includes Iris propositions $iProp$ and the camera parameter M ; as well as the later modality and a guarded fixpoint operator:

$$P ::= \dots \mid \triangleright P \mid \mu x : \tau. P$$

where every occurrence of x in P must appear under a \triangleright modality; and a number of modalities and atomic propositions to reason about resources:

$$P ::= \dots \mid \text{Own}(a) \mid \mathcal{V}(a) \mid \Box P \mid \boxplus P \mid \ulcorner \Phi \urcorner$$

where a in an element of M , and $\Phi \in \mathbb{P}$ is a proposition in the ambient logic. The formula $\text{Own}(a)$ expresses the ownership of the resource a , and $\mathcal{V}(a)$ expresses that it is valid. The formula $\Box P$ (read *persistently* P) asserts that the persistent part of P holds, and $\ulcorner \Phi \urcorner$ states that the proposition Φ is true.

Finally, the update modality $\boxplus P$ is related to the viewshift connector \Rightarrow which we have seen before:

$$P \Rightarrow Q \quad := \quad P \multimap \boxplus Q$$

The formula $\boxplus P$ means intuitively that the current world (an element of M) can be updated so that the new world satisfies P . The condition that this update must not disturb the rest of the world is formalized using the notion of **frame-preserving update**: Given $a \in M$ and $B \subseteq M$, there is a frame-preserving update from a to B , written $a \rightsquigarrow B$ if:

$$\forall n, a_f^?, n \in \mathcal{V}(a \cdot a_f^?) \Rightarrow \exists b \in B, n \in \mathcal{V}(b \cdot a_f^?)$$

where $a_f^?$, which represents the rest of the world, or, in other words, what is owned by the environment, is an element of $M^? := M \uplus \perp$. One way to interpret it is that, whatever the frame $a_f^?$, which is assumed to be compatible with a , we can find an element $b \in B$ which can replace a without conflicting with the environment.

One example which we have already encountered is the allocation of a new location in the partial monoid **LogState** of logical states in Part II. The fact that there is always a free location in a logical state can be expressed as the following frame preserving update:

$$\emptyset \rightsquigarrow \{[\ell := (v, 1)] \mid \ell \in \text{Loc}\}$$

The fact that full ownership of a memory location allows one to update its contents corresponds to the following frame-preserving update:

$$[\ell := (v_1, 1)] \rightsquigarrow [\ell := (v_2, 1)]$$

where we write $a \rightsquigarrow b$ for $a \rightsquigarrow \{b\}$.

As we will see, these two frame-preserving updates corresponds, respectively, to the two laws in the logic:

$$\top \multimap \text{true} \Leftrightarrow \exists \ell. \ell \dot{\vdash} v \qquad \ell \dot{\vdash} v_1 \multimap \text{true} \Leftrightarrow \ell \dot{\vdash} v_2$$

where we define $\ell \dot{\vdash}^p v$ as $\text{Own}([\ell := (v, p)])$.

4.2.4 Interpretation of Iris predicates

Iris predicates P are interpreted as monotone maps $P : M \rightarrow SProp$, where the order on M is given by divisibility, and the order on $SProp$ is given by containment. Hence, a predicate P is upward-closed with respect to the order on M (which corresponds to the fact that Iris is an affine logic) and downward-closed with respect to step-indices (from the definition of $SProp$).

Writing $a, n \models P$ for $n \in P(a)$, we define the interpretations of the main constructions of the base logic:

$$\begin{aligned} a, n \models P * Q &\Leftrightarrow \exists b_1, b_2, a \stackrel{n}{=} b_1 \cdot b_2 \wedge b_1, n \models P \wedge b_2, n \models Q \\ a, n \models P \multimap Q &\Leftrightarrow \forall m, b, m \leq n \wedge m \in \mathcal{V}(a \cdot b) \wedge b, m \models P \Rightarrow a \cdot b, m \models Q \\ a, n \models \Box P &\Leftrightarrow |a|, n \models P \\ a, n \models \triangleright P &\Leftrightarrow n = 0 \vee a, n - 1 \models P \\ a, n \models \text{true} P &\Leftrightarrow \forall m, a', m \leq n \wedge m \in \mathcal{V}(a \cdot a') \Rightarrow \exists b, m \in \mathcal{V}(b \cdot a') \wedge b, m \models P \\ a, n \models \ulcorner \Phi \urcorner &\Leftrightarrow \Phi \\ a, n \models \text{Own}(b) &\Leftrightarrow b \stackrel{n}{\leq} a \\ a, n \models \mathcal{V}(b) &\Leftrightarrow n \in \mathcal{V}(b) \end{aligned}$$

The guarded fixed-points are interpreted using a fixed point operator defined on **contractive** maps in the category **OFE**, as introduced by America and Rutten (1989) and generalized by Lars Birkedal, Støvring, and Thamsborg (2010).

Definition 4.2.4. A predicate P is valid, which we write $\models P$ if it is valid for all element of the **camera** and all the step-indices:

$$\models P \quad \Leftrightarrow \quad \forall a \in M, n \in \mathbb{N}, a, n \models P.$$

4.2.5 Logical rules

This model satisfies a number of logical properties. We write here a few of them, and refer the reader to (Iris Team, 2021) for an exhaustive list.

The separating conjunction reflects the operation of the `camera` M , and an owned element is always valid:

$$\text{Own}(a) * \text{Own}(b) \Leftrightarrow \text{Own}(a \cdot b) \qquad \text{Own}(a) \Rightarrow \mathcal{V}(a)$$

The update modality behaves like a strong monad

$$P \Rightarrow \Vdash P \qquad \Vdash \Vdash P \Rightarrow \Vdash P \qquad Q * \Vdash P \Rightarrow \Vdash Q * P$$

and it reflects [frame-preserving updates](#):

$$\frac{a \rightsquigarrow B}{\text{Own}(a) \Rightarrow \Vdash \exists b, \text{Own}(b)}$$

The persistence modality behaves like a comonad and corresponds to the core operation

$$\Box P \Rightarrow P \qquad \Box P \Rightarrow \Box \Box P \qquad \text{Own}(a) \Rightarrow \text{Own}(|a|)$$

and it makes predicate duplicable:

$$\Box P \Leftrightarrow \Box P * \Box P$$

A predicate P is called *persistent* if it is equivalent to $\Box P$. Note that this is stronger than being duplicable: the predicate $\exists p, \ell \stackrel{p}{\rightarrow} 1$ is duplicable but not persistent. Of course, the laws of higher-order bunched implications also hold, such as \multimap being a right adjoint to $*$, etc.

The later modality allows to define recursive predicates using guarded recursion. Such predicates come with a powerful induction principle which allows reasoning about them, called *Löb induction*:

$$\frac{\triangleright P \Rightarrow P}{P}$$

This corresponds to an induction on the step-index in the semantics of the predicates. In the setting of a program logic, this means that the induction hypothesis can be used as soon as the program reduces by one step, because the definition of wp contains a later modality, tying logical steps with program steps. This allows for eliminating a later modality guarding an hypothesis using the rule:

$$\frac{P \Rightarrow Q}{\triangleright P \Rightarrow \triangleright Q}$$

In particular, the proof of a recursive function is able to use its induction hypothesis to reason about recursive calls because the execution of a recursive function always starts with the beta reduction step which substitute its argument.

Adequacy of the base logic

The adequacy of the Iris base logic states that the modalities of the Iris base logic do not allow to prove false statements in the meta-logic.

Theorem 4.2.5 (Adequacy). *Suppose $\models m_1 \cdots m_k \ulcorner \Phi \urcorner$, where each m_i is an Iris modality such as $\triangleright, \Rightarrow, \square, \dots$, and $\Phi \in \mathbb{P}$ is a proposition, then Φ holds.*

4.3 High level logic

The logic we have described so far is not usable to prove programs. For instance, it is parameterized by a single **camera**, whereas programs most often need several kinds of resources to be specified and proved. Another missing feature are invariants, which are the way concurrent separation logics allow reasoning about how threads communicate and synchronize with each others. These features can be defined using the base logic above. We sketch in this section how they are defined and how they can be used.

4.3.1 Combining cameras

The mechanism with which several **cameras** can be used is simple: Given a list M_1, \dots, M_n of cameras, we define a new camera

$$Res := \prod_{i \in \{1, \dots, n\}} GName \rightarrow_{fn} M_i$$

where $GName := \mathbb{N}$ are the type of ghost (or logical) variables. The camera structure induced by the product is straightforward, and the product of two partial finite maps $f_1, f_2 : X \rightarrow_{fn} M$ is defined as:

$$(f_1 \cdot f_2)(x) := \begin{cases} f_1(x) & \text{if } x \in \text{dom}(f_1) \setminus \text{dom}(f_2) \\ f_2(x) & \text{if } x \in \text{dom}(f_2) \setminus \text{dom}(f_1) \\ f_1(x) \cdot f_2(x) & \text{if } x \in \text{dom}(f_1) \cap \text{dom}(f_2) \end{cases}$$

and a finite partial map is valid iff all its component are valid. This definition is the same as the product of logical states from Definition 2.5.3 in Part I for $X := \mathbf{Loc}$ and $M := \mathbf{Val} \times \mathbf{Perm}$.

We can define the predicate

$$\boxed{m : M_i}^\gamma := \text{Own}((\emptyset, \dots, \emptyset, [\gamma := m], \emptyset, \dots, \emptyset))$$

which expresses the ownership of an n -tuple of empty partial finite maps, except for the k th component which contains the singleton finite map associating the element m of M_i to the ghost variable γ .

This new predicate admits a rule for allocation, and the same rules as the primitive Own predicate for updating and combining:

$$\frac{a \in \mathcal{V}_{M_i}}{\models \exists \gamma. \boxed{a : M_i}^\gamma} \quad \frac{a \rightsquigarrow B}{\boxed{a : M_i}^\gamma \models \exists b \in B. \boxed{b : M_i}^\gamma} \quad \boxed{a}^\gamma * \boxed{b}^\gamma \Leftrightarrow \boxed{a \cdot b}^\gamma$$

4.3.2 Invariants

We will not explain how invariants are constructed, but the gist of it is that, as we sketched at the beginning of this chapter, Iris predicates can be endowed with a **camera** structure. To be useful, invariants cannot always hold, there needs to be a mechanism for opening invariants and then requiring they be reestablished.

This control is achieved in Iris using the so-called **fancy update** modalities $\varepsilon_1 \models \varepsilon_2$, which are defined using the update modality \models . This modality is parameterized by a masks \mathcal{E}_1 and \mathcal{E}_2 which denote sets of invariant names. The predicate $\varepsilon_1 \models \varepsilon_2 P$ means that we can update from a world where every invariant in \mathcal{E}_1 holds to a world where every invariant in \mathcal{E}_2 hold and where P is verified. For instance $\top \models^\emptyset P$ means that one is able to prove P while opening all the invariants (the mask \top denotes the set of all invariant names).

The existence of an invariant is claimed using the predicate $\text{inv}^\mathcal{N}(P)$, where \mathcal{N} is the name of the invariant. It is a persistent predicate, which means in particular that it is duplicable, and that it can be shared between threads. The rules for creating and for opening invariants are the following:

$$\triangleright P \Rightarrow \varepsilon \models^{\mathcal{E} \cup \mathcal{N}} \text{inv}^\mathcal{N}(P) \quad \frac{\mathcal{N} \subseteq \mathcal{E}}{\text{inv}^\mathcal{N}(P) \multimap \varepsilon \models^{\mathcal{E} \setminus \mathcal{N}} \triangleright P * \square(\triangleright P \multimap \varepsilon \setminus \mathcal{N} \models^\mathcal{E} \top)}$$

The second rule deserves an explanation: If the invariant \mathcal{N} contains predicate P , and if it currently holds ($\mathcal{N} \subseteq \mathcal{E}$), then is it possible to update to a world where the invariant \mathcal{N} no longer hold but where $\triangleright P$ holds. To close the invariant and go back to a situation where \mathcal{N} holds, one can use the implication which consumes the resource $\triangleright P$. The later modalities \triangleright in these rules come from the later operator in the domain equation at the beginning of Section (4.2).

4.3.3 The standard Iris weakest precondition

We now have the tools to define the standard Iris weakest precondition, which guarantees the safety of programs. It differs from the one sketched at the beginning of the section by the masks and the **fancy updates** which need to be added to be able to use invariants when proving programs.

$$\begin{aligned}
 \text{wp}_\varepsilon e \{Q\} &:= \mu \text{wp}_0. \\
 &(\ulcorner e \in \text{Val} \urcorner \wedge \varepsilon \Vdash^\varepsilon Q e) \vee \\
 &(\exists s, \text{state_interp } s \multimap \varepsilon \Vdash^\emptyset \ulcorner \text{reducible } e \urcorner s \urcorner \wedge \\
 &\forall e', s', \vec{e}_f. \ulcorner (e, s) \rightarrow (e', s', \vec{e}_f) \urcorner \multimap \emptyset \Vdash^\emptyset \triangleright \emptyset \Vdash^\varepsilon \text{state_interp } s' \urcorner * \\
 &\text{wp}_{0 \ \varepsilon} e' \{Q\} * \bigotimes_{e_f \in \vec{e}_f} \text{wp}_{0 \ \top} e_f \{\top\})
 \end{aligned}$$

All the expected rules about this weakest precondition are proved using the rules of Iris logic. The soundness theorem states that if we can prove $\text{wp}_\top e \{v. \ulcorner \Phi v \urcorner\}$ from any world where the initial machine state is s , then all executions which start from s are safe, and if they terminate, their final value satisfies Φ . This is stated formally as follows:

Theorem 4.3.1 (Soundness of the standard weakest precondition). *Given a program e , a machine state s and a postcondition $\Phi : \mathbf{Val} \rightarrow \mathbb{P}$, if*

$$\models \top \Vdash^\top \text{state_interp } s \urcorner * \text{wp}_\top e \{v. \ulcorner \Phi v \urcorner\}$$

then for all execution paths $([e], s) \rightarrow_{\text{tp}}^ (tp', s')$, for all $e' \in tp'$, either e' is a value, or (e', s') is reducible. If moreover the main thread has terminated (that is, $e' = \text{tp}'[0]$ is a value), then $\Phi e'$ holds.*

The proof proceeds by Löb induction and uses the adequacy theorem of the Iris base logic.

The soundness theorem above does not allow to prove intensional properties of the program; the next chapter presents a new weakest precondition which allows to specify more accurately the behavior of the program.

5 A program logic to relate traces

Program logics describe program behaviors through their soundness theorems: the fact that some predicate about a program is provable inside the program logic implies that the semantics of this program satisfies some property in the ambient logic. This means that the expressiveness of the logic is limited by its soundness theorem.

5.1 Introduction

The Soundness Theorem 4.3.1 for the standard Iris weakest precondition is, essentially, able to express that all vertices in the graph of all possible executions of a program e starting at machine state s satisfy some property, which corresponds to invariants being always satisfied.

To see why one would want a finer characterization of the behavior of programs, let us consider a program `hanoi` which solves the Tower of Hanoi problem. Recall that, in the Tower of Hanoi problem, there are three rods, A , B and C and a number n of disks of different diameters $\{1, \dots, n\}$. Each disk has a hole in its center so that it can slide onto any rod A , B or C . In this initial state, all disks are on rod A , in increasing diameter from the top. The goal is to move this stack from rod A to rod B , with the following constraints:

1. a disk must never lie on a smaller disk;
2. at most one disk can be missing from the stacks at any time;
3. each move consists in moving one disk from the top of a stack to the top of another stack (including the empty stack).

It is a standard exercise for students learning recursive programming to write a program which solves the Hanoi problem. One way to encode the Towers of Hanoi as a program specification is as follows: each disk is represented as its diameter, and each stack is represented as a linked list of the diameters of the disks (the head of the list contains the disk on the top of the stack.)

Conditions 1 and 2 can be expressed as invariants of the program logic: at each instant, each linked list is sorted, and the multiset of all elements of every list contains all disks except for at most one (which is being held in the hands of the player of the game) and no disk has a multiplicity greater than 1.

The problem with this specification is that the following program satisfies it:

$$\text{hanoi } \ell_A \ell_B \ell_C := \text{swap } \ell_A \ell_B$$

indeed, at the end of the execution, all the disks have been moved to the second rod, and the invariant described above was never invalidated. The problem is that we are not able to express the third condition above, which restricts the *steps* the program can take.

5.1.1 Relating STSs to programs in Iris

More generally, the Towers of Hanoi problem described above is an instance of a more general pattern: there is a state transition system (STS) $(\mathcal{M}, \rightarrow)$ which is an abstract specification of the behavior the program should follow. It can be seen as a specification of the problem to solve as in the example above, or it can represent the algorithm the program should implement.

In the case of the Towers of Hanoi, $\mathcal{M} \subseteq \mathbb{N}^* \times \mathbb{N}^* \times \mathbb{N}^*$ represents the legal states of the three stacks of disks, and the transition relation $\rightarrow \subseteq \mathcal{M} \times \mathcal{M}$ expresses the allowed transitions of the game.

A common strategy, used for example by Robbert Krebbers, Timany, and Lars Birkedal (2017) to encode binary logical relations in Iris, is to state the following invariant

$$\text{inv}^{\mathcal{N}} (\exists s \in \mathcal{M}. \text{model_is } s * \ulcorner s_0 \rightarrow^* s \urcorner * I s)$$

where s_0 is the initial state of the STS—for Hanoi, $s_0 = ([1, \dots, n], \varepsilon, \varepsilon)$ —and model_is is a predicate which states that the “current state” of the STS is equal to s . The Iris predicate $I : \mathcal{M} \rightarrow iProp$ expresses that the memory of the program corresponds to the state $s \in \mathcal{M}$ of the STS: In the Towers of Hanoi case, it states that the three linked lists contains the same integer sequences as the components of the model.

Using the soundness theorem of the logic, we can prove that at every step k of the execution of the program, there is a state s_k of the STS which is reachable from the initial state and which matches the state of the program according to the predicate I . What is lacking is that the sequence of witnesses s_1, \dots, s_k is a valid run of the STS \mathcal{M} . The technical objective of this work is to construct a program logic using the Iris base logic which does provide such a guarantee.

5.1.2 Simulations

The formal notion of correspondence between executions of the program and the runs of the STS which we use is *simulation*. This makes sense because a small step semantics of a programming language in the sense of Definition 4.1.1 defines an STS \mathcal{L} over the states $(\text{tp}, \mathfrak{s})$ where tp is a threadpool and \mathfrak{s} is a machine state. The transition relation is \rightarrow_{tp} , ignoring the label. A simulation is a relation between states of two STSs which is, in a sense, stable under transitions:

Definition 5.1.1. Given two STSs $(\mathcal{M}_1, \rightarrow_1)$ and $(\mathcal{M}_2, \rightarrow_2)$, a *simulation* is a relation $\phi \subseteq \mathcal{M}_1 \times \mathcal{M}_2$ which satisfies the following property. For all $(s_1, s_2) \in \mathcal{M}_1 \times \mathcal{M}_2$, if $\phi(s_1, s_2)$, and for all $s'_1 \in \mathcal{M}_1$ such that $s_1 \rightarrow s'_1$, there exists s'_2 such that $s_2 \rightarrow s'_2$ and $\phi(s'_1, s'_2)$.

Any STS induces another STS whose states are its runs:

Definition 5.1.2. Given an STS $\mathcal{M} = (\mathcal{M}, \rightarrow)$, a *run, execution or trace* of \mathcal{M} is a finite or infinite sequence $t = s_1, \dots, s_n, \dots \in \mathcal{M}$ of states such that, for all $1 \leq i$, $s_i \rightarrow s_{i+1}$. We write $\mathbf{Runs}(\mathcal{M})$ the set of runs of \mathcal{M} , and $\mathbf{FRuns}(\mathcal{M})$ for finite runs. If t is a finite run of length n with the notations above, and if s_{n+1} is such that $s_n \rightarrow s_{n+1}$, then $t \cdot s_{n+1}$ denotes the run s_1, \dots, s_{n+1} .

Given an STS \mathcal{M} , the STS $\mathfrak{R}(\mathcal{M})$ of runs of \mathcal{M} is the STS $\mathfrak{R}(\mathcal{M}) = (\mathbf{FRuns}(\mathcal{M}), \rightarrow)$ where

$$s_1 \cdots s_n \rightarrow s_1 \cdots s_n s_{n+1} \quad \Leftrightarrow \quad s_n \rightarrow s_{n+1}$$

This allows us to define the notion of *history-sensitive simulation* as follows: a history-sensitive simulation between two STSs \mathcal{M}_1 and \mathcal{M}_2 is a *simulation* between $\mathfrak{R}(\mathcal{M}_1)$ and $\mathfrak{R}(\mathcal{M}_2)$. Such a simulation relates pairs of “executions”, and may examine the history of those computations.

This notion of simulation is too strict for our purposes however: each step of the program would need to be related to a step of the STS \mathcal{M} , but, because \mathcal{M} is meant to be an abstract description of a computational process, it is natural to expect that it will take fewer steps than the program. For example, for the Towers of Hanoi, all the steps of the program which do not change the state of the memory are not matched by steps in the STS representing the rules of the Hanoi game.

The solution, of course, is to allow the simulation to stutter. In this chapter, we allow infinite stuttering, and define

Definition 5.1.3. Given two state transition systems \mathcal{M}_1 and \mathcal{M}_2 , a *(history-sensitive termination-insensitive) stuttering simulation* is a *history-sensitive simulation* between \mathcal{M}_1 and $\mathcal{M}_2^{\bar{\bar{}}}$, where $\mathcal{M}_2^{\bar{\bar{}}}$ is the reflexive closure of \mathcal{M}_2 .

In this chapter, we will define a program logic which, given a proof of $\text{trwp } e \{Q\}$, allows for easily proving that some user-chosen relation ϕ contains a **stuttering simulation** which relates an abstract model \mathcal{M} and a program e , seen as an STS.

More precisely, we define

Definition 5.1.4. Given two STSs \mathcal{M}_1 and \mathcal{M}_2 and a relation $\phi \subseteq \mathcal{M}_1 \times \mathcal{M}_2$, we define **ϕ -similarity** as the greatest simulation included in ϕ . Two states $s_1 \in \mathcal{M}_1$ and $s_2 \in \mathcal{M}_2$ are called **ϕ -similar** if they are related by the ϕ -similarity relation.

This is well defined because simulations are closed under union, and so are the relations $\subseteq \phi$. Intuitively, s_1 and s_2 are ϕ -similar when they are related by ϕ and for any step of s_1 , there exists a step of s_2 so that the reducts are related by ϕ . This can be made formal with the following characterization of ϕ -similarity as the greatest fixpoint of the following monotone map over relations in $\wp(\mathcal{M}_1 \times \mathcal{M}_2)$:

$$\mathcal{R} \mapsto \{(x_1, x_2) \mid \phi(x_1, x_2) \wedge \forall x'_1, x_1 \rightarrow x'_1 \Rightarrow \exists x'_2, x_2 \rightarrow x'_2 \wedge \mathcal{R}(x'_1, x'_2)\}$$

This induces the notion of **ϕ -history sensitive similarity** and of **ϕ -stuttering similarity** since the corresponding notions of simulation are defined as **simulations** on transformed STSs.

5.1.3 Motivation

This goal of this logic is to be able to state formally that some program module e implements some algorithm \mathcal{A} , described as a state transition system \mathcal{M} . The usual program logic of Iris allows, in a sense, to prove that e satisfies properties which are expected of \mathcal{A} by specifying e with Hoare triples. The relevance of these Hoare triples is supported by writing clients which use these specifications and whose safety depend on the program module e having the properties expected of it.

Practically speaking, one advantage of our approach is that one can use an off-the-shelf formal description of the algorithm \mathcal{A} . For example, we took a TLA+ specification, written by experts, of the single decree Paxos algorithm of L. Lamport (2001), and we proved that some program implements this algorithm. This also allows to split the proof of this subtle program in two independent parts: first we prove that e implements \mathcal{M} dealing with issues related to the implementation, and then we prove that \mathcal{M} does solve the consensus problem in an idealized setting. See (Timany, Gregersen, et al., 2021) for more details.

As we will see in the next chapter, this technique can be adapted to prove termination of concurrent programs under fair schedulers by building a **fair simulation** between the program and a well-behaved abstract model.

5.2 The trace weakest-precondition

The new weakest precondition trwp which we introduce in this chapter relates executions of the program with traces in the model. Therefore, we generalize the state_interp predicate to take both the memory state \mathfrak{s} and the model state s as parameter. Using the so-called authoritative **camera** construction which we will not detail here, the state_interp and model_is predicates can be defined in such a way that they follow the law:

$$\text{state_interp } \mathfrak{s} \ s \ * \ \text{model_is } s' \quad \dashv \! \! \dashv \quad \ulcorner s = s' \urcorner$$

Only the model_is predicate is seen by the user of the logic during the proof of the program; the state_interp predicate is only used in the definition of the weakest precondition and in the proof of lemmas about trwp and model_is .

The logic is parameterized by a predicate $\text{state_evol } \mathfrak{s}_1 \ s_1 \ \mathfrak{s}_2 \ s_2$ which expresses which transitions in the model are allowed, a notion which can depend on the old memory state \mathfrak{s}_1 and the new memory state \mathfrak{s}_2 of the program.

To make trwp a strict generalization of the standard Iris weakest precondition wp , in that every rule which is valid for the latter is also valid for the former, it suffices to choose a state_evol predicate which satisfies:

$$\forall \mathfrak{s}_1, \mathfrak{s}_2, s, \quad \text{state_evol } \mathfrak{s}_1 \ s \ \mathfrak{s}_2 \ s.$$

The standard definition for this predicate, given a model STS \mathcal{M} , is:

$$\text{state_evol } \mathfrak{s}_1 \ s_1 \ \mathfrak{s}_2 \ s_2 \quad := \quad \ulcorner s_1 \rightarrow s_2 \vee s_1 = s_2 \urcorner$$

which corresponds to ignoring the machine states and asking that each step of the program correspond to a step in $\mathcal{M}^=$, the reflexive closure of the model STS \mathcal{M} .

The new weakest precondition is defined as the following guarded fixpoint:

$$\begin{aligned} \text{trwp}_{\mathcal{E}} e_1 \{Q\} &:= (\ulcorner e_1 \in \mathbf{Val} \urcorner \wedge \mathcal{E} \Vdash^{\mathcal{E}} Q \ e_1) \vee (\ulcorner e \notin \mathbf{Val} \urcorner \wedge \\ &\forall \mathfrak{s}_1, s_1, \text{state_interp } \mathfrak{s}_1 \ s_1 \dashv \! \! \dashv \mathcal{E} \Vdash^{\emptyset} \ulcorner \text{reducible } e_1 \ \mathfrak{s}_1 \urcorner * \\ &\forall e_2, \mathfrak{s}_2, \vec{e}_f. \ulcorner (e_1, \mathfrak{s}_1) \rightarrow (e_2, \mathfrak{s}_2, \vec{e}_f) \urcorner \dashv \! \! \dashv \emptyset \Vdash^{\emptyset} \triangleright \emptyset \Vdash^{\mathcal{E}} \exists s_2, \\ &\ulcorner \text{state_evol } \mathfrak{s}_1 \ s_1 \ \mathfrak{s}_2 \ s_2 \urcorner * \text{state_interp } \mathfrak{s}_2 \ s_2 * \\ &\text{trwp}_{\mathcal{E}} e_2 \{Q\} * \bigotimes_{e \in \vec{e}_f} \text{trwp}_{\top} e \{\top\}) \end{aligned}$$

The weakest precondition takes as input the “current” machine state \mathfrak{s}_1 and the current state s_1 of the STS \mathcal{M} . If the program e_1 is not a value, it must be reducible in state \mathfrak{s} ,

and for all possible reduction steps of the program $(e_1, s_1) \rightarrow (e_2, s_2, \vec{e}_f)$, there must exist a corresponding state s_2 of the STS such that the transition from s_1, s_1 to s_2, s_2 is allowed, and such that $\text{trwp}_\varepsilon e_2 \{Q\}$ holds in a world where the new states are s_2 and s_2 and where the new threads are safe.

Given that guarded recursion is closely related to condinduction, this definition is reminiscent of *similarity*.

The predicate trwp satisfies all the rules expected of a weakest precondition. In addition, it is possible to change the state of the STS while executing an atomic expression (by definition, an expression which reduces to a value in one step). The following rule allows to take a step in the model:

$$\frac{\ulcorner s \rightarrow s' \urcorner * \text{model_is } s * \text{trwp}_\varepsilon e \{v. \text{model_is } s' \dashv^* \varepsilon \Rightarrow^\varepsilon Q v\}}{\text{trwp}_\varepsilon e \{Q\}}$$

In other words, if we own the fact that the current models of the STS is s and that $s \rightarrow s'$, then we can assume that, after e has executed, the new state of the STS is s' .

5.2.1 Soundness of trwp

The soundness theorem provides a criterion, given a relation ϕ and a proof of $\text{trwp}_\top e \{Q\}$ in a logical state where the current program and model states are respectively s_0 and s_0 , for $([e], s_0)$ and s_0 to be ϕ -*history sensitive similar*. This criterion states —if we ignore histories— that $\phi s s$ holds if $\ulcorner \text{state_interp } s s \urcorner$ does.

The reason why this criterion is easy to use is that this it needs to be satisfied inside of Iris, while being able to use every invariant; and in practice the invariant and ϕ state essentially the same property, the first from inside the Iris logic and the second from the outside. For example, in the case of the Towers of Hanoi, the invariant states that the three linked lists contain the same elements as the three components of the abstract model using points-to predicates, and ϕ states the same fact, in the form of a predicate over the memory and the state of the STS.

We need some notations to state the soundness theorem. Given a run t of an STS, we write

$$t = x \cdots y$$

to mean that the first state of the run t is x and the last one is y . Note that if t is of length 1, x and y are the same state. We write $|t|$ for the length of the run t . Finally, if $t = x_1 \cdots x_{n-1} x_n$ is a run of length > 1 , we write $t_{[: -1]}$ for the trace $x_1 \cdots x_{n-1}$.

As we will explain in the next section, the following technical condition is required by the fact that the Iris logic is step-indexed.

Definition 5.2.1. Given relation ϕ between finite executions of \mathcal{L} and finite runs of \mathcal{M} , the STS \mathcal{M} is called **ϕ -finitary** if, for all s, t_1, t_2 , the the following set is finite:

$$\{s' \mid s \xrightarrow{\ell} s' \wedge \phi(t_1 \cdot s)(t_2 \cdot s')\}$$

Theorem 5.2.2. Let ϕ be a relation between finite executions of \mathcal{L} and finite runs of \mathcal{M} . Given a program e , an initial memory s and a model STS \mathcal{M} with an initial state s_0 , if the Iris formula

$$\begin{aligned} & \top \Vdash \top \text{state_interp } s_0 \ s_0 * \text{trwp}_\top e \{Q\} * \\ & (\forall t_1, t_2, \text{tp}, s, s, \top t_1 = ([e], s_0) \cdots (\text{tp}, s) \wedge t_2 = s_0 \cdots s \top \dashv^* \\ & (\top |t_1| > 0 \wedge |t_2| > 0 \Rightarrow \phi t_{1[: -1]} t_{2[: -1]} \top) \dashv^* \\ & \text{state_interp } s \ s \dashv^* \\ & \top \Vdash \emptyset \top \phi t_1 t_2 \top) \end{aligned}$$

is valid, and if the STS \mathcal{M} is **ϕ -finitary**, then the singleton traces $[[e], s_0]$ of the program and $[s_0]$ of the STS are **ϕ -history sensitive similar**.

The implication

$$|t'_1| > 0 \wedge |t'_2| > 0 \Rightarrow \phi t_{1[: -1]} t_{2[: -1]}$$

is a terse way of stating that the user needs to establish that ϕ relates the initial machine and STS states, and that afterwards they can assume that ϕ relates the previous states.

In the standard instantiation explained above with the reflexive closure $\mathcal{M} := \mathcal{M}_0^\top$ of a STS \mathcal{M}_0 , the conclusion states that the singleton executions $[[e], s_0]$ and $[s_0]$ are **ϕ -history sensitive similar** in the STSs \mathcal{M}^\top and \mathcal{L} . By definition of **stuttering simulation**, this means that $[[e], s_0]$ and $[s_0]$ are **ϕ -stuttering similar**.

Remark 5.2.3. We have not presented the soundness theorem of the logic in its full generality for clarity, instead we have presented it as it looks when instantiated for some programming language. We refer the reader to the Coq development for the more general definition, where the predicates `state_interp` and `state_evol` are existentially quantified, and where there are more hypothesis available to prove $\top \phi t_1 t_2 \top$.

5.2.2 Proof of the soundness theorem

Because **ϕ -similarity** can be defined as a coinductive predicate, the proof of the soundness theorem boils down to the fact that a guarded fixed point `trwp` in the Iris logic implies a coinductive predicate in the ambient logic.

This proof proceeds by considering an intermediate “pure” predicate gsim which sits in between those two predicates. It states the definition of ϕ -similarity expressed as a guarded fixpoint which is independent of resources. Writing $\text{Hyp}(\phi, et, mt)$ for the hypothesis in Theorem 5.2.2, where et and mt are singleton program trace and STS run respectively, the proof proceeds by proving the following two implications:

1. $\models \text{Hyp}(\phi, et, mt) \multimap \text{gsim}_\phi(et, mt)$
2. $(\models \text{gsim}_\phi(et, mt)) \implies et \text{ and } mt \text{ are } \phi\text{-similar}$

Before we describe the proofs of these two implications in turn, let us define the $\text{gsim}_\phi(et, mt)$ predicate as the Iris guarded fixpoint:

$$\begin{aligned} \text{gsim}_\phi(et, mt) &:= \triangleright \ulcorner \phi \text{ et } mt \urcorner \wedge \forall tp, s, tp', s'. \\ &\quad \ulcorner et = \dots (tp, s) \wedge (tp, s) \rightarrow_{tp} (tp', s') \urcorner \rightarrow \\ &\quad \triangleright \triangleright (\exists s'. \ulcorner \text{state_evol } s \ s' \ s' \urcorner \wedge \text{gsim}_\phi(et \cdot (tp', s'), mt \cdot s')) \end{aligned}$$

where the first \triangleright spans over the whole formula. Notice that all the atomic predicates used in this definition are of the form $\ulcorner \Phi \urcorner$. This implies that it does not depend on resources: for all resources a and a'

$$a, n \models \text{gsim}_\phi(et, mt) \iff a', n \models \text{gsim}_\phi(et, mt)$$

The three later modalities in the definition of gsim are necessary to eliminate the later and the fancy updates (which are defined using a later modality) in the definition of trwp during the proof of the first implication $\text{Hyp}(\phi, et, mt) \multimap \text{gsim}_\phi(et, mt)$.

The proof of this implication proceeds by Löb induction and it fairly similar to the proof of soundness of the usual Iris weakest precondition.

Extracting a trace from a guarded predicate

The more interesting part of the proof is the second implication:

$$(\models \text{gsim}_\phi(et, mt)) \implies et \text{ and } mt \text{ are } \phi\text{-similar}$$

To emphasize the resemblance between the gsim predicate and ϕ -similarity, we express the latter in a language close to the Coq syntax for coinductive predicates:

$$\begin{aligned} \text{Coinductive } \text{gsim}_\phi(et, mt) &:= \\ &\phi \text{ et } mt \wedge \forall tp, s, tp', s'. \\ &\quad et = \dots (tp, s) \wedge (tp, s) \rightarrow_{tp} (tp', s') \rightarrow \\ &\quad \exists s'. \text{state_evol } s \ s' \ s' \wedge \text{gsim}_\phi(et \cdot (tp', s'), mt \cdot s') \end{aligned}$$

They only differ syntactically by the modalities and the different nature of their fixpoint. The proof of this implication is not trivial however: the STS needs to be ϕ -finitary for it to hold. The reason is that the following result does not hold in general:

$$\models \exists x : A. P x \quad \Rightarrow \quad \exists x : A, \models P x \quad (5.1)$$

where the predicate P does not depend on resources. The reason is that, unfolding the definition of $\models \exists x : A. P x$ and of validity, this amounts to proving that

$$(\forall i : \mathbb{N}, \exists x_i : A, i \models P x_i) \quad \Rightarrow \quad \exists x : A, \forall i : \mathbb{N}, i \models P x$$

which, not surprisingly, does not hold for all A (we omit the resource a since validity of these formulas do not depend on it). If A is finite however, if the LHS of the implication holds, by the pigeonhole principle and the axiom of choice, there exists x_∞ such that for infinitely many $i \in \mathbb{N}, i \models P x_\infty$.

Because, by definition, predicates are downward closed for step indices ($i \models P$ implies that, for any $j < i, j \models P$), this witness x_∞ holds for every step-index i , and the implication (5.1) holds.

The case of the universal quantification and conjunction is not problematic because this amounts to permuting two universal quantifiers. Disjunction is equivalent to existential quantification over Booleans which means that “disjunction distributes over \models ” (though we do not need it here).

Using (5.1) to extract the witness s' , the proof proceeds by a simple coinduction.

This technique was used by Tassarotti, Jung, and Harper (2017) and Spies et al. (2021) for similar purposes.

Remark 5.2.4. The condition that the STS \mathcal{M} is finitely branching which we presented here is stronger than necessary. In the next chapter, we need the weaker but more complicated condition, that for all program configuration (tp, \mathfrak{s}) and machine state \mathfrak{s}' , for all STS state $s \in \mathcal{M}$,

$$(\exists et, mt, et = \dots (tp, \mathfrak{s}) \wedge mt = \dots s \wedge \phi et mt) \Rightarrow |\{s' \mid \text{state_evol } \mathfrak{s} s' s'\}| < \infty$$

5.2.3 Infinite traces

The soundness theorem of this logic is a proof method to prove that a singleton program trace $[(tp, \mathfrak{s})]$ and a singleton run $[s]$ of the STS \mathcal{M} are ϕ -similar for some binary relation ϕ .

This implies in particular that given a possibly infinite trace et of the form $et = (tp, s) \dots$, there exists a possibly infinite run mt of \mathcal{M} of the same length as et which begins by s , and such that:

$$\forall n < |et|, \phi \text{ } et_{[:n]} \text{ } mt_{[:n]}$$

where we write $et_{[:n]}$ for the prefix of et which consists in the first n -transitions.

The fact that this relates infinite traces will allow us to transport liveness properties from the STS run to the execution trace in the next chapter.

5.3 Related works

There have been several works which aim at proving refinement of programs using a concurrent separation logic such as Iris, often using logical relations: (Robbert Krebbers, Timany, and Lars Birkedal, 2017) explains how to define binary logical relations and prove them sound using Iris, (Timany, Léo Stefanescu, et al., 2017) defines a binary logical relation to prove that the ST monad preserves observational purity of the language, and ReLoC is a relational logic encoded in Iris (Dan Frumin, Robbert Krebbers, and Lars Birkedal, 2018). The goal of these logics is to prove contextual refinements, whereas our goal is to prove more intentional notions of refinements at the level of traces. As a consequence, they can use the usual Iris weakest precondition, and put the state of the program being refined in the logical state. They state in an Iris invariant the fact the current state of the refined program is reachable from its initial state, as we explained at the beginning of this chapter.

The goal of the Polaris logic, by Tassarotti and Harper (2019), is to reason about programs written in a probabilistic concurrent shared memory programming language. The nondeterminism of the scheduler mixed with the probabilistic behavior of the program is a difficult combination to reason about. Their solution is to relate such a program e with a simple probabilistic program which denotes an element of a monad for mixing non-determinism and probabilities defined by Varacca and Winskel (2006). Our method is similar since we relate a program with an abstract specification of its behavior, which is easier to study. Their weakest precondition is similar to ours, specialized to their notion of abstract model.

There are works aiming at proving non-interference properties of programs using program logics. The work closest to ours is SeLoC (D. Frumin, R. Krebbers, and L. Birkedal, 2021) a logic on top of Iris to prove strong low-bisimulations. Bisimulations are easier in some respect to deal with because they do not have existential quantifiers interacting with the step indexing, and they do not deal with stuttering.

6 Fair termination in Iris

Most concurrent programs need a fair scheduler to make progress. This is especially true when programs cannot rely on an operating system to put to sleep threads waiting for a lock to become unlocked and to wake them up when another thread unlocks this lock.

This means that whenever a thread A is waiting on another thread B to be able to make progress, for instance the following program which uses a lock ℓ :

$$\text{acquire}(\ell); \text{release}(\ell) \quad \parallel \quad \text{acquire}(\ell); \text{release}(\ell)$$

an unfair scheduler could first execute the left hand-side acquire operation, followed by only executing the right hand-side thread. Because the program cannot go to sleep, the implementation of acquire must spin and try to acquire the lock in a loop until it succeeds. This means that this program has infinite executions.

In most circumstances, it is not a reasonable behavior for the scheduler, and we only consider traces where the scheduler is behaving fairly:

Definition 6.0.1. Consider an execution trace of a program, of the form:

$$(\text{tp}_1, \mathfrak{s}_1) \xrightarrow{\text{tp}}_{\text{tid}_1} (\text{tp}_2, \mathfrak{s}_2) \xrightarrow{\text{tp}}_{\text{tid}_2} \dots$$

This trace is called *fair* if, either the trace is finite, or, for all $n \in \mathbb{N} \setminus \{0\}$, if the thread at index tid of tp_n is enabled, then there exists $m \in \mathbb{N}$ such that either:

1. the thread at index tid is not enabled in tp_{n+m} , or
2. the thread tid is executed in step $n + m$, in that: $\text{tid}_{n+m} = \text{tid}$.

In words and in the language of linear temporal logic, it is *always* the case that if a thread is enabled, *eventually* it is disabled or it runs.

Note that, in our particular setting, a thread stops being enabled only after it takes at least one step, so the first disjunct of the definition could be removed. We choose however to use the usual definition.

It is easy to convince oneself that every **fair execution** of the simple program above is finite. By definition, this means that it is **fairly terminating**.

Our goal in this chapter is to use the framework described in the previous chapter to prove the **fair termination** of concurrent programs in the Iris logic.

6.1 Termination and fairness preserving simulation

Proving termination of a program “directly” in Iris is made very difficult by the step-indexing underlying it: the validity of an Iris predicate states that it holds at every step-index. If, as it is the case for most weakest preconditions, each step of the program corresponds to at least one later modality \triangleright in the weakest precondition, then we are limited to expressing properties about all finite prefixes for executions of the program. Termination being a liveness property, finite prefixes are not enough.

One solution, which is part of the Iris development, is to define a weakest precondition without later modalities. Such a weakest precondition guarantees termination, but does not allow to eliminate later modalities. As a consequence, not all predicates can be used in invariants, and Löb induction cannot be used.

6.1.1 Fairness models

Our approach is instead to relate the program to an abstract model expressed as an STS in such a way that if the model is fairly terminating, so is the program. This means that the traces of the STSs we consider need to come with a notion of fairness, which is not the case of the simple unlabeled STSs of the previous chapter. Instead, we use the following notion:

Definition 6.1.1. A **fairness model** \mathcal{F} is the data of a set \mathcal{F} of states, a set \mathcal{R} of roles, and a transition relation $\rightarrow \subseteq \mathcal{F} \times \mathcal{R} \times \mathcal{F}$ labeled by roles. Moreover, it is equipped with a map $\text{enabled_roles} : \mathcal{F} \rightarrow \wp_{\text{fin}}(\mathcal{R})$ which associates a finite set of *enabled roles* to each state $s \in \mathcal{F}$. It must approximate the set of outgoing roles:

$$\forall \rho \in \mathcal{R}, \forall s, s' \in \mathcal{F}, s \xrightarrow{\rho} s' \implies \rho \in \text{enabled_roles } s.$$

and it must not disable other roles: if $s \xrightarrow{\rho} s'$, then

$$\forall \rho' \neq \rho, \rho' \in \text{enabled_roles } s \implies \rho' \in \text{enabled_roles } s'$$

Finally, \mathcal{F} comprises a map $\text{fuel_limit} : \mathcal{F} \rightarrow \mathbb{N}$ which will be useful to ensure finite branching conditions.

A run of \mathcal{F} is a non-empty finite or infinite sequence of the form:

$$s_1 \xrightarrow{\rho_1} s_2 \xrightarrow{\rho_2} \dots$$

The notion of role plays the same role as thread identifiers in the STS of programs: they allow to talk about “who” is taking steps. In particular, this notion allows to define a notion of fair run for a **fairness model**:

Definition 6.1.2. Consider a run of a **fairness model** \mathcal{F} of the form:

$$s_1 \xrightarrow{\rho_1} s_2 \xrightarrow{\rho_2} \dots$$

This run is called **fair** if it is finite, or if, for all $n \in \mathbb{N} \setminus \{0\}$, for all $\rho \in \text{enabled_roles } s_n$, there exists $m \in \mathbb{N}$ such that either $\rho \notin \text{enabled_roles } s_{n+m}$, or $\rho = \rho_{n+m}$.

The fairness model \mathcal{F} is called **fairly terminating** if all its fair runs are finite.

6.1.2 Running example

The running example in this chapter is a classical example of a simple program which fairly terminates. It has two threads which each flip k times a shared Boolean variable b . Thread Yes flips it from true to false, and thread No flips it from false to true. In code:

```
Yes () := (if CAS(b, true, false) then n := !n - 1); if !n > 0 then Yes ()
No () := (if CAS(b, false, true) then m := !m - 1); if !m > 0 then No ()
```

where we use ML notations for references: $!n$ reads the contents of the reference n , and $n := v$ updates it with a new value v . Note that only b is shared, the variables n and m are local to each thread, and initially hold k .

The $\text{CAS}(\ell, v_1, v_2)$, which stands for compare-and-swap, or compare-and-set, is an atomic primitive which updates the contents of ℓ with v_2 if it currently contains v_1 , in which case it returns true; otherwise, if ℓ contains another value than v_1 , it does not modify ℓ and returns false.


This program **fairly terminates** because, intuitively, at each stage of the computation, exactly one of the two threads can make progress. In a **fair execution**, each thread eventually gets executed, which means that the program makes progress eventually. This means that the contents of n and m always eventually decrease, and since they are always positive, the program terminates.

The model we will use, in a sense, formalizes the above reasoning by only counting the relevant steps, which are the CAS operations, which are less than one twentieth of the reduction steps in this simple example.

To define a small model, it is useful to note that the contents of n and of m are almost equal just before executing the CAS operation:

- if b is true, n and m are equal,
- if b is false, $m = n + 1$

this means that, in a way, the state of the program can be summarized by the value of m and the value of b . Hence, the states of the **fairness model** \mathcal{F}_{YN} for the example above are the pairs (m, b) of a natural number $m \in \mathbb{N}$ and a Boolean $b \in \mathbb{2}$. There are two roles, Yes and No, which represent the actions of the corresponding threads of the program. The transitions are the following, labeled with their role:



The loops represent the failed CASes, and the horizontal transitions represent the successful CASes which make the program progress. For now, let us say that both roles are enabled in every state, that is, `enabled_roles` is the constant map equal to $\{\text{Yes}, \text{No}\}$. We will see in the sequel that, for rather unfortunate technical reasons, the model needs to be slightly modified.

Since our strategy to prove that the parallel composition of Yes and of No is **fairly terminating** is to reduce it to the fact that \mathcal{F}_{YN} is a **fairly terminating fairness model**, it is natural to ask how does one prove the latter.

6.1.3 A local criterion for the fair termination of models

Proving directly that a simple model such as \mathcal{F}_{YN} is fairly terminating is very tedious, as one needs to do a lexicographic induction on the pairs which define the states, but also an induction on the “deadline” m which Definition 6.1.2 supplies.

Instead, one can use a local criterion which formalizes the intuition that at each step, there is a certain role ρ which will make progress when it gets to run. Given a fair run, we know that this role is eventually going to run, and therefore that progress will be made. The notion of progress is defined using the usual notion of well-founded order:

Definition 6.1.3. An ordered set (A, \leq) is a *well-founded ordered set* if there does not exist any infinite decreasing sequences of elements of A .

The states of the fairness model $\mathcal{F}_{\mathbb{N}}$ equipped with the lexicographic order over $\mathbb{N} \times \mathbb{2}$ are well-founded. We can now state our criterion:

Definition 6.1.4. A *fairness model* \mathcal{F} is called *locally fairly terminating* if there exists a well founded order \leq over \mathcal{F} and a map $\pi : \mathcal{F} \rightarrow \mathcal{R}$ from states to roles which satisfies the following conditions:

1. for all transitions $s \xrightarrow{\rho} s', s' \leq s$;
2. for all states $s \in \mathcal{F}$ which are not dead ends, $\pi(s) \in \text{enabled_roles } s$ and for all s' such that $s \xrightarrow{\pi(s)} s', s' < s$;
3. for all transitions $s \xrightarrow{\rho} s',$ if $\rho \neq \pi(s)$, then $\pi(s') = \pi(s)$;

where a state s is called a dead end if there are no outgoing transitions from it.

We call this criterion *local* because it can be checked for each transition independently, without any reference to traces. This criterion is correct:

Lemma 6.1.5. *If a fairness model \mathcal{F} is locally fairly terminating, then it is fairly terminating.*

Proof. The proof proceeds by proving the following statement by well-founded induction over \bar{s} : *For any run mt whose first state is $\leq \bar{s}$, if it is fair, then it is finite.*

Let \bar{s} be a state and mt be such a fair model run; write s_k its k th state and ρ_k the role of its k th transition, pose $\rho := \pi(s_1)$. By definition of locally fair model, $\rho \in \text{enabled_roles } s_1$, which means, since mt is fair, that there exists $m \in \mathbb{N}$ such that

$$\rho \notin \text{enabled_roles } s_{1+m} \vee \rho = \rho_{1+m} \tag{6.1}$$

where we see s_k and ρ_k as depending on the trace mt . We proceed by induction over natural numbers m .

For $m = 0$, the left disjunct is in contradiction with $\rho \in \text{enabled_roles } s_1$, therefore $\rho = \rho_1$ and, by definition of local fairness, $s_2 < s_1$, and we conclude with the “outer” induction hypothesis applied to the tail of the run mt .

For $m > 0$, we examine the first transition of mt :

1. If the trace is a singleton it is finite and we are done.

2. If it is of the form $s_1 \xrightarrow{\rho} s_2$, we proceed as above with the “outer” induction hypothesis.
3. Otherwise, $\rho_1 \neq \rho$, so $\pi(s_2) = \pi(s_1)$, and the tail of mt satisfies (6.1) for $m - 1$ instead of m , and so we can conclude using the “inner” induction hypothesis.

□

The criterion is enough to prove that \mathcal{F}_{YN} is **fairly terminating**: define $\pi((m, b))$ to be Yes for $b = 1$ and No otherwise. Checking that this defines a **locally fairly terminating** model is immediate.

6.1.4 Termination and fairness preserving refinements

Now that we have explained how to establish the **fair termination** of a **fairness model** \mathcal{F} , it remains to define a notion of refinement *between traces* \lesssim which preserves the properties we need: Given a program execution et and a fairness model run mt such that $et \lesssim mt$,

1. if et is a **fair execution**, then mt is a **fair run**,
2. if mt is a finite run, then et is a finite execution.

Given such a notion of refinement, to prove that a program e **fairly terminates**, it is sufficient to find a **fairly terminating fairness model** \mathcal{F} such that, for all execution trace et of e , there exists a run mt of \mathcal{F} such that $et \lesssim mt$. The existence of such a run will follow in great part from the corollary of the soundness of trwp which we stated in Section 5.2.3.

The main difficulty in defining the trace refinement relation \lesssim is to enforce fairness preservation. The intuition is the following. If the execution trace et is fair, this means that “resources” are equitably allocated to each thread; and such a trace must only be related to run of \mathcal{F} where resources are allocated equitably to each role. Keeping in mind that we want our proof method to be thread-local, we ask that each role ρ be the responsibility of a thread $\mathfrak{T}(\rho) \in \mathbb{N}$, and each thread tid needs to allocate the resources it gets from the scheduler fairly among the set $\mathfrak{T}^{-1}(tid)$ of roles it is in charge of.

More concretely, any infinite sequence of transitions of thread tid in et must correspond to an infinite number of transitions in mt of *every* role in $\mathfrak{T}^{-1}(tid)$. This is achieved by associating a “fuel amount” $\mathfrak{F}(\rho) \in \mathbb{N}$ to each role ρ . Each step of a thread tid must decrease (strictly) the amount of fuel of every role in $\mathfrak{T}^{-1}(tid)$, unless this step corresponds, according to \lesssim , to a step labeled with a role $\rho \in \mathfrak{T}^{-1}(tid)$, in which case

the fuels of the roles in $\mathfrak{I}^{-1}(tid) \setminus \{\rho\}$ must decrease, and the fuel of ρ can be filled up.

Remark 6.1.6. This notion of fuel is, in a sense, dual to step-indexing: if the step-index reaches “-1”, Prover wins the game (that is, $0 \models \triangleright \perp$ is true), in other words, Verifier needs to find a contradiction in a finite number of moves. Here, on the other hand, Prover needs to produce a witness before the fuel runs out.

This relation \lesssim over traces needs to keep track of these two mappings \mathfrak{I} and \mathfrak{F} which need to evolve at each step. To manage this complexity, we define \lesssim as the composition of two relations: $\lesssim := \lesssim_s \circ \lesssim_f$ which relate the former traces to an existentially quantifier trace of an STS $\mathcal{Vive}(\mathcal{F})$. In other words, for program execution et and a model run mt ,

$$et \lesssim mt \quad \Leftrightarrow \quad \exists t \in \mathbf{Runs}(\mathcal{Vive}(\mathcal{F})), \quad et \lesssim_f t \wedge t \lesssim_s mt$$

The relation \lesssim_f is the predicate on traces which is induced by a mild generalization of the notion of [simulation](#) of Definition 5.1.1, and \lesssim_s is a very simple as well. Hence, the whole complexity of the accounting of fuel and of the correspondence between threads and roles (which can change dynamically) is confined in the definition of $\mathcal{Vive}(\mathcal{F})$.

6.1.5 The \mathcal{Vive} construction

Given a [fairness model](#) \mathcal{F} , we define a (labeled) STS $\mathcal{Vive}(\mathcal{F})$ which keeps track of mapping between roles and threads, and of fuels, as suggested above.

Definition of the labeled STS

A state of $\mathcal{Vive}(\mathcal{F})$ is a triple $(s, \mathfrak{F}, \mathfrak{I})$ of a state $s \in \mathcal{F}$, together with two maps $\mathfrak{F} : \text{enabled_roles } s \rightarrow \mathbb{N}$ and $\mathfrak{I} : \text{enabled_roles } s \rightarrow \mathbb{N}$ which associate a fuel amount and a thread ID to each role ρ which is enabled in the current underlying state s .

The set of labels of $\mathcal{Vive}(\mathcal{F})$ is

$$\{\text{Step } \rho \text{ } tid \mid \rho \in \mathcal{R}, tid \in \mathbb{N}\} \quad \cup \quad \{\text{Silent } tid \mid tid \in \mathbb{N}\}$$

(recall that \mathcal{R} is the set of roles of the [fairness model](#) \mathcal{F}) The intuition is that a step labeled by $\text{Step } \rho \text{ } tid$ corresponds to the situation where the thread tid takes a step in the program, and one of the roles ρ under its responsibility takes a step in the fairness model. A step labeled $\text{Silent } tid$, on the other hand, corresponds to a step in the program which does not correspond to a step in the fairness model, in other words, a stuttering step.

We now describe the transitions in the labeled STS $\mathcal{L}\text{ive}(\mathcal{F})$. The idea is that there are two kinds of transitions

- A thread tid can take a step in $\mathcal{L}\text{ive}(M)$ of the form

$$(s, \mathfrak{F}, \mathfrak{I}) \xrightarrow{\text{Silent } tid} (s, \mathfrak{F}', \mathfrak{I}')$$

which does not corresponds to a step in the underlying fairness model \mathcal{F} in exchange of consuming fuel: for *every* $\rho \in \mathfrak{I}^{-1}(tid)$ which the thread tid is in charge of, the corresponding fuel $\mathfrak{F}(\rho)$ must decrease strictly, in that $\mathfrak{F}'(\rho) < \mathfrak{F}(\rho)$.

- A thread tid can also take a step in the underlying model which corresponds to a role $\rho \in \mathfrak{I}^{-1}(tid)$

$$(s, \mathfrak{F}, \mathfrak{I}) \xrightarrow{\text{Step } \rho \text{ } tid} (s', \mathfrak{F}', \mathfrak{I}')$$

This allows to refill the fuel of ρ up to the limit $\text{fuel_limit } s'$, that is, $\mathfrak{F}'(\rho) \leq \text{fuel_limit } s'$; this is required to keep the STS finitely branching. All the other roles which are associated with tid must decrease:

$$\forall \rho' \in \mathfrak{I}(tid) \setminus \{\rho\}, \quad \mathfrak{F}'(\rho') < \mathfrak{F}(\rho')$$

Roles which appear between s and s' can have any fuel $\leq \text{fuel_limit } s'$ in \mathfrak{F}' . Of course, we also require there be a step

$$s \xrightarrow{\rho} s'$$

in the fairness model \mathcal{F} .

In addition to the two constraints above, in both cases, the fuel of the roles which are not associated with the thread tid must not increase:

$$\forall \rho \in \mathcal{R} \setminus \mathfrak{I}^{-1}(tid), \quad \mathfrak{F}'(\rho) \leq \mathfrak{F}(\rho)$$

and roles which change owners ($\mathfrak{I}'(\rho) \neq \mathfrak{I}(\rho)$) must have their fuel decrease strictly.

The trace refinement predicates \lesssim_f and \lesssim_s

Given two labeled STS \mathcal{M}_1 and \mathcal{M}_2 with states S_1 and S_2 and labels L_1 and L_2 respectively, and relations $\phi_S \subseteq S_1 \times S_2$ and $\phi_L \subseteq L_1 \times L_2$, we say that a trace et of \mathcal{M}_1 and a trace mt of \mathcal{M}_2 are **related by ϕ_S and ϕ_L** if: they have the same size, and, for all n , the n th state of et and the n th state of mt are related by ϕ_S , and similarly ϕ_L relates the respective n th labels of the traces.

The refinement predicate \lesssim_f is defined as follows:

$$et \lesssim_f t \iff et \text{ and } t \text{ are related by } \text{enabled_tids} \text{ and } \text{labels_match}$$

where those two relations are defined as follow. Given a label tid of the labeled STS \mathcal{L} of the programming language, and a label ℓ of the labeled STS $\mathcal{Qive}(\mathcal{F})$, $\text{labels_match } tid \ell$ holds iff

$$\ell = \text{Silent } tid \quad \vee \quad \ell = \text{Step } tid \rho \quad (\text{for some } \rho)$$

The predicate over states of the two STSs enforces that thread IDs which appear in the states of $\mathcal{Qive}(\mathcal{F})$ correspond to threads in the corresponding state of the program, and that no role is the responsibility of a thread which is not enabled. Hence we define $\text{enabled_tids}((tp, s), (s, \mathfrak{F}, \mathfrak{I}))$ as

$$\begin{aligned} & (\forall \rho, tid, \mathfrak{I}(\rho) = tid \Rightarrow tid < |tp|) \wedge \\ & (\forall tid, tp[tid] \in \mathbf{Val} \Rightarrow (\forall \rho, \mathfrak{I}(\rho) \neq tid)) \end{aligned} \tag{6.2}$$

The second trace refinement relation \lesssim_s is simpler. There is forgetful map \mathcal{D} from the states of $\mathcal{Qive}(\mathcal{F})$ to the sets of \mathcal{F} :

$$(s, \mathfrak{F}, \mathfrak{I}) \mapsto s$$

and a partial map \mathcal{L} from the labels of of \mathcal{Qive} to the sets of \mathcal{F} which maps the labels of the form $\text{Step } tid \rho$ to ρ and is undefined otherwise. This defines a map \mathcal{D} from the traces of $\mathcal{Qive}(\mathcal{F})$ to the traces of \mathcal{F} by disregarding transitions which are not in the domain of the label partial map above; this corresponds to destuttering of the trace.

This map is very well behaved:

Lemma 6.1.7. *If the image under \mathcal{L} of the label of transition (trace of length 2) $(s, \mathfrak{F}, \mathfrak{I}) \xrightarrow{\ell} (s', \mathfrak{F}', \mathfrak{I}')$ is not defined, then $s = s'$.*

Moreover, there exists a map Ψ from the states of $\mathcal{Qive}(\mathcal{F})$ to \mathbb{N} such that, for transitions as above:

$$\Psi(s', \mathfrak{F}', \mathfrak{I}') < \Psi(s, \mathfrak{F}, \mathfrak{I})$$

Proof. Take $\Psi(s, \mathfrak{F}, \mathfrak{I}) := \sum_{\rho \in \text{enabled_roles } s} \mathfrak{F}(\rho)$. □

The existence of the decreasing potential Ψ means in particular that the destuttering map \mathcal{D} can be defined in Coq. Finally, we define \lesssim_s as the graph of \mathcal{D} :

$$t \lesssim_f mt \iff \mathcal{D}(t) = mt$$

Preservation of fairness and termination

There is a natural notion of fairness for runs of the STS $\mathcal{L}ive(\mathcal{F})$:


Definition 6.1.8. Consider a run of $\mathcal{L}ive(\mathcal{F})$, for a **fairness model** \mathcal{F} , of the form:

$$(s_1, \mathfrak{F}_1, \mathfrak{I}_1) \xrightarrow{\ell_1} (s_2, \mathfrak{F}_2, \mathfrak{I}_2) \xrightarrow{\ell_2} \dots$$

This run is called **fair** if it is finite, or for all $n \in \mathbb{N} \setminus \{0\}$, for all $\rho \in \text{enabled_roles } s_n$, then there exists $m \in \mathbb{N}$ such that either: $\rho \notin \text{enabled_roles } s_{n+m}$, or $\ell_{n+m} = \text{Step } tid \rho$ for some $tid \in \mathbb{N}$.

This amounts to defining a fair run of $\mathcal{L}ive(\mathcal{F})$ as being in the preimage under \mathcal{D} of a **fair run** of \mathcal{F} .

We are going to prove the following:



where, for instance, the arrow from et to t means that if et is fair, then t is fair. The right hand-side arrows are easy to prove, because \simeq_s relates traces which are equal up to finite stuttering. The implication that, if t is finite, then et is also finite is obvious because of the definition of \simeq_f states that they must have the same length.

The part which is more difficult to prove is that \simeq_f preserves fairness.

Lemma 6.1.9. *Given a execution et and a run of $\mathcal{L}ive(\mathcal{F})$, if et is a **fair execution**, then t is a **fair run**.*

Idea of the proof. We write $(s_n, \mathfrak{F}_n, \mathfrak{I}_n) \xrightarrow{\ell_n} (s_{n+1}, \mathfrak{F}_{n+1}, \mathfrak{I}_{n+1})$ the n -th transition of t , and, similarly, we write $(tp_n, \mathfrak{s}_n) \xrightarrow{tid_n} (tp_{n+1}, \mathfrak{s}_{n+1})$ the n th transition of et .

Let ρ be a role, and assume without loss of generality that it is enabled in the first state of t . We need to prove that there exists a natural number $m \in \mathbb{N}$ such that $\rho \notin \text{enabled_roles } s_{1+m}$ or such that $\ell_{1+m} = \text{Step } \rho \text{ } tid$ for some tid .

Let $tid = \mathfrak{I}_1(\rho)$, which is defined since we assume the role ρ is enabled at the beginning. Because et is fair, we know that there exists N such that $tid_N = tid$. We proceed by lexicographic induction on the pair $(\mathfrak{I}_1(\rho), N)$ of the fuel of ρ at the beginning of t and this bound N . Let us consider the first transition of t .

1. If $\ell_1 = \text{Step } \rho \text{ } tid$, we are done.
2. If $\ell_1 = \text{Step } \rho' \text{ } tid'$ or $\text{Silent } tid'$, with $tid' \neq tid$, we use the induction hypothesis with the tails of t and of et , with $N - 1$ and the same fuel.
3. If $\ell_1 = \text{Step } \rho' \text{ } tid$ with $\rho' \neq \rho$, then the definition of the transition in $\mathcal{L}\text{ive}(\mathcal{F})$ states that the fuel of ρ must decrease, and fairness of et gives us a new bound N before which thread tid will take a step, and we conclude using the induction hypothesis.

□

We have now proved everything we need about the trace refinement relations \lesssim_f and \lesssim_s . It remains to explain how to prove that, given a program execution et , there exists runs t of $\mathcal{L}\text{ive}(\mathcal{F})$ and mt of \mathcal{F} such that $et \lesssim_f t \lesssim_s mt$. It is easy to see that it suffices to find a trace t , since we can then define $mt := \mathcal{D}(t)$. The next section explains how the program logic presented in the previous chapter can be used to prove the existence of t .

6.2 A logic for proving fairness and termination preserving refinements

We use the logic we described in the previous chapter instantiated with the labeled STS $\mathcal{L}\text{ive}(\mathcal{F})$ which depends on a user chosen **fairness model** \mathcal{F} .

6.2.1 Changes to the logic

To accommodate the change from unlabeled to labeled state transition systems, we need to slightly change the definition of trwp . Another modification is that because we need to keep track of which thread takes which transition in the fairness model, we need to annotate the weakest precondition with the thread ID of the thread running the code. This new weakest precondition is defined as follows, with the changes highlighted

in green:

$$\begin{aligned}
 \text{fwp}_\varepsilon e_1 @tid\{Q\} &:= \\
 &(\ulcorner e_1 \in \mathbf{Val} \urcorner \wedge \varepsilon \Vdash^\varepsilon Q e_1) \vee (\ulcorner e \notin \mathbf{Val} \urcorner \wedge \\
 &\forall s_1, s_1, n, \text{state_interp } s_1 s_1 n \multimap \varepsilon \Vdash^\emptyset \ulcorner \text{reducible } e_1 s_1 \urcorner * \\
 &\forall e_2, s_2, \vec{e}_f. \ulcorner (e_1, s_1) \rightarrow (e_2, s_2, \vec{e}_f) \urcorner \multimap \emptyset \Vdash^\emptyset \triangleright \emptyset \Vdash^\varepsilon \exists s_2, \ell. \\
 &\ulcorner \text{state_evol } s_1 s_1 tid \ell (n, |\vec{e}_f|) s_2 s_2 \urcorner \multimap \text{state_interp } s_2 s_2 (n + |\vec{e}_f|) * \\
 &\text{fwp}_\varepsilon e_2 @tid\{Q\} * \bigotimes_{k \mapsto e \in \vec{e}_f} \text{fwp}_\top e @n+k \{\text{fork_post } n+k\}
 \end{aligned}$$

In addition to the new *tid* parameter to *fwp*, the main change is that the *state_evol* predicate takes labels and the thread count into account. The pair $(n, |\vec{e}_f|)$ which is passed to it corresponds to the thread count before the reduction step is taken, and the number of threads created during the reduction step of the program. The total number of threads is retrieved using a new parameter to *state_interp* which corresponds to the total number of threads. Finally, threads which are spawned can have a non-trivial postcondition *fork_post tid* which may depend on their thread ID.

In the case where the labeled STS is $\mathcal{Q}\text{ive}(\mathcal{F})$, the predicate *state_evol* is instantiated as follows:

$$\begin{aligned}
 \text{state_evol } s_1 s_1 tid \ell (n, k) s_2 s_2 &:= \\
 \text{labels_match } tid \ell \wedge s_1 \xrightarrow{\ell} s_2 \wedge (\text{tids_le } n s_1 \Rightarrow \text{tids_le } (n + k) s_1)
 \end{aligned}$$

where *tids_le* $n s$ states that every *tid* in the \mathcal{T} component of s is $\leq n$, and where the *labels_match* predicate is defined in Section 6.1.5, page 149.

The soundness theorem of the logic can also be slightly adapted to fit the labeled setting without any difficulty.

All that remains to do is to deduce a trace refinement from the ϕ -history sensitive similarity predicate which the soundness theorem provides.

Lemma 6.2.1. *Let \mathcal{M}_1 and \mathcal{M}_2 be two labeled STSs, and let $\phi \subseteq \mathbf{FRuns}(\mathcal{M}_1) \times \mathbf{FRuns}(\mathcal{M}_2)$ a relation between finite runs. Let $\phi_S \subseteq \mathcal{M}_1 \times \mathcal{M}_2$ be a relation between their states, and ϕ_L be a relation between their labels. Then, if $s_1 \in \mathcal{M}_1$ and $s_2 \in \mathcal{M}_2$ are ϕ -history sensitive similar, then, for all finite or infinite trace $t_1 \in \mathbf{Runs}(\mathcal{M}_1)$ which starts with the state s_1 , there exists a trace $t_2 \in \mathbf{Runs}(\mathcal{M}_2)$ such that t_1 and t_2 are related by ϕ_S and ϕ_L .*

Using this lemma, we can deduce the existence, for every execution et of a program from an initial state which has been proved with fwp , of a trace t of $\mathcal{L}\text{ive}(\mathcal{F})$ such that $et \preceq_f t$. We now finish by explaining how the logical resources are setup to provide Iris predicates and reasoning rules which are easy to manipulate.

6.2.2 Logical resources

As expected, the Iris logical state contains Iris resources for each the three components of the states of $\mathcal{L}\text{ive}(\mathcal{F})$, and they are tied to the “raw” model state in the state_interp predicate.

First, as in the previous chapter, there is a predicate $\text{model_is } s$ which state that the current state of the underlying **fairness model** \mathcal{F} is equal to s . As before, it must agree with the state_interp predicate:

$$\text{state_interp } s (s, \mathfrak{F}, \mathfrak{I}) n * \text{model_is } s' \quad \dashv * \quad \ulcorner s = s' \urcorner$$

For each role ρ which is enabled in the current state of the **fairness model** \mathcal{F} , there is a predicate $\rho \mapsto_F \mathfrak{f}$ which states that the fuel amount associated to ρ is \mathfrak{f} . It satisfies the rule:

$$\text{state_interp } s (s, \mathfrak{F}, \mathfrak{I}) n * \rho \mapsto_F \mathfrak{f} \quad \dashv * \quad \ulcorner \mathfrak{F}(\rho) = \mathfrak{f} \urcorner$$

The way we reflect mapping is less straightforward: instead of maps from roles to thread IDs, they are represented as maps from thread IDs to sets of roles. Formally, we represent the map $\mathfrak{I} : R \rightarrow \mathbb{N}$ as a map $M' : \mathbb{N} \rightarrow \wp_{\text{fin}}(R)$ such that:

$$M \rho = \text{tid} \quad \Leftrightarrow \quad \rho \in M' \text{tid}$$

where $R \subseteq \mathcal{R}$ is a set of roles of \mathcal{F} . We put an exclusive structure on the finite sets of roles in the codomain of the map. That is, the **camera** operation on $\wp_{\text{fin}}(R)$ which we use is never defined. We write $\text{tid} \mapsto_T \vec{\rho}$ to denote the exclusive ownership of the singleton map $[\text{tid} := \vec{\rho}]$, where $\vec{\rho} \subseteq R$. The exclusive structure on sets of roles means that $\text{tid} \mapsto_T \vec{\rho} * \text{tid}' \mapsto_T \vec{\rho}'$ implies that $\text{tid} \neq \text{tid}'$.

The law which relates it to the states interpretation is the following:

$$\text{state_interp } s (s, \mathfrak{F}, \mathfrak{I}) n * \text{tid} \mapsto_T \vec{\rho} \quad \dashv * \quad \ulcorner \forall \rho, \rho \in \vec{\rho} \Leftrightarrow \mathfrak{I}(\rho) = \text{tid} \urcorner \quad (6.3)$$

Remark 6.2.2 (Why exclusive?). The reason we need an exclusive structure on the set of roles above is that the resource corresponds to an *obligation*: when the thread tid makes a transition, the set of roles whose fuel decreases must include the set of roles

associated with the thread tid . A simple way of achieving this is to associate tid with exactly its set of roles in the model.

If we had used a more natural camera structure on these sets, such as disjoint union, the reasoning principle (6.3) would have been too weak, and the equivalence on the right would have been a simple implication. Another way to see this is that the set of roles $\vec{\rho}$ in $tid \mapsto_T \vec{\rho}$ is *linear*. Of course the whole $tid \mapsto_T \vec{\rho}$ can be dropped since Iris is an affine logic, but we cannot drop one single $\rho \in \vec{\rho}$ like we would have been able to do if we had put a disjoint union structure on $\wp_{\text{fin}}(\mathcal{R})$.

The resources for the fuel and the mapping are often used in concert since they contain all the information about some thread. We define

$$\text{tid_has_fuels } tid \ \mathfrak{f}s := tid \mapsto_T \text{dom}(\mathfrak{f}s) * \bigotimes_{\rho \mapsto \mathfrak{f} \in \mathfrak{f}s} \rho \mapsto_F \mathfrak{f}$$

where $\mathfrak{f}s$ is a finite partial map from roles ρ to fuel amounts \mathfrak{f} . This predicate states that the thread tid is associated to the roles which are in the domain of $\mathfrak{f}s$ and each of these roles ρ has fuel $\mathfrak{f}s(\rho)$.

6.2.3 Inference rules

Unlike in the previous chapter, where the standard instantiation of trwp allows for infinite stutter by choosing as STS the reflexive closure of some user-chosen STS \mathcal{M} , which allows the logic to be conservative over the usual Iris weakest precondition wp , fwp is *not* conservative of wp . Indeed, each “internal” step of of each thread tid which does not correspond to a step in the **fairness model** \mathcal{F} consumes the fuel of the roles which are under its responsibility.

Silent steps

Iris has a useful concept of a pure reduction step of a program. Intuitively, for ML-style languages, this corresponds to steps of a pure λ -calculus.

Definition 6.2.3. A *pure reduction step* $e \rightarrow e'$ between two expressions is the data of two expressions which correspond to a reduction step which ignores the machine state

$$\forall s, (e, s) \rightarrow (e', s)$$

and which is deterministic:

$$\forall s, s', e'' (e, s) \rightarrow (e'', s') \Rightarrow e'' = e' \wedge s' = s$$

This definition allows to define a rule for pure reductions: if $e \rightarrow e'$ is a pure reduction, then:

$$\frac{\text{tid_has_fuels } tid \ (\text{fs} + 1) * (\text{tid_has_fuels } tid \ \text{fs} \multimap \text{fwp}_\varepsilon e' @tid \{Q\}) \quad \text{fs} \neq \emptyset}{\text{fwp}_\varepsilon e @tid \{Q\}}$$

where $\text{fs} + 1$ adds 1 to each fuel in the finite partial map fs . This simply says that reducing by one step “costs” one unit of fuel for every role associated to tid . The standard instantiation of trwp in the previous chapter and the standard Iris weakest precondition wp , in contrast, have the simpler rule:

$$\frac{\text{wp}_\varepsilon e' \{Q\}}{\text{wp}_\varepsilon e \{Q\}}$$

The rule for fwp can be automated in a proof assistant so that the fuel is automatically decreased and the program e automatically reduced until it reaches an expression which is not purely reducible. (Currently, this tactic only supports the common case where fs is a singleton, but this should not be too difficult to fix)

The other rules of the logic, for reading and writing references or for CAS operations, including the invariant opening, admit inference rules which are the same as for wp except for the fuel being decremented using the pattern above.

Steps in the fairness model

For the steps of a thread tid which correspond to a step in the fairness model \mathcal{F} , the rules are a little more complicated because the role ρ of the step in \mathcal{F} is distinguished from the other roles associated to the thread tid . We give as an example the rule for a successful CAS, since it is a common point for externally visible actions, and with a role associated to the thread tid . The rule is expressed as a Hoare triple, which is defined using fwp as explained in Chapter 4.

The Hoare triple

$$\frac{\{\triangleright \ell \mapsto v_1 * \triangleright \text{model_is } s_1 * \triangleright \text{tid_has_fuels } tid \ [\rho := f_1]\}}{\text{CAS}(\ell, v_1, v_2) @tid} \frac{}{\{\ell \mapsto v_2 * \text{model_is } s_2 * \text{tid_has_fuels } tid \ [\rho := f_2]\}}$$

holds when

- there is a transition $s_1 \xrightarrow{\rho} s_2$ in the underlying fairness model \mathcal{F} ;
- the role ρ is still enabled in s_2 : $\rho \in \text{enabled_roles } s_2$;

- the new fuel associated to ρ is under the limit: $\mathfrak{f}_2 \leq \text{fuel_limit } s_2$.

If the second condition had been false, the postcondition would have contained the predicate $\text{tid} \mapsto_T \emptyset$ instead of $\text{tid_has_fuels } \text{tid} [\rho := \mathfrak{f}_2]$.

Postcondition for forks

We need to maintain the invariant `enabled_tids` which is defined in equation (6.2), which states that only non-finished threads are associated to roles. To do this, we require the post-condition of every forked thread to be `fork_post tid := tid \mapsto_M \emptyset`. The soundness theorem of the logic will also require the main thread to have this postcondition.

For this reason, in the example with the two threads Yes and No, the final step of each thread needs to take a step in the fairness model \mathcal{F}_{YN} to move to a state where the corresponding role is not enabled. This has to be done at the last step because a thread which is associated to no role does not have the right to take any step. We describe how we handle this requirement in the sequel when we describe the proof of the example.

6.2.4 Soundness theorem

We are finally ready to state the soundness theorem of the logic when instantiated with a labeled STS of the form $\mathcal{L}\text{ive}(\mathcal{F})$ for some *fairness model* \mathcal{F} .

Theorem 6.2.4. *Given a program e , a finitely branching *fairly terminating fairness model* \mathcal{F} , a state s_0 in \mathcal{F} , if the following Iris predicate is valid*

$$\text{model_is } \sigma_0 \text{ } * \text{ tid_has_fuels } 0 \text{ } \text{initial_fuel} \text{ } * \text{ } \mathbb{T} \Vdash \mathbb{T} \text{ fwp}_{\top} e @0 \{0 \mapsto_T \emptyset\}$$

*then e is *fairly terminating*.*

In the Iris predicate above, 0 is the thread ID of the main thread, and `initial_fuel` is the constant map defined on `enabled_roles` s_0 equal to `fuel_limit` s_0 .

Proof. The proof is a straightforward application of the soundness theorem of `fwp` and of the results in Section 6.1. One needs in addition to prove that, if \mathcal{F} is finitely branching, then so is $\mathcal{L}\text{ive}(\mathcal{F})$, assuming a bound on the thread IDs which is given by the fact that `tids_le` holds.

Given a state $(s, \mathfrak{F}, \mathfrak{X})$, a successor state $(s', \mathfrak{F}', \mathfrak{X}')$ is the data of s' such that $s \rightarrow s'$, of which we assume there is a finite number, and of two maps defined on a fixed finite set `enabled_roles` s' with finite ranges: for $\rho \in \text{enabled_roles } s'$,

$$\mathfrak{F}'(\rho) \leq \max(\max(\mathfrak{F}), \text{fuel_limit } s') \quad \mathfrak{X}'(\rho) \leq \text{maxtid}$$

where $maxtid$ is the bound given by $tids_le$, and $\max(\mathfrak{F})$ is the maximal fuel value which appears in \mathfrak{F} . \square

6.2.5 Back to the example

Recall that we considered two threads

```
Yes () := (if CAS(b, true, false) then n := !n - 1); if !n > 0 then Yes ()
No () := (if CAS(b, false, true) then m := !m - 1); if !m > 0 then No ()
```

flipping a Boolean variable b back and forth a finite number of time, each corresponding to a role in the fairness model \mathcal{F}_{YN} . They do not form a complete program, since they rely on references to be allocated and initialized. We use a program Start which allocates the references, spawns the two threads and stops immediately. Roughly (ignoring lexical scoping issues):

```
Start () :=
  b := alloc(true); n := alloc(k); m := alloc(k);
  fork(Yes); fork(No)
```

We do not need to give a role in \mathcal{F}_{YN} to Start , instead, it is initially responsible of both Yes and No , until the first fork operation, after which it is only responsible of No , and finally it terminates and is responsible of no role. We are lucky, in that there is no operation after it has delegated its last role to the thread No .

This is not the case for the threads Yes and No , because they must decrement their variable after their last CAS. The solution is to add flags to the states of \mathcal{F}_{YN} : a state is of the form

$$(m, b, ye, ne)$$

where the new flags $ye, ne \in \mathbb{Z}$ mean that Yes , respectively No , is still executing. The enabled_roles predicate is defined using these flags using

$$\text{Yes} \in \text{enabled_roles}(m, b, ye, ne) \Leftrightarrow ye = 1$$

and symmetrically for No . The transitions express that threads can shutdown only after their counter has reached zero.

This difficulty surmounted, the proof is fairly straightforward, using an invariant to related the values of b and m in \mathcal{F}_{YN} to the values of the three references.

6.3 Related works

The most closely related work is a paper by Tassarotti, Jung, and Harper (2017) which uses Iris to construct a logic which allows to prove the existence of a fairness preserving termination preserving refinement between two programs. In their case, their goal is to prove that a simple compiler which replaces abstract channel operations by shared memory implementations preserves fair termination. Their soundness theorem is similar as well, in that it states that each fair infinite trace of the target program is refined by a fair trace of the source program.

They have a second example where they prove that a Craig-Landin-Hagersten queue lock refines a ticket-lock. The statement is made formal by implementing a compiler which translates a program using the first into one using the second.

The main technical difference between the two approaches is that they need to change the base model of Iris in two ways. First, to be able to talk about traces, they add a new piece of data to the definition of a **camera**: Each camera M has a step-indexed transition relation $(\rightsquigarrow_n)_{n \in \mathbb{N}}$. For every camera except for the camera which holds the state of the source program, this relation is the total relation.

The second modification to the Iris model is the addition of linear predicates, which they need in order to prevent the prover from dropping the obligation to take a step in the source language. The interpretation of an Iris predicate becomes a **non-expansive** map

$$M \times M \rightarrow \mathbb{P}$$

which is upward closed only with respect to its first parameter, which corresponds to the usual Iris affine predicates; the second parameter corresponds to linear predicates. They have an “affinely” modality to express rules which only hold for affine propositions such as framing and weakening. Linear propositions cannot be put inside an invariant.

Our intuition for why we do not need the predicate $tid \mapsto_{\mathcal{T}} \vec{\rho}$ to be linear is that the weakest precondition takes as input a resource of the form $tid \mapsto_{\mathcal{T}} \vec{\rho}'$, which means that dropping this resource is not a risk. Moreover, as we mentioned in Remark 6.2.2, the exclusive structure on $\vec{\rho}$ makes the predicate behave somewhat linearly.

We think the main advantage of our approach, beyond the fact that changing the base logic of Iris is a fair amount of work, is that putting the role-to-thread mapping and the fuel accounting in the **Live** construction clarifies the reasoning. Their construction of the camera which holds the state of the source program is fairly involved and takes around 7000 lines of Coq (for comparison the **Live** construction and the proofs of Section 6.1 take 1350 lines).

Less closely related, there are several program logics to reason about fairness of concurrent programs, such as the Rely-Guarantee logic LiLi by Liang and Feng (2016), and the concurrent separation logic Tada Live described in a draft by D’Oswaldo et al. (2019).

There also has been a lot of work on the topic of fairness preserving simulations from the model checking community, for example in the context of LTL (Kesten, Manna, and Pnueli, 1994), TLA+ (L. Lamport, 1994), games (Henzinger, Kupferman, and Rajamani, 1997), etc.

Another related line of work are the time credit-base logics by Guéneau, Charguéraud, and François Pottier (2018) and Mével, Jourdan, and François Pottier (2019). Their goal is to prove worst case complexity bounds using a program logic. They add a predicate $\$k$ to the logic to state that the program must finish in less than k steps. This is related to our $tid \mapsto_F \dagger$ predicate. The main difference is that they do not keep track of the identity of the thread, and credits always decrease.

Conclusion

This thesis contains two works I was involved in during my PhD. This last chapter concludes this thesis with remarks about the two works which were presented, and future directions which I think are interesting to pursue.

Semantic models of CSL

In Part I of this thesis, I presented a semantic model of a simple concurrent separation logic and its associated programming language. Our goal was to give a semantic model which gives a semantic interpretation to the proof of CSL themselves, in the tradition of Curry-Howard, and which takes on data races head on.

The simplicity of the language, and the fact that the logic is strongly associated to the language compared to Iris which can be instantiated with any language written in the format of Definition 4.1.1, had the benefit, besides simplifying the model, to bring to light structures in the logic. For example, the change of lock operations on the interpretations of the proofs make sense because the logic follows the same lexical scoping of locks as the language.

Choice of the language

Looking back, I believe however that a better choice of programming language would have been a first order λ -calculus with ML-style references. The first advantage is that variables would have been immutable which can be substituted into in the specifications, instead of using the $\text{Own}(x)$ predicates.

The other advantage is that it would have clarified a question we had as to what is the right shape of the cobordisms we use to interpret programs. In this thesis, the “input interface” I needs to be isomorphic to $\mathbb{Z}[0, i]$, but not the output interface, because we had not found this constraint to have been useful. In a language where programs reduce to values however, the natural solution is to interpret a program of type $C : A \rightarrow B$,

where A and B denote sets if the language is first order, as a cobordism roughly of the form:

$$\begin{array}{ccccc}
 A \times \mathbb{Z}[0] & \longrightarrow & \llbracket C \rrbracket_{support} & \longleftarrow & B' \times \mathbb{Z}[0] \\
 \pi_2 \downarrow & & \downarrow & & \downarrow \pi_2 \\
 \mathbb{Z}[0] & \longrightarrow & \mathbb{Z}[1] & \longleftarrow & \mathbb{Z}[0].
 \end{array}$$

where $B' \subseteq B$, following the intuition that the input and output are “covering spaces” above base space $\mathbb{Z}[0]$ such that the fiber above any memory state is the set of possible values.

Since, in the language we consider, every command can be considered as having the type $\mathbb{1} \rightarrow \mathbb{1}$, this suggests that the right definition is to ask that the map $\llbracket C \rrbracket_{out} \rightarrow \mathbb{Z}[0]$ be an iso. Alas, this remark came to me too late.

Higher order

The variant of CSL we considered cannot reason about higher-order functions. There are two possible ways to scale our approach to this setting.

One possibility would be to do a construction similar to the Int construction of Joyal, Street, and Verity (1996), which constructs a compact-closed category from a traced symmetric monoidal category, which $\mathbf{Cob}(\mathbb{Z})$ is, at least intuitively. The issue is that the resulting double category would not be closed with respect to the parallel product, but the disjoint union. I believe this approach could be made to work for a sequential separation logic.

Another possibility is to switch from our cospan-based semantics to the span-based template games of Melliès (2019). The approach would be to create a game model of the logic of bunched implications which would be parameterized by a partial commutative monoid (PCM). In his setting, the template \mathbb{Z} is an internal category which describes the polarities of the games and the strategies. The idea would be to annotate these polarities with elements of the PCM which would represent the resources owned by each of the players.

Weak memory

Another natural extension is giving a semantics to a concurrent separation logic for weak memory. This is natural because axiomatic semantics of weak memory models are expressed as partial orders over the events of computations, and asynchronous graphs are true concurrency structures which induce partial orders over events.

A counter argument is that most successful concurrent separation logics, eg. (Kaiser et al., 2017), for weak memory use the promising semantics of Kang et al. (2017), whose distinctive feature is that it is, essentially, expressed as a normal state transition system. The complexity of weak memory is largely put inside the state, which is a set of timestamped messages instead of a simple partial map from locations to values. Part of this success can be explained by the fact that separation logic was designed to deal with complex state rather than complex reduction semantics (for instance, separation logics for probabilistic languages are an active area of research).

Fairness in Iris

Part II of this thesis develops a framework on top of Iris for relating the behavior of a program and of an abstract model of computation in an intensional manner. This approach is quite flexible: From a single soundness theorem which asserts that the initial state of the program and of the model are similar in some sense, using constructions on the model we are able to, on the one hand, prove that distributed programs implement classical algorithms, and on the other hand, that concurrent programs fairly terminate.

Finite but unbounded stuttering

The genesis of Chapter 6 is a previous failed attempt at proving fair termination of programs. The goal was to avoid the need for a fixed (in our case, given by the current state of the fairness model) bound to impose finite stuttering. Inspired by the fact that finite stutter can be expressed as a nested coinductive-inductive predicate, our plan was to define a weakest precondition as a *least fixpoint* which would look schematically like this:

$$\begin{aligned} \text{fwp } e \{Q\} &:= \text{“}e \text{ is a value”} \vee \\ &(\forall e \rightarrow e', \text{fwp } e' \{Q\}) \vee (\forall e \rightarrow e', \text{“take a step in } \mathcal{M}\text{”} \wedge \triangleright \text{fwp } e' \{Q\}) \end{aligned}$$

Note the lack of later modality in the disjunct corresponding to not taking a step in the model. Because this is a least fixpoint, the program e would either eventually reduce to a value, or the proof of $\text{fwp } e \{Q\}$ would eventually take a step in the model.

The issue is that we did not manage to define an Iris predicate similar to gsim which does not depend on the resources and which expresses that *each thread* always eventually takes a step in the model without quantifying over ordinals larger than the indexing ordinal, which is needed to get a simulation in the ambient logic out.

This method works for single-threaded programs, as it was successfully used in (Spies et al., 2021) in version of Iris with transfinite step-indexing, and I think it would improve the usability of the logic significantly if this obstacle was cleared.

Liveness of distributed systems

I think the objective of proving the liveness of distributed systems, under fairness assumptions about the network, is reachable. The idea would be that, for each message which must be received, there is a role in an abstract model, which becomes disabled when the message is received. Locally, this means that each node must send it infinitely often until it receives an acknowledgment.

Liveness is more of an issue in distributed systems compared to typical shared memory algorithms because the nodes and the network are unreliable, so this would be a good case study.

Compositionality of models

More generally, the main limitation of Part II is that the whole program must be proved using one single abstract model \mathcal{M} . In the future, we need to develop ways of combining such models as well as hiding them.

Bibliography

- America, Pierre and Jan Rutten (1989). “Solving reflexive domain equations in a category of complete metric spaces.” In: *Journal of Computer and System Sciences* (cit. on p. 126).
- Appel, Andrew W., Paul-André Melliès, Christopher D. Richards, and Jérôme Vouillon (2007). “A very modal model of a modern, major, general type system.” In: *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007, Nice, France, January 17-19, 2007*. Ed. by Martin Hofmann and Matthias Felleisen (cit. on p. 121).
- Balabonski, Thibaut, François Pottier, and Jonathan Protzenko (2014). “Type Soundness and Race Freedom for Mezzo.” In: *FLOPS* (cit. on p. 37).
- Bénabou, Jean (1967). “Introduction to bicategories.” In: *Reports of the Midwest Category Seminar*. Berlin, Heidelberg, pp. 1–77. ISBN: 978-3-540-35545-8 (cit. on pp. 68, 69).
- Birkedal, L., R. Møgelberg, J. Schwinghammer, and K. Støvring (2012). “First steps in synthetic guarded domain theory: step-indexing in the topos of trees.” In: *Logical Methods in Computer Science* (cit. on p. 123).
- Birkedal, Lars, Kristian Støvring, and Jacob Thamsborg (2010). “The category-theoretic solution of recursive metric-space equations.” In: *Theoretical Computer Science* (cit. on p. 126).
- Bodin, Martin, Arthur Chargueraud, Daniele Filaretti, Philippa Gardner, Sergio Maffei, Daiva Naudziuniene, Alan Schmitt, and Gareth Smith (2014). “A Trusted Mechanised JavaScript Specification.” In: *POPL ’14*. San Diego, California, USA (cit. on p. 8).
- Bornat, Richard, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson (2005). “Permission Accounting in Separation Logic.” In: *POPL* (cit. on pp. 24, 54).
- Bornat, Richard, Cristiano Calcagno, and Hongseok Yang (2006). “Variables as Resource in Separation Logic.” In: *ENTCS 155* (cit. on pp. 32, 54).
- Boyland, John (2003). “Checking Interference with Fractional Permissions.” In: *Static Analysis*. Ed. by Radhia Cousot. Springer Berlin Heidelberg (cit. on pp. 24, 54).
- Brookes, Stephen (2004). “A semantics for concurrent separation logic.” In: *CONCUR* (cit. on pp. 21, 29, 30, 32, 41, 54).
- Brookes, Stephen (2011). “A Revisionist History of Concurrent Separation Logic.” In: *Mathematical Foundations of Programming Semantics, MFPS*. Ed. by Michael W. Mislove and Joël Ouaknine (cit. on p. 29).

- Courser, Kenny Allen (2020). “Open Systems: A Double Categorical perspective.” PhD thesis. University of California Riverside (cit. on pp. 65, 68).
- D’Osualdo, Emanuele, Azadeh Farzan, Philippa Gardner, and Julian Sutherland (2019). “TaDA Live: Compositional Reasoning for Termination of Fine-grained Concurrent Programs.” In: *CoRR* (cit. on p. 159).
- da Rocha Pinto, Pedro, Thomas Dinsdale-Young, and Philippa Gardner (2014). “TaDA: A Logic for Time and Data Abstraction.” In: *Proceedings of the 28th European Conference on Object-Oriented Programming (ECOOP’14)*. Ed. by Richard E. Jones (cit. on p. 33).
- Day, Brian and Ross Street (1997). “Monoidal Bicategories and Hopf Algebroids.” In: *Advances in Mathematics* 129.1, pp. 99–157. ISSN: 0001-8708 (cit. on p. 81).
- Dinsdale-Young, Thomas, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang (2013). “Views: compositional reasoning for concurrent programs.” In: *POPL* (cit. on p. 33).
- Dodds, Mike, Suresh Jagannathan, Matthew J. Parkinson, Kasper Svendsen, and Lars Birkedal (2016). “Verifying Custom Synchronization Constructs Using Higher-Order Separation Logic.” In: *ACM Trans. Program. Lang. Syst.* 38.2 (cit. on p. 36).
- Ehresmann, Charles (1963). “Catégories structurées.” In: *Annales scientifiques de l’École Normale Supérieure* 80.4, pp. 349–426 (cit. on pp. 64, 66).
- Filliâtre, Jean-Christophe and Andrei Paskevich (2013). “Why3 – Where Programs Meet Provers.” In: *Proceedings of the 22nd European Symposium on Programming*. Lecture Notes in Computer Science (cit. on p. 12).
- Floyd, Robert (1967). “Mathematical Aspects of Computer Science.” In: *Proceedings of Symposium on Applied Mathematics*. American Mathematical Society. Chap. Assigning Meanings to Programs (cit. on p. 7).
- Frumin, D., R. Krebbers, and L. Birkedal (2021). “Compositional Non-Interference for Fine-Grained Concurrent Programs.” In: *Proceedings of Security and Privacy* (cit. on p. 140).
- Frumin, Dan, Robbert Krebbers, and Lars Birkedal (2018). “ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency.” In: *LICS ’18* (cit. on p. 140).
- Georges, Aïna Linn, Armaël Guéneau, Thomas Van Strydonck, Amin Timany, Alix Trieu, Sander Huyghebaert, Dominique Devriese, and Lars Birkedal (2021). “Efficient and Provable Local Capability Revocation Using Uninitialized Capabilities.” In: *Proc. ACM Program. Lang.* 5.POPL (cit. on p. 24).
- Giarrusso, Paolo G., Léo Stefanescu, Amin Timany, Lars Birkedal, and Robbert Krebbers (2020). “Scala Step-by-Step: Soundness for DOT with Step-Indexed Logical Relations in Iris.” In: *Proc. ACM Program. Lang.* 4.ICFP (cit. on p. 26).
- Gondelman, Léon, Simon Oddershede Gregersen, Abel Nieto, Amin Timany, and Lars Birkedal (2021). “Distributed Causal Memory: Modular Specification and Verification in Higher-Order Distributed Separation Logic.” In: *Proc. ACM Program. Lang.* POPL (cit. on p. 24).

- Gotsman, Alexey, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv (2007a). “Local Reasoning for Storable Locks and Threads.” In: *Programming Languages and Systems*. Ed. by Zhong Shao (cit. on p. 37).
- Gotsman, Alexey, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv (2007b). “Local Reasoning for Storable Locks and Threads.” In: *Programming Languages and Systems, 5th Asian Symposium, APLAS 2007, Singapore, November 29–December 1, 2007, Proceedings*. Ed. by Zhong Shao. Vol. 4807. Lecture Notes in Computer Science. Springer, pp. 19–37. DOI: 10.1007/978-3-540-76637-7_3. URL: https://doi.org/10.1007/978-3-540-76637-7_3 (cit. on p. 23).
- Guéneau, Armaël, Arthur Charguéraud, and François Pottier (2018). “A Fistful of Dollars: Formalizing Asymptotic Complexity Claims via Deductive Program Verification.” In: *ESOP*. Ed. by Amal Ahmed (cit. on p. 159).
- Hayman, Jonathan and Glynn Winskel (2008). “Independence and Concurrent Separation Logic.” In: *LMCS* (cit. on p. 37).
- Henzinger, Thomas A., Orna Kupferman, and Sriram K. Rajamani (1997). “Fair Simulation.” In: *CONCUR ’97*. Ed. by Antoni W. Mazurkiewicz and Józef Winkowski (cit. on p. 159).
- Hoare, C. A. R. (1969). “An Axiomatic Basis for Computer Programming.” In: *Commun. ACM* 12.10 (cit. on pp. 1, 2, 7, 9).
- Iris Team (2021). *The Iris 3.4 Documentation*. <https://plv.mpi-sws.org/iris/appendix-3.4.pdf> (cit. on pp. 124, 127).
- Ishtiaq, Samin S. and Peter W. O’Hearn (2001). “BI as an Assertion Language for Mutable Data Structures.” In: *Conference Record of POPL 2001*. Ed. by Chris Hankin and Dave Schmidt. ACM (cit. on p. 17).
- Jones, Cliff B. (1983). “Specification and Design of (Parallel) Programs.” In: *Information Processing 83, Proceedings of the IFIP 9th World Computer Congress, Paris, France, September 19–23, 1983*. Ed. by R. E. A. Mason. North-Holland/IFIP, pp. 321–332 (cit. on p. 15).
- Joyal, André, Ross Street, and Dominic Verity (1996). “Traced monoidal categories.” In: *Mathematical Proceedings of the Cambridge Philosophical Society*, pp. 447–468 (cit. on p. 161).
- Jung, Ralf, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer (2017). “RustBelt: Securing the Foundations of the Rust Programming Language.” In: *2.POPL* (cit. on pp. 8, 26, 36).
- Jung, Ralf, Robbert Krebbers, Lars Birkedal, and Derek Dreyer (2016). “Higher-Order Ghost State.” In: *51.9* (cit. on pp. 25, 36).
- Jung, Ralf, Robbert Krebbers, Jacques-Henri Jourdan, Aleš Bizjak, Lars Birkedal, and Derek Dreyer (2018). “Iris from the ground up: A modular foundation for higher-order concurrent separation logic.” In: *Journal of Functional Programming* 28. DOI: 10.1017/S0956796818000151 (cit. on p. 36).

- Jung, Ralf, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer (2015). “Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning.” In: 50.1 (cit. on pp. 25, 36, 118).
- Kaiser, Jan-Oliver, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis (2017). “Strong Logic for Weak Memory: Reasoning About Release-Acquire Consistency in Iris.” In: *31st European Conference on Object-Oriented Programming (ECOOP 2017)*. Ed. by Peter Müller (cit. on p. 162).
- Kang, Jeehoon, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer (2017). “A Promising Semantics for Relaxed-Memory Concurrency.” In: *SIGPLAN Not.* (Cit. on p. 162).
- Kesten, Yonit, Zohar Manna, and Amir Pnueli (1994). “Temporal verification of simulation and refinement.” In: *A Decade of Concurrency Reflections and Perspectives*. Ed. by J. W. de Bakker, W. -P. de Roever, and G. Rozenberg (cit. on p. 159).
- Krebbers, Robbert, Ralf Jung, Aleš Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal (2017). “The Essence of Higher-Order Concurrent Separation Logic.” In: *Programming Languages and Systems*. Ed. by Hongseok Yang. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 696–723 (cit. on pp. 25, 36, 118).
- Krebbers, Robbert, Amin Timany, and Lars Birkedal (2017). “Interactive Proofs in Higher-Order Concurrent Separation Logic.” In: *SIGPLAN Not.* (Cit. on pp. 26, 132, 140).
- Krogh-Jespersen, Morten, Amin Timany, Marit Edna Ohlenbusch, Simon Oddershede Gregersen, and Lars Birkedal (2020). “Aneris: A Mechanised Logic for Modular Reasoning about Distributed Systems.” In: *ESOP*. Ed. by Peter Müller (cit. on p. 24).
- Lack, Stephen and Paweł Sobociński (2004). “Adhesive Categories.” In: *Foundations of Software Science and Computation Structures*. Ed. by Igor Walukiewicz (cit. on p. 103).
- Lampert (1979). “How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs.” In: *IEEE Transactions on Computers C-28.9*, pp. 690–691 (cit. on p. 23).
- Lampert, Leslie (1994). “The Temporal Logic of Actions.” In: *ACM Trans. Program. Lang. Syst.* (Cit. on pp. 7, 26, 159).
- Lampert, Leslie (2001). “Paxos Made Simple.” In: *ACM SIGACT News (Distributed Computing Column)*, pp. 51–58 (cit. on p. 134).
- Leino, Rustan (2010). “Dafny: An Automatic Program Verifier for Functional Correctness.” In: *16th International Conference, LPAR-16, Dakar, Senegal* (cit. on p. 12).
- Leroy, Xavier (2009). “Formal verification of a realistic compiler.” In: *Communications of the ACM* 52.7, pp. 107–115 (cit. on p. 8).
- Ley-Wild, Ruy and Aleksandar Nanevski (2013). “Subjective Auxiliary State for Coarse-Grained Concurrency.” In: *SIGPLAN Not.* 48.1, pp. 561–574 (cit. on p. 33).
- Liang, Hongjin and Xinyu Feng (2016). “A Program Logic for Concurrent Objects under Fair Scheduling.” In: *SIGPLAN Not.* (Cit. on p. 159).

- McCusker, Guy and David Pym (2007). “A Games Model of Bunched Implications.” In: *Computer Science Logic*. Ed. by Jacques Duparc and Thomas A. Henzinger (cit. on p. 20).
- Melliès, Paul-André (2019). “Categorical Combinatorics of Scheduling and Synchronization in Game Semantics.” In: *Proc. ACM Program. Lang.* 3.POPL. ISSN: 2475-1421 (cit. on pp. 42, 54, 63, 81, 161).
- Melliès, Paul-André and Léo Stefanescu (2017). “A Game Semantics for Concurrent Separation Logic.” In: *MFPS* (cit. on p. 26).
- Melliès, Paul-André and Léo Stefanescu (2018). “An Asynchronous Soundness Theorem for Concurrent Separation Logic.” In: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018* (cit. on p. 26).
- Melliès, Paul-André and Léo Stefanescu (2020). “Concurrent Separation Logic Meets Template Games.” In: *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*. LICS '20 (cit. on p. 26).
- Melliès, Paul-André and Noam Zeilberger (2015). “Functors are Type Refinement Systems.” In: *POPL 2015*. ACM, pp. 3–16 (cit. on p. 39).
- Mellor-Crummey, John M. and Michael L. Scott (1991). “Algorithms for Scalable Synchronization on Shared-Memory Multiprocessors.” In: *ACM Trans. Comput. Syst.* 9.1 (cit. on p. 8).
- Mével, Glen, Jacques-Henri Jourdan, and François Pottier (2019). “Time credits and time receipts in Iris.” In: *ESOP*. Ed. by Luis Caires (cit. on p. 159).
- Nakano, Hiroshi (2000). “A Modality for Recursion.” In: *15th Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, June 26-29, 2000* (cit. on p. 121).
- O’Hearn, Peter (2007). “Resources, Concurrency, and Local Reasoning.” In: *TCS* 375 (cit. on p. 54).
- O’Hearn, Peter W. (2004). “Resources, Concurrency and Local Reasoning.” In: *CONCUR*. Ed. by Philippa Gardner and Nobuko Yoshida. Lecture Notes in Computer Science (cit. on p. 21).
- O’Hearn, Peter W. and David J. Pym (1999). “The logic of bunched implications.” In: *Bull. Symb. Log.* (Cit. on p. 19).
- O’Hearn, Peter W., John C. Reynolds, and Hongseok Yang (2001). “Local Reasoning about Programs that Alter Data Structures.” In: *Computer Science Logic, 15th International Workshop, CSL 2001. 10th Annual Conference of the EACSL, Paris, France, September 10-13, 2001, Proceedings*. Ed. by Laurent Fribourg. Lecture Notes in Computer Science (cit. on p. 17).
- Owicki, Susan and David Gries (1978). “An Axiomatic Proof Technique for Parallel Programs.” In: *Programming Methodology: A Collection of Articles by Members of IFIP WG2.3*. Ed. by David Gries. New York, NY: Springer New York, pp. 130–152.

- ISBN: 978-1-4612-6315-9. DOI: 10.1007/978-1-4612-6315-9_12. URL: https://doi.org/10.1007/978-1-4612-6315-9_12 (cit. on p. 15).
- Parkinson, Matthew, Richard Bornat, and Peter O’Hearn (2007). “Modular Verification of a Non-Blocking Stack.” In: *POPL ’07*. ISBN: 1595935754 (cit. on p. 23).
- Pottier, François (2013). “Syntactic soundness proof of a type-and-capability system with hidden state.” In: *J. Funct. Program.* (Cit. on p. 37).
- Reddy, Uday S. (1996). “Global State Considered Unnecessary: An Introduction to Object-Based Semantics.” In: *LISP Symb. Comput.* 9.1, pp. 7–76 (cit. on p. 40).
- Reynolds, John C. (1997). “The Essence of Algol.” In: *Algol-like Languages*. Ed. by Peter W. O’Hearn and Robert D. Tennent. Boston, MA: Birkhäuser Boston (cit. on p. 87).
- Reynolds, John C. (2000). “Intuitionistic Reasoning about Shared Mutable Data Structure.” In: *Millennial Perspectives in Computer Science*. Palgrave (cit. on p. 17).
- Reynolds, John C. (2002). “Separation Logic: A Logic for Shared Mutable Data Structures.” In: *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings* (cit. on p. 17).
- Sergey, Ilya, James R. Wilcox, and Zachary Tatlock (2017). “Programming and Proving with Distributed Protocols.” In: *2.POPL* (cit. on p. 24).
- Shulman, Michael (2008). “Framed bicategories and monoidal fibrations.” In: *Theory and Applications of Categories [electronic only]* 20, pp. 650–738 (cit. on p. 67).
- Spies, Simon, Lennard Gäher, Daniel Gratzer, Joseph Tassarotti, Robbert Krebbers, Derek Dreyer, and Lars Birkedal (2021). “Transfinite Iris: resolving an existential dilemma of step-indexed separation logic.” In: *PLDI*. Ed. by Stephen N. Freund and Eran Yahav (cit. on pp. 139, 163).
- Svendsen, Kasper and Lars Birkedal (2014). “Impredicative Concurrent Abstract Predicates.” In: *Programming Languages and Systems*. Ed. by Zhong Shao (cit. on p. 36).
- Tassarotti, Joseph and Robert Harper (2019). “A Separation Logic for Concurrent Randomized Programs.” In: *3.POPL* (cit. on p. 140).
- Tassarotti, Joseph, Ralf Jung, and Robert Harper (2017). “A Higher-Order Logic for Concurrent Termination-Preserving Refinement.” In: *ESOP*. Ed. by Hongseok Yang. Lecture Notes in Computer Science (cit. on pp. 139, 158).
- Timany, Amin, Simon Oddershede Gregersen, Léo Stefanescu, Léon Gondelman, Abel Nieto, and Lars Birkedal (2021). *Trillium: Unifying Refinement and Higher-Order Distributed Separation Logic*. arXiv: 2109.07863 [cs.PL] (cit. on pp. 27, 134).
- Timany, Amin, Léo Stefanescu, Morten Krogh-Jespersen, and Lars Birkedal (2017). “A Logical Relation for Monadic Encapsulation of State: Proving Contextual Equivalences in the Presence of RunST.” In: *Proc. ACM Program. Lang.* POPL (cit. on pp. 26, 140).
- Turing, Alan (1949). “Checking a Large Routine.” In: *Report of a Conference on High Speed Automatic Calculating Machines* (cit. on p. 7).

- Turon, Aaron, Viktor Vafeiadis, and Derek Dreyer (2014). “GPS: Navigating Weak Memory with Ghosts, Protocols, and Separation.” In: (cit. on p. 24).
- Vafeiadis, Viktor (2011). “Concurrent Separation Logic and Operational Semantics.” In: *ENTCS 276* (cit. on pp. 32–34, 41, 54).
- Varacca, Daniele and Glynn Winskel (2006). “Distributing probability over non-determinism.” In: *Math. Struct. Comput. Sci.* (Cit. on p. 140).
- Xia, Li-yao, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic (2020). “Interaction trees: representing recursive and impure programs in Coq.” In: *Proc. ACM Program. Lang.* 4.POPL, 51:1–51:32 (cit. on p. 33).
- Zakowski, Yannick, Paul He, Chung-Kil Hur, and Steve Zdancewic (2020). “An equational theory for weak bisimulation via generalized parameterized coinduction.” In: *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP*, pp. 71–84 (cit. on p. 33).